



User Guide

Amazon Detective



Amazon Detective: User Guide

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

Table of Contents

What is Detective?	1
Features of Amazon Detective	1
Accessing Amazon Detective	3
Pricing for Amazon Detective	4
How does Detective work?	5
Who uses Detective?	6
Related services	6
Concepts and terminology	8
Getting started	13
Setting up	13
Sign up for an AWS account	13
Create a user with administrative access	14
Prerequisites	15
Granting the required Detective permissions	15
Supported AWS Command Line Interface version	16
Recommendations	16
Recommended alignment with GuardDuty and AWS Security Hub	16
Recommended update to the GuardDuty CloudWatch notification frequency	17
Enabling Detective	17
Checking that data is being extracted	19
Data in a behavior graph	21
How Detective populates a behavior graph	21
How Detective processes source data	22
Detective extraction	22
Detective analytics	23
Training period for new behavior graphs	23
Overview of the behavior graph data structure	24
Types of elements in the behavior graph data structure	24
Types of entities in the behavior graph data structure	24
Source data used in a behavior graph	30
Types of core data sources in Detective	31
Types of optional data sources in Detective	31
Amazon EKS audit logs	32
AWS security findings	33

How Detective ingests and stores source data	34
How Detective enforces the data volume quota for behavior graphs	35
Summary dashboard	36
Investigations	36
Newly observed geolocations	37
Active finding groups in the last 7 days	37
Roles and users with the most API call volume	38
EC2 instances with the most traffic volume	38
Container clusters with the most Kubernetes pods	39
Approximate value notification	39
How Detective is used for investigation	41
Investigation phases	41
Starting points for a Detective Investigation	42
Findings detected by GuardDuty	42
AWS security findings aggregated by Security Hub	42
Entities extracted from Detective source data	42
Detective Investigation flow	43
Detective Investigation	44
Running a Detective Investigation	45
Reviewing investigations reports	47
Understanding a Detective Investigations report	48
Investigations report summary	50
Downloading an investigation report	50
Archiving an investigation report	51
Analyzing findings	53
Finding overview	53
Scope time used for the finding overview	54
Finding details	54
Related entities	54
Troubleshooting 'Page not found'	54
Finding groups	55
Understanding the finding groups page	56
Informational findings in finding groups	58
Finding group profiles	59
Finding group visualization	60
Finding group summary	63

Reviewing finding group summary	63
Disabling finding group summary	65
Enabling finding group summary	65
Supported Regions	65
Archiving a GuardDuty finding	66
Analyzing entities	67
Using entity profiles	67
Scope time for an entity profile	68
Entity identifier and type	68
Involved findings	68
Finding groups involving this entity	68
Profile panels containing entity details and analytics results	68
Navigating in an entity profile	69
Profile panels	69
Types of information on a profile panel	70
Types of profile panel visualizations	73
Other notes on profile panel content	77
Preferences for profile panels	78
Navigating to an entity profile	79
Pivoting from another console	79
Navigating using a URL	82
Adding Detective URLs for findings to Splunk	85
Pivoting to another console	86
Pivoting to another entity profile	86
Exploring activity details	86
Overall API call volume	87
Geolocations	94
Overall VPC flow volume	98
Overall Kubernetes API call volume	102
Managing the scope time	107
Setting specific start and end dates and times	107
Edit the length of time for the scope time	108
Setting the scope time to a finding time window	108
Setting the scope time on the summary page	109
Viewing findings for an entity	109
High-volume entities	110

What is a high-volume entity?	110
Viewing the high-volume entity notification on a profile	111
Viewing the list of high-volume entities for the current scope time	111
Searching for a finding or entity	113
Completing the search	113
Using the search results	115
Troubleshooting the search	115
Managing accounts	117
Restrictions and recommendations	118
Maximum number of member accounts	118
Accounts and Regions	118
Alignment of administrator accounts with Security Hub and GuardDuty	118
Granting the required permissions for administrator accounts	118
Reflecting organization updates in Detective	119
Using Organizations to manage behavior graph accounts	119
Designate a Detective administrator account for your organization	120
Enable organization accounts as member accounts	120
Designating the Detective administrator account	121
Designating a Detective administrator	123
Removing the Detective administrator account	126
Available actions for accounts	128
Viewing the list of accounts	130
Listing accounts (Console)	131
Listing your member accounts (Detective API, AWS CLI)	133
Managing organization member accounts	134
Enabling new organization accounts	134
Enabling organization accounts as member accounts	136
Disassociating organization accounts	138
Managing invited member accounts	139
Inviting individual accounts to a behavior graph	141
Inviting a list of member accounts to a behavior graph	143
Enabling a member account that is Not enabled	144
Removing member accounts	146
For member accounts: Managing invitations and memberships	147
IAM policy for a member account	148
Viewing behavior graph invitations	149

Responding to a behavior graph invitation	150
Removing your account from a behavior graph	152
Effect of account actions	153
Detective disabled	153
Member account removed from the behavior graph	153
Member account leaves the organization	153
AWS account suspended	154
AWS account closed	154
Amazon Detective Python scripts	155
Overview of the enableDetective.py script	155
Overview of the disableDetective.py script	156
Required permissions for the scripts	156
Setting up the run environment for the Python scripts	157
Creating a .csv list of member accounts to add or remove	159
Running enableDetective.py	160
Running disableDetective.py	161
Integration with Amazon Security Lake	163
Enabling the integration	163
Before you begin	164
Step 1: Creating a Security Lake subscriber	165
Step 2: Adding the required IAM permissions to your account	166
Step 3: Accept the Resource Share ARN invitation and enable the integration	169
Changing the integration configuration	175
Supported AWS Regions	176
Querying raw logs in Detective	178
Query raw logs for an AWS role	181
Query raw logs for an Amazon EKS cluster	182
Query raw logs for an Amazon EC2 instance	182
Disabling the integration	183
Deleting a CloudFormation stack	183
Forecasting and monitoring costs	185
About the free trial for behavior graphs	186
Free trial for optional data sources	187
Administrator account usage and cost	187
Volume of data ingested for each account	187
Projected costs for the behavior graph	188

Projected cost for the behavior graph	188
Volume of data ingested by source packages	189
Member account usage tracking	189
Ingested volume for each behavior graph	189
Projected cost across behavior graphs	190
How Detective calculates projected cost	190
Security	192
Data protection	193
Key management	194
Identity and access management	194
Audience	194
Authenticating With Identities	195
Managing Access Using Policies	198
How Amazon Detective works with IAM	200
Identity-based policy examples	207
AWS managed policies	212
Using service-linked roles	223
Troubleshooting identity and access	225
Compliance validation	227
Resilience	227
Infrastructure security	228
Security best practices	228
Best practices for Detective administrator accounts	228
Best practices for member accounts	229
Logging API calls	230
Detective information in CloudTrail	230
Understanding Detective log file entries	231
Regions and quotas	233
Detective Regions and endpoints	233
Detective quotas	233
Internet Explorer 11 not supported	234
Managing tags	235
Viewing the tags for a behavior graph	235
Adding tags to a behavior graph	236
Removing tags from a behavior graph	237
Disabling Amazon Detective	238

Disabling Detective (Console)	238
Disabling Detective (Detective API, AWS CLI)	238
Disabling Detective across Regions (Python script on GitHub)	239
Document history	240

What is Amazon Detective?

Amazon Detective helps you analyze, investigate, and quickly identify the root cause of security findings or suspicious activities. Detective automatically collects log data from your AWS resources. It then uses machine learning, statistical analysis, and graph theory to generate visualizations that help you to conduct faster and more efficient security investigations. The Detective prebuilt data aggregations, summaries, and context help you to quickly analyze and determine the nature and extent of possible security issues.

With Detective, you can access up to a year of historical event data. This data is available through a set of visualizations that show changes in the type and volume of activity over a selected time window. Detective links these changes to GuardDuty findings. For more information on source data in Detective, see [the section called "Source data used in a behavior graph"](#).

By automatically aggregating data and providing visual tools, Amazon Detective lets you to conduct faster, more efficient security investigations. You can quickly analyze potential issues and determine the scope of security threats.

Topics

- [Features of Amazon Detective](#)
- [Accessing Amazon Detective](#)
- [Pricing for Amazon Detective](#)
- [How does Detective work?](#)
- [Who uses Detective?](#)
- [Related services](#)

Features of Amazon Detective

Here are some of the key ways that Amazon Detective is helpful for investigating suspicious activity in your AWS environment and analyze resources to identify the root cause of security issues.

Detective finding groups

[Detective finding groups](#) lets you examine multiple activities as they relate to a potential security event. You can analyze the root cause for high severity GuardDuty findings using

finding groups. If a threat actor is attempting to compromise your AWS environment, they typically perform a sequence of actions that generate multiple security findings and unusual behaviors.

The finding groups page in Detective displays all the related finding groups extracted from your behavior graph. For more information about how you can leverage finding groups to analyze the root cause of security findings, see [Analyzing finding groups in Detective](#).

Detective provides an interactive visualization of each finding group to help you investigate security issues faster and more thoroughly. The visualization is designed to display entities and findings involved in a security incident, making it easier to understand connections and root causes. help you investigate issues faster and more thoroughly with less effort. The [Finding group Visualization](#) panel displays the findings and entities involved in a finding group.

Detective Investigation to triage findings

With [Detective Investigation](#) you can investigate IAM users and IAM roles using indicators of compromise, which can help you determine if a resource is involved in a security incident. An indicator of compromise (IOC) is an artifact observed in or on a network, system, or environment that can (with a high level of confidence) identify malicious activity or a security incident. With Detective investigations, you can maximize efficiency, focus on the security threats, and strengthen incidence response capabilities.

Detective Investigation uses machine learning models and threat intelligence to surface only the most critical, suspicious issues, allowing you to focus on high-level investigations. It automatically analyzes resources in your AWS environment to identify potential indicators of compromise or suspicious activity. This lets you identify patterns and comprehend which resources are impacted by security events, offering a proactive approach to threat identification and mitigation.

You can use start a Detective Investigation from the Detective console by [Running a Detective Investigation](#). To run an investigation programmatically, use the [StartInvestigation](#) operation of the Detective API. If you're using the AWS Command Line Interface (AWS CLI) run the [start-investigation](#) command.

Detective integration with Amazon Security Lake

[Detective integrates with Amazon Security Lake](#), which means that you can query and retrieve the raw log data stored by Security Lake. With this integration, you can collect logs and events from the following sources which Security Lake natively supports.

- AWS CloudTrail management events version 1.0 and after
- Amazon Virtual Private Cloud (Amazon VPC) Flow Logs version 1.0 and after
- Amazon Elastic Kubernetes Service (Amazon EKS) Audit Log version 2.0

After you integrate Detective with Security Lake, Detective begins pulling raw logs from Security Lake related to AWS CloudTrail management events and Amazon VPC Flow Logs. You can [query raw logs](#) to view the logs and events in Detective.

Investigate VPC flow volume

With Detective you can interactively examine the [activity details of the virtual private cloud \(VPC\) network flows](#) of your Amazon Elastic Compute Cloud (Amazon EC2) instances and Kubernetes pods. Detective automatically collects VPC flow logs from your monitored accounts, aggregates them by EC2 instance, and presents visual summaries and analytics about these network flows.

For an EC2 instance, the activity details for Overall VPC flow volume show the interactions between the EC2 instance and IP addresses during a selected time range.

For a Kubernetes pod, Overall VPC flow volume displays the overall volume of bytes into and out of the Kubernetes pod's assigned IP address for all destination IP addresses.

Accessing Amazon Detective

Amazon Detective is available in most AWS Regions. For a list of Regions where Detective is currently available, see [Amazon Detective endpoints and quotas](#) in the *AWS General Reference*. For information about managing AWS Regions for your AWS account, see [Specifying which AWS Regions your account can use](#) in the *AWS Account Management Reference Guide*.

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

AWS Management Console

The AWS Management Console is a browser-based interface that you can use to create and manage AWS resources. As part of that console, the Amazon Detective console provides access to your Detective account, data, and resources. You can perform any Detective task by using the Detective console—review potential security threats and analyze, investigate, and identify the root cause of security findings.

AWS command line tools

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

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

AWS SDKs

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

Amazon Detective REST API

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

Pricing for Amazon Detective

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

Detective pricing is based on several dimensions—and charges a tiered flat rate per GB for all data regardless of the source. For more information, see [Amazon Detective pricing](#).

To help you understand and forecast the cost of using Detective, Detective provides estimated usage costs for your account. You can [review these estimates](#) on the Amazon Detective console and access them with the Amazon Detective API. Depending on how you use the service, you might

incur additional costs for using other AWS services in combination with certain Detective features, such as Security Lake integration and Detective Investigations.

When you enable Detective for the first time, your AWS account is automatically enrolled in the 30-day free trial of Detective. This includes individual accounts that are enabled as part of an organization in AWS Organizations. During the free trial, there's no charge for using Detective in the applicable AWS Region.

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

How does Detective work?

Detective automatically extracts time-based events such as login attempts, API calls, and network traffic from AWS CloudTrail and Amazon VPC flow logs. It also ingests findings detected by GuardDuty.

From those events, Detective uses machine learning and visualization to create a unified, interactive view of your resource behaviors and the interactions between them over time. You can explore this behavior graph to examine disparate actions such as failed logon attempts or suspicious API calls. You can also see how these actions affect resources such as AWS accounts and Amazon EC2 instances. You can adjust the behavior graph's scope and timeline for a variety of tasks:

- Rapidly investigate any activity that falls outside the norm.
- Identify patterns that may indicate a security issue.
- Understand all of the resources affected by a finding.

Detective tailored visualizations provide a baseline for and summarize the account information. These findings can help answer questions such as "Is this an unusual API call for this role?" Or "Is this spike in traffic from this instance expected?"

With Detective, you don't have to organize any data or develop, configure, or tune your own queries and algorithms. There are no upfront costs and you pay only for the events analyzed, with no additional software to deploy or other feeds to subscribe to.

Who uses Detective?

When an account enables Detective, it becomes the administrator account for a behavior graph. A behavior graph is a linked set of extracted and analyzed data from one or more AWS accounts. Administrator accounts invite member accounts to contribute their data to the administrator account's behavior graph.

Detective is also integrated with AWS Organizations. Your organization management account designates a Detective administrator account for the organization. The Detective administrator account enables organization accounts as member accounts in the organization behavior graph.

For information about how Detective uses source data from behavior graph accounts, see [the section called "Source data used in a behavior graph"](#).

For information on how administrator accounts manage behavior graphs, see [Managing accounts](#). For information on how member accounts manage their behavior graph invitations and memberships, see [the section called "For member accounts: Managing invitations and memberships"](#).

The administrator account uses the analytics and visualizations generated from the behavior graph to investigate AWS resources and GuardDuty findings. Using the Detective integrations with GuardDuty and AWS Security Hub, you can pivot from a GuardDuty finding in these services directly into the Detective console.

A Detective investigation focuses on the activity that is connected to the involved AWS resources. For an overview of the investigation process in Detective, see [How Amazon Detective is used for investigation](#) in *Detective User Guide*.

Related services

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

AWS Security Hub

AWS Security Hub gives you a comprehensive view of the security state of your AWS resources and helps you check your AWS environment against security industry standards and best practices. It does this partly by consuming, aggregating, organizing, and prioritizing your security findings from multiple AWS services (including Detective) and supported AWS Partner

Network (APN) products. Security Hub helps you analyze your security trends and identify the highest priority security issues across your AWS environment.

To learn more about Security Hub, see the [AWS Security Hub User Guide](#).

Amazon GuardDuty

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

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

Amazon Security Lake

Amazon Security Lake is a fully managed security data lake service. You can use Security Lake to automatically centralize security data from AWS environments, SaaS providers, on-premises sources, cloud sources, and third-party sources into a purpose-built data lake that's stored in your AWS account. Security Lake helps you analyze security data, so you can get a more complete understanding of your security posture across your entire organization. With Security Lake, you can also improve the protection of your workloads, applications, and data.

To learn more about Security Lake, see the [Amazon Security Lake User Guide](#). To learn more about using Detective and Security Lake together, see [Integration with Amazon Security Lake](#).

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

Amazon Detective concepts and terminology

The following terms and concepts are important for understanding Amazon Detective and how it works.

Administrator account

The AWS account that owns a behavior graph and that uses the behavior graph for investigation.

The administrator account invites member accounts to contribute their data to the behavior graph. For more information, see [the section called "Managing invited member accounts"](#).

For the organization behavior graph, the administrator account is the Detective administrator account that the organization management account designates. For more information, see [the section called "Designating the Detective administrator account"](#). The Detective administrator account can enable any organization account as a member account in the organization behavior graph. For more information, see [the section called "Managing organization member accounts"](#).

Administrator accounts can also view data usage for the behavior graph, and remove member accounts from the behavior graph.

Autonomous System Organization (ASO)

The titled organization which is assigned an autonomous system. This autonomous system is a heterogenous network or a set of networks using similar routing logic and policies.

Behavior graph

A linked set of data generated from incoming source data that is associated with one or more AWS accounts.

Each behavior graph uses the same structure of findings, entities, and relationships.

Delegated administrator account (AWS Organizations)

In Organizations, the delegated administrator account for a service is able to manage the use of a service for the organization.

In Detective, the Detective administrator account is also the delegated administrator account, unless the Detective administrator account is the organization management account. The organization management account cannot be a delegated administrator account.

In Detective, self-delegation is allowed. An organization management account can delegate their own account to be the delegated administrator of Detective but this would be registered or remembered only in the scope of Detective and not organizations.

Detective administrator account

The account designated by the organization management account to be the administrator account for the organization behavior graph in a Region. For more information, see [the section called “Designating the Detective administrator account”](#).

Detective recommends that the organization management account chooses an account other than their account.

If the account is not the organization management account, then the Detective administrator account is also the delegated administrator account for Detective in Organizations.

Detective source data

Processed, structured versions of information from the following types of feeds:

- Logs from AWS services, such as AWS CloudTrail logs and Amazon VPC Flow Logs
- GuardDuty findings

Detective uses the Detective source data to populate the behavior graph. Detective also stores copies of the Detective source data to support its analytics.

Entity

An item extracted from the ingested data.

Each entity has a type, which identifies the type of object it represents. Examples of entity types include IP addresses, Amazon EC2 instances, and AWS users.

Entities can be AWS resources that you manage, or external IP addresses that have interacted with your resources.

For each entity, the source data is also used to populate entity properties. Property values can be extracted directly from source records or aggregated across multiple records.

Finding

A security issue detected by Amazon GuardDuty.

Finding group

A collection of related findings, entities, and evidence that may be related to the same event or security issue. Detective generates finding groups based on a built-in machine learning model.

Detective evidence

Detective identifies additional evidence related to a finding group based on data in your behavior graph collected within the last 45 days. This evidence is presented as a finding with the severity value of **Informational**. Evidence provides supporting information that highlights an unusual activity or unknown behavior that is potentially suspicious when viewed within a finding group. An example of this might be newly observed geolocations or API calls observed within the scope time of a finding. At this time, these findings are only viewable in Detective and not sent to Security Hub.

Finding overview

A single page that provides a summary of information about a finding.

A finding overview contains the list of involved entities for the findings. From the list, you can pivot to the profile for an entity.

A finding overview also contains a details panel that contains the finding attributes.

High-volume entity

An entity that has connections to or from a large number of other entities during a time interval. For example, an EC2 instance might have connections from millions of IP addresses. The number of connections exceeds the threshold that Detective can accommodate.

When the current scope time contains a high-volume time interval, Detective notifies the user.

For more information, see [Viewing details for high-volume entities](#) in the *Amazon Detective User Guide*.

Investigation

The process of triaging suspicious or interesting activity, determining its scope, getting to its underlying source or cause, and then determining how to proceed.

Member account

An AWS account that an administrator account invited to contribute data to a behavior graph. In the organization behavior graph, a member account can be an organization account that the Detective administrator account enabled as a member account.

Member accounts that are invited can respond to the behavior graph invitation and remove their account from the behavior graph. For more information, see [the section called “For member accounts: Managing invitations and memberships”](#).

Organization accounts cannot change their membership in the organization behavior graph.

All member accounts can also view usage information for their account across the behavior graphs that they contribute data to.

They have no other access to the behavior graph.

Organization behavior graph

The behavior graph that is owned by the Detective administrator account. The organization management account designates the Detective administrator account. For more information, see [the section called “Designating the Detective administrator account”](#).

In the organization behavior graph, the Detective administrator account controls whether an organization account is a member account. An organization account cannot remove itself from the organization behavior graph.

The Detective administrator account can also invite other accounts to the organization behavior graph.

Profile

A single page that provides a collection of data visualizations related to activity for an entity.

For findings, profiles help analysts to determine whether the finding is of genuine concern or a false positive.

Profiles provide information to support an investigation into a finding or for a general hunt for suspicious activity.

Profile panel

A single visualization on a profile. Each profile panel is intended to help answer a specific question or questions to assist an analyst in an investigation.

Profile panels can contain key-value pairs, tables, timelines, bar charts, or geolocation charts.

Relationship

Activity that occurs between individual entities. Relationships are also extracted from the incoming source data.

Similar to an entity, a relationship has a type, which identifies the types of entities involved and the direction of the connection. An example of a relationship type is an IP address connecting to an Amazon EC2 instance.

Scope time

The time window that is used to scope the data displayed on profiles.

The default scope time for a finding reflects the first and last times when the suspicious activity was observed.

The default scope time for an entity profile is the previous 24 hours.

Getting started with Amazon Detective

This tutorial provides an introduction to Amazon Detective. You'll learn how to enable Detective for your AWS account. You'll also learn how to verify that Detective has begun to ingest and extract data from your AWS account into your behavior graph.

When you enable Amazon Detective, Detective creates a Region-specific behavior graph that has your account as its administrator account. This is initially the only account in the behavior graph. The administrator account can then invite other AWS accounts to contribute their data to the behavior graph. See [Managing accounts](#).

Enabling Detective in a Region for the first time also begins a 30-day free trial for the behavior graph. If the account disables Detective and then enables it again, no free trial is available. See [the section called "About the free trial for behavior graphs"](#).

After the free trial, each account in the behavior graph is billed for the data they contribute to it. The administrator account can track the usage and see the total projected cost for a typical 30-day period for their entire behavior graph. For more information, see [the section called "Administrator account usage and cost"](#). Member accounts can track the usage and projected cost for the behavior graphs that they belong to. For more information, see [the section called "Member account usage tracking"](#).

Topics

- [Setting up your AWS account](#)
- [Prerequisites](#)
- [Recommendations](#)
- [Enabling Amazon Detective](#)

Setting up your AWS account

Before you can enable Amazon Detective, you must have an AWS account. If you do not have an AWS account, complete the following steps to create one.

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform [tasks that require root user access](#).

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <https://aws.amazon.com/> and choosing **My Account**.

Create a user with administrative access

After you sign up for an AWS account, secure your AWS account root user, enable AWS IAM Identity Center, and create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the [AWS Management Console](#) as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see [Signing in as the root user](#) in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see [Enable a virtual MFA device for your AWS account root user \(console\)](#) in the *IAM User Guide*.

Create a user with administrative access

1. Enable IAM Identity Center.

For instructions, see [Enabling AWS IAM Identity Center](#) in the *AWS IAM Identity Center User Guide*.

2. In IAM Identity Center, grant administrative access to a user.

For a tutorial about using the IAM Identity Center directory as your identity source, see [Configure user access with the default IAM Identity Center directory](#) in the *AWS IAM Identity Center User Guide*.

Sign in as the user with administrative access

- To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see [Signing in to the AWS access portal](#) in the *AWS Sign-In User Guide*.

Assign access to additional users

1. In IAM Identity Center, create a permission set that follows the best practice of applying least-privilege permissions.

For instructions, see [Create a permission set](#) in the *AWS IAM Identity Center User Guide*.

2. Assign users to a group, and then assign single sign-on access to the group.

For instructions, see [Add groups](#) in the *AWS IAM Identity Center User Guide*.

Prerequisites

Make sure that the following requirements are met before enabling Detective.

Granting the required Detective permissions

Before you can enable Detective, you must make sure that your IAM principal has the required Detective permissions. The principal can be an existing user or role that you are already using, or you can create a new user or role to use for Detective.

When you sign up for Amazon Web Services (AWS), your account is automatically signed up for all AWS services, including Amazon Detective. However, to enable and use Detective, you first have to set up permissions that allow you to access the Amazon Detective console and API operations. You or your administrator can do this by using AWS Identity and Access Management (IAM) to attach the [AmazonDetectiveFullAccess managed policy](#) to your IAM principal, which grants access to all Detective actions.

Supported AWS Command Line Interface version

To use the AWS CLI to perform Detective tasks, the minimum required version is 1.16.303.

Recommendations

Consider following these recommendations before enabling Detective

Recommended alignment with GuardDuty and AWS Security Hub

If you are enrolled in GuardDuty and AWS Security Hub, we recommend that your account be an administrator account for those services. If the administrator accounts are the same for all three services, then the following integration points work seamlessly.

- In GuardDuty or Security Hub, when viewing details for a GuardDuty finding, you can pivot from the finding details to the Detective finding profile.
- In Detective, when investigating a GuardDuty finding, you can choose the option to archive that finding.

If you have different administrator accounts for GuardDuty and Security Hub, we recommend that you align the administrator accounts based on the service you use more frequently.

- If you use GuardDuty more frequently, then enable Detective using the GuardDuty administrator account.

If you use AWS Organizations to manage accounts, designate the GuardDuty administrator account as the Detective administrator account for the organization.

- If you use Security Hub more frequently, then enable Detective using the Security Hub administrator account.

If you use Organizations to manage accounts, designate the Security Hub administrator account as the Detective administrator account for the organization.

If you cannot use the same administrator accounts across all of the services, then after you enable Detective, you can optionally create a cross-account role. This role grants an administrator account access to other accounts.

For information about how IAM supports this type of role, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.

Recommended update to the GuardDuty CloudWatch notification frequency

In GuardDuty, detectors are configured with an Amazon CloudWatch notification frequency for reporting subsequent occurrences of a finding. This includes sending notifications to Detective.

By default, the frequency is six hours. This means that even if a finding recurs many times, the new occurrences are not reflected in Detective until up to six hours later.

To reduce the amount of time it takes for Detective to receive these updates, we recommend that the GuardDuty administrator account changes the setting on their detectors to 15 minutes. Note that changing the configuration has no effect on the cost of using GuardDuty.

For information about setting the notification frequency, see [Monitoring GuardDuty Findings with Amazon CloudWatch Events](#) in the *Amazon GuardDuty User Guide*.

Enabling Amazon Detective

You can enable Detective from the Detective console, the Detective API, or the AWS Command Line Interface.

You can only enable Detective once in each Region. If you already are the administrator account for a behavior graph in the Region, then you cannot enable Detective again in that Region.

Console

To enable Detective (console)

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. Choose **Get started**.
3. On the **Enable Amazon Detective** page, **Align administrator accounts (recommended)** explains the recommendation to align the administrator accounts between Detective and Amazon GuardDuty and AWS Security Hub. See [the section called "Recommended alignment with GuardDuty and AWS Security Hub"](#).
4. The **Attach IAM policy** button takes you directly to the IAM console and opens up the recommended policy, You have the option to attach the recommended policy to the principal you use for Detective. If you do not have permissions to operate in the IAM console, within the **Required permissions** you can copy the policy Amazon Resource Name (ARN) to provide it to your IAM administrator. They can attach the policy on your behalf.

Confirm that the required IAM policy is in place.

5. The **Add tags** section allows you to add tags to the behavior graph.

To add a tag, do the following:

- a. Choose **Add new tag**.
- b. For **Key**, enter the name of the tag.
- c. For **Value**, enter the value of the tag.

To remove a tag, choose the **Remove** option for that tag.

6. Choose **Enable Amazon Detective**.
7. After you enable Detective, you can invite member accounts to your behavior graph.

To navigate to the **Account management** page, choose **Add members now**. For information about inviting member accounts, see [the section called "Managing invited member accounts"](#).

Detective API, AWS CLI

You can enable Amazon Detective from the Detective API or the AWS Command Line Interface.

To enable Detective (Detective API, AWS CLI)

- **Detective API:** Use the [CreateGraph](#) operation.
- **AWS CLI:** At the command line, run the [create-graph](#) command.

```
aws detective create-graph --tags '{"tagName": "tagValue"}
```

The following command enables Detective and sets the value of the Department tag to Security.

```
aws detective create-graph --tags '{"Department": "Security"}
```

Python script on GitHub

You can enable Detective across Regions using the Detective Python script on GitHub. Detective provides an open-source script in GitHub that does the following:

- Enables Detective for an administrator account in a specified list of Regions
- Adds a provided list of member accounts to each of the resulting behavior graphs
- Sends invitation emails to the member accounts
- Automatically accepts the invitations for the member accounts

For information about how to configure and use the GitHub scripts, see [the section called “Amazon Detective Python scripts”](#).

Checking that data is being extracted

After you enable Detective, it begins to ingest and extract data from your AWS account into your behavior graph.

For the initial extraction, data usually becomes available in the behavior graph within 2 hours.

One way to check that Detective is extracting data is to look for example values on the Detective **Search** page.

To check for example values on the Search page

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Search**.
3. From the **Select type** menu, choose a type of item.

Examples from your data contains a sample set of identifiers of the selected type that are in your behavior graph data.

If you can see example values, then you know that data is being ingested and extracted into your behavior graph.

Data in a Detective behavior graph

In Amazon Detective, you conduct investigations using data from a Detective behavior graph. In this section you can learn about the core data sources used in a Detective behavior graph and how Detective uses the source data to populate it.

A behavior graph is a linked set of data generated from the Detective source data that is ingested from one or more Amazon Web Services (AWS) accounts.

The behavior graph uses the source data to do the following.

- Generate an overall picture of your systems, users, and the interactions among them over time
- Perform more detailed analysis of specific activity to help you answer questions that arise as you conduct investigations
- Correlate collections of findings, entities, and evidence that may be related to the same event or security issue.

Note that all extraction, modeling, and analytics of behavior graph data occurs within the context of each individual behavior graph.

Each behavior graph contains data from one or more accounts. When an account enables Detective, it becomes the administrator account for the behavior graph, and it chooses the member accounts for the behavior graph. A behavior graph can have up to 1,200 member accounts. For information about how an administrator account manages the member accounts in a behavior graph, see [Managing accounts in Detective](#).

Contents

- [How Detective populates a behavior graph](#)
- [Training period for new Detective behavior graphs](#)
- [Overview of the behavior graph data structure](#)
- [Source data used in a Detective behavior graph](#)

How Detective populates a behavior graph


To provide the raw data for investigations, Detective brings together data from across your AWS environment and beyond, including the following:

- Log data, including Amazon Virtual Private Cloud (Amazon VPC) and AWS CloudTrail
- Findings from Amazon GuardDuty
- Findings from AWS Security Hub

To learn more about the source data used in a behavior graph, see [Source data used in a behavior graph](#).

How Detective processes source data

As new data comes in, Detective uses a combination of extraction and analytics to populate the behavior graph.



Detective extraction

Extraction is based on configured mapping rules. A mapping rule basically says, "Whenever you see this piece of data, use it in this specific way to update behavior graph data."

For example, an incoming Detective source data record might include an IP address. If it does, Detective uses the information in that record to create a new IP address entity or update an existing IP address entity.

Detective analytics

Analytics are more complex algorithms that analyze the data to provide insight into activity that is associated with entities.

For example, one type of Detective analytic analyzes how often activity occurs by running algorithms. For entities that make API calls, the algorithm looks for API calls that the entity doesn't normally use. The algorithm also looks for a large spike in the number of API calls.

Analytic insights support investigations by providing answers to key analyst questions and are frequently used to populate finding and entity profile panels.

Training period for new Detective behavior graphs

One avenue of investigation for a finding is to compare the activity during the finding scope time to activity that occurred before the finding was detected. Activity that has not been seen before might be more likely to be suspicious.

Some Amazon Detective profile panels highlight activity that was not observed during the time period before the finding. Several profile panels also display a baseline value to show the average activity during the 45 days before the scope time. Scope time is the summary of activity of an entity over time.

As more data is extracted into your behavior graph, Detective develops a more accurate picture of what activity is normal in your organization and what activity is unusual.

However, to create this picture, Detective needs access to at least two weeks of data. The maturity of the Detective analysis also increases with the number of accounts in the behavior graph.

The first two weeks after you activate Detective are considered a training period. During this period, profile panels that compare scope time activity to earlier activity display a message that Detective is in a training period.

During the trial period, Detective recommends that you add as many member accounts as you can to the behavior graph. This provides Detective with a larger pool of data, which allows it to generate a more accurate picture of the normal activity for your organization.

Overview of the behavior graph data structure

The behavior graph data structure defines the structure of the extracted and analyzed data. It also defines how the source data is mapped to the behavior graph.

Types of elements in the behavior graph data structure

The behavior graph data structure is made up of the following information elements.

Entity

An entity represents an item extracted from the Detective source data.

Each entity has a type, which identifies the type of object it represents. Examples of entity types include IP addresses, Amazon EC2 instances, and AWS users.

For each entity, the source data is also used to populate entity properties. Property values might be extracted directly from source records or aggregated across multiple records.

Some properties consist of a single scalar or aggregated value. For example, for an EC2 instance, Detective tracks the type of instance and the total number of bytes processed.

Time series properties track activity over time. For example, for an EC2 instance, Detective tracks over time the unique ports that it used.

Relationships

A relationship represents activity occurring between individual entities. Relationships are also extracted from the Detective source data.

Similar to an entity, a relationship has a type, which identifies the types of entities involved and the direction of the connection. An example of a relationship type is IP addresses connecting to EC2 instances.

For each individual relationship, such as a specific IP address connecting to a specific instance, Detective tracks the occurrences over time.

Types of entities in the behavior graph data structure

The behavior graph data structure consists of entity and relationship types that do the following:

- Track the servers, IP addresses, and user agents being used

- Track the AWS users, roles, and accounts being used
- Track the network connections and authorizations that occur in your AWS environment

The behavior graph data structure contains the following entity types.

AWS account

AWS accounts that are present in the Detective source data.

For each account, Detective answers several questions:

- What API calls has the account used?
- What user agents has the account used?
- What autonomous system organizations (ASOs) has the account used?
- In what geographic locations has the account been active?

AWS role

AWS roles that are present in the Detective source data.

For each role, Detective answers several questions:

- What API calls has the role used?
- What user agents has the role used?
- What ASOs has the role used?
- In what geographic locations has the role been active?
- What resources have assumed this role?
- What roles has this role assumed?
- What role sessions have involved this role?

AWS user

AWS users that are present in the Detective source data.

For each user, Detective answers several questions:

- What API calls has the user used?
- What user agents has the user used?
- In what geographic locations has the user been active?
- What roles has this user assumed?

- What role sessions have involved this user?

Federated user

Instances of a federated user. Examples of federated users include the following:

- An identity that logs in using Security Assertion Markup Language (SAML)
- An identity that logs in using web identity federation

For each federated user, Detective answers these questions:

- What identity provider did the federated user authenticate with?
- What was the audience of the federated user? The audience identifies the application that requested the web identity token of the federated user.
- In what geographic locations has the federated user been active?
- What user agents has the federated user used?
- What ASOs has the federated user used?
- What roles has this federated user assumed?
- What role sessions have involved this federated user?

EC2 instance

EC2 instances that are present in the Detective source data.

For EC2 instances, Detective answers several questions:

- What IP addresses have communicated with the instance?
- What ports have been used to communicate with the instance?
- What volume of data has been sent to and from the instance?
- What VPC contains the instance?
- What API calls has the EC2 instance used?
- What user agents has the EC2 instance used?
- What ASOs has the EC2 instance used?
- In what geographic locations has the EC2 instance been active?
- What roles has the EC2 instance assumed?

Role session

Instances of a resource that is assuming a role. Each role session is identified by the role identifier and a session name.

For each role, Detective answers several questions:

- What resources were involved in this role session? In other words, what role was assumed, and what resource assumed the role?

Note that for cross-account role assumption, Detective cannot identify the resource that assumed the role.

- What API calls has the role session used?
- What user agents has the role session used?
- What ASOs has the role session used?
- In what geographic locations has the role session been active?
- What user or role started this role session?
- What role sessions started from this role session?

Finding

Findings uncovered by Amazon GuardDuty that are fed into the Detective source data.

For each finding, Detective tracks the finding type, origin, and the time window for the finding activity.

It also stores information specific to the finding, such as roles or IP addresses that are involved in the detected activity.

IP address

IP addresses that are present in the Detective source data.

For each IP address, Detective answers several questions:

- What API calls has the address used?
- What ports has the address used?
- What users and user agents have used the IP address?
- In what geographic locations has the IP address been active?
- What EC2 instances has this IP address been assigned to and communicated with?

S3 bucket

S3 buckets that are in the Detective source data.

For each S3 bucket, Detective answers these questions:

- What principals interacted with the S3 bucket?
- What API calls were made to the S3 bucket?
- From what geographic locations did principals make API calls to the S3 bucket?
- What user agents were used to interact with the S3 bucket?
- What ASOs were used to interact with the S3 bucket?

You can delete an S3 bucket and then create a new bucket with the same name. Because Detective uses the S3 bucket name to identify the S3 bucket, it treats these as a single S3 bucket entity. On the entity profile, **Creation time** is the first creation time. **Deletion time** is the most recent deletion time.

To view all of the creation and deletion events, set the scope time to start with the creation time and end with the deletion time. On the **Overall API call volume** profile panel, display the activity details for the scope time. Filter the API methods to show Create and Delete methods. See [the section called “Overall API call volume”](#).

User agent

User agents that are present in the Detective source data.

For each user agent, Detective answers questions such as the following:

- What API calls has the user agent used?
- What users and roles have used the user agent?
- What IP addresses have used the user agent?

EKS Cluster

EKS clusters that are present in the Detective source data.

Note

To see complete details for this entity type the optional EKS audit logs data source must be enabled. For more info see [Optional data sources](#)

For each EKS cluster, Detective answers questions such as the following:

- What Kubernetes API calls have been run in this cluster?
- What Kubernetes users and service accounts (subjects) are active in this cluster?
- What containers have been launched in this cluster?

- What images are used to launch containers in this cluster?

Kubernetes Pod

Kubernetes pods that are present in the Detective source data.

Note

To see complete details for this entity type the optional EKS audit logs data source must be enabled. For more info see [Optional data sources](#)

For each pod, Detective answers questions such as the following:

- What container images in this pod are common in my accounts?
- What activity has been directed at this pod?
- What containers run in this pod?
- Are registries from containers in this pod common in my accounts?
- What other containers are running in the other pods of the workload?
- Are there any anomalous containers in this pod that are not in the other pods of the workload?

Container Image

Container images that are present in the Detective source data.

Note

To see complete details for this entity type the optional EKS audit logs data source must be enabled. For more info see [Optional data sources](#)

For each container image, Detective answers questions such as the following:

- What other images in my environment share the same repository or registry with this image?
- How many copies of this image are running in my environment?

Kubernetes Subject

Kubernetes subjects that are present in the Detective source data. A Kubernetes subject is a user or service account.

Note

To see complete details for this entity type the optional EKS audit logs data source must be enabled. For more info see [Optional data sources](#)


For each subject, Detective answers questions such as the following:

- What IAM principals have authenticated as this subject?
- What findings are associated with this subject?
- What IP addresses is the subject using?

Source data used in a Detective behavior graph

To populate a behavior graph, Amazon Detective uses source data from the behavior graph administrator account and member accounts.

With Detective, you can access up to a year of historical event data. This data is available through a set of visualizations that show changes in the type and volume of activity over a selected time window. Detective links these changes to GuardDuty findings.



For details about the behavior graph data structure, see [Overview of the behavior graph data structure](#) in *Detective User Guide*.

Types of core data sources in Detective

Detective ingests data from these types of AWS logs:

- AWS CloudTrail logs
- Amazon Virtual Private Cloud (Amazon VPC) flow logs
 - Ingests both IPv4 and IPv6 records, but not MAC records produced by Elastic Fabric Adapters.
 - Ingests log records when the value of the `log-status` field is in OK state. For more information, see [Flow log records](#) in the Amazon VPC User Guide.
 - Ingests flow logs produced by Amazon Elastic Compute Cloud instances running in those VPCs only. No other resources, such as NAT gateways, RDS instances, or Fargate clusters are used.
 - Ingests both accepted and rejected traffic.
- For accounts that are enrolled in GuardDuty, Detective also ingests GuardDuty findings.

Detective consumes CloudTrail and VPC flow log events using independent and duplicative streams of CloudTrail and VPC flow logs. These processes do not affect or use your existing CloudTrail and VPC flow log configurations. They also do not affect the performance of or increase your costs for these services.

Types of optional data sources in Detective

Detective offers optional source packages in addition to the three data sources offered in the Detective core package (the core package includes AWS CloudTrail logs, VPC Flow logs, and GuardDuty findings). An optional data source package can be started or stopped for a behavior graph at any time.

Detective provides a 30-day free trial for all core and optional source packages per Region.

Note

Detective retains all data received from each data source package for up to 1 year.

Currently the following optional source packages are available:

- **EKS audit logs**

This optional data source package allows Detective to ingest detailed information on EKS clusters in your environment and adds that data to your behavior graph. Detective correlates user activities with AWS CloudTrail Management events and network activity with Amazon VPC Flow Logs without the need for you to enable or store these logs manually. See [Amazon EKS audit logs](#) for details.

- **AWS security findings**

This optional data source package allows Detective to ingest data from Security Hub and adds that data to your behavior graph. See [AWS security findings](#) for details.

Starting or stopping an optional data source:

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. From the navigation panel under **Settings**, choose **General**.
3. Under **Optional source packages**, select **Update**. Then select the data source you wish to enable or deselect a box for an already enabled data source and choose **Update** to change which data source packages are enabled.

Note

If you stop and then restart an optional data source you will see a gap in the data displayed on some entity profiles. This gap will be noted in the console display and represent the period of time when the data source was stopped. When a data source is restarted Detective does not retroactively ingest data.

Amazon EKS audit logs

Amazon EKS audit logs is an optional data source package that can be added to your Detective behavior graph. You can view the available optional source packages, and their status in your account, from the **Settings** page in the console or through the Detective API.

A 30 day free trial is provided for this data source. To learn more see [Free trial for optional data sources](#).

Enabling Amazon EKS audit logs allows Detective to add in-depth information about resources created with Amazon EKS to your behavior graph. This data source enhances the information provided about the following entity types: EKS Cluster, Kubernetes Pod, Container Image and Kubernetes subject.

Additionally, if you have enabled EKS audit logs as a data source in Amazon GuardDuty you will be able to see details for Kubernetes findings from GuardDuty. For more info on enabling this data source in GuardDuty see [Kubernetes protection in Amazon GuardDuty](#).

Note

This data source is enabled by default for new behavior graphs created after July 26, 2022. For behavior graphs created before July 26, 2022 it must be enabled manually.

Adding or removing Amazon EKS audit logs as an optional data source:

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. From the navigation panel under **Settings**, choose **General**.
3. Under **Source packages**, select **EKS audit logs** to enable this data source. If it is already enabled, select it again to stop ingesting **EKS audit logs** into your behavior graph.

AWS security findings

AWS security findings is an optional data source package that can be added to your Detective behavior graph.

You can view the available optional source packages, and their status in your account, from the Settings page in the console or through the Detective API.

A 30 day free trial is provided for this data source. To learn more see [Free trial for optional data sources](#).

Enabling **AWS security findings** allows Detective to use the findings from Security Hub aggregated by Security Hub from upstream services in a standard findings format called the AWS Security Format (ASFF), which eliminates the need for time-consuming data conversion efforts. Then it correlates ingested findings across products to prioritize the most important ones.

Adding or removing AWS security findings as an optional data source:

Note

The AWS security findings data source is enabled by default for new behavior graphs created after May 16, 2023. For behavior graphs created before May 16, 2023 it must be enabled manually.

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. From the navigation panel under **Settings**, choose **General**.
3. Under **Source packages**, select AWS security findings to enable this data source. If it is already enabled, select it again to stop ingesting AWS Security Finding Format (ASFF) findings into your behavior graph.

Currently supported findings

Detective ingests all ASFF findings in Security Hub from services that are owned by Amazon or AWS.

- To see the list of supported service integrations, see [Available AWS service integrations](#) in the AWS Security Hub User Guide.
- For the list of supported resources, see [Resources](#) in the AWS Security Hub User Guide.
- AWS Service Findings with a Compliance status not set to FAILED and cross-Region aggregated findings are not ingested.

How Detective ingests and stores source data

When Detective is enabled, Detective begins ingesting source data from the behavior graph administrator account. As member accounts are added to the behavior graph, Detective also begins using the data from those member accounts.

Detective source data consists of structured and processed versions of the original feeds. To support Detective analytics, Detective stores copies of the Detective source data.

The Detective ingest process feeds data into Amazon Simple Storage Service (Amazon S3) buckets in the Detective source data store. As new source data arrives, other Detective components pick up

the data and start the extraction and analytics processes. For more information, see [How Detective uses source data to populate a behavior graph](#) in *Detective User Guide*.

How Detective enforces the data volume quota for behavior graphs

Detective has strict quotas on the volume of data it allows in each behavior graph. The data volume is the amount of data per day that flows into the Detective behavior graph.

Detective enforces these quotas when an administrator account enables Detective, and when a member account accepts an invitation to contribute to a behavior graph.

- If the data volume for an administrator account exceeds 10 TB per day, then the administrator account cannot enable Detective.
- If the added data volume from a member account would cause the behavior graph to exceed 10 TB per day, the member account cannot be enabled.

The data volume for a behavior graph also can grow naturally over time. Detective checks the behavior graph data volume each day to make sure that it does not exceed the quota.

If the behavior graph data volume is approaching the quota, Detective displays a warning message on the console. To avoid exceeding the quota, you can remove member accounts.

If the behavior graph data volume exceeds 10 TB per day, then you cannot add a new member account to the behavior graph.

If the behavior graph data volume exceeds 15 TB per day, then Detective stops ingesting data into the behavior graph. The 15 TB per day quota reflects both normal data volume and spikes in the data volume. When this quota is reached, no new data is ingested into the behavior graph, but existing data is not removed. You can still use that historical data for investigation. The console displays a message to indicate that the data ingest is suspended for the behavior graph.

If the data ingest is suspended, you must work with AWS Support to get it re-enabled. If possible, before you contact AWS Support, try to remove member accounts to get the data volume below the quota. This makes it easier to re-enable the data ingest for the behavior graph.

Using the Summary dashboard

Use the Summary dashboard in Amazon Detective to identify entities to investigate the origin of activity during the previous 24 hours. The Amazon Detective Summary dashboard helps you to identify entities that are associated with specific types of unusual activity. It is one of several possible starting points for an investigation.

To display the **Summary** dashboard, in the Detective navigation pane, choose **Summary**. The **Summary** dashboard is also displayed by default when you first open the Detective console.

From the **Summary** dashboard, you can identify entities that meet the following criteria:

- Investigations that show potential security events identified by Detective
- Entities involved in activity that occurred in newly observed geolocations
- Entities that made the largest number of API calls
- EC2 instances that had the largest volume of traffic
- Container clusters that had the largest number of containers

From each **Summary** dashboard panel, you can pivot to the profile for a selected entity.

As you review the **Summary** dashboard, you can adjust the **Scope time** to view the activity for any 24-hour time frame in the previous 365 days. When you change the **Start date and time**, the **End date and time** is automatically updated to 24 hours after your chosen start time.

With Detective, you can access up to a year of historical event data. This data is available through a set of visualizations that show changes in the type and volume of activity over a selected time window. Detective links these changes to GuardDuty findings.

For more information about source data in Detective, see [Source data used in a behavior graph](#).

Investigations

Investigations shows you the potential security events identified by Detective. On the Investigations panel, you can view Critical investigations and the corresponding AWS roles and users that were impacted by security events over a set period of time. Investigations groups together indicators of compromise to help determine if a AWS resource is involved in unusual activity that could indicate malicious behavior and its impact.

Select **View all investigations** to review findings, triage finding groups, and resource details to accelerate your security investigation. Investigations are displayed depending on the selected Scope time. You can adjust the scope time to view investigations in a 24-hour time frame in the previous 365 days. You can pivot directly to **Critical investigations** to see a detailed investigation report.

If you identify a AWS role or user that seems to have suspicious activity, you can pivot directly from the **Investigations** panel to the role or user to continue your investigation. Pivot to a role or user and click **Run investigation** to generate an investigations report. Once you run an investigation on a role or user, the role or user is moved to the **Investigated** tab.

Newly observed geolocations

Newly observed geolocations highlights geographic locations that were the origin of activity during the previous 24 hours, but that were not seen during the baseline time period before that.

The panel includes up to 100 geolocations. The locations are marked on the map and listed in the table below the map.

For each geolocation, the table displays the number of failed and successful API calls made from that geolocation during the previous 24 hours.

You can expand each geolocation to display the list of users and roles that made API calls from that geolocation. For each principal, the table lists the type and the associated AWS account.

If you identify a user or role that seems suspicious, then you can pivot directly from the panel to the user or role profile to continue your investigation. To pivot to a profile, choose the user or role identifier.

Detective determines the location of requests using MaxMind GeoIP databases. MaxMind reports very high accuracy of their data at the country level, although accuracy varies according to factors such as country and type of IP. For more information about MaxMind, see [MaxMind IP Geolocation](#). If you think any of the GeoIP data is incorrect, you can submit a correction request to Maxmind at [MaxMind Correct GeolP2 Data](#).

Active finding groups in the last 7 days

Active finding groups in the last 7 days shows you correlated groupings of Detective findings, entities, and evidence in your environment that occurred over a set period of time. These groupings correlate unusual activity that could indicate malicious behavior. The summary dashboard shows

up to five groups sorted by the groups containing the most critical findings that have been active in the last week.

You can select values in the **Tactic**, **Account**, **Resource**, and **Findings** content to see more details.

Findings groups are generated on a daily basis. If you identify a finding group of interest, you can select the title to move to a detailed view of a group profile to continue your investigation.

Roles and users with the most API call volume

Roles and users with the most API call volume identifies the users and roles that have made the largest number of API calls during the previous 24 hours.

The panel can include up to 100 users and roles. For each user or role, you can see the type (user or role) and the associated account. You can also see the number of API calls issued by that user or role during the previous 24 hours.

By default, service-linked roles are displayed. Service-linked roles can produce large volumes of AWS CloudTrail activity, which displaces the principals that you want to investigate further. You can choose to turn off **Show service-linked roles**, to filter out service-linked roles from the summary dashboard view.

You can export a comma-separated values (.csv) file that contains the data in this panel. .

There is also a timeline of the API call volume for the previous 7 days. The timeline can help you to determine whether the volume of API calls is unusual for that principal.

If you identify a user or role for which the API call volume seems suspicious, then you can pivot directly from the panel to the user or role profile to continue your investigation. You can also view the profile of the account associated with the user or role. To view a profile, choose the user, role, or account identifier.

EC2 instances with the most traffic volume

EC2 instances with the most traffic volume identifies the EC2 instances that have had the largest total volume of traffic during the previous 24 hours.

The panel can include up to 100 EC2 instances. For each EC2 instance, you can see the associated account and the number of inbound bytes, outbound bytes, and total bytes from the previous 24 hours.

You can export a comma-separated values (.csv) file that contains the data in this panel.

You can also see a timeline showing the inbound and outbound traffic over the previous 7 days. The timeline can help determine whether the volume of traffic is unusual for that EC2 instance.

If you identify an EC2 instance that has suspicious traffic volume, then you can go directly from the panel to the EC2 instance profile to continue your investigation. You can also view the profile of the account that owns the EC2 instance. To view a profile, choose the EC2 instance or account identifier.

Container clusters with the most Kubernetes pods

Container clusters with the most Kubernetes pods created identifies the clusters that have had the most containers running during the previous 24 hours.

This panel includes up to 100 clusters organized by which clusters had the most findings associated with them. For each cluster you can see the associated account, the current number of containers in that cluster, and the number of findings associated with that cluster over the last 24 hours. You can export a comma-separated values (.csv) file that contains the data in this panel.

If you identify a cluster with recent findings you can pivot directly from the panel to the cluster profile to continue your investigation. You can also pivot to the profile of the account that owns the cluster. To pivot to a profile, choose the cluster name or account identifier.

Approximate value notification

On **Roles and users with the most API call volume** and **EC2 instances with the most traffic volume**, if a value is followed by an asterisk (*), it means that the value is an approximation. The true value is either equal to or greater than the displayed value.

This occurs because of the method that Detective uses to calculate the volume for each time interval. On the **Summary** page, the time interval is an hour.

For each hour, Detective calculates the total volume for the 1,000 users, roles, or EC2 instances with the largest volume. It excludes the data for the remaining users, roles, or EC2 instances.

If a resource was sometimes in the top 1,000 and sometimes not, then the calculated volume for that resource might not include all of the data. The data for the time intervals where it was not in the top 1,000 is excluded.

Note that this only applies to the **Summary** page. The profile for the user, role, or EC2 instance provides precise details.

How Amazon Detective is used for investigation

Amazon Detective makes it easy to analyze, investigate, and quickly identify the root cause of security findings or suspicious activity. Detective provides tools to support the overall investigation process. An investigation in Detective can start from a finding, a finding group, or an entity.

Investigation phases

Any Detective investigation process involves the following phases:

Triage

The investigation process starts when you are notified about a suspected instance of malicious or high-risk activity. For example, you are assigned to look into findings or alerts uncovered by services such as Amazon GuardDuty and Amazon Inspector.

In the triage phase, you determine whether you believe the activity is a true positive (genuine malicious activity) or false positive (not malicious or high-risk activity). Detective profiles support the triage process by providing insight into the activity for the involved entity.

For true positive instances, you continue to the next phase.

Scoping

During the scoping phase, analysts determine the extent of the malicious or high-risk activity and the underlying cause.

Scoping answers the following types of questions:

- What systems and users were compromised?
- Where did the attack originate?
- How long has the attack been going on?
- Is there other related activity to uncover? For example, if an attacker is extracting data from your system, how did they obtain it?

Detective visualizations can help you to identify other entities that were involved or affected.

Response

The final step is to respond to the attack in order to stop the attack, minimize the damage, and prevent a similar attack from happening again.

Starting points for a Detective Investigation

Every investigation in Detective has an essential starting point. For example, you might be assigned an Amazon GuardDuty or AWS Security Hub finding to investigate. Or you might have a concern about unusual activity for a specific IP address.

Typical starting points for an investigation include findings detected by GuardDuty and entities extracted from Detective source data.

Findings detected by GuardDuty

GuardDuty uses your log data to uncover suspected instances of malicious or high-risk activity. Detective provides resources that help you investigate these findings.

For each finding, Detective provides the associated finding details. Detective also shows the entities, such as IP addresses and AWS accounts, that are connected to the finding.

You can then explore the activity for the involved entities to determine whether the detected activity from the finding is a genuine cause for concern.

For more information, see [the section called "Finding overview"](#).

AWS security findings aggregated by Security Hub

AWS Security Hub aggregates security findings from various findings providers in a single place, and provides you with a comprehensive view of your security state in AWS. Security Hub eliminates the complexity of addressing large volumes of findings from multiple providers. It reduces the effort required to manage and improve the security of all of your AWS accounts, resources, and workloads. Detective provides resources that help you investigate these findings.

For each finding, Detective provides the associated finding details. Detective also shows the entities, such as IP addresses and AWS accounts, that are connected to the finding.

For more information, see [the section called "Finding overview"](#).

Entities extracted from Detective source data

From the ingested Detective source data, Detective extracts entities such as IP addresses and AWS users. You can use one of these as an investigation starting point.


Detective provides general details about the entity, such as the IP address or user name. It also provides details on activity history. For example, Detective can report what other IP addresses an entity has connected to, been connected to, or used.

For more information, see [Analyzing entities](#).

Amazon Detective Investigation flow

You can use Amazon Detective to investigate an entity such as an EC2 instance or an AWS user. You can also investigate security findings.

At a high level, the following image shows the process for a Detective Investigation.



Step 1: Select the entity to investigate

When looking at a finding in GuardDuty, analysts can choose to investigate an associated entity in Detective. See [the section called “Pivoting from another console”](#).

Selecting the entity takes you to the entity profile in Detective.

Step 2: Analyze visualizations on profiles

Each entity profile contains a set of visualizations that are generated from the behavior graph. The behavior graph is created from the log files and other data that are fed into Detective.

The visualizations show activity that is related to an entity. You use these visualizations to answer questions to determine whether the entity activity is unusual. See [Analyzing entities](#).

To help guide the investigation, you can use the Detective guidance provided for each visualization. The guidance outlines the displayed information, suggests questions for you to ask, and proposes next steps based on the answers. See [the section called “Using profile panel guidance”](#).

Each profile contains a list of associated findings. You can view the details for a finding, and view the finding overview. See [the section called “Viewing findings for an entity”](#).

From an entity profile, you can pivot to other entity and finding profiles, to investigate further into activity for related assets.

Step 3: Take action

Based on the results of your investigation, take the appropriate action.

For a finding that is a false positive, you can archive the finding. From Detective, you can archive GuardDuty findings. For more details, see [Archiving an Amazon GuardDuty finding](#).

Otherwise, you take the appropriate action to address the vulnerability and mitigate damage. For example, you might need to update the configuration of a resource.

Detective Investigation

You can use Amazon Detective Investigation to investigate IAM users and IAM roles using indicators of compromise, which can help you determine if a resource is involved in a security incident. An indicator of compromise (IOC) is an artifact observed in or on a network, system, or environment that can (with a high level of confidence) identify malicious activity or a security incident. With Detective Investigations you can maximize efficiency, focus on the security threats, and strengthen incidence response capabilities.

Detective Investigation uses machine learning models and threat intelligence to automatically analyze resources in your AWS environment to identify potential security incidents. It lets you proactively, effectively, and efficiently use automation built on top of Detective’s behavioral graph to improve security operations. Using Detective Investigation you can investigate attack tactics, impossible travel, flagged IP addresses, and finding groups. It performs initial security investigation steps and generates a report highlighting the risks identified by Detective, to help you understand security events and respond to potential incidents.

Topics

- [Running a Detective Investigation](#)
- [Reviewing investigations reports](#)
- [Understanding a Detective Investigations report](#)
- [Investigations report summary](#)
- [Downloading an investigation report](#)
- [Archiving an investigation report](#)

Running a Detective Investigation

Use **Run investigation** to analyze resources such as IAM users and IAM roles and to generate an investigation report. The generated report details anomalous behavior that indicates potential compromise.

Console

Follow these steps to run a Detective Investigation from the **Investigations page** using the Amazon Detective console.

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Investigations**.
3. In the **Investigations** page, choose **Run investigation** in the top right corner.
4. In the **Select resource** section, you have three ways to run an investigation. You can choose to run the investigation for a resource recommended by Detective. You can run the investigation for a specific resource. You can also investigate a resource from the Detective Search page.

1. Choose a recommended resource – Detective recommends resources based on its activity in findings and finding groups. To run the investigation for a resource recommended by Detective, in the **Recommended resources** table, select a resource to investigate.

The Recommended resources table provides the following details:

- **Resource ARN** – The Amazon Resource Name (ARN) of the AWS resource.

- **Reason to investigate** – Displays the key reason(s) to investigate the resource. The reasons for which Detective recommends to investigate a resource are as follows:
 - If a resource was involved in a High Severity finding in the last 24 hours.
 - If a resource was involved in a finding group observed in the last 7 days. Detective finding groups let you examine multiple activities as they relate to a potential security event. For more details, see [the section called “Finding groups”](#).
 - If a resource was involved in a finding in the last 7 days.
 - **Latest finding** – Latest findings are prioritized on top of the list.
 - **Resource type** – Identifies the type of resource. For example, an AWS user or AWS role.
2. Specify an AWS role or user with an ARN – You can select an AWS role or AWS user and run an investigation for the specific resource.

Follow these steps to investigate a specific resource type.

- a. From the **Select resource type** drop-down list, choose AWS role or AWS user.
 - b. Enter the **Resource ARN** of the IAM resource. For more details about Resource ARNs, see [Amazon Resource Names \(ARNs\)](#) in the IAM User Guide.
3. Find a resource to investigate from the Search page – You can search all of your IAM resources from the Detective **Search** page.

Follow these steps to investigate a resource from the Search page.

- a. In the navigation pane, choose **Search**.
 - b. In the Search page, search for an IAM resource.
 - c. Navigate to the profile page of the resource and run investigation from there.
5. In the **Investigation scope time** section, choose the **Scope time** for the investigation to assess the selected resource's activity. You can select a **Start date** and **Start time**; and **End date** and **End time** in UTC format. The selected scope time window can be between at a minimum of 3 hours and a maximum of 30 days.
 6. Choose **Run investigation**.

API

To run an investigation programmatically, use the [StartInvestigation](#) operation of the Detective API. If you're using the AWS Command Line Interface (AWS CLI) run the [start-investigation](#) command.

In your request, use these parameters to run an investigation in Detective:

- `GraphArn` – Specify the Amazon Resource Name (ARN) of the behavior graph.
- `EntityArn` – Specify the unique Amazon Resource Name (ARN) of the IAM user and IAM role.
- `ScopeStartTime` – Optionally, specify the data and time from which the investigation should begin. The value is an UTC ISO8601 formatted string. For example, `2021-08-18T16:35:56.284Z`.
- `ScopeEndTime` – Optionally, specify the data and time when the investigation should end. The value is an UTC ISO8601 formatted string. For example, `2021-08-18T16:35:56.284Z`.

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

```
aws detective start-investigation \  
--graph-arn arn:aws:detective:us-  
east-1:123456789123:graph:fdac8011456e4e6182facb26dfceade0  
--entity-arn arn:aws:iam::123456789123:role/rolename --scope-start-  
time 2023-09-27T20:00:00.00Z  
--scope-end-time 2023-09-28T22:00:00.00Z
```

You can also run an investigation from the following pages in Detective:

- An IAM user or IAM role profile page in Detective.
- Graph visualization pane of a finding group.
- Actions column of an involved resource.
- IAM user or IAM role on a finding page.

After Detective runs the investigation for a resource, an investigation report is generated. To access the report, go to **Investigations** from the navigation pane.

Reviewing investigations reports

Investigations reports lets you review the generated **Reports** for investigations that you have run previously in Detective.

To review investigations reports

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Investigations**.

Take note of the following attributes from an investigations report.

- **ID** – The generated identifier of the investigations report. You can choose this **ID** to read a summary of the investigation report, which has the details of the investigation.
- **Status** – Each investigation is associated with a **Status** based on the completion status of the investigation. Status values can be **In progress**, **Succeeded**, or **Failed**.
- **Severity** – Each investigation is assigned a **Severity**. Detective automatically assigns a severity to the finding.

A severity represents the disposition as analyzed by the investigation of a single resource at a given scope time. A severity reported by an investigation doesn't imply or otherwise indicate the criticality or importance that an affected resource might have for your organization.

Investigation severity values can be **Critical**, **High**, **Medium**, **Low**, or **Informational** from most to least severe.

Investigations that are assigned a Critical or High severity value should be prioritized for further inspection, as they are more likely to represent high-impact security issues identified by Detective.

- **Entity** – The **Entity** column contains details on the specific entities detected in the investigation. Some entities are AWS accounts, such as user and role.
- **Status** – The **Creation** date column contains details on the date and time the investigation report was first created.

Understanding a Detective Investigations report

A Detective Investigations report lists a summary of the uncommon behavior or malicious activity that indicates compromise. It also lists the recommendations that Detective suggests to mitigate the security risk.

To view an investigations report for a specific investigation ID.

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Investigations**.
3. In the **Reports** table, select an investigation **ID**.

Admin report summary Info High

We observed anomalous behavior for the role from [redacted] indicating potential compromise. The role invoked CloudTrail management actions mapped to Impact MITRE tactic(s). The role was also involved in Findings that map to the MITRE tactic(s) Discovery, as well as other tactic(s). The role was also involved in 10 findings, 1 finding group, 170 impossible travels, 3 new geolocations, and 5 new user agents.

<p>Scope time</p> <p>05/25/2023 13:00 UTC - 05/31/2023 19:00 UTC</p> <p>role</p> <p>[redacted]</p>	<p>Indicators of compromise</p> <p>5 Tactics</p> <p>0 Flagged IP</p> <p>170 Impossible travel</p> <p>1 Finding group</p>	<p>Recommendation</p> <p>Based on our investigation, we recommend you take action to mitigate what we've found on AWS role Admin. Please review Security Best Practices in IAM to secure your AWS resource.</p>
--	--	---

Detective generates the report for the selected **Scope** time and **User**. The report contains an **Indicators of Compromise** section that includes details regarding one or more of the indicators of compromise listed below. As you review each indicator of compromise, optionally choose an item to drill down and review its details.

- **Tactics, Techniques, and Procedures** – Identifies tactics, techniques, and procedures (TTPs) used in a potential security event. The MITRE ATT&CK framework is used to understand the TTPs. Tactics are based on the [MITRE ATT&CK matrix for Enterprise](#).
- **Threat Intelligence Flagged IP Addresses** – Suspicious IP addresses are flagged and identified as critical or severe threats based on Detective threat intelligence.
- **Impossible Travel** – Detects and identifies unusual and impossible user activity for an account. For example, this indicator lists a drastic change between source to destination location of a user within a short time span.
- **Related Finding Group** – Shows multiple activities as they relate to a potential security event. Detective uses graph analysis techniques that infers relationships between findings and entities, and groups them together as a finding group.
- **Related Findings** – Related activities associated with a potential security event. Lists all distinct categories of evidence that are connected to the resource or the finding group.

- **New Geolocations** – Identifies new geolocations used either at the resource or account level. For example, this indicator lists an observed geolocation that is an infrequent or unused location based on previous user activity.
- **New User Agents** – Identifies new user agents used either at the resource or account level.
- **New ASOs** – Identifies new Autonomous System Organizations (ASOs) used either at the resource or account level. For example, this indicator lists a new organization assigned as an ASO.

Investigations report summary

Investigations summary highlights anomalous indicators that require attention, for the selected scope time. Using the summary, you can more quickly identify the root cause of potential security issues, identify patterns, and understand the resources impacted by security events.

In the detailed investigations report summary, you can view the following details.

Investigations overview

In the **Overview** panel, you can see a visualization of IPs with high severity activity, which can give more context on the pathway of an attacker.

Detective highlights **Unusual activity** in the investigation, for example impossible travel from a source to a faraway destination by the IAM user.

Detective maps the investigations to tactics, techniques, and procedures (TTPs) used in a potential security event. The MITRE ATT&CK framework is used to understand the TTPs. Tactics are based on the [MITRE ATT&CK matrix for Enterprise](#).

Investigations indicators

You can use the information in the **Indicators** pane, to determine if an AWS resource is involved in unusual activity that could indicate malicious behavior and its impact. An indicator of compromise (IOC) is an artifact observed in or on a network, system, or environment that can (with a high level of confidence) identify malicious activity or a security incident.

Downloading an investigation report

You can download the Detective Investigations report in JSON format, to analyze it further or store it to your preferred storage solution such as an Amazon S3 bucket.

To download an investigations report from the Reports table.

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Investigations**.
3. Select an investigation, from the **Reports** table, and choose **Download**.

To download an investigations report from the summary page.

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Investigations**.
3. Select an investigation, from the **Reports** table.
4. In the investigations summary page, choose **Download**.

Archiving an investigation report

When you complete your investigation in Amazon Detective, you can **Archive** the investigation report. An archived investigation indicates you have completed reviewing the investigation.

You can archive or unarchive an investigation only if you are a Detective Administrator. Detective will store your archived investigations for 90 days.

To archive an investigations report from the Reports table.

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Investigations**.
3. Select an investigation, from the **Reports** table, and choose **Archive**.

To archive an investigations report from the summary page.

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Investigations**.
3. Select an investigation, from the **Reports** table.

4. In the investigations summary page, choose **Archive**.

Analyzing findings in Amazon Detective

A finding is an instance of potentially malicious activity or other risk that was detected. Amazon GuardDuty and AWS security findings are loaded into Amazon Detective so that you can use Detective to investigate the activity associated with the involved entities. GuardDuty findings are part of the Detective core package and are ingested by default. All other AWS security findings that are aggregated by Security Hub are ingested as an optional data source. See [Source data used in a behavior graph](#) for more details.

A Detective finding overview provides detailed information about the finding. It also displays a summary of the involved entities, with links to the associated entity profiles.

If a finding is correlated to a larger activity, Detective notifies you to **Go to finding group**. We recommend using finding groups to continue your investigation, as finding groups enable you to examine multiple activities that relate to a potential security event. See [the section called "Finding groups"](#).

Amazon Detective provides an interactive visualization of finding groups. This visualization is designed to help you investigate issues faster and more thoroughly with less effort. The finding group **Visualization** panel displays the findings and entities involved in a finding group. You can use this interactive visualization to analyze, understand, and triage the impact of the finding group. This panel helps visualize the information presented in the **Involved entities** and **Involved findings** table. From the visual presentation, you can select findings or entities for further analysis. See [Finding group visualization](#).

Contents

- [Analyzing a finding overview](#)
- [Analyzing finding groups](#)
- [Finding group summary powered by generative AI](#)
- [Archiving an Amazon GuardDuty finding](#)

Analyzing a finding overview

A Detective finding overview provides detailed information about the finding. It also displays a summary of the involved entities, with links to the associated entity profiles.

Scope time used for the finding overview

The scope time for a finding overview is set to the finding time window. The finding time window reflects the first and last time that the finding activity was observed.

Finding details

The panel at the right contains the details for the finding. These are the details provided by the finding provider.

From the finding details, you can also archive the finding. For more details, see [Archiving an Amazon GuardDuty finding](#).

Related entities

The finding overview contains a list of entities that are involved in the finding. For each entity, the list provides overview information about the entity. This information reflects the information on the entity details profile panel on the corresponding entity profile.

You can filter the list based on entity type. You can also filter the list based on text in the entity identifier.

To pivot to the profile for an entity, choose **See profile**. When you pivot to the entity profile, the following occurs:

- The scope time is set to the finding time window.
- On the **Associated findings** panel for the entity, the finding is selected. The finding details remain displayed at the right of the entity profile.

Troubleshooting 'Page not found'

When you navigate to an entity or a finding in Detective, you may see a **Page not found** error message.

To resolve this, do one of the following:

- Make sure that the entity or finding belongs to one of your member accounts. For information on how to review member accounts, see [Viewing the list of accounts](#).

- Make sure your administrator account is aligned with GuardDuty and/or Security Hub to pivot to Detective from these services. For the recommendations, see [Recommended alignment with GuardDuty and Security Hub](#).
- Verify that the finding occurred after the member account accepted your invitation.
- Verify the Detective behavior graph is ingesting data from an optional data source package. For more information about source data used in Detective behavior graphs, see [Source data used in a behavior graph](#).
- To allow Detective to ingest data from Security Hub and add that data to your behavior graph, you must enable Detective for AWS security findings as a data source package. For more information, see [AWS security findings](#).
- If you are navigating to an entity profile or finding overview in Detective, make sure that the URL is in the right format. For details on the formation of a profile URL, see [Navigating to an entity profile or finding overview using URL](#).

Analyzing finding groups

Amazon Detective finding groups let you examine multiple activities as they relate to a potential security event. A finding group in Amazon Detective is created when Detective detects a pattern or relationship among multiple findings that suggest they are related to the same potential security incident. This grouping helps in managing and investigating related findings more efficiently.

You can analyze the root cause for high severity GuardDuty findings using finding groups. If a threat actor is attempting to compromise your AWS environment, they typically perform a sequence of actions that lead to multiple security findings and unusual behaviors. These actions are often spread across time and entities. When security findings are investigated in isolation, it can lead to a misinterpretation of their significance, and difficulty in finding the root cause. Amazon Detective addresses this problem by applying a graph analysis technique that infers relationships between findings and entities, and groups them together. We recommend treating finding groups as the starting point for investigating the involved entities and findings.

Detective analyzes data from findings and groups them with other findings that are likely to be related based on resources they share. For example, findings related to actions taken by the same IAM role sessions or originating from the same IP address are very likely to be part of the

same underlying activity. It's valuable to investigate findings and evidence as a group, even if the associations made by Detective aren't related.

Finding groups are created based on the following criteria.

- **Temporal Proximity** – Findings that occur within a close time frame are often grouped together, as they are likely related to the same incident.
- **Common Entities** – Findings involving the same entities, such as IP addresses, users, or resources, are grouped together. This helps in understanding the scope of the incident across different parts of the environment.
- **Patterns and Behaviors** – Detective analyzes patterns and behaviors in the findings, such as similar types of attacks or suspicious activities, to determine relationships and group them accordingly.
- **Tactics, Techniques, and Procedures (TTPs)** – Findings that share similar TTPs, as described in frameworks like MITRE ATT&CK, are grouped together to highlight potential coordinated attacks.

These criteria help streamline the investigation process so you can focus on correlated findings that likely represent the same security incident.

In addition to findings, each group includes entities involved in the findings. The entities can include resources outside of AWS such as IP Addresses or user agents.

Note

After an initial GuardDuty finding occurs that is related to another finding, the finding group with all related findings and all involved entities is created within 48 hours.

Understanding the finding groups page

The finding groups page lists all the finding groups collected by Amazon Detective from your behavior graph. take note of the following attributes of finding groups:

Severity of a group

Each finding group is assigned a severity based on the AWS Security Finding Format (ASFF) severity of the associated findings. ASFF finding severity values are **Critical, High, Medium,**

Low, or **Informational** from most to least severe. The severity of a grouping is equal to the highest severity finding among the findings in that grouping.

Groups that consist of **Critical** or **High** severity findings that impact a large number of entities should be prioritized for investigations, as they are more likely to represent high-impact security issues.

Group title

In the **Title** column, each group has a unique ID and a non-unique title. These are based on the ASFF type namespace for the group and the number of findings within that namespace in the cluster. For example, if a grouping has the title: Group with: **TTP (2), Effect (1), and Unusual behavior (2)** it includes five total findings consisting of two findings in the **TTP** namespace, one finding in the **Effect** namespace, and two findings in the **Unusual Behavior** namespace. For a complete list of namespaces, see [Types taxonomy for ASFF](#).

Tactics in a group

The **Tactics** column in a group details which tactics category the activity falls into. The tactics, techniques, and procedures categories in the following list align to the [MITRE ATT&CK matrix](#).

You can select a tactic on the chain to see a description of the tactic. Following the chain is a list of the tactics detected within the group. These categories and the activities they typically represent are as follows:

- **Initial Access** – An adversary is trying to get into someone else’s network.
- **Execution** – An adversary is trying to get into someone else’s network.
- **Persistence** – An adversary is trying to maintain their foothold.
- **Privilege Escalation** – An adversary is trying to gain higher-level permissions.
- **Defense Evasion** – An adversary is trying to avoid being detected.
- **Credential Access** – An adversary is trying to steal account names and passwords.
- **Discovery** – An adversary is trying to understand and learn about an environment.
- **Lateral Movement** – An adversary is trying to move through an environment.
- **Collection** – An adversary is trying to gather data of interest to their goal.
- **Command and Control** – An adversary is trying to get into someone else’s network.
- **Exfiltration** – An adversary is trying to steal data.
- **Impact** – An adversary is trying to manipulate, interrupt, or destroy your systems and data.

- **Other** – Indicates activity from a finding that does not align with tactics listed in the matrix.

Entities within a group

The **Entities** column contains details on the specific entities detected within this grouping. Select this value for a breakdown of entities based on the categories: **Identity**, **Network**, **Storage**, and **Compute**. Examples of entities in each category are:

- **Identity** – IAM principals and AWS accounts, such as user and role
- **Network** – IP address or other networking and VPC entities
- **Storage** – Amazon S3 buckets or DDBs
- **Compute** Amazon EC2 instances or Kubernetes containers

Accounts within a group

The **Accounts** column tells you what AWS accounts own entities involved with the findings in the group. The AWS Accounts are listed by name and AWS ID so you can prioritize investigations of activity involving critical accounts.

Findings within a group

The **Findings** column has a lists the entities within a group by severity. The findings include Amazon GuardDuty findings, Amazon Inspector findings, AWS security findings, and evidence from Detective. You can select the graph to see an exact count of findings by severity.

GuardDuty findings are part of the Detective core package and are ingested by default. All other AWS security findings that are aggregated by Security Hub are ingested as an optional data source. See [Source data used in a behavior graph](#) for more details.

Informational findings in finding groups

Amazon Detective identifies additional information related to a finding group based on data in your behavior graph collected within the last 45 days. Detective presents this information as a finding with the **Informational** severity. Evidence provides supporting information that highlights an unusual activity or unknown behavior that is potentially suspicious when viewed within a finding group. This might include newly observed geolocations or API calls observed within the scope time of a finding. Evidence findings are only viewable in Detective and are not sent to AWS Security Hub.

Detective determines the location of requests using MaxMind GeoIP databases. MaxMind reports very high accuracy of their data at the country level, although accuracy varies according to factors

such as country and type of IP. For more information about MaxMind, see [MaxMind IP Geolocation](#). If you think any of the GeoIP data is incorrect, you can submit a correction request to Maxmind at [MaxMind Correct GeoIP2 Data](#).

You can observe evidence for different principal types (such as IAM user or IAM role). For some evidence types, you can observe evidence for **all accounts**. This means evidences affect your entire behavior graph. If an evidence finding is observed for all accounts, you will also see at least one additional informational evidence finding of the same type for an individual IAM role. For example, if you see a **New geolocation observed for all accounts** finding, you will see another for **New geolocation observed for a principal**.

Types of evidence in finding groups

- New geolocation observed
- New Autonomous System Organization (ASO) observed
- New user agent observed
- New API call issued
- New geolocation observed for all accounts
- New IAM principal observed for all accounts

Finding group profiles

When you select a group title, a finding group profile opens with additional details about that group. The details panel in the finding groups profile page supports the display of up to 1000 entities and findings for finding groups parent and children.

The group profile page displays the set **Scope time** of the group. This is the date and time from the earliest finding or evidence included in the group to the most recently updated finding or evidence in a group. You can also see the **Finding group severity**, which is equal to the highest severity category among findings in the group. Other details within this profile panel include:

- The **Involved tactics** chain shows you which tactics, are attributed to the findings in the group. Tactics are based on the [MITRE ATT&CK Matrix for Enterprise](#). The tactics are shown as a chain of colored dots that represents the typical progression of an attack from the earliest to latest stages. This means the leftmost circles on the chain typically represent less severe activities where an adversary is trying to gain or maintain access your environment. Conversely, activities toward the right are the most severe and can include data tampering or destruction.

- The relationships that this group has with other groups. Occasionally, one or more previously unconnected groups of findings could be merged into a new group based on a newly discovered link, for example, a finding that involves entities from the existing groups. In this case, Amazon Detective deactivates the parent groups and creates a child group. You can trace the lineage for any group back to its parent groups. Groups can have the following relationships:
 - **Child finding group** – A finding group created when a finding involved in two other finding groups is involved in a new finding. The parent groups of the finding are listed for any child group.
 - **Parent finding group** – A finding group is a parent when a child group has been created from it. If a finding group is a parent, the related children are listed with it. A parent group's status becomes **Inactive** when it's merged into an **Active** child group.

There are two information tabs that open profile panels. Using the **Involved entities** and **Involved findings** tabs, you can view further details about the group.

Use **Run investigation** to generate an investigation report. The generated report details anomalous behavior that indicates compromise. .

Profile panels within groups

Involved entities

Focuses on the entities in the finding group, including what findings within the group each entity is linked to. The tags attached to each entity are also displayed so you can quickly identify important entities based on tagging. Select an entity to view its entity profile.

Involved findings

Has details about each finding, including finding severity, each entity involved, and when that finding was first and last seen. Select a finding type in the list to open a finding details panel with additional information about that finding. As part of the **Involved findings** panel, you may see **Informational** findings based on Detective evidence from your behavior graph.

Finding group visualization

Amazon Detective provides an interactive visualization of finding groups. This visualization is designed to help you investigate issues faster and more thoroughly with less effort. The finding group **Visualization** panel displays the findings and entities involved in a finding group. You can

use this interactive visualization to analyze, understand, and triage the impact of the finding group. This panel helps visualize the information presented in the **Involved entities** and **Involved findings** table. From the visual presentation, you can select findings or entities for further analysis.

Detective finding groups with aggregated findings are a cluster of findings that are connected to the same type of resource. With aggregated findings, you can quickly assess the makeup of a finding group and interpret security issues faster. In the finding groups details panel, similar findings are combined and you can expand the findings to view relatively similar findings together. For example, an evidence node, which has informational findings and medium findings of the same type are aggregated. Currently, you can view the title, source, type, and severity of finding groups with aggregated findings.

From this interactive panel, you can:


- Use **Run investigation** to generate an investigation report. The generated report details anomalous behavior that indicates compromise. For more details, see [Detective Investigations](#).
- View more details on finding groups with aggregated findings to analyze the involved evidence, entities, and findings.
- View the labels for the entities and findings to identify the affected entities with potential security issues. You can toggle off the **Label**.
- Rearrange the entities and findings to better understand their interconnectedness. Isolate entities and findings from a group by moving the selected item in the finding group.
- Select the evidences, entities, and findings to view more details about them. To select multiple items, choose **command/control** and either choose the items, or drag and drop them using your pointer.
- Adjust the layout to fit all entities and findings into the finding group window. View what entity types are prevalent in a finding group.

 **Note**

The finding group **Visualization** panel supports the display of finding groups with up to 100 entities and findings.

You can use the drop-down to view the findings and entities in a **Radial**, **Circle**, **Force-directed**, or **Grid** layout. The **Radial** layout provides improved visualization for easier data interpretation.

The **Force-directed** layout positions the entities and findings so that links are a consistent length between items and the links are distributed evenly. This helps to reduce overlapping. The layout that you select defines the placement of findings in the **Visualization** panel.



You can use the following keyboard shortcuts to interact with the finding group Visualization panel:

- Click – Selects a single node, deselects all other nodes, deselects all nodes if white space is clicked.
- Ctrl + Click – Selects a single node, does not deselect other nodes.
- Drag – Pans the view.
- Ctrl + Drag – Marquee selects, does not deselect other nodes.
- Shift + Drag – Marquee selects, deselects all other nodes.
- Arrow keys – Changes the focus between nodes.
- Ctrl + Space – Selects or deselects the currently focused node.
- Shift + Arrow keys – Changes the focus between nodes and selects them.

The dynamic **Legend** changes based on the entities and findings in your current graph. It helps you identify what each visual element represents.

Finding group summary powered by generative AI

By default, Amazon Detective automatically provides summaries of an individual finding group. The summaries are powered by generative artificial intelligence (generative AI) models hosted on [Amazon Bedrock](#).

By using finding groups, you can examine multiple security findings, as they relate to a potential security event, and identify potential threat actors. Finding group summaries for finding groups builds upon these capabilities. Finding group summaries consume the data for a finding group, rapidly analyze relationships between the findings and affected resources, and then summarize potential threats in natural language. You can leverage these summaries to identify larger security threats, improve investigation efficiency, and shorten the response timelines.

Note

Finding group summaries powered by generative AI may and not always provide completely accurate information. See [AWS Responsible AI Policy](#) for more information.

Reviewing finding group summary


The finding group summary for a finding group gives you a clear, detailed explanation of a security event. In natural language, the explanation includes a succinct title, a summary of the resources involved, and curated information about those resources.

To review a finding group summary

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Finding groups**.
3. In the **Finding groups** table, choose the finding group that you want to display a summary of. A details page appears.

On the details page, you can use the **Summary** pane to review a generated, descriptive summary of the top findings in the finding group. You can also review an analysis of the top threat events in the finding group, which you can then investigate further. To add the generated summary to your notes or a ticketing system, choose the copy icon in the pane. This copies the summary to your clipboard. You can also share your feedback about the finding group summary output in the

summary, which can provide a better experience in the future. To share your feedback, choose the thumbs up or thumbs down icon, depending on the nature of your feedback.

 **Note**

If you provide feedback about the finding group summary, your feedback is not used for model tuning. We use it only to help facilitate that the prompts in Detective are crafted effectively.



Summary - new [Info](#)

Credentials exfiltration from i-0e5f7e596391b28eb using role privilegedRole

Instance i-0e5f7e596391b28eb had newly observed API calls and user agents for role privilegedRole.

Credentials for role privilegedRole on i-0e5f7e596391b28eb were exfiltrated and used from account [REDACTED] and IP [REDACTED].

The exfiltrated credentials were used to access S3 bucket private-bucket-[REDACTED].

i-0e5f7e596391b28eb was vulnerable to CVE-2021-44228 and CVE-2021-45046.



Disabling finding group summary

By default, finding group summary is enabled for finding groups. You can disable finding group summary at any time. If you disable, you can enable them again later.

To disable finding group summary

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Preferences**.
3. Under **Finding group summary**, choose **Edit**.
4. Turn off **Enabled**.
5. Choose **Save**.

Enabling finding group summary

If you previously disabled finding group summary for finding groups, you can enable them again at any time.

To enable finding group summary

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Preferences**.
3. Under **Finding group summary**, choose **Edit**.
4. Turn on **Enabled**.
5. Choose **Save**.

Supported Regions

Finding group summary is available in the following AWS Regions.

- US East (N. Virginia)
- US West (Oregon)
- Asia Pacific (Tokyo)
- Europe (Frankfurt)

Archiving an Amazon GuardDuty finding

When you complete your investigation of an Amazon GuardDuty finding, you can archive the finding from Amazon Detective. This saves you the trouble of having to return to GuardDuty to make the update. Archiving a finding indicates that you have finished your investigation of it.

You can only archive a GuardDuty finding from within Detective if you are also the GuardDuty administrator account for the account associated with the finding. If you are not a GuardDuty administrator account and you attempt to archive a finding, GuardDuty displays an error.

To archive a GuardDuty finding

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective console, in the finding details panel, choose **Archive finding**.
3. When prompted to confirm, choose **Archive**.

You can view archived GuardDuty findings in the GuardDuty console. The archived finding is stored in GuardDuty for 90-days and can be viewed at any time during that period. You can view suppressed findings in the GuardDuty console by selecting Archived from the findings table, or through the GuardDuty API using the [ListFindings API](#) with a findingCriteria criterion of service.archived equal to true. To learn more, see [Suppression Rules](#) in the *Amazon GuardDuty User Guide*.

Analyzing entities in Amazon Detective

An entity is a single object extracted from the source data. Examples include a specific IP address, Amazon EC2 instance, or AWS account. For a list of entity types, see [the section called “Types of entities in the behavior graph data structure”](#).

An Amazon Detective entity profile is a single page that provides detailed information about the entity and its activity. You might use an entity profile to get supporting details for an investigation into a finding or as part of a general hunt for suspicious activity.

Contents

- [Using entity profiles](#)
- [Viewing and interacting with profile panels](#)
- [Navigating directly to an entity profile or finding overview](#)
- [Pivoting from a profile panel to another console](#)
- [Exploring activity details on a profile panel](#)
- [Managing the scope time](#)
- [Viewing details for associated findings](#)
- [Viewing details for high-volume entities](#)

Using entity profiles

An entity profile appears when you perform one of the following actions:

- From the Amazon GuardDuty console, choose the option to investigate an entity that is related to a selected finding.

See [the section called “Pivoting from another console”](#).

- Go to the Detective URL for the entity profile.

See [the section called “Navigating using a URL”](#).

- Use the Detective search in the Detective console to look up an entity.
- Choose a link to the entity profile from another entity profile or from a finding overview.

Scope time for an entity profile

When you navigate directly to an entity profile without providing the scope time, the scope time is set to the previous 24 hours.

When you navigate to an entity profile from another entity profile, the currently selected scope time remains in place.

When you navigate to an entity profile from a finding overview, the scope time is set to the finding time window.

For information on customizing the scope time to limit the data displayed on entity profiles, see [Managing the scope time](#).

Entity identifier and type

At the top of the profile are the entity identifier and the entity type. Each entity type has a corresponding icon, to provide a visual indicator of the type of profile.

Involved findings

Each profile contains a list of findings that the entity was involved in during the scope time.

You can see the details for each finding, change the scope time to reflect the finding time window, and go to the finding overview to look for other involved resources.

See [the section called "Viewing findings for an entity"](#).

Finding groups involving this entity

Each profile contains a list of finding groups that an entity is included in.

A finding group is made up of findings, entities, and evidence that Detective collects into a group to provide more context on possible security issues.

For more information on finding groups, see [the section called "Finding groups"](#).

Profile panels containing entity details and analytics results

Each entity profile contains a set of one or more tabs. Each tab contains one or more profile panels. Each profile panel contains text and visualizations that are generated from the behavior graph data. The specific tabs and profile panels are tailored to the entity type.

For most entities, the panel at the top of the first tab provides high-level summary information about the entity.

Other profile panels highlight different types of activity. For an entity that is involved with a finding, the information on the entity profile panels can provide additional supporting evidence to help complete an investigation. Each profile panel provides access to guidance on how to use the information. For more information, see [the section called “Using profile panel guidance”](#).


For more details about profile panels, the types of data they contain, and available options for interacting with them, see [the section called “Profile panels”](#).

Navigating in an entity profile

An entity profile contains a set of one or more tabs. Each tab contains one or more profile panels. Each profile panel contains text and visualizations that are generated from the behavior graph data.

As you scroll down through a profile tab, the following information remains visible at the top of the profile:

- Entity type
- Entity identifier
- Scope time



Viewing and interacting with profile panels

Each entity profile on the Amazon Detective console consists of a set of profile panels. A profile panel is a visualization that provides general details or highlights specific activity associated

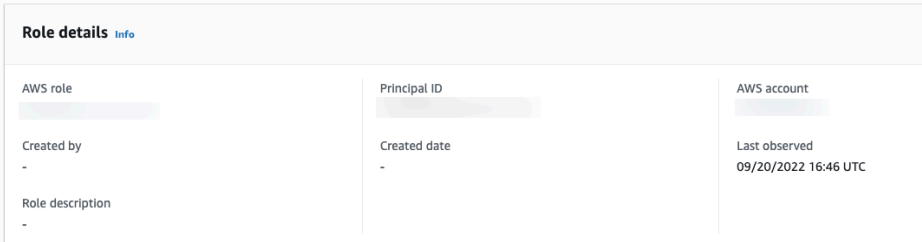
with an entity. Profile panels use different types of visualizations to present different types of information. They can also provide links to additional details or to other profiles.

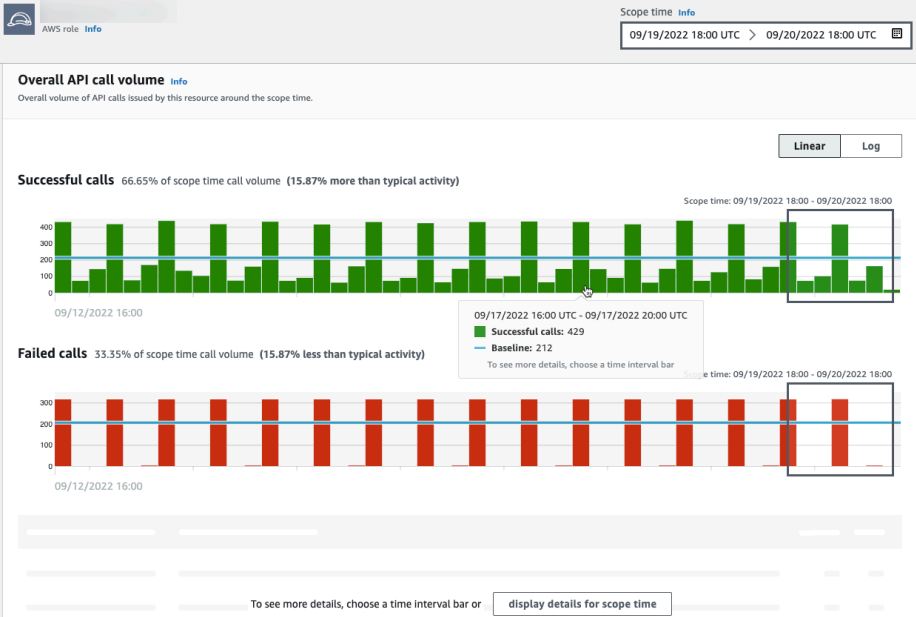
Each profile panel is intended to help analysts find answers to specific questions about entities and their associated activity. The answers to those questions help lead to a conclusion about whether the activity represents a genuine threat.

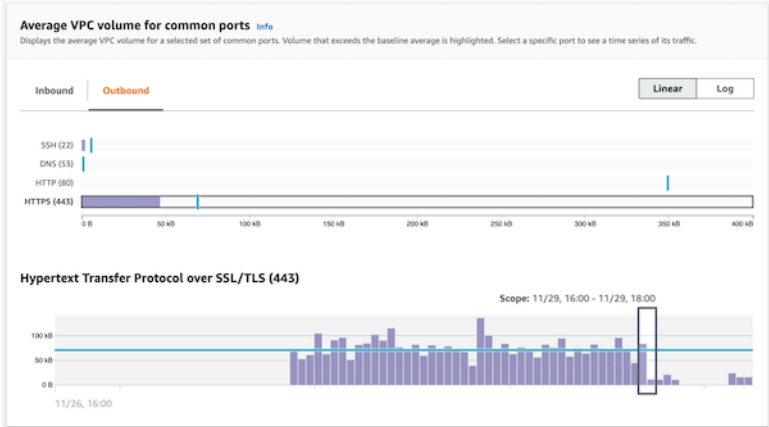
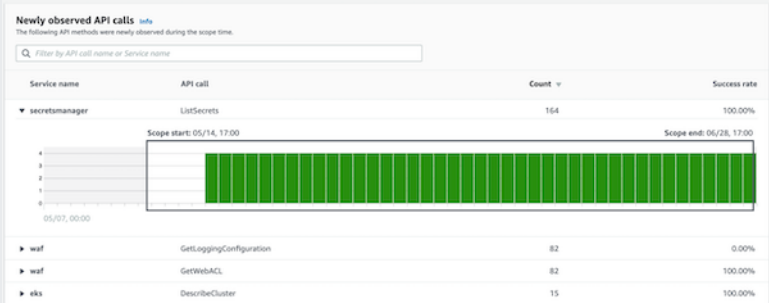
Profile panels use different types of visualizations to present different types of information.

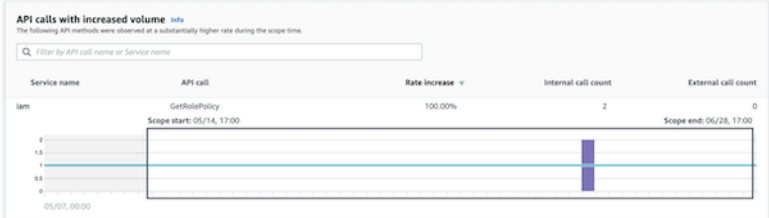
Types of information on a profile panel

Profile panels typically provide the following types of data.

Panel data type	Description
<p>High-level information about a finding or entity</p>	<p>The simplest type of panel provides some basic information about an entity.</p> <p>Examples of information included on an information panel include the identifier, name, type, and creation date.</p>  <p>Most entity profiles contain an information panel for that entity.</p>
<p>General summary of activity over time</p>	<p>Displays a summary of activity for an entity over time.</p> <p>This type of panel provides an overall view of how an entity is behaving during the scope time.</p>

Panel data type	Description
	 <p>Here are some examples of summary data provided on Detective profile panels:</p> <ul style="list-style-type: none">• Failed and successful API calls• Inbound and outbound VPC volume

Panel data type	Description
<p>Summary of activity grouped by values</p>	<p>Displays a summary of activity for an entity, grouped by specific values.</p> <p>You can see this type of profile panel on the profile for an EC2 instance. The profile panel shows the average volume of VPC flow log data to and from an EC2 instance for common ports that are associated with specific types of services.</p> 
<p>Activity that only started during the scope time</p>	<p>During an investigation, it is valuable to see what activity only began to occur during a specific time frame.</p> <p>For example, are there API calls, geographic locations, or user agents that were not seen before?</p>  <p>If the behavior graph is still in training mode, the profile panel displays a notification message. The message is removed when the behavior graph has accumulated at least two weeks of data. For more information about training mode, see the section called “Training period for new behavior graphs”.</p>

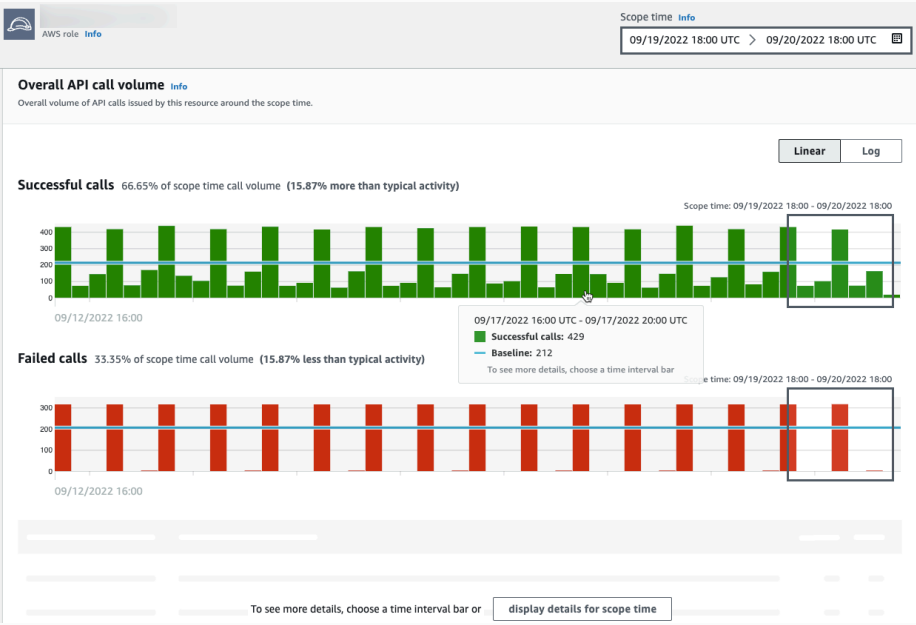
Panel data type	Description
Activity that changed significantly during the scope time	<p>Similar to the new activity panels, profile panels can also display activity that changed significantly during the scope time.</p> <p>For example, a user might regularly issue a certain API call a few times a week. If the same user suddenly issues the same call multiple times in a single day, that might be evidence of malicious activity.</p>  <p>If the behavior graph is still in training mode, the profile panel displays a notification message. The message is removed when the behavior graph has accumulated at least two weeks of data. For more information about training mode, see the section called "Training period for new behavior graphs".</p>

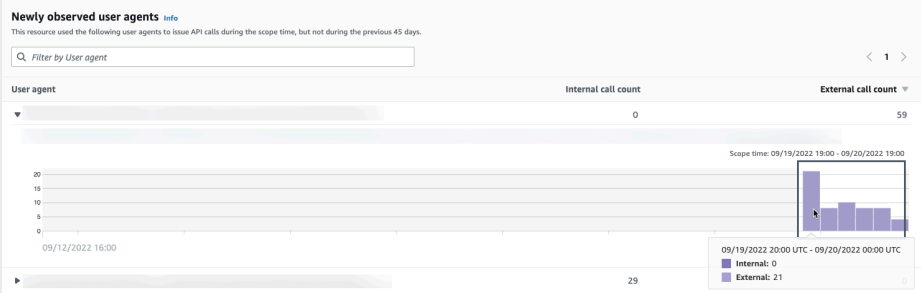
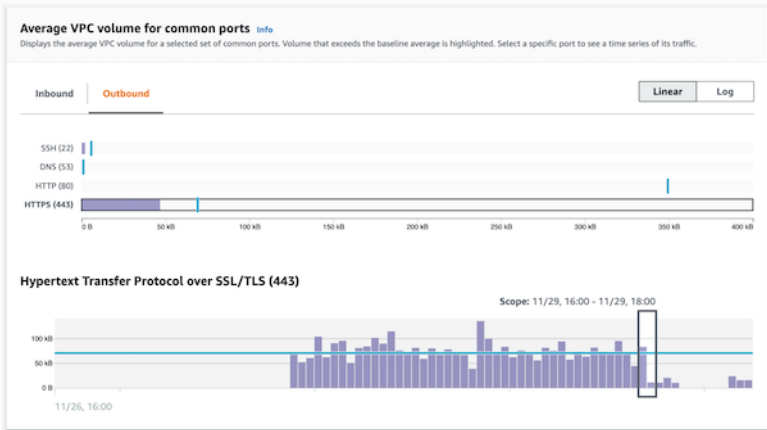
Types of profile panel visualizations

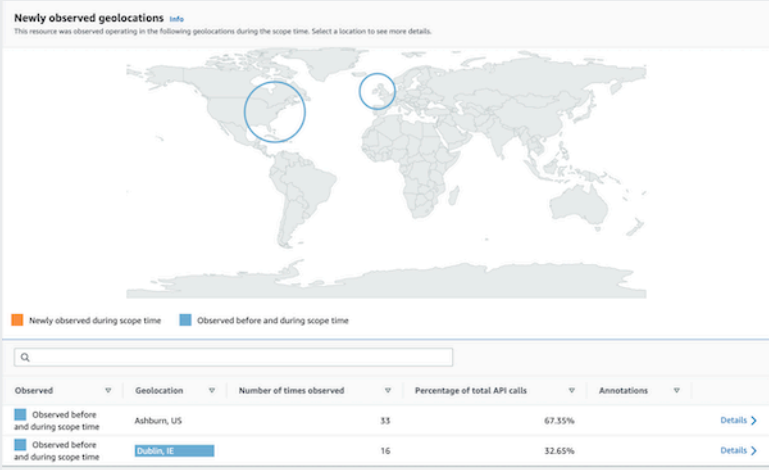
Profile panel content can take one of the following forms.

Visualization type	Description
Key-value pairs	<p>The simplest type of visualization is a set of key-value pairs.</p> <p>A finding or entity information panel is the most common example of a key-value pair panel.</p>

Visualization type	Description															
	<div data-bbox="592 214 1507 453"> <p>Role details Info</p> <table border="1"> <tr> <td>AWS role</td> <td>Principal ID</td> <td>AWS account</td> </tr> <tr> <td>Created by</td> <td>Created date</td> <td>Last observed</td> </tr> <tr> <td>-</td> <td>-</td> <td>09/20/2022 16:46 UTC</td> </tr> <tr> <td>Role description</td> <td></td> <td></td> </tr> <tr> <td>-</td> <td></td> <td></td> </tr> </table> </div> <p>Key-value pairs can also be used to add additional information to other types of panels.</p> <p>From a key-value pair panel, if a value is an identifier of an entity, then you can pivot to its profile.</p>	AWS role	Principal ID	AWS account	Created by	Created date	Last observed	-	-	09/20/2022 16:46 UTC	Role description			-		
AWS role	Principal ID	AWS account														
Created by	Created date	Last observed														
-	-	09/20/2022 16:46 UTC														
Role description																
-																
Table	<p>A table is a simple multiple-column list of items.</p> <div data-bbox="592 810 1507 966"> <p>Observed IP address assignments based on VPC Flow</p> <p>These IP addresses were assigned to this EC2 instance and also had traffic with the instance</p> <p>Q Filter by IP CIDR</p> <table border="1"> <thead> <tr> <th>IP address</th> <th>First observed</th> <th>Last observed</th> </tr> </thead> <tbody> <tr> <td>10.101.0.119</td> <td>04/27/2021 15:19 UTC</td> <td>09/20/2022 17:45 UTC</td> </tr> </tbody> </table> </div> <p>You can sort, filter, and page through the table.</p> <p>You can change the number of entries to display on each page. See the section called "Preferences for profile panels".</p> <p>If a value in the table is an identifier of an entity, then you can pivot to its profile.</p>	IP address	First observed	Last observed	10.101.0.119	04/27/2021 15:19 UTC	09/20/2022 17:45 UTC									
IP address	First observed	Last observed														
10.101.0.119	04/27/2021 15:19 UTC	09/20/2022 17:45 UTC														

Visualization type	Description
Timeline	<p>A timeline visualization shows an aggregated value for defined intervals over time.</p>  <p>The timeline highlights the current scope time, and includes additional peripheral time before and after the scope time. The peripheral time provides context for the activity in the scope time.</p> <p>Hover over a time interval to display a summary of the data for that time interval.</p>

Visualization type	Description
<p>Expandable table</p>	<p>An expandable table combines tables and timelines.</p>  <p>The visualization starts as a table.</p> <p>You can sort, filter, and page through the table.</p> <p>You can change the number of entries to display on each page. See the section called "Preferences for profile panels".</p> <p>You can then expand each row to show a timeline visualization specific to that row.</p>
<p>Bar chart</p>	<p>A bar chart shows values based on groupings.</p> <p>Depending on the chart, you might be able to choose a bar to display a timeline of related activity.</p> 

Visualization type	Description
Geolocation chart	<p>A geolocation chart displays a map that is marked to highlight data based on geographic location. It may be followed by a table containing details about individual geolocations.</p>  <p>Note that when processing incoming geographic data, Detective rounds the latitude and longitude values to a single decimal point.</p>

Other notes on profile panel content

When viewing the content of a profile panel, be aware of the following items:

Approximate count data warning

This warning indicates that items with extremely low counts do not appear due to the volume of applicable data.

To ensure a completely accurate count, reduce the amount of data. The simplest way to do that is to reduce the length of the scope time. See [the section called “Managing the scope time”](#).

Rounding for geographic locations

Detective rounds all latitude and longitude values to a single decimal point.

Changes to how Detective represents API calls

Beginning on July 14, 2021, Detective tracks the service that made each API call. Whenever Detective displays an API method, it also displays the associated service. On profile panels that

display information about API calls, the calls are always grouped by the service. For data that Detective ingested before that date, the service name is listed as **Unknown service**.

Also beginning on July 14, 2021, for accounts and roles, the activity details for the **Overall API call volume** profile panel no longer show the AKID of the resource that issued the call. For accounts, Detective displays the identifier of the principal (user or role) that issued the call. For roles, Detective displays the identifier of the role session. For data that Detective ingested before July 14, 2021, the identifier is listed as **Unknown resource**.

For profile panels that display a list of API calls, the associated timeline highlights the period of time during which this transition occurred. The highlight starts on July 14, 2021, and ends when the update was fully propagated in Detective.

Setting the preferences for a profile panel

For profile panels, you can customize the number of rows that appear on each page on profile panels by and configure the timestamp format preference.

Setting the table length

For profile panels that contain tables or expandable tables, you can configure the number of rows to display on each page.

Set your preference for the number of entries on each page.

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, under **Settings**, choose **Preferences**.
3. On the **Preferences** page, under **Table length**, click **Edit**.
4. Choose the number of table rows you want to display on each page.
5. Choose **Save**.

Setting the timestamp format

For profile panels, you can configure the timestamp format preference that will be applied to all timestamps for each IAM user or IAM role in Detective.

Note

The timestamp format preference is not applied across the entire AWS account.

Set the preference for timestamp.

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, under **Settings**, choose **Preferences**.
3. On the **Preferences** page, under **Timestamp preferences**, view and change the preferred display for all timestamps.
4. By default, the timestamp format is set to UTC. Click **Edit** to choose your local timezone.

Example:

Example

UTC - 09/20/22 16:39 UTC

Local - 09/20/2022 9:39 (UTC-07:00)

5. Choose **Save**.

Navigating directly to an entity profile or finding overview

To navigate directly to an entity profile or finding overview in Amazon Detective, you can use one of these options.

- From Amazon GuardDuty or AWS Security Hub, you can pivot from a GuardDuty finding to the corresponding Detective finding profile.
- You can assemble a Detective URL that identifies a finding or entity and sets the scope time to use.

Pivoting to an entity profile or finding overview from Amazon GuardDuty or AWS Security Hub

From the Amazon GuardDuty console, you can navigate to the entity profile for an entity that is related to a finding.

From the GuardDuty and AWS Security Hub consoles, you can also navigate to a finding overview. This also provides links to the entity profiles for the involved entities.

These links can help to streamline the investigation process. You can quickly use Detective to see the associated entity activity and determine next steps. You can then archive a finding if it is a false positive or explore further to determine the scope of the problem.

How to pivot to the Amazon Detective console

The investigation links are available for all GuardDuty findings. GuardDuty also allows you to choose whether to navigate to an entity profile or to the finding overview.

To pivot to Detective from the GuardDuty console

1. Open the GuardDuty console at <https://console.aws.amazon.com/guardduty/>.
2. If necessary, choose **Findings** in the left navigation pane.
3. On the GuardDuty **Findings** page, choose the finding.

The finding details pane displays to the right of the finding list.

4. On the finding details pane, choose **Investigate in Detective**.

GuardDuty displays a list of available items to investigate in Detective.

The list contains both the related entities, such as IP addresses or EC2 instances, and the finding.

5. Choose an entity or the finding.

The Detective console opens in a new tab. The console opens to the entity or finding profile.

If you have not enabled Detective, then the console opens to a landing page that provides an overview of Detective. From there, you can choose to enable Detective.

To pivot to Detective from the Security Hub console

1. Open the AWS Security Hub console at <https://console.aws.amazon.com/securityhub/>.
2. If necessary, choose **Findings** in the left navigation pane.
3. On the Security Hub **Findings** page, choose a GuardDuty finding.
4. In the details pane, choose **Investigate in Detective** and then choose **Investigate finding**.

When you choose **Investigate finding**, the Detective console opens in a new tab. The console opens to the finding overview.

The Detective console always opens to the Region where the finding originated, even if you pivot from your aggregation Region. For more information about finding aggregation, see [Aggregating findings across Regions](#) in the *AWS Security Hub User Guide*.

If you have not enabled Detective, the console opens to the Detective landing page. From there, you can enable Detective.

Troubleshooting the pivot

To use the pivot, one of the following must be true:

- Your account must be an administrator account for both Detective and the service you are pivoting from.
- You have assumed a cross-account role that grants you administrator account access to the behavior graph.

For more information about the recommendation to align administrator accounts, see [Recommended alignment with Amazon GuardDuty and AWS Security Hub](#).

If the pivot does not work, check the following.

- **Does the finding belong to an enabled member account in your behavior graph?** If the associated account was not invited to the behavior graph as a member account, then the behavior graph does not contain data for that account.

If an invited member account did not accept the invitation, then the behavior graph does not contain data for that account.

- **Is the finding archived?** Detective does not receive archived findings from GuardDuty.
- **Did the finding occur before Detective began to ingest data into your behavior graph?** If the finding is not present in the data that Detective ingests, then the behavior graph does not contain data for it.
- **Is the finding from the correct Region?** Each behavior graph is specific to a Region. A behavior graph does not contain data from other Regions.

Navigating to an entity profile or finding overview using a URL

To navigate to an entity profile or finding overview in Amazon Detective, you can use a URL that provides a direct link to it. The URL identifies the finding or entity. It can also specify the scope time to use on the profile. Detective maintains up to a year of historical event data.

Format of a profile URL

Note

If you are using the old URL format, Detective will automatically redirect to the new URL. The old format of the URL was:
`https://console.aws.amazon.com/detective/home?region=Region#type/namespace/instanceID?parameters`

The new format of the profile URL is as follows:

- For entities - `https://console.aws.amazon.com/detective/home?region=Region#entities/namespace/instanceID?parameters`
- For findings - `https://console.aws.amazon.com/detective/home?region=Region#findings/instanceID?parameters`

The URL requires the following values.

Region

The Region that you want to use.

type

The type of item for the profile that you are navigating to.

- `entities` - Indicates that you are navigating to an entity profile
- `findings` - Indicates that you are navigating to a finding overview

namespace

For entities, the namespace is the name of the entity type.

- `AwsAccount`

- `AwsRole`
- `AwsRoleSession`
- `AwsUser`
- `Ec2Instance`
- `FederatedUser`
- `IpAddress`
- `S3Bucket`
- `UserAgent`
- `FindingGroup`
- `KubernetesSubject`
- `ContainerPod`
- `ContainerCluster`
- `ContainerImage`

instanceID

The instance identifier of the finding or entity.

- For a GuardDuty finding, the GuardDuty finding identifier.
- For an AWS account, the account ID.
- For AWS roles and users, the principal ID of the role or of the user.
- For federated users, the principal ID of the federated user. The principal ID is either `<identityProvider>:<username>` or `<identityProvider>:<audience>:<username>`.
- For IP addresses, the IP address.
- For user agents, the user agent name.
- For EC2 instances, the instance ID.
- For role sessions, the session identifier. The session identifier uses the format `<rolePrincipalID>:<sessionName>`.
- For S3 buckets, the bucket name.
- For FindingGroups, a UUID. for example, `ca6104bc-a315-4b15-bf88-1c1e60998f83`
- For EKS resources, use the following formats:
 - EKS cluster: `<clusterName>~<accountId>~EKS`

- Kubernetes Pod: `<podUid>~<clusterName>~<accountId>~EKS`
- Kubernetes Subject: `<subjectName>~<clusterName>~<accountId>`
- Container image: `<registry>/<repository>:<tag>@<digest>`

The finding or entity must be associated with an enabled account in your behavior graph.

The URL can also include the following optional parameters, which are used to set the scope time. For more information about scope time and how it is used on profiles, see [the section called "Managing the scope time"](#).

scopeStart

Start time for the scope time to use on the profile. Start time must be within the last 365 days.

The value is the epoch timestamp.

If you provide a start time but no end time, then the scope time ends at the current time.

scopeEnd

End time for the scope time to use on the profile.

The value is the epoch timestamp.

If you provide an end time, but no start time, then the scope time includes all time before the end time.

If you don't specify the scope time, then the default scope time is used.

- For findings, the default scope time uses the first and last times that the finding activity was observed.
- For entities, the default scope time is the previous 24 hours.

Here is an example of a Detective URL:

```
https://console.aws.amazon.com/detective/home?region=us-east-1#entities/IpAddress/192.168.1.1?scopeStart=1552867200&scopeEnd=1552910400
```

This example URL provides the following instructions.

- Display the entity profile for the IP address 192.168.1.
- Use a scope time that starts Monday, March 18, 2019 12:00:00 AM GMT and that ends Monday, March 18, 2019 12:00:00 PM GMT.

Troubleshooting a URL

If the URL does not display the expected profile, first check that the URL uses the correct format and that you have provided the correct values.

- Did you start with the correct URL (findings or entities)?
- Did you specify the correct namespace?
- Did you provide the correct identifier?

If the values are correct, then you can also check the following.

- **Does the finding or entity belong to an enabled member account in your behavior graph?** If the associated account was not invited to the behavior graph as a member account, then the behavior graph does not contain data for that account.

If an invited member account did not accept the invitation, then the behavior graph does not contain data for that account.

- **For a finding, is the finding archived?** Detective does not receive archived findings from Amazon GuardDuty.
- **Did the finding or entity occur before Detective began to ingest data into your behavior graph?** If the finding or entity is not present in the data that Detective ingests, then the behavior graph does not contain data for it.
- **Is the finding or entity from the correct Region?** Each behavior graph is specific to a Region. A behavior graph does not contain data from other Regions.

Adding Detective URLs for findings to Splunk

The Splunk Trumpet project allows you send data from AWS services to Splunk.

You can configure the Trumpet project to generate Detective URLs for Amazon GuardDuty findings. You can then use these URLs to pivot directly from Splunk to the corresponding Detective finding profiles.

The Trumpet project is available from GitHub at <https://github.com/splunk/splunk-aws-project-trumpet>.

On the configuration page for the Trumpet project, from **AWS CloudWatch Events**, choose **Detective GuardDuty URLs**.

Pivoting from a profile panel to another console

For EC2 instances, IAM users, and IAM roles, you can navigate directly from the details profile panel to the corresponding console. The information available from the console can provide additional input for your investigation.

On the **EC2 instance details** profile panel, the EC2 instance identifier is linked to the Amazon EC2 console.

On the **User details** profile panel, the user name is linked to the IAM console.

On the **Role details** profile panel, the role name is linked to the IAM console.

Pivoting from a profile panel to another entity profile

When a profile panel contains an identifier of a different entity, it is usually a link to that entity profile. The exceptions are the links to the Amazon EC2 and IAM consoles on the EC2 instance, IAM users, and IAM roles profiles. See [the section called "Pivoting to another console"](#).

For example, from a list of IP addresses, you might be able to display the profile for a specific IP address. That way you can see if there is any other information available to help you to complete your investigation.

Exploring activity details on a profile panel

During an investigation, you might want to investigate further into the pattern of activity for an entity.

On the following profile panels, you can display a summary of the activity details:

- **Overall API call volume**, except for the profile panel on the user agent profile
- **Newly observed geolocations**

- **Overall VPC flow volume**
- **VPC flow volume to and from the finding IP address**, for findings that are associated with a single IP address
- **Container details**
- **VPC flow volume** for clusters
- **Overall Kubernetes API activity**

The activity details can answer these types of questions:

- Which IP addresses were used?
- Where were those IP addresses located?
- Which API calls did each IP address make, and from which services did they make those calls?
- Which principals or access key identifiers (AKIDs) were used to make the calls?
- What resources were used to make those calls?
- How many calls were made? How many succeeded and failed?
- What volume of VPC flow log data was sent to or from each IP address?
- What containers were active for a given cluster, image, or pod?

Topics

- [Activity details for Overall API call volume](#)
- [Activity details for a geolocation](#)
- [Activity details for overall VPC flow volume](#)
- [Overall Kubernetes API activity involving EKS cluster](#)

Activity details for Overall API call volume

The activity details for **Overall API call volume** show the API calls that were issued during a selected time range.

To display the activity details for a single time interval, choose the time interval on the chart.

To display the activity details for the current scope time, choose **Display details for scope time**.

Note that Detective began to store and display the service name for API calls as of July 14, 2021. That date is highlighted on the profile panel timeline. For activity that occurs before that date, the service name is **Unknown service**.

Content of the activity details (users, roles, accounts, role sessions, EC2 instances, S3 buckets)

For IAM users, IAM roles, accounts, role sessions, EC2 instances, and S3 buckets, the activity details contain the following information:

- Each tab provides information about the set of API calls that were issued during the selected time range.

For S3 buckets, the information reflects API calls that were made to the S3 bucket.

The API calls are grouped by the services that called them. For S3 buckets, the service is always Amazon S3. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

- For each entry, the activity details show the number of successful and failed calls. The **Observed IP addresses** tab also shows the location of each IP address.
- Each entry shows information about who made the calls. For accounts, the activity details identify the users or roles. For roles, the activity details identify the role sessions. For users and role sessions, the activity details identify the access key identifiers (AKIDs).

Note that as of July 14, 2021, for account profiles, the activity details show users or roles instead of AKIDs. For role profiles, the activity details show role sessions instead of AKIDs. For activity that occurs before July 14, 2021, the caller is listed as **Unknown resource**.

The activity details contain the following tabs:

Observed IP addresses

Initially displays the list of IP addresses used to issue API calls.

You can expand each IP address to display the list of API calls that were issued from that IP address. The API calls are grouped by the services that called them. For S3 buckets, the service is always Amazon S3. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

You can then expand each API call to display the list of callers from that IP address. Depending on the profile, the caller might be a user, role, role session, or AKID.

Showing activity: 07/08/2021, 00:00 UTC - 07/08/2021, 04:00 UTC Edit

Observed IP addresses | API method by service | Resource

Filter by IP CIDR, Service name, API Method name, or Resource string

IP address	Successful calls	Failed calls	Location
...	421	311	-
s3	316	311	
config	61	0	
kms	15	0	
DescribeKey	14	0	
Role session (.....)	14	0	
ListKeys	1	0	
rds	7	0	
ec2	4	0	
autoscaling	3	0	
secretsmanager	2	0	
guardduty	2	0	
es	2	0	

API method by service

Initially displays the list of API calls that were issued. The API calls are grouped by the services that issued the calls. For S3 buckets, the service is always Amazon S3. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

You can expand each API method to display the list of IP addresses from which the calls were issued.

You can then expand each IP address to display the list of AKIDs that issued that API call from that IP address.

Showing activity: 07/08/2021, 00:00 UTC - 07/08/2021, 04:00 UTC Edit

Observed IP addresses | API method by service | Resource

Filter by IP CIDR, Service name, API Method name, or Resource string

API method	Successful calls	Failed calls
s3	316	311
config	61	0
kms	15	0
DescribeKey	14	0
Role session (.....)	14	0
ListKeys	1	0
rds	7	0
ec2	4	0
autoscaling	3	0

Resource or Access Key ID

Initially displays the list of users, roles, role sessions, or AKIDs that were used to issue API calls.

You can expand each caller to display the list of IP addresses from which the caller issued API calls.

You can then expand each IP address to display the list of API calls that were issued from that IP address by that caller. The API calls are grouped by the services that issued the calls. For S3 buckets, the service is always Amazon S3. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

Showing activity: 07/08/2021, 00:00 UTC - 07/08/2021, 04:00 UTC

Observed IP addresses | API method by service | **Resource**

Filter by IP CIDR, Service name, API Method name, or Resource string

Resource	Successful calls	Failed calls
Role session	322	310
Role session	91	0
Role session	91	0
config	61	0
kms	15	0
DescribeKey	14	0
ListKeys	1	0
ec2	3	0
secretsmanager	2	0
guardduty	2	0
...

Content of the activity details (IP addresses)

For IP addresses, the activity details contain the following information:

- Each tab provides information about the set of API calls that were issued during the selected time range. The API calls are grouped by the services that issued the calls. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.
- For each entry, the activity details show the number of successful and failed calls.

The activity details contain the following tabs:

Resource

Initially displays the list of resources that issued API calls from the IP address.

For each resource, the list includes the resource name, the type, and the AWS account.

You can expand each resource to display the list of API calls that the resource issued from the IP address. The API calls are grouped by the services that issued the calls. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

Showing activity: 07/08/2021, 00:00 UTC - 07/08/2021, 04:00 UTC Edit

Resource | API method by service

Filter by Resource string, Service name or API Method name

Resource	Successful calls	Failed calls	Account ID
<ul style="list-style-type: none"> <ul style="list-style-type: none"> DescribeComplianceByConfigRule PutEvaluations SelectResourceConfig DescribeDeliveryChannelStatus DescribeConfigurationRecorderSta... DescribeConfigurationRecorders ec2 shield waf-regional 	3,520	0	
<ul style="list-style-type: none"> config 	1,754	0	
<ul style="list-style-type: none"> <ul style="list-style-type: none"> DescribeComplianceByConfigRule PutEvaluations SelectResourceConfig DescribeDeliveryChannelStatus DescribeConfigurationRecorderSta... DescribeConfigurationRecorders 	1,408	0	
<ul style="list-style-type: none"> PutEvaluations 	244	0	
<ul style="list-style-type: none"> SelectResourceConfig 	78	0	
<ul style="list-style-type: none"> DescribeDeliveryChannelStatus 	8	0	
<ul style="list-style-type: none"> DescribeConfigurationRecorderSta... 	8	0	
<ul style="list-style-type: none"> DescribeConfigurationRecorders 	8	0	
<ul style="list-style-type: none"> ec2 	1,690	0	
<ul style="list-style-type: none"> shield 	50	0	
<ul style="list-style-type: none"> waf-regional 	26	0	
<ul style="list-style-type: none"> <ul style="list-style-type: none"> DescribeComplianceByConfigRule PutEvaluations SelectResourceConfig DescribeDeliveryChannelStatus DescribeConfigurationRecorderSta... DescribeConfigurationRecorders 	1,715	0	
<ul style="list-style-type: none"> <ul style="list-style-type: none"> DescribeComplianceByConfigRule PutEvaluations SelectResourceConfig DescribeDeliveryChannelStatus DescribeConfigurationRecorderSta... DescribeConfigurationRecorders 	504	480	

API method by service

Initially displays the list of API calls that were issued. The API calls are grouped by the services that issued the calls. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

You can expand each API call to display the list of resources that issued the API call from the IP address during the selected time period.

Showing activity: 07/08/2021, 00:00 UTC - 07/08/2021, 04:00 UTC Edit

Resource | API method by service

Filter by Resource string, Service name or API Method name

API method	Successful calls	Failed calls
<ul style="list-style-type: none"> config 	3,787	0
<ul style="list-style-type: none"> ec2 	2,538	0
<ul style="list-style-type: none"> s3 	1,269	1,016
<ul style="list-style-type: none"> <ul style="list-style-type: none"> ListCommands <ul style="list-style-type: none"> AWS role AWS role SendCommand 	481	16
<ul style="list-style-type: none"> ListCommands 	392	0
<ul style="list-style-type: none"> <ul style="list-style-type: none"> AWS role 	222	0
<ul style="list-style-type: none"> <ul style="list-style-type: none"> AWS role 	170	0
<ul style="list-style-type: none"> SendCommand 	89	16
<ul style="list-style-type: none"> logs 	165	0
<ul style="list-style-type: none"> sts 	149	0
<ul style="list-style-type: none"> iam 	149	12

Sorting the activity details

You can sort the activity details by any of the list columns.

When you sort using the first column, only the top-level list is sorted. The lower-level lists are always sorted by the count of successful API calls.

Filtering the activity details

You can use the filtering options to focus on specific subsets or aspects of the activity represented in the activity details.

On all of the tabs, you can filter the list by any of the values in the first column.

To add a filter

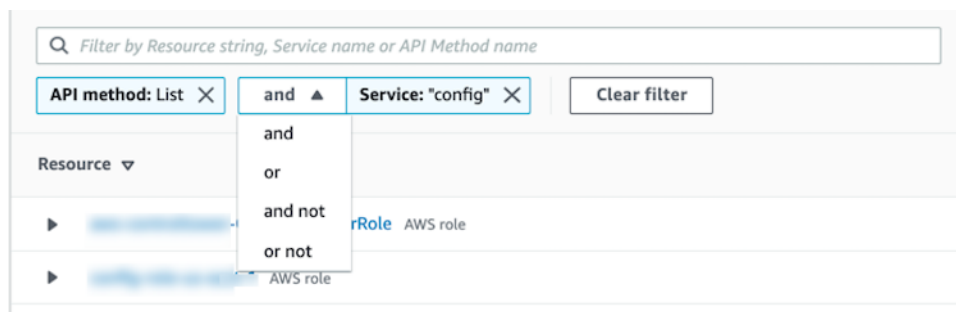
1. Choose the filter box.
2. From **Properties**, choose the property to use for the filtering.
3. Provide the value to use for the filtering. The filter supports partial values. For example, when you filter by API method, if you filter by **Instance**, the results include any API operation that has Instance in its name. So both ListInstanceAssociations and UpdateInstanceInformation would match.

For service names, API methods, and IP addresses, you can either specify a value or choose a built-in filter.

For **Common API substrings**, choose the substring that represents the type of operation, such as List, Create, or Delete. Each API method name starts with the operation type.

For **CIDR patterns**, you can choose to include only public IP addresses, private IP addresses, or IP addresses that match a specific CIDR pattern.

4. If you have multiple filters, choose a Boolean option to set how those filters are connected.



5. To remove a filter, choose the x icon in the top-right corner.
6. To clear all of the filters, choose **Clear filter**.

Selecting the time range for the activity details

When you first display the activity details, the time range is either the scope time or a selected time interval. You can change the time range for the activity details.


To change the time range for the activity details

1. Choose **Edit**.
2. On **Edit time window**, choose the start and end time to use.

To set the time window to the default scope time for the profile, choose **Set to default scope time**.

3. Choose **Update time window**.

The time range for the activity details is highlighted on the profile panel charts.



Querying raw logs

Amazon Detective integrates with Amazon Security Lake, which means that you can query and retrieve the raw log data stored by Security Lake. For more details about this integration, see [Integration with Amazon Security Lake](#).

Using this integration, you can collect and query logs and events from the following sources which Security Lake natively supports.

- AWS CloudTrail management events version 1.0 and after
- Amazon Virtual Private Cloud (Amazon VPC) Flow Logs version 1.0 and after
- Amazon Elastic Kubernetes Service (Amazon EKS) Audit Log version 2.0

Note

There are no additional charges to query raw data logs in Detective. Usage charges for other AWS Services, including Amazon Athena, still apply at published rates.

To query raw logs

1. Choose **display details for scope time**.
2. From here, you can start to **Query raw logs**.
3. In the **Raw log preview** table, you can view the logs and events retrieved by querying data from Security Lake. For more details about the raw event logs, you can view the data displayed in Amazon Athena.

From the Query raw logs table, you can **Cancel query request**, **See results in Amazon Athena**, and **Download results** as a comma-separated values (.csv) file.

If you see logs in Detective, but the query returned no results, it could happen because of the following reasons.

- Raw logs may become available in Detective before showing up in Security Lake log tables. Try again later.
- Logs may be missing from Security Lake. If you waited for an extended period of time, it indicates that logs are missing from Security Lake. Contact your Security Lake administrator to resolve the issue.

Activity details for a geolocation

The activity details for **Newly observed geolocations** show the API calls that were issued from a geolocation during the scope time. The API calls include all calls issued from the geolocation. They are not limited to calls that used the finding or profile entity. For S3 buckets, the activity calls are API calls made to the S3 bucket.

Detective determines the location of requests using MaxMind GeoIP databases. MaxMind reports very high accuracy of their data at the country level, although accuracy varies according to factors such as country and type of IP. For more information about MaxMind, see [MaxMind IP Geolocation](#).

If you think any of the GeolP data is incorrect, you can submit a correction request to Maxmind at [MaxMind Correct GeolP2 Data](#).

The API calls are grouped by the services that issued the calls. For S3 buckets, the service is always Amazon S3. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

To display the activity details, do one of the following:

- On the map, choose a geolocation.
- In the list, choose **Details** for a geolocation.

The activity details replace the geolocation list. To return to the geolocation list, choose **Return to all results**.

Note that Detective began to store and display the service name for API calls as of July 14, 2021. For activity that occurs before that date, the service name is **Unknown service**.

If Detective cannot determine the credentials that issued a call, the call is listed as an **Unknown resource**. You can expand the unknown resource to see the IP addresses and API calls that were made by it.

Content of the activity details

Each tab provides information about all of the API calls that were issued from the geolocation during the scope time.

For each IP address, resource, and API method, the list shows the number of successful and failed API calls.

The activity details contain the following tabs:

Observed IP addresses

Initially displays the list of IP addresses that were used to issue API calls from the selected geolocation.

You can expand each IP address to display the resources that issued API calls from that IP address. The list displays the resource name. To see the principal ID, hover over the name.

You can then expand each resource to display the specific API calls that were issued from that IP address by that resource. The API calls are grouped by the services that issued the calls. For S3

buckets, the service is always Amazon S3. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

Ashburn, US from 05/14/2021 - 06/28/2021

Observed IP addresses | Resource

Filter by IP CIDR, API Method name, or Resource string

IP address	Successful calls	Failed calls
10.0.0.0/24	27,564	2,453
10.0.0.0/24 AWS role (1000000000000)	27,564	2,453
ssm	25,111	0
UpdateInstanceInformation	13,066	0
ListInstanceAssociations	6,482	0
PutInventory	2,544	0
GetDeployablePatchSnapshotForIns...	2,453	0
UpdateInstanceAssociationStatus	466	0
PutComplianceItems	98	0
GetDocument	2	0
sts	2,453	0
s3	0	2,453
10.0.0.0/24	24,635	1,512
10.0.0.0/24	24,632	1,511

Resource

Initially displays the list of resources that issued API calls from the selected geolocation. The list displays the resource name. To see the principal ID, pause on the name. For each resource, the **Resource** tab also displays the associated AWS account.

You can expand each user or role to display the list of API calls that were issued by that resource. The API calls are grouped by the services that issued the calls. For S3 buckets, the service is always Amazon S3. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

You can then expand each API call to display the list of IP addresses from which the resource issued the API call.

Ashburn, US from 05/14/2021 - 06/28/2021

Observed IP addresses | Resource

Filter by IP CIDR, API Method name, or Resource string

Resource	Successful calls	Failed calls	Account ID
10.0.0.0/24 AWS role (1000000000000)	189,097	17	1000000000000
10.0.0.0/24 AWS role (1000000000000)	49,267	3,023	1000000000000
ssm	46,254	0	
UpdateInstanceInformation	25,932	0	
10.0.0.0/24	12,968	0	
10.0.0.0/24	12,964	0	
ListInstanceAssociations	12,964	0	
PutInventory	3,194	0	
GetDeployablePatchSnapshotForIns...	3,011	0	
UpdateInstanceAssociationStatus	949	0	
PutComplianceItems	199	0	
GetDocument	5	0	
sts	3,013	0	
s3	0	3,023	

Sorting the activity details

You can sort the activity details by any of the list columns.

When you sort using the first column, only the top-level list is sorted. The lower-level lists are always sorted by the count of successful API calls.

Filtering the activity details

You can use the filtering options to focus on specific subsets or aspects of the activity represented in the activity details.

On all of the tabs, you can filter the list by any of the values in the first column.

To add a filter


1. Choose the filter box.
2. From **Properties**, choose the property to use for the filtering.
3. Provide the value to use for the filtering. The filter supports partial values. For example, when you filter by API method, if you filter by **Instance**, the results include any API operation that has Instance in its name. So both `ListInstanceAssociations` and `UpdateInstanceInformation` would match.

For service names, API methods, and IP addresses, you can either specify a value or choose a built-in filter.

For **Common API substrings**, choose the substring that represents the type of operation, such as `List`, `Create`, or `Delete`. Each API method name starts with the operation type.

For **CIDR patterns**, you can choose to include only public IP addresses, private IP addresses, or IP addresses that match a specific CIDR pattern.

4. If you have multiple filters, choose a Boolean option to set how those filters are connected.



5. To remove a filter, choose the **x** icon in the top-right corner.
6. To clear all of the filters, choose **Clear filter**.

Activity details for overall VPC flow volume

For an EC2 instance, the activity details for **Overall VPC flow volume** show the interactions between the EC2 instance and IP addresses during a selected time range.

For a Kubernetes pod, **Overall VPC flow volume** displays the overall volume of bytes into and out of the Kubernetes pod's assigned IP address for all destination IP addresses. The Kubernetes pod's IP address is not unique when `hostNetwork: true`. In this case, the panel shows traffic to other pods with the same configuration and the node hosting them.

For an IP address, the activity details for **Overall VPC flow volume** show the interactions between the IP address and EC2 instances during a selected time range.

To display the activity details for a single time interval, choose the time interval on the chart.

To display the activity details for the current scope time, choose **display details for scope time**.

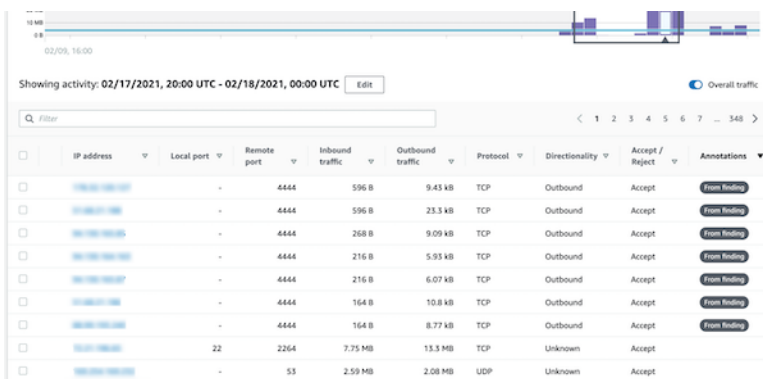
Content of the activity details

The content reflects the activity during the selected time range.

For an EC2 instance, the activity details contain an entry for each unique combination of IP address, local port, remote port, protocol, and direction.

For an IP address, the activity details contain an entry for each unique combination of EC2 instance, local port, remote port, protocol, and direction.

Each entry displays the volume of inbound traffic, the volume of outbound traffic, and whether the access request was accepted or rejected. On finding profiles, the **Annotations** column indicates when an IP address is related to the current finding.



Showing activity: 02/17/2021, 20:00 UTC - 02/18/2021, 00:00 UTC Edit Overall traffic

<input type="checkbox"/>	IP address	Local port	Remote port	Inbound traffic	Outbound traffic	Protocol	Directionality	Accept / Reject	Annotations
<input type="checkbox"/>	10.0.0.1	-	4444	596 B	9.43 kB	TCP	Outbound	Accept	From Finding
<input type="checkbox"/>	10.0.0.1	-	4444	596 B	23.5 kB	TCP	Outbound	Accept	From Finding
<input type="checkbox"/>	10.0.0.1	-	4444	268 B	9.09 kB	TCP	Outbound	Accept	From Finding
<input type="checkbox"/>	10.0.0.1	-	4444	216 B	5.93 kB	TCP	Outbound	Accept	From Finding
<input type="checkbox"/>	10.0.0.1	-	4444	216 B	6.07 kB	TCP	Outbound	Accept	From Finding
<input type="checkbox"/>	10.0.0.1	-	4444	164 B	10.8 kB	TCP	Outbound	Accept	From Finding
<input type="checkbox"/>	10.0.0.1	-	4444	164 B	8.77 kB	TCP	Outbound	Accept	From Finding
<input type="checkbox"/>	10.0.0.1	22	2264	7.75 MB	13.5 MB	TCP	Unknown	Accept	
<input type="checkbox"/>	10.0.0.1	-	53	2.59 MB	2.08 MB	UDP	Unknown	Accept	

Sorting the activity details

You can sort the activity details by any of the columns in the table.

By default, the activity details are sorted first by the annotations, then by the inbound traffic.

Filtering the activity details

To focus on specific activity, you can filter the activity details by the following values:

- IP address or EC2 instance
- Local or remote port
- Direction
- Protocol
- Whether the request was accepted or rejected


To add and remove filters

1. Choose the filter box.
2. From **Properties**, choose the property to use for the filtering.
3. Provide the value to use for the filtering. The filter supports partial values.

To filter by IP address, you can either specify a value or choose a built-in filter.

For **CIDR patterns**, you can choose to include only public IP addresses, private IP addresses, or IP addresses that match a specific CIDR pattern.

4. If you have multiple filters, choose a Boolean option to set how those filters are connected.



5. To remove a filter, choose the x icon in the top-right corner.
6. To clear all of the filters, choose **Clear filter**.

Selecting the time range for the activity details

When you first display the activity details, the time range is either the scope time or a selected time interval. You can change the time range for the activity details.


To change the time range for the activity details

1. Choose **Edit**.
2. On **Edit time window**, choose the start and end time to use.

To set the time window to the default scope time for the profile, choose **Set to default scope time**.

3. Choose **Update time window**.


The time range for the activity details is highlighted on the profile panel charts.




Displaying the volume of traffic for selected rows

When you identify rows that are of interest, you can display on the main charts the volume of traffic over time for those rows.

For each row to add to the charts, select the check box. For each selected row, the volume is displayed as a line on the inbound or outbound charts.



To focus on the traffic volume for the selected entries, you can hide the overall volume. To show or hide the overall traffic volume, toggle **Overall traffic**.



Displaying the VPC flow traffic for EKS clusters

Detective has visibility into your Amazon Virtual Private Cloud (Amazon VPC) flow logs, which represent the traffic that traverses your Amazon Elastic Kubernetes Service (Amazon EKS) clusters. For Kubernetes resources, the content of the VPC flow logs depends on the Container Network Interface (CNI) deployed in the EKS cluster.

An EKS cluster with a default configuration uses the Amazon VPC CNI plugin. For more details, see [Managing VPC CNI](#) in the Amazon EKS User Guide. The Amazon VPC CNI plugin sends internal traffic with the IP address of the pod and translates the source IP address to the IP address of the node for external communication. Detective can capture and correlate internal traffic to the correct pod but it can't do the same for external traffic.

If you want Detective to have visibility into the external traffic of your pods, enable External Source Network Address Translation (SNAT). Enabling SNAT comes with limitations and drawbacks. For more details, see [SNAT for pods](#) in the Amazon EKS User Guide.

If you use a different CNI plugin, Detective has limited visibility to pods with `hostNetwork: true`. For these pods, the **VPC Flow** panel displays all traffic to the IP address of the pod. This includes the traffic to the host node and any pod on the node with the `hostNetwork: true` configuration.

Detective displays traffic in the **VPC flow** panel of an EKS pod for the following EKS cluster configurations:

- In a cluster with the Amazon VPC CNI plugin, any pod with the configuration `hostNetwork: false` sending traffic inside the VPC of the cluster.

- In a cluster with the Amazon VPC CNI plugin and the configuration `AWS_VPC_K8S_CNI_EXTERNALSNAT=true`, any pod with `hostNetwork:false` sending traffic outside the VPC of the cluster.
- Any pod with the configuration `hostNetwork:true`. Traffic from the node is mixed with traffic from other pods that have the configuration `hostNetwork:true`.

Detective does not display traffic in the **VPC flow** panel for:

- In a cluster with the Amazon VPC CNI plugin and the configuration `AWS_VPC_K8S_CNI_EXTERNALSNAT=false`, any pod with the configuration `hostNetwork:false` sending traffic outside the VPC of the cluster.
- In a cluster without the Amazon VPC CNI plugin for Kubernetes, any pod with the configuration `hostNetwork:false`.
- Any pod sending traffic to another pod that is hosted in the same node.

Displaying the VPC flow traffic for shared Amazon VPCs

Detective has visibility into your Amazon Virtual Private Cloud (Amazon VPC) flow logs for shared VPCs:

- If a Detective member account has a shared Amazon VPC and there are other non-Detective accounts using the shared VPC, Detective monitors all traffic from that VPC, and provides visualization on all the traffic flow within the VPC.
- If you have an Amazon EC2 instance inside a shared Amazon VPC and the shared VPC owner is not a Detective member, Detective will not monitor any traffic from the VPC. If you want to view the traffic flow within the VPC, you must add the Amazon VPC owner as a member of your Detective graph.

Overall Kubernetes API activity involving EKS cluster

The activity details for **Overall Kubernetes API activity involving EKS cluster** show the number of successful and failed Kubernetes API calls that were issued during a selected time range.

To display the activity details for a single time interval, choose the time interval on the chart.

To display the activity details for the current scope time, choose **Display details for scope time**.

Content of the activity details (Cluster, pod, user, role, role session)

For a cluster, pod, user, role, or role session, the activity details contain the following information:

- Each tab provides information about the set of API calls that were issued during the selected time range.

For clusters, the API calls occurred inside the cluster.

For pods, the API calls targeted the pod.

For users, roles, and role sessions, the API calls were issued by Kubernetes users that authenticated as that user, role, or role session.

- For each entry, the activity details show the number of successful, failed, unauthorized, and forbidden calls.
- The information includes the IP address, the type of Kubernetes call, the entity that was affected by the call, and the subject (service account or user) that made the call. From the activity details, you can pivot to the profiles for the IP address, subject, and the affected entity.

The activity details contain the following tabs:

Subject

Initially displays the list of service accounts and users that were used to make API calls.

You can expand each service account and user to display the list of IP addresses from which the account or user made API calls.

You can then expand each IP address to show the Kubernetes API calls that were made by that account or user from that IP address.

Expand the Kubernetes API call to see the `requestURI` to identify the action that was done.

Showing activity: 05/09/2022, 23:00 UTC - 05/10/2022, 23:00 UTC [Edit](#)

Subject | IP address | Kubernetes API call

Filter by Kubernetes subject, IP CIDR, API verb, or API method name

Subject	Success	Failure	Unauthorized	Forbidden
<ul style="list-style-type: none"> cloud-controller-manager Kubernetes user <ul style="list-style-type: none"> 10.0.100.200 IP address <ul style="list-style-type: none"> update 80,343 get 80,343 watch 720 10.0.100.55 IP address 25,245 	186,651	1	0	0

IP Address

Initially displays the list of IP addresses from which the API calls were made.

You can expand each call to display the list of Kubernetes subjects (service accounts and users) that made the call.

You can then expand each subject to a list of API call types made by the subject during the scope time.

Expand the API call type to see the requestURI to identify the action that was done.

Showing activity: 05/09/2022, 23:00 UTC - 05/10/2022, 23:00 UTC [Edit](#)

Subject | **IP address** | Kubernetes API call

Filter by Kubernetes subject, IP CIDR, API verb, or API method name

IP address	Success	Failure	Unauthorized	Forbidden	Location
<ul style="list-style-type: none"> 10.0.100.200 IP address <ul style="list-style-type: none"> cloud-controller-manager Kubernetes user <ul style="list-style-type: none"> update <ul style="list-style-type: none"> /apis/coordination.k8s.io/v1/namespaces/kube-system/leases/cloud-provider-extraction-migration 40,172 /apis/coordination.k8s.io/v1/namespaces/kube-system/leases/cloud-controller-manager 40,171 	599,250	2,706	0	0	-

Kubernetes API call

Initially displays the list of Kubernetes API call verbs.

You can expand each API verb to display the requestURIs associated with that action.

You can then expand each requestURI to see Kubernetes subject (service accounts and users) that made the API call.

Expand the subject to see which IPs that subject used to make the API call.

Resource	Successful calls	Failed calls
Role session	322	310
Role session	91	0
Role session	91	0
config	61	0
kms	15	0
DescribeKey	14	0
ListKeys	1	0
ec2	3	0
secretsmanager	2	0
guardduty	2	0
...

Sorting the activity details

You can sort the activity details by any of the list columns.

When you sort using the first column, only the top-level list is sorted. The lower-level lists are always sorted by the count of successful API calls.

Filtering the activity details

You can use the filtering options to focus on specific subsets or aspects of the activity represented in the activity details.

On all of the tabs, you can filter the list by any of the values in the first column.

Selecting the time range for the activity details

When you first display the activity details, the time range is either the scope time or a selected time interval. You can change the time range for the activity details.


To change the time range for the activity details

1. Choose **Edit**.
2. On **Edit time window**, choose the start and end time to use.

To set the time window to the default scope time for the profile, choose **Set to default scope time**.

3. Choose **Update time window**.

The time range for the activity details is highlighted on the profile panel charts.



Using profile panel guidance during an investigation

Each profile panel is designed to provide answers to specific questions that arise as you conduct an investigation and analyze the activity for the related entities.

The guidance provided for each profile panel helps you find these answers.

Profile panel guidance starts with a single sentence on the panel itself. This guidance provides a brief explanation of the data presented on the panel.

To display more detailed guidance for a panel, choose **More info** from the panel heading. This extended guidance appears in the help pane.

The guidance can provide these types of information:

- An overview of the panel content
- How to use the panel to answer the relevant questions
- Suggested next steps based on the answers

Managing the scope time

Customize the scope time used to limit the data displayed on entity profiles.

The charts, timelines, and other data displayed on entity profiles are all based on the current scope time. Scope time is the summary of activity for an entity over time. This appears at the top right of each profile in the Amazon Detective console. The data displayed on those charts, timelines, and other visualizations is based on the scope time. For some profile panels, additional time is added before and after the scope time to provide context. In Detective, all timestamps are displayed in UTC by default. You can select your local time zone by changing the **Timestamp preferences**. To update the **Timestamp preference**, see [the section called "Setting the timestamp format"](#).

Detective analytics uses the scope time when checking for unusual activity. The analytics process gets the activity during the scope time, then compares it to the activity during the 45 days before the scope time. It also uses that 45-day timeframe to generate baselines of activity.

On a finding overview, the scope time reflects the first and last time the finding was observed. For more information about finding overview, see [the section called "Finding overview"](#).

As you work through an investigation, you can adjust the scope time. For example, if the original analysis was based on activity from a single day, you might want to expand that to a week or a month. The expanded period could help you get a better sense of whether the activity fits a normal pattern or is unusual.

You can also set the scope time to match an associated finding for the current entity.

When you change the scope time, Detective repeats its analysis and updates the displayed data based on the new scope time.

The scope time cannot be shorter than one hour and not longer than one year. The start and end time must be on an hour.

Setting specific start and end dates and times

You can set the scope time start and end dates from the Detective console.

To set specific start and end times for the new scope time

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.

2. On an entity profile, choose the scope time.
3. On the **Edit scope time** panel, under **Start**, choose the new start date and time for the scope time. For the new start time, you choose the hour only.
4. Under **End**, choose the new end date and time for the scope time. For the new end time, you choose the hour only. The end time must be at least an hour later than the start time.
5. When you're finished editing, to save the changes and update the displayed data, choose **Update scope time**.

Edit the length of time for the scope time

When you set a scope time length, Detective sets the scope time to that amount of time from the current time.

To edit the length of time for the scope time

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. On an entity profile, choose the scope time.
3. On the **Edit scope time** panel, next to **Historical**, choose the length of time for the scope time.

Specifying a time range updates the **Start** and **End** settings.

4. When you're finished editing, to save the changes and update the displayed data, choose **Update scope time**.

Setting the scope time to a finding time window

Each finding has an associated time window, which reflects the first and last times the finding was observed. When you view a finding overview, the scope time changes to the finding time window.

From an entity profile, you can align the scope time to the time window for an associated finding. This allows you to investigate the activity that occurred during that time.

To align the scope time to a finding time window, on the **Associated findings** panel, choose the finding that you want to use.

Detective populates the finding details and sets the scope time to the finding time window.

Setting the scope time on the summary page

As you review the **Summary** page, you can adjust the Scope time to view the activity for any 24-hour time frame in the previous 365 days.

To set the scope time on the Summary page

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Summary**.
3. On the **Scope time** panel, next to **Summary**, you can change the **Start date and time**. Start time must be within the last 365 days.

When you change the **Start date and time**, the **End date and time** is automatically updated to 24 hours after your chosen start time.

Note

With Detective, you can access up to a year of historical event data. For more information on source data in Detective, see [Source data used in a behavior graph](#).

4. When you're finished editing, to save the changes and update the displayed data, choose **Update scope time**.

Viewing details for associated findings

Each entity profile contains an associated findings panel that lists the findings that involved the entity during the current scope time. One indication that an entity has been compromised is its involvement in multiple findings. The types of findings can also provide insight into the type of activity to be concerned about.

The associated findings panel is displayed immediately below the entity details profile panel.

For each finding, the table includes the following information:

- The finding title, which is also a link to the finding overview.
- The AWS account associated with the finding, which is also a link to the account profile
- The finding type

- The earliest time that the finding was observed
- The most recent time that the finding was observed
- The finding severity

To display the finding details for a finding, choose the radio button for the finding. Detective populates the finding details panel at the right of the page. Detective also changes the scope time to be the finding time window. This allows you to focus on activity that occurred during that time.

If you navigated to the entity profile from a finding overview, then that finding is selected automatically and the details for the finding are displayed.

From the finding details, to navigate back to the finding overview, choose **See all related entities**.

You can also archive the finding. For more details, see [Archiving an Amazon GuardDuty finding](#).

Viewing details for high-volume entities

In the [behavior graph](#), Amazon Detective tracks relationships between entities. For example, each behavior graph tracks when an AWS user creates an AWS role and when an EC2 instance connects to an IP address.

When an entity has too many relationships during a time period, Detective cannot store all of the relationships. When this occurs during the current scope time, Detective notifies you. Detective also provides a list of occurrences of high-volume entities.

What is a high-volume entity?

During a given time interval, an entity might be the origin or destination of an extremely large number of connections. For example, an EC2 instance may have connections from millions of IP addresses.

Detective maintains a limit on the number of connections that it can accommodate during each time interval. If an entity exceeds that limit, then Detective discards the connections for that time interval.

For example, assume that the limit is 100,000,000 connections per time interval. If an EC2 instance is connected to by more than 100,000,000 IP addresses during a time interval, then Detective discards the connections from that time interval.

However, you might be able to analyze that activity based on the entity at the other end of the relationship. To continue the example, while an EC2 instance might be connected to from millions of IP addresses, a single IP address connects to far fewer EC2 instances. Each IP address profile provides details about the EC2 instances that the IP address connected to.

Viewing the high-volume entity notification on a profile

Detective displays a notice at the top of a finding or entity profile if the scope time includes a time interval where the entity is high-volume. For finding profiles, the notice is for the involved entity.

The notice includes the list of relationships that have high-volume time intervals. Each list entry contains a description of the relationship and the start of the high-volume time interval.

A high-volume time interval might be an indicator of suspicious activity. To understand what other activity occurred at the same time, you can focus your investigation on a high-volume time interval. The high-volume entity notice includes an option to set the scope time to that time interval.

To set the scope time to a high-volume time interval

1. In the high-volume entity notice, choose the time interval.
2. On the pop-up menu, choose **Apply scope time**.

Viewing the list of high-volume entities for the current scope time

The **High-volume entities** page contains a list of high-volume time intervals and entities during the current scope time.

To display the High-volume entities page

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **High-volume entities**.

Each entry in the list contains the following information:

- The start of the high-volume time interval
- The identifier and type of the entity
- The description of the relationship, such as "EC2 instance connected from IP address"

You can filter and sort the list by any of the columns. You can also navigate to the entity profile for an involved entity.

To navigate to the profile for an entity

1. In the **High-volume entities** list, choose the row to navigate from.
2. Choose **View profile with high-volume scope time**.

When you use this option to navigate to an entity profile, the scope time is set as follows:

- The scope time starts 30 days before the high-volume time interval.
- The scope time ends at the end of the high-volume time interval.

Searching for a finding or entity

With the Amazon Detective search function, you can search for a finding or entity. From the search results, you can navigate to an entity profile or a finding overview. If your search returns more than 10,000 results, only the top 10,000 results are displayed. Changing the sorting order changes the returned results.

You can export your search results to a comma-separated values (.csv) file. This file contains the data returned in the search page. The data is exported in comma-separated values (CSV) format. The file name of the exported data follows the pattern `detective-page-panel-yyyy-mm-dd.csv` format. You can enrich your security investigations by manipulating the data using other AWS services, third-party applications, or spreadsheet programs that support CSV import.

Note

If an export is currently in progress, wait until the export is complete before you try to export additional data.

Completing the search

To complete the search, choose the type of entity to search for. Then provide the exact identifier or identifier with wildcard characters * or ?. To search for a range of IP addresses, you can also use CIDR or dot notations. See the following example search strings.

For IP addresses:

- `1.0.*.*`
- `1.0.133.*`
- `1.0.0.0/16`
- `0.239.48.198/31`

For all other types of entities:

- `Admin`
- `ad*`

- ad*n
- ad*n*
- adm?n
- a?m*
- *min

For each entity type, the following identifiers are supported:

- For Findings, the finding identifier or finding Amazon Resource Name (ARN).
- For AWS accounts, the account ID.
- For AWS roles and AWS users, either the principal ID, the name, or the ARN.
- For Container clusters, the cluster name or ARN.
- For Container images, the repository or the full digest of the container image.
- For container Pods or Tasks, the pod name or the UID of the pod.
- For EC2 instances, the instance identifier or the ARN.
- For Finding group, the finding group identifier.
- For IP addresses, the address in CIDR or dot notation.
- For Kubernetes subjects (service accounts or users), the name.
- For a role session, you can use any of the following values to search:
 - Role session identifier.

The role session identifier uses the format *<rolePrincipalID>:<sessionName>*.

Here is an example: AROA12345678910111213:MySession.

- Role session ARN
- Session name
- Principal ID of the role that was assumed
- Name of the role that was assumed
- For S3 buckets, the bucket name or bucket ARN.
- For federated users, the principal ID or the user name. The principal ID is either *<identityProvider>:<username>* or *<identityProvider>:<audience>:<username>*.
- For user agents, the user agent name.

To search for a finding or entity

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Search**.
3. From the **Choose type** menu, choose the type of item you're looking for.

Note that when you choose **User**, you can search for either an AWS user or a federated user.

Examples from your data contains a sample set of identifiers of the selected type that are in your behavior graph data. To display the profile for one of the examples, choose its identifier.

4. Enter the exact identifier or an identifier with wildcard characters to search for.

The search is case insensitive.

5. Choose **Search** or press **Enter**.

Using the search results

When you complete the search, Detective displays a list of up to 10,000 matching results. For searches that use a unique identifier, there is only one matching result.

From the results, to navigate to the entity profile or finding overview, choose the identifier.

For findings, roles, users, and EC2 instances, the search results include the associated account. To navigate to the profile for the account, choose the account identifier.

Troubleshooting the search

If Detective does not find the finding or entity, first check that you entered the correct identifier. If the identifier is correct, you can also check the following.

- **Does the finding or entity belong to an enabled member account in your behavior graph?**
If the associated account was not invited to the behavior graph as a member account, then the behavior graph does not contain data for that account.

If an invited member account did not accept the invitation, then the behavior graph does not contain data for that account.

- **For a finding, is the finding archived?** Detective does not receive archived findings from Amazon GuardDuty.
- **Did the finding or entity occur before Detective began to ingest data into your behavior graph?** If the finding or entity is not present in the data that Detective ingests, then the behavior graph does not contain data for it.
- **Is the finding or entity from the correct Region?** Each behavior graph is specific to an AWS Region. A behavior graph does not contain data from other Regions.

Managing accounts in Detective

When an account enables Detective, it becomes the administrator account for the behavior graph, and it chooses the member accounts for the behavior graph. An administrator account can invite accounts to join a behavior graph. When the account accepts the invitation, Detective enables the account as a member account. Member accounts that are added by invitation can remove themselves from the behavior graph.

When an account is enabled as a member account, Detective begins to ingest and extract the member account's data into that behavior graph.

Each behavior graph contains data from one or more accounts. A behavior graph can have up to 1,200 member accounts.

If you are integrated with AWS Organizations, then the organization management account designates the Detective administrator account for the organization. That Detective administrator account then becomes the administrator account for the organization behavior graph. The Detective administrator account can enable any organization account as a member account in the organization behavior graph. Organization accounts cannot remove themselves from the organization behavior graph.

Detective charges each account for the data that it contributes to each behavior graph. For information on tracking the volume of data for each account in a behavior graph, see [Forecasting and monitoring Amazon Detective costs](#).

Contents

- [Account restrictions and recommendations in Detective](#)
- [Using Organizations to manage behavior graph accounts](#)
- [Designating the Detective administrator for an organization](#)
- [Available actions for accounts](#)
- [Viewing the list of accounts](#)
- [Managing organization accounts as Detective member accounts](#)
- [Managing invited member accounts in Detective](#)
- [For member accounts: Managing behavior graph invitations and memberships](#)
- [Effect of account actions on behavior graphs](#)
- [Using Detective Python scripts to manage accounts](#)

Account restrictions and recommendations in Detective

When managing accounts in Amazon Detective, be aware of the following restrictions and recommendations.

Maximum number of member accounts

Detective allows up to 1,200 member accounts in each behavior graph.

Accounts and Regions

If you use AWS Organizations to manage accounts, the organization management account designates a Detective administrator account for the organization. The Detective administrator account becomes the administrator account for the organization behavior graph.

The Detective administrator account must be the same in all Regions. The organization management account designates the Detective administrator account separately in each Region. The Detective administrator account also manages the organization behavior graphs and member accounts separately in each Region.

For member accounts created by invitation, the administrator-member association is created only in the Region that the invitation is sent from. The administrator account must enable Detective in each Region, and has a separate behavior graph in each Region. The administrator account then invites each account to associate as a member account in that Region.

An account can be a member account of multiple behavior graphs in the same Region. An account can only be the administrator account of one behavior graph per Region. An account can be an administrator account in different Regions.

Alignment of administrator accounts with Security Hub and GuardDuty

To ensure that the integrations with AWS Security Hub and Amazon GuardDuty work smoothly, we recommend that the same account is the administrator account in all of these services.

See [the section called "Recommended alignment with GuardDuty and AWS Security Hub"](#).

Granting the required permissions for administrator accounts

To ensure that an administrator account has the required permissions to manage its behavior graph, attach the [AmazonDetectiveFullAccess managed policy](#) to the IAM principal.

Reflecting organization updates in Detective

Changes to an organization are not immediately reflected in Detective.

For most changes, such as new and removed organization accounts, it can take up to an hour for Detective to be notified.

A change to the designated Detective administrator account in Organizations takes less time to propagate.

Using Organizations to manage behavior graph accounts

You might have an existing behavior graph with member accounts that accepted a manual invitation. If you are enrolled in AWS Organizations, use the following steps to use Organizations to enable and manage member accounts instead of using the manual invitation process:


1. [Designate the Detective administrator account for your organization.](#) This creates the organization behavior graph.

If the Detective administrator account already has a behavior graph, then that behavior graph becomes the organization behavior graph.

2. [Enable organization accounts as member accounts in the organization behavior graph.](#)

If the organization behavior graph has existing member accounts that are organization accounts, those accounts are enabled automatically.

The following diagram shows an overview of a behavior graph structure before the transition, the configuration in Organizations, and the behavior graph account structure after the transition.



Designate a Detective administrator account for your organization

Your organization management account designates the Detective administrator account from your organization. See [the section called "Designating the Detective administrator account"](#).

To make the transition simpler, Detective recommends that you choose a current administrator account as the Detective administrator account for the organization.

If there is a delegated administrator account for Detective in Organizations, then you must use either that account or the organization management account as the Detective administrator account.

Otherwise, the first time you designate a Detective administrator account that is not the organization management account, Detective calls Organizations to make that account the delegated administrator account for Detective.

Enable organization accounts as member accounts

The Detective administrator account is the administrator account for the organization behavior graph. The Detective administrator account chooses the organization accounts to enable

as member accounts in the organization behavior graph. See [the section called “Managing organization member accounts”](#).

On the **Accounts** page, the Detective administrator account sees all of the accounts in the organization.

If the Detective administrator account was already the administrator account for a behavior graph, then that behavior graph becomes the organization behavior graph. Organization accounts that were already member accounts in that behavior graph are enabled as member accounts automatically. Other organization accounts have a status of **Not a member**.

Organization accounts have a type of **By organization**, even if they were previously member accounts by invitation.

Member accounts that do not belong to the organization have a type of **By invitation**.

The **Account management** page also provides an option, **Automatically enable new organization accounts**, to automatically enable new accounts as they are added to an organization. See [the section called “Enabling new organization accounts”](#). The option is initially turned off.

When the Detective administrator account first displays the **Account management** page, it displays a message that contains an **Enable all organization accounts** button. When you choose **Enable all organization accounts**, Detective performs the following actions:

- Enables all of the current organization accounts as member accounts.
- Turns on the option to automatically enable new organization accounts.

There is also an **Enable all organization accounts** option on the member account list.

Designating the Detective administrator for an organization

In the organization behavior graph, the Detective administrator account manages the behavior graph membership for all organization accounts.

How the Detective administrator account is managed – The organization management account designates the Detective administrator account for the organization in each AWS Region.

Setting the Detective administrator account as the delegated administrator account – The Detective administrator account also becomes the delegated administrator account for Detective in


AWS Organizations. The exception is if the organization management account designates itself as the Detective administrator account. The organization management account cannot be a delegated administrator in Organizations.

After the delegated administrator account is set in Organizations, the organization management account can only choose either the delegated administrator account or their own account as the Detective administrator account. We recommend that you choose the delegated administrator account in all Regions.

Creating and managing the organization behavior graph – When the organization management account chooses a Detective administrator account, Detective creates a new behavior graph for that account. That behavior graph is the organization behavior graph.

If the Detective administrator account is an administrator account for an existing behavior graph, then that behavior graph becomes the organization behavior graph.

The Detective administrator account chooses organization accounts to enable as member accounts in the organization behavior graph.



The Detective administrator account can also send invitations to accounts that do not belong to the organization. For more information, see [the section called “Managing organization member accounts”](#) and [the section called “Managing invited member accounts”](#).

Required permissions to configure the Detective administrator account – To ensure that the organization management account is able to configure the Detective administrator account, you

can attach the [AmazonDetectiveOrganizationsAccess managed policy](#) to your AWS Identity and Access Management (IAM) entities.

Designating a Detective administrator

The organization management account can use the Detective console to designate the Detective administrator account.

You do not need to enable Detective in order to manage the Detective administrator account. You can manage the Detective administrator account from the **Enable Detective** page.

Enable Detective page (Console)

To designate a Detective administrator from the **Enable Detective** page, follow these steps.

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. Choose **Get started**.
3. In the **Required permissions for administrator accounts** panel, grant necessary the permissions to the account you choose so that they can operate as a Detective administrator with full access to all actions in Detective. To operate as an administrator, We recommend attaching the AmazonDetectiveFullAccess policy to the principal.
4. Choose **Attach policy from IAM** to view the recommended policy directly in the IAM console.
5. Depending on whether you have permissions in the IAM console, proceed as follows:
 - If you have permissions to operate in the IAM console, attach the recommended policy to the principal you use for Detective.
 - If you don't have permissions to operate in the IAM console, copy the Amazon Resource Name (ARN) of the policy and provide it to your IAM administrator. They can then attach the policy on your behalf.
6. Under **Delegated administrator**, choose the Detective administrator account.

The available options depend on whether you have a delegated administrator account for Detective in Organizations.

- If you do not have a delegated administrator account for Detective in Organizations, then enter the account identifier of the account to designate it as the Detective administrator account.

You might have an existing administrator account and behavior graph from the manual invitation process. If so, we recommend that you designate that account as the Detective administrator account.

If you have a delegated administrator account in Organizations for Amazon GuardDuty, AWS Security Hub, or Amazon Macie, then Detective prompts you to select one of those accounts. You can also enter a different account.

- If you do have a delegated administrator account for Detective in Organizations, then you are prompted to choose either that account or your account. We recommend that you choose the delegated administrator account in all Regions.

7. Choose **Delegate**.

If you have Detective enabled, or are a member account in an existing behavior graph, then you can designate the Detective administrator account from the **General** page.

General page (Console)

To designate a Detective administrator from the **General** page, follow these steps.

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, under **Settings**, choose **General**.
3. In the **Managed policies** panel, you can learn more about all the managed policies Detective supports. You can grant necessary permissions to an account depending on the actions you want users to perform in Detective. To operate as an administrator, We recommend attaching the `AmazonDetectiveFullAccess` policy to the principal.
4. Depending on whether you have permissions in the IAM console, proceed as follows:
 - If you have permissions to operate in the IAM console, attach the recommended policy to the principal you use for Detective.
 - If you don't have permissions to operate in the IAM console, copy the Amazon Resource Name (ARN) of the policy and provide it to your IAM administrator. They can then attach the policy on your behalf.

The available options depend on whether you have a delegated administrator account for Detective in Organizations.

- If you do not have a delegated administrator account for Detective in Organizations, then enter the account identifier of the account to designate it as the Detective administrator account.

You might have an existing administrator account and behavior graph from the manual invitation process. If so, then we recommend that you designate that account as the Detective administrator account.

If you have a delegated administrator account in Organizations for Amazon GuardDuty, AWS Security Hub, or Amazon Macie, then Detective prompts you to select one of those accounts. You can also enter a different account.

- If you do have a delegated administrator account for Detective in Organizations, then you are prompted to choose either that account or your account. We recommend that you choose the delegated administrator account in all Regions.

5. Choose **Delegate**.

Detective API, AWS CLI

To designate the Detective administrator account, you can use an API call or the AWS Command Line Interface. You must use the organization management account credentials.

If you already have a delegated administrator account for Detective in organizations, then you must choose either that account or your account we recommend that you choose the delegated administrator account.

To designate the Detective administrator account (Detective API, AWS CLI)

- **Detective API:** Use the [EnableOrganizationAdminAccount](#) operation. You must provide the AWS account identifier of the Detective administrator account. To obtain the account identifier, use the [ListOrganizationAdminAccounts](#) operation.
- **AWS CLI:** At the command line, run the [enable-organization-admin-account](#) command.

```
aws detective enable-organization-admin-account --account-id <admin account ID>
```

Example

```
aws detective enable-organization-admin-account --account-id 777788889999
```

Removing the Detective administrator account

The organization management account can remove the current Detective administrator account in a Region. When you remove the Detective administrator account, Detective only removes it from the current Region. It does not change the delegated administrator account in Organizations.

When the organization management account removes the Detective administrator account in a Region, Detective deletes the organization behavior graph. Detective is disabled for the removed Detective administrator account.

To remove the current delegated administrator account for Detective, you use the Organizations API. When you remove the delegated administrator account for Detective in Organizations, Detective deletes all of the organization behavior graphs where the delegated administrator account is the Detective administrator account. Organization behavior graphs that have the organization management account as the Detective administrator account are not affected.

Console

From the Detective console, you can remove the Detective administrator account.

When you remove the Detective administrator account, Detective is disabled for the account, and the organization behavior graph is deleted. The Detective administrator account is only removed in the current Region.

Important

Removing a Detective administrator account does not affect the delegated administrator account in Organizations.

To remove the Detective administrator account (Enable Detective page)

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. Choose **Get started**.
3. Under **Delegated Administrator**, choose **Disable Amazon Detective**.

4. On the confirmation dialog box, enter **disable**, then choose **Disable Amazon Detective**.

To remove a Detective administrator account (General page)

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, under **Settings**, choose **General**.
3. Under **Delegated Administrator**, choose **Disable Amazon Detective**.
4. On the confirmation dialog box, enter **disable**, then choose **Disable Amazon Detective**.

Detective API, AWS CLI

To remove the Detective administrator account, you can use an API call or the AWS CLI. You must use the organization management account credentials.

When you remove the Detective administrator account, Detective is disabled for the account, and the organization behavior graph is deleted.

Important

Removing a Detective administrator account does not affect the delegated administrator account in Organizations.

To remove the Detective administrator account (Detective API, AWS CLI)

- **Detective API:** Use the [DisableOrganizationAdminAccount](#) operation.

When you use the Detective API to remove the Detective administrator account, it is only removed in the Region where the API call or command was issued.

- **AWS CLI:** At the command line, run the [disable-organization-admin-account](#) command.

```
aws detective disable-organization-admin-account
```


Removing the delegated administrator account

Removing the Detective administrator account does not automatically remove the delegated administrator account in Organizations. To remove the delegated administrator account for Detective, you can use the Organizations API.

When you remove the delegated administrator account, this deletes all organization behavior graphs where the delegated administrator account is the Detective administrator account. It also disables Detective for the account in those Regions.

To remove the delegated administrator account (Organizations API, AWS CLI)

- **Organizations API:** Use the [DeregisterDelegatedAdministrator](#) operation. You must provide the account identifier of the Detective administrator account, and the service principal for Detective, which is `detective.amazonaws.com`.
- **AWS CLI:** At the command line, run the [deregister-delegated-administrator](#) command.

```
aws organizations deregister-delegated-administrator --account-id <Detective administrator account ID> --service-principal <Detective service principal>
```

Example

```
aws organizations deregister-delegated-administrator --account-id 777788889999 --service-principal detective.amazonaws.com
```

Available actions for accounts

Administrator and member accounts have access to the following Detective actions. In the table, the values have the following meanings:

- **Any** – The account can perform the action for all of the accounts under the same Detective administrator account.
- **Self** – The account can only perform the action on their own account.
- Dash (–) – The account cannot perform the action.

In the organization behavior graph, the Detective administrator account determines which organization accounts to enable as member accounts. They can configure Detective to enable new

organization accounts as member accounts automatically, or they can enable organization accounts manually.

An administrator account can invite accounts to be member accounts in the behavior graph. When a member account accepts the invitation and is enabled, Amazon Detective begins to ingest and extract the member account's data into that behavior graph.

For behavior graphs other than the organization behavior graph, all of the member accounts are invited accounts.

The following table reflects the default permissions for administrator and member accounts. You can use custom IAM policies to restrict access further to Detective features and functions.

Action	Administrator account (Organization)	Administrator account (Invitation)	Member (Organization)	Member (Invitation)
View accounts	Any	Any	Self (View administrator accounts)	Self (View administrator accounts)
Remove member account	Any Invited accounts are removed Organization accounts are disassociated	Any	–	Self
Add or remove optional data source packages	Any (Setting applies to all member accounts)	Any (Setting applies to all member accounts)	–	–
Disable Detective	Self	Self	–	–
View behavior graph data	Any	Any	–	–

Action	Administrator account (Organization)	Administrator account (Invitation)	Member (Organization)	Member (Invitation)
Enable or disable optional data source packages	All	All	–	–

Viewing the list of accounts

The administrator account can use the Detective console or API to view a list of accounts. The list can include:

- Accounts that the administrator account invited to join the behavior graph. These accounts have a type of **By invitation**.
- For the organization behavior graph, all of the accounts in the organization. These accounts have a type of **By organization**.

The results do not include invited member accounts that declined an invitation or that the administrator account removed from the behavior graph. It only includes accounts with the following statuses.

Verification in progress

For invited accounts, Detective is verifying the account email address before it sends the invitation.

For organization accounts, Detective is verifying that the account belongs to the organization. Detective also verifies that it was the Detective administrator account that enabled the account.

Verification failed

The verification failed. The invitation was not sent, or the organization account was not enabled as a member.

Invited

For invited accounts. The invitation was sent, but the member account has not yet responded.

Not a member

For organization accounts in the organization behavior graph. The organization account is not currently a member account. It does not contribute data to the organization behavior graph.

Enabled

For invited accounts, the member account accepted the invitation and contributes data to the behavior graph.

For organization accounts in the organization behavior graph, the Detective administrator account enabled the account as a member account. The account contributes data to the organization behavior graph.

Not enabled

For invited accounts, the member account accepted the invitation, but cannot be enabled.

For organization accounts in the organization behavior graph, the Detective administrator account tried to enable the account, but the account cannot be enabled.

For invited accounts, Detective checks the number of member accounts. The maximum number of member accounts for a behavior graph is 1,200. If the behavior graph already contains 1,200 member accounts, then new accounts cannot be enabled.

Detective checks whether your data volume is within the Detective quota. The volume of data flowing into a behavior graph must be less than the maximum allowed by Detective. If the current volume ingested is above the limit of 10 TB per day for Behavior graph data volume, then Detective will not allow you to add additional member accounts.

Listing accounts (Console)

You can use the AWS Management Console to see and filter your list of accounts.

To display the list of accounts (console)

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.

The member account list contains the following accounts:

- Your account
- Accounts that you invited to contribute data to the behavior graph
- In the organization behavior graph, all of the organization accounts

For each account, the list displays the following information:

- The AWS account identifier.
- For organization accounts, the account name.
- The account type (**By invitation** or **By organization**).
- For invited accounts, the account root user email address.
- The account status.
- The daily data volume for the account. Detective cannot retrieve the data volume for accounts that are not enabled as member accounts.
- The date when the account status was last updated.

You can use the tabs at the top of the table to filter the list based on the member account status. Each tab shows the number of matching member accounts.

- Choose **All** to view all of the member accounts.
- Choose **Enabled** to view accounts that have a status of **Enabled**.
- Choose **Not enabled** to view accounts that have a status other than **Enabled**.

You also can add other filters to the member account list.

To add a filter to the list of accounts in the behavior graph (console)

1. Choose the filter box.
2. Choose the column that you want to use to filter the list.
3. For the specified column, choose the value to use for the filter.
4. To remove a filter, choose the x icon at the top right.
5. To update the list with the most recent status information, choose the refresh icon at the top right.

Listing your member accounts (Detective API, AWS CLI)

You can use an API call or the AWS Command Line Interface to view a list of member accounts in your behavior graph.

To get the ARN of your behavior graph to use in the request, use the [ListGraphs](#) operation.

To retrieve a list of member accounts (Detective API, AWS CLI)

- **Detective API:** Use the [ListMembers](#) operation. To identify the intended behavior graph, specify the behavior graph ARN.

Note that for the organization behavior graph, [ListMembers](#) does not return organization accounts that you did not enable as member accounts or that you disassociated from the behavior graph.

- **AWS CLI:** At the command line, run the [list-members](#) command.

```
aws detective list-members --graph-arn <behavior graph ARN>
```

Example:

```
aws detective list-members --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```

To retrieve details about specific member accounts in your behavior graph (Detective API, AWS CLI)

- **Detective API:** Use the [GetMembers](#) operation. Specify the behavior graph ARN and the list of account identifiers for the member accounts.
- **AWS CLI:** At the command line, run the [get-members](#) command.

```
aws detective get-members --account-ids <member account IDs> --graph-arn <behavior graph ARN>
```

Example:

```
aws detective get-members --account-ids 444455556666 123456789012 --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```

Managing organization accounts as Detective member accounts

In the organization behavior graph, the Detective administrator account determines which organization accounts to enable as member accounts. By default, new organization accounts are not enabled as member accounts. Their status is **Not a member**. The Detective administrator account can configure Detective to automatically enable new organization accounts as member accounts in the organization behavior graph.

The Detective administrator can configure Detective to enable new organization accounts as member accounts automatically. When you choose to enable organization accounts automatically, then Detective begins to enable new accounts as member accounts as they are added to the organization. Detective does not enable existing organization accounts that are not yet enabled.

The Detective can enable organization accounts as member accounts manually, if you do not want to automatically enable new organization accounts. They can also manually enable disassociated organization accounts. The Detective administrator cannot enable an organization account as a member account if the organization behavior graph already has the maximum 1,200 enabled accounts. In this case, the organization account status remains **Not a member**.

The Detective administrator also can disassociate organization accounts from the organization behavior graph. To stop ingesting data from an organization account in the organization behavior graph, you can disassociate the account. Existing data for that account remains in the behavior graph.

Contents

- [Enabling new organization accounts as member accounts](#)
- [Enabling organization accounts as member accounts](#)
- [Disassociating organization accounts as member accounts](#)

Enabling new organization accounts as member accounts

The Detective administrator account can configure Detective to automatically enable new organization accounts as member accounts in the organization behavior graph.

When new accounts are added to your organization, they are added to the list on the **Account management** page. For organization accounts, **Type** is **By organization**.

By default, new organization accounts are not enabled as member accounts. Their status is **Not a member**.

When you choose to enable organization accounts automatically, then Detective begins to enable new accounts as member accounts as they are added to the organization. Detective does not enable existing organization accounts that are not yet enabled.

Detective can enable organization accounts as member accounts only if the maximum number of member accounts for a behavior graph is 1,200. If your behavior graph already contains 1,200 member accounts, then new accounts cannot be enabled.

Enabling new organization accounts automatically (console)

On the **Account management** page, the **Automatically enable new organization accounts** setting determines whether to automatically enable accounts as they are added to an organization.

To automatically enable new organization accounts as member accounts

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.
3. Toggle **Automatically enable new organization accounts** to the on position.

Enabling new organization accounts automatically (Detective API, AWS CLI)

To determine whether to automatically enable new organization accounts as member accounts, the administrator account can use the Detective API or the AWS Command Line Interface.

To view and manage the configuration, you must provide the behavior graph ARN. To obtain the ARN, use the [ListGraphs](#) operation.

To view the current configuration for automatically enabling organization accounts

- **Detective API:** Use the [DescribeOrganizationConfiguration](#) operation.

In the response, if new organization accounts are enabled automatically, then `AutoEnable` is `true`.

- **AWS CLI:** At the command line, run the [describe-organization-configuration](#) command.


```
aws detective describe-organization-configuration --graph-arn <behavior graph ARN>
```

Example

```
aws detective describe-organization-configuration --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```

To automatically enable new organization accounts

- **Detective API:** Use the [UpdateOrganizationConfiguration](#) operation. To automatically enable new organization accounts, set `AutoEnable` to `true`.
- **AWS CLI:** At the command line, run the [update-organization-configuration](#) command.

```
aws detective update-organization-configuration --graph-arn <behavior graph ARN> --auto-enable | --no-auto-enable
```

Example

```
aws detective update-organization-configuration --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234 --auto-enable
```

Enabling organization accounts as member accounts

If you do not automatically enable new organization accounts, then you can enable those accounts manually. You must also manually enable accounts that you disassociated.

Determining whether an account can be enabled

You cannot enable an organization account as a member account if the organization behavior graph already has the maximum 1,200 enabled accounts. In this case, the organization account status remains **Not a member**. The account does not contribute data to the behavior graph.

As soon as the member account can be enabled, Detective automatically changes the member account status to **Enabled**. For example, the member account status changes to **Enabled** if the administrator account removes other member accounts to make space for an account.

Enabling organization accounts as member accounts (console)

From the **Account management** page, you can enable organization accounts as member accounts.

To enable organization accounts as member accounts

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.
3. To view the list of accounts that are not currently enabled, choose **Not enabled**.
4. You can either select specific organization accounts, or enable all organization accounts.

To enable selected organization accounts:

- a. Select each organization account that you want to enable.
- b. Choose **Enable accounts**.

To enable all organization accounts, choose **Enable all organization accounts**.

Enabling organization accounts as member accounts (Detective API, AWS CLI)

You can use the Detective API or the AWS Command Line Interface to enable organization accounts as member accounts in the organization behavior graph. To get the ARN of your behavior graph to use in the request, use the [ListGraphs](#) operation.

To enable organization accounts as member accounts (Detective API, AWS CLI)

- **Detective API:** Use the [CreateMembers](#) operation. You must provide the graph ARN.

For each account, specify the account identifier. Organization accounts in the organization behavior graph do not receive an invitation. You do not need to provide an email address or other invitation information.

- **AWS CLI:** At the command line, run the [create-members](#) command.

```
aws detective create-members --accounts AccountId=<AWS account ID> --graph-arn <behavior graph ARN>
```

Example

```
aws detective create-members --accounts AccountId=444455556666 AccountId=123456789012
--graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```

Disassociating organization accounts as member accounts

To stop ingesting data from an organization account in the organization behavior graph, you can disassociate the account. Existing data for that account remains in the behavior graph.

When you disassociate an organization account, the status changes to **Not a member**. Detective stops ingesting data from that account, but the account remains in the list.

Disassociating organization accounts (console)

From the **Account management** page, you can disassociate organization accounts as member accounts.

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.
3. To display the list of enabled accounts, choose **Enabled**.
4. Select the check box for each account to disassociate.
5. Choose **Actions**. Then choose **Disable accounts**.

The account status for the disassociated accounts changes to **Not a member**.

Disassociating organization accounts (Detective API, AWS CLI)

You can use the Detective API or the AWS Command Line Interface to disassociate organization accounts as member accounts in your behavior graph.

To get the ARN of your behavior graph to use in the request, use the [ListGraphs](#) operation.

To disassociate organization accounts from the organization behavior graph (Detective API, AWS CLI)

- **Detective API:** Use the [DeleteMembers](#) operation. Specify the graph ARN and the list of account identifiers for the member accounts to disassociate.
- **AWS CLI:** At the command line, run the [delete-members](#) command.

```
aws detective delete-members --account-ids <account ID list> --graph-arn <behavior graph ARN>
```

Example

```
aws detective delete-members --account-ids 444455556666 123456789012 --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```

Managing invited member accounts in Detective

A Detective administrator account can invite accounts to be member accounts in their behavior graph. A behavior graph can contain up to 1,200 member accounts. When a member account accepts the invitation and is enabled, Amazon Detective begins to ingest and extract the member account's data into that behavior graph.

To invite individual accounts, you can manually specify the member accounts to invite to contribute their data to a behavior graph. If you want to add a list of member accounts, you can choose to provide a .csv file containing a list of member accounts to invite to your behavior graph.

For behavior graphs other than the organization behavior graph, all of the member accounts are invited accounts. The Detective administrator account can also invite accounts that are not organization accounts to the organization behavior graph.

At a high level, the process for inviting accounts to contribute to a behavior graph is as follows.

1. For each member account to add, the administrator account provides the AWS account identifier and the root user email address.
2. Detective validates that the email address is the root user email address for the account. If the account information is valid, Detective sends the invitation to the member account.

Detective does not perform this validation or sends email invitations to member accounts in these Regions:

- AWS GovCloud (US-East) Region
- AWS GovCloud (US-West) Region

For other Regions, you can `DisableEmailNotification` using the [CreateMembers](#) operation of the Detective API. If `DisableEmailNotification` is set to true, then Detective will not

send invitations to the member accounts. This is a useful setting for accounts that are managed centrally.

3. The member account accepts or declines the invitation.

Even if the administrator account does not send invitation emails, the member account still must respond to the invitation.

4. After the member account accepts the invitation, Detective begins to ingest data from the member account into the behavior graph.

5. As soon as the member account is eligible to be enabled, Detective automatically changes the member account status to **Enabled**.

For example, the member account status changes to **Enabled** if the administrator account removes other member accounts to make space for an account.

If more than one account is **Not enabled**, then Detective enables the accounts in the order in which they were invited. The process to check whether to enable any **Not enabled** accounts runs every hour.

The administrator account also can enable accounts manually, instead of waiting for the automatic process. For example, the administrator account might want to select the accounts to enable. For information on how to enable a member account, see [the section called “Enabling a member account that is Not enabled”](#).

Note that Detective began to automatically enable accounts that are **Not enabled** on May 12, 2021. Accounts that were **Not enabled** before then are not enabled automatically. The administrator account must enable them manually.

The administrator account can remove invited member accounts from the behavior graph. Detective does not remove any existing data from the behavior graph, which aggregates data across member accounts.

Contents

- [Inviting individual accounts to a behavior graph](#)
- [Inviting a list of member accounts to a behavior graph](#)
- [Enabling a member account that is Not enabled](#)
- [Removing member accounts from a behavior graph](#)

Inviting individual accounts to a behavior graph

You can manually specify the member accounts to invite to contribute their data to a behavior graph.

Console

To manually select the member accounts to invite using the Detective console.

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.
3. Choose **Actions**. Then choose **Invite accounts**.
4. Under **Add accounts**, choose **Add individual accounts**.
5. To add a member account to the invitation list, perform the following steps.
 - a. Choose **Add account**.
 - b. For **AWS Account ID**, enter the AWS account ID.
 - c. For **Email address**, enter the root user email address for the account.
6. To remove an account from the list, choose **Remove** for that account.
7. Under **Personalize invitation email**, add customized content to include in the invitation email.

For example, you can use this area to provide contact information. Or use it to remind the member account that they need to attach the required IAM policy to their user or role before they can accept the invitation.

8. **Member account IAM policy** contains the text of the required IAM policy for member accounts. The email invitation includes this policy text. To copy the policy text, choose **Copy**.
9. Choose **Invite**.

Detective API/AWS CLI

You can use the Detective API or the AWS Command Line Interface to invite member accounts to contribute their data to a behavior graph. To get the ARN of your behavior graph to use in the request, use the [ListGraphs](#) operation.

To invite member accounts to a behavior graph (Detective API, AWS CLI)

- **Detective API:** Use the [CreateMembers](#) operation. You must provide the graph ARN. For each account, specify the account identifier and the root user email address.

To not send invitation emails to the member accounts, set `DisableEmailNotification` to true. By default, `DisableEmailNotification` is false .

If you do send invitation emails, you can optionally provide custom text to add to the invitation email.

- **AWS CLI:** At the command line, run the `create-members` command.

```
aws detective create-members --accounts AccountId=<AWS account ID>,EmailAddress=<root user email address> --graph-arn <behavior graph ARN> --message "<Custom message text>"
```

Example

```
aws detective create-members --accounts AccountId=444455556666,EmailAddress=mmajor@example.com AccountId=123456789012,EmailAddress=jstiles@example.com --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234 --message "This is Paul Santos. I need to add your account to the data we use for security investigation in Amazon Detective. If you have any questions, contact me at psantos@example.com."
```

To indicate to not send invitation emails to the member accounts, include `--disable-email-notification`.

```
aws detective create-members --accounts AccountId=<AWS account ID>,EmailAddress=<root user email address> --graph-arn <behavior graph ARN> --disable-email-notification
```

Example

```
aws detective create-members --accounts AccountId=444455556666,EmailAddress=mmajor@example.com AccountId=123456789012,EmailAddress=jstiles@example.com --graph-arn
```

```
arn:aws:detective:us-east-1:111122223333:graph:123412341234 --disable-email-notification
```

Inviting a list of member accounts to a behavior graph

From the Detective console, you can provide a .csv file containing a list of member accounts to invite to your behavior graph.

The first line in the file is the header row. Each account is then listed on a separate line. Each member account entry contains the AWS account ID and the account's root user email address.

Example:

```
Account ID,Email address
111122223333,srodriguez@example.com
444455556666,rroe@example.com
```

When Detective processes the file, it ignores accounts that were already invited, unless the account status is **Verification failed**. That status indicates that the email address provided for the account did not match the account's root user email address. In that case, Detective deletes the original invitation and tries again to verify the email address and send the invitation.

This option also provides a template that you can use to create the list of accounts.

To invite member accounts from a .csv list (console)

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.
3. Choose **Actions**. Then choose **Invite accounts**.
4. Under **Add accounts**, choose **Add from .csv**.
5. To download a template file to work from, choose **Download .csv template**.
6. To select the file containing the list of accounts, choose **Choose .csv file**.
7. Under **Review member accounts**, verify the list of member accounts that Detective found in the file.
8. Under **Personalize invitation email**, add customized content to include in the invitation email.

For example, you can provide contact information, or remind the member account about the required IAM policy.

9. **Member account IAM policy** contains the text of the required IAM policy for member accounts. The email invitation includes this policy text. To copy the policy text, choose **Copy**.
10. Choose **Invite**.

Adding a list of member accounts across Regions

Detective provides an open-source Python script in GitHub that allows you to do the following:

- Add a specified list of member accounts to an administrator account's behavior graphs across a specified list of Regions.
- If the administrator account does not have a behavior graph in a Region, then the script also enables Detective and creates the behavior graph in that Region.
- Send invitation emails to the member accounts.
- Automatically accept the invitations for the member accounts.

For information on how to configure and use the GitHub scripts, see [the section called "Amazon Detective Python scripts"](#).

Enabling a member account that is Not enabled

After a member account accepts an invitation, Amazon Detective checks the number of member accounts. The maximum number of member accounts for a behavior graph is 1,200. If your behavior graph already contains 1,200 member accounts, then new accounts cannot be enabled. If Detective cannot enable the member account, then it sets the member account status to **Not enabled**.

Member accounts that are **Not enabled** do not contribute data to the behavior graph.

Detective automatically enables accounts as the behavior graph can accommodate them.

You can also try to enable member accounts manually that are **Not enabled** member accounts. For example, you might remove existing member accounts to reduce the data volume. Instead of waiting for the automatic process to enable accounts, you can try to enable **Not enabled** member accounts.

Console

The member account list includes an option to enable selected member accounts that are **Not enabled**.

To enable a member account that is Not enabled

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.
3. Under **My member accounts**, select the check box for each member account to enable.

You can only enable member accounts that have a status of **Not enabled**.

4. Choose **Enable accounts**.

Detective determines whether the member account can be enabled. If the member account can be enabled, the status changes to **Enabled**.

Detective API/CLI

You can use an API call or the AWS Command Line Interface to enable a single member account that is **Not enabled**. To get the ARN of your behavior graph to use in the request, use the [ListGraphs](#) operation.

To enable a member account that is Not enabled

- **Detective API:** Use the [StartMonitoringMember](#) API operation. You must provide the behavior graph ARN. To identify the member account, use the AWS account identifier.
- **AWS CLI:** At the command line, run the [start-monitoring-member](#) command:

```
start-monitoring-member --graph-arn <behavior graph ARN> --account-id <AWS account ID>
```

For example:

```
start-monitoring-member --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234 --account-id 444455556666
```

Removing member accounts from a behavior graph

The administrator account can remove invited member accounts from a behavior graph at any time.

Detective automatically removes member accounts that are terminated in AWS, except in the AWS GovCloud (US-East) and AWS GovCloud (US-West) Regions.

When an invited member account is removed from a behavior graph, the following occurs.

- The member account is removed from **My member accounts**.
- Amazon Detective stops ingesting data from the removed account.

Detective does not remove any existing data from the behavior graph, which aggregates data across member accounts.

Console

You can use the AWS Management Console to remove invited member accounts from your behavior graph.

To remove member accounts (console)

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.
3. In the account list, select the check box for each member account to remove.

You cannot remove your own account from the list.

4. Choose **Actions**. Then choose **Disable accounts**.

Detective API/CLI

You can use the Detective API or the AWS Command Line Interface to remove invited member accounts from your behavior graph. To get the ARN of your behavior graph to use in the request, use the [ListGraphs](#) operation.

To remove invited member accounts from your behavior graph (Detective API, AWS CLI)

- **Detective API:** Use the [DeleteMembers](#) operation. Specify the graph ARN and the list of account identifiers for the member accounts to remove.
- **AWS CLI:** At the command line, run the [delete-members](#) command.

```
aws detective delete-members --account-ids <account ID list> --graph-arn <behavior graph ARN>
```

Example:

```
aws detective delete-members --account-ids 444455556666 123456789012 --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```

Python script

Detective provides an open-source script in GitHub. You can use this script to remove a specified list of member accounts from an administrator account's behavior graphs across a specified list of Regions.

For information on how to configure and use the GitHub scripts, see [the section called "Amazon Detective Python scripts"](#).

For member accounts: Managing behavior graph invitations and memberships

Amazon Detective charges each member account for the ingested data for each behavior graph that it contributes to.

The **Account management** page allows member accounts to see the administrator accounts for the behavior graphs they are a member of.

Member accounts that are invited to a behavior graph can view and respond to their invitations. They can also remove their account from the behavior graph.

For the organization behavior graph, organization accounts do not control whether their account is a member account. The Detective administrator account chooses the organization accounts to enable or disable as member accounts.

Contents

- [Required IAM policy for a member account](#)
- [Viewing your list of behavior graph invitations](#)
- [Responding to a behavior graph invitation](#)
- [Removing your account from a behavior graph](#)

Required IAM policy for a member account

Before a member account can view and manage invitations, the required IAM policy must be attached to their principal. The principal can be an existing user or role, or you can create a new user or role to use for Detective.

Ideally, the administrator account has their IAM administrator attach the required policy.

The member account IAM policy grants access to member account actions in Amazon Detective. The email invitation to contribute to a behavior graph includes the text of that IAM policy.

To use this policy, replace *<behavior graph ARN>* with the graph ARN.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "detective:AcceptInvitation",
        "detective:DisassociateMembership",
        "detective:RejectInvitation"
      ],
      "Resource": "<behavior graph ARN>"
    },
    {
      "Effect": "Allow",
      "Action": [
        "detective:BatchGetMembershipDatasources",
        "detective:GetFreeTrialEligibility",
        "detective:GetPricingInformation",
        "detective:GetUsageInformation",
        "detective:ListInvitations"
      ]
    }
  ]
}
```

```
"Resource": "*"
}
]
}
```

Note that organization accounts in the organization behavior graph do not receive invitations and cannot disassociate their account from the organization behavior graph. If they do not belong to other behavior graphs, then they only require the `ListInvitations` permission. `ListInvitations` allows them to see the administrator account for the behavior graph. The permissions to manage invitations and disassociate memberships only apply to memberships by invitation.

Viewing your list of behavior graph invitations

From the Amazon Detective console, Detective API, or AWS Command Line Interface, a member account can see their behavior graph invitations.

Viewing behavior graph invitations (console)

You can view behavior graph invitations from the AWS Management Console.

To view behavior graph invitations (console)

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.

On the **Account management** page, **My administrator accounts** contains your open and accepted behavior graph invitations in the current Region. For an organization account, **My administrator accounts** also contains the organization behavior graph.

If your account is currently in the free trial period, the page also displays the number of days remaining in your free trial.

The list does not contain invitations that you declined, memberships that you resigned, or memberships that the administrator account removed.

Each invitation shows the administrator account number, the date that the invitation was accepted, and the current status of the invitation.

- For invitations that you have not responded to, the status is **Invited**.
- For invitations that you accepted, the status is either **Enabled** or **Not enabled**.

If the status is **Enabled**, then your account contributes data to the behavior graph.

If the status is **Not enabled**, then your account does not contribute data to the behavior graph.

Your account status is set initially to **Not enabled** while Detective checks whether you have GuardDuty enabled, and if so, whether your account would cause the data volume for the behavior graph to exceed the Detective quota.

If your account would not cause the behavior graph to exceed the quota, Detective updates your account status to **Enabled**. Otherwise, the status remains **Not enabled**.

When the behavior graph is able to accommodate the data volume for your account, Detective automatically updates it to **Enabled**. For example, the administrator account might remove other member accounts so that your account can be enabled. The administrator account can also enable your account manually.

Viewing behavior graph invitations (Detective API, AWS CLI)

You can list behavior graph invitations from the Detective API or the AWS Command Line Interface.

To retrieve a list of open and accepted invitations to behavior graphs (Detective API, AWS CLI)

- **Detective API:** Use the [ListInvitations](#) operation.
- **AWS CLI:** At the command line, run the [list-invitations](#) command.

```
aws detective list-invitations
```

Responding to a behavior graph invitation

After you accept an invitation, Detective checks the number of member accounts. The maximum number of member accounts for a behavior graph is 1,200. If your behavior graph already contains 1,200 member accounts, then new accounts cannot be enabled.

After you accept the invitation, Detective is enabled in your account. Detective checks whether your data volume is within the Detective quota. The volume of data flowing into a behavior graph must

be less than the maximum allowed by Detective. If the current volume ingested is above the limit of 10 TB per day, you cannot add more accounts and Detective will disable further ingestion of data. The Detective console displays a notification to indicate that data volume is too large and the status remains **Not enabled**.

If you decline the invitation, then it is removed from your list of invitations, and Detective does not use your account data in the behavior graph.

Responding to a behavior graph invitation (console)

You can use the AWS Management Console to respond to the email invitation, which includes a link to the Detective console. You can only respond to an invitation that has a status of **Invited**.

To respond to a behavior graph invitation (console)

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.
3. Under **My administrator accounts**, to accept the invitation and begin contributing data to the behavior graph, choose **Accept invitation**.

To decline the invitation and remove it from the list, choose **Decline**.

Responding to a behavior graph invitation (Detective API, AWS CLI)

You can respond to behavior graph invitations from the Detective API or the AWS Command Line Interface.

To accept a behavior graph invitation (Detective API, AWS CLI)

- **Detective API:** Use the [AcceptInvitation](#) operation. You must specify the graph ARN.
- **AWS CLI:** At the command line, run the [accept-invitation](#) command.

```
aws detective accept-invitation --graph-arn <behavior graph ARN>
```

Example:

```
aws detective accept-invitation --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```


To decline a behavior graph invitation (Detective API, AWS CLI)

- **Detective API:** Use the [RejectInvitation](#) operation. You must specify the graph ARN.
- **AWS CLI:** At the command line, run the [reject-invitation](#) command.

```
aws detective reject-invitation --graph-arn <behavior graph ARN>
```

Example:

```
aws detective reject-invitation --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```

Removing your account from a behavior graph

After you accept an invitation, you can remove your account from a behavior graph at any time. When you remove your account from a behavior graph, Amazon Detective stops ingesting data from your account into the behavior graph. Existing data remains in the behavior graph.

Only invited accounts can remove their account from a behavior graph. Organization accounts cannot remove their account from the organization behavior graph.

Removing your account from a behavior graph (Console)

You can use the AWS Management Console to remove your account from a behavior graph.

To remove your account from a behavior graph (console)

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, choose **Account management**.
3. Under **My administrator accounts**, for the behavior graph you want to resign from, choose **Resign**.

Removing your account from a behavior graph (Detective API, AWS CLI)

You can use the Detective API or the AWS Command Line Interface to remove your account from a behavior graph.

To remove your account from a behavior graph (Detective API, AWS CLI)

- **Detective API:** Use the [DisassociateMembership](#) operation. You must specify the graph ARN.
- **AWS CLI:** At the command line, run the [disassociate-membership](#) command.

```
aws detective disassociate-membership --graph-arn <behavior graph ARN>
```

Example:

```
aws detective disassociate-membership --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```

Effect of account actions on behavior graphs

These actions have the following effects on Amazon Detective data and access.

Detective disabled

When an administrator account disables Detective, the following occurs:

- The behavior graph is removed.
- Detective stops ingesting data from the administrator account and the member accounts for that behavior graph.

Member account removed from the behavior graph

When a member account is removed from a behavior graph, Detective stops ingesting data from that account.

Existing data in the behavior graph is not affected.

For invited accounts, the account is removed from the **My member accounts** list.

For organization accounts in the organization behavior graph, the account status changes to **Not a member**.

Member account leaves the organization

When a member account leaves an organization, the following occurs:

- The account is removed from the **My member accounts** list for the organization behavior graph.
- Detective stops ingesting data from that account.

Existing data in the behavior graph is not affected.

AWS account suspended

When an administrator account is suspended in AWS, the account loses permission to view the behavior graph in Detective. Detective stops ingesting data into the behavior graph.

When a member account is suspended in AWS, Detective stops ingesting data for that account.

After 90 days, the account is either terminated or reactivated. When an administrator account is reactivated, its Detective permissions are restored. Detective resumes the ingest of data from the account. When a member account is reactivated, Detective resumes the ingest of data from the account.

AWS account closed

When an AWS account is closed, Detective responds to the closure as follows.

- For an administrator account, Detective deletes the behavior graph.
- For a member account, Detective removes the account from the behavior graph.

AWS retains the policy data for the account for 90 days from the effective date of the administrator account closure. At the end of the 90 day period, AWS permanently deletes all policy data for the account.

- To retain findings for more than 90 days, you can archive the policies. You can also use a custom action with an EventBridge rule to store the findings in an S3 bucket.
- As long as AWS retains the policy data, when you reopen the closed account, AWS reassigns the account as the service administrator and recovers the service policy data for the account.
- For more information, see [Closing an account](#).

Important

For customers in the AWS GovCloud (US) Regions:

- Before closing your account, back up and then delete account resources. You will no longer have access to them after you close the account.

Using Detective Python scripts to manage accounts

Amazon Detective provides a set of open-source Python scripts in the GitHub repository [amazon-detective-multiaccount-scripts](#). The scripts require Python 3.

You can use these to perform the following tasks:

- Enable Detective for an administrator account across Regions.

When you enable Detective, you can assign tag values to the behavior graph.

- Add member accounts to an administrator account's behavior graphs across Regions.
- Optionally send invitation emails to the member accounts. You can also configure the request to not send invitation emails.
- Remove member accounts from an administrator account's behavior graphs across Regions.
- Disable Detective for an administrator account across Regions. When an administrator account disables Detective, the administrator account's behavior graph in each Region is disabled.

Overview of the `enableDetective.py` script

The `enableDetective.py` script does the following:

1. Enables Detective in for an administrator account in each specified Region, if the administrator account does not already have Detective enabled in that Region.

When you use the script to enable Detective, you can assign tag values to the behavior graph.

2. Optionally sends invitations from the administrator account to the specified member accounts for each behavior graph.

The invitation email messages use the default message content and cannot be customized.

You can also configure the request to not send invitation emails.

3. Automatically accepts the invitations for the member accounts.

Because the script automatically accepts the invitations, member accounts can ignore these messages.

We recommend reaching out directly to the member accounts to notify them that the invitations are accepted automatically.

Overview of the `disableDetective.py` script

The `disableDetective.py` script deletes the specified member accounts from the administrator account's behavior graphs across the specified Regions.

It also provides an option to disable Detective for the administrator account across the specified Regions.

Required permissions for the scripts

The scripts require a preexisting AWS role in the administrator account and in all of the member accounts that you add or remove.

Note

The role name must be the same in all of the accounts.

IAM policy [recommended best practices](#) are to use least scoped roles. To execute the script's workflow of [creating a graph](#), [creating members](#), and [adding members to the graph](#) the required permissions are:

- `detective:CreateGraph`
- `detective:CreateMembers`
- `detective>DeleteGraph`
- `detective>DeleteMembers`
- `detective:ListGraphs`
- `detective:ListMembers`
- `detective:AcceptInvitation`

Role trust relationship

The role trust relationship must allow your instance or local credentials to assume the role.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::<ACCOUNTID>:user/<USERNAME>"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

If you do not have a common role that includes the required permissions, you must create a role with at least those permissions in each member account. You must also create the role in the administrator account.

When you create the role, make sure that you do the following:

- Use the same role name in every account.
- Add the required permissions above (recommended) or select the [AmazonDetectiveFullAccess](#) managed policy.
- Add role trust relationship block as discussed above.

To automate this process, you can use the `EnableDetective.yaml` AWS CloudFormation template. Because the template creates only global resources, it can be run in any Region.

Setting up the run environment for the Python scripts

You can run the scripts from either an EC2 instance or from a local machine.

Launching and configuring an EC2 instance

One option for running the scripts is to run them from an EC2 instance.

To launch and configure an EC2 instance

1. Launch an EC2 instance in your administrator account. For details on how to launch an EC2 instance, see [Getting Started with Amazon EC2 Linux Instances](#) in the *Amazon EC2 User Guide*.
2. Attach to the instance an IAM role that has permissions to allow the instance to call `AssumeRole` within the administrator account.

If you used the `EnableDetective.yaml` AWS CloudFormation template, then an instance role with a profile named `EnableDetective` was created.

Otherwise, for information on creating an instance role, see the blog post [Easily Replace or Attach an IAM Role to an Existing EC2 Instance by Using the EC2 Console](#).

3. Install the required software:
 - **APT:** `sudo apt-get -y install python3-pip python3 git`
 - **RPM:** `sudo yum -y install python3-pip python3 git`
 - **Boto (minimum version 1.15):** `sudo pip install boto3`
4. Clone the repository to the EC2 instance.

```
git clone https://github.com/aws-samples/amazon-detective-multiaccount-scripts.git
```

Configuring a local machine to run the scripts

You can also run the scripts from your local machine.

To configure a local machine to run the scripts

1. Make sure that you have set up on your local machine credentials for your administrator account that have permission to call `AssumeRole`.
2. Install the required software:
 - Python 3
 - Boto (minimum version 1.15)
 - GitHub scripts

Platform	Setup instructions
Windows	<ol style="list-style-type: none"> 1. Install Python 3 (https://www.python.org/downloads/windows/). 2. Open a command prompt. 3. To install Boto, run: <code>pip install boto3</code> 4. Download the script source code from GitHub (https://github.com/aws-samples/amazon-detective-multiaccount-scripts).
Mac	<ol style="list-style-type: none"> 1. Install Python 3 (https://www.python.org/downloads/mac-osx/). 2. Open a command prompt. 3. To install Boto, run: <code>pip install boto3</code> 4. Download the script source code from GitHub (https://github.com/aws-samples/amazon-detective-multiaccount-scripts).
Linux	<ol style="list-style-type: none"> 1. To install Python 3, run one of the following: <ul style="list-style-type: none"> • <code>sudo apt-get -y install python3-pip python3 git</code> • <code>sudo yum install git python</code> 2. To install Boto, run: <code>sudo pip install boto3</code> 3. Clone the script source code from https://github.com/aws-samples/amazon-detective-multiaccount-scripts.

Creating a .csv list of member accounts to add or remove

To identify the member accounts to add to or remove from the behavior graphs, you provide a .csv file that contains the list of accounts.

List each account on a separate line. Each member account entry contains the AWS account ID and the account's root user email address.

See the following example:

```
111122223333,srodriguez@example.com  
444455556666,rroe@example.com
```

Running enableDetective.py

You can run the `enableDetective.py` script from an EC2 instance or your local machine.

To run enableDetective.py

1. Copy the `.csv` file to the `amazon-detective-multiaccount-scripts` directory on your EC2 instance or local machine.
2. Change to the `amazon-detective-multiaccount-scripts` directory.
3. Run the `enableDetective.py` script.

```
enableDetective.py --master_account administratorAccountID --assume_role roleName  
  --input_file inputFileName --tags tagValueList --enabled_regions regionList --  
  disable_email
```

When you run the script, replace the following values:

administratorAccountID

The AWS account ID for the administrator account.

roleName

The name of the AWS role to assume in the administrator account and each member account.

inputFileName

The name of the `.csv` file containing the list of member accounts to add to the administrator account's behavior graphs.

tagValueList

(Optional) A comma-separated list of tag values to assign to a new behavior graph.

For each tag value, the format is *key=value*. For example:

```
--tags Department=Finance,Geo=Americas
```

regionList

(Optional) A comma-separated list of Regions in which to add the member accounts to the administrator account's behavior graph. For example:

```
--enabled_regions us-east-1,us-east-2,us-west-2
```

The administrator account might not already have Detective enabled in a Region. In that case, the script enables Detective and creates a new behavior graph for the administrator account.

If you do not provide a list of Regions, then the script acts across all Regions that Detective supports.

`--disable_email`

(Optional) If included, Detective does not send invitation emails to the member accounts.

Running `disableDetective.py`

You can run the `disableDetective.py` script from an EC2 instance or your local machine.

To run `disableDetective.py`

1. Copy the `.csv` file to the `amazon-detective-multiaccount-scripts` directory.
2. To use the `.csv` file to delete the listed member accounts from the administrator account's behavior graphs across a specified list of Regions, run the `disableDetective.py` script as follows:

```
disabledetective.py --master_account administratorAccountID --assume_role roleName  
--input_file inputFileName --disabled_regions regionList
```

3. To disable Detective for the administrator account across all Regions, run the `disableDetective.py` script with the `--delete-master` flag.

```
disabledetective.py --master_account administratorAccountID --assume_role roleName  
--input_file inputFileName --disabled_regions regionList --delete_master
```

When you run the script, replace the following values:

administratorAccountID

The AWS account ID for the administrator account.

roleName

The name of the AWS role to assume in the administrator account and each member account.

inputFileName

The name of the .csv file containing the list of member accounts to remove from the administrator account's behavior graphs.

You must provide a .csv file even if you are disabling Detective.

regionList

(Optional) A comma-separated list of Regions in which to do one of the following:

- Remove the member accounts from the administrator account's behavior graphs.
- If the `--delete-master` flag is included, disable Detective.

For example:

```
--disabled_regions us-east-1,us-east-2,us-west-2
```

If you do not provide a list of Regions, then the script acts across all Regions that Detective supports.

Integration with Amazon Security Lake

Amazon Security Lake is a fully managed security data lake service. You can use Security Lake to automatically centralize security data from AWS environments, SaaS providers, on-premises sources, cloud sources, and third-party sources into a purpose-built data lake that's stored in your AWS account. Security Lake helps you analyze security data, so you can get a more complete understanding of your security posture across your entire organization. With Security Lake, you can also improve the protection of your workloads, applications, and data.

Amazon Detective integrates with Amazon Security Lake, which means that you can query and retrieve the raw log data stored by Security Lake.

Using this integration, you can collect logs and events from the following sources which Security Lake natively supports. Detective supports up to source version 2 (OCSF 1.1.0).

- AWS CloudTrail management events version 1.0 and after
- Amazon Virtual Private Cloud (Amazon VPC) Flow Logs version 1.0 and after
- Amazon Elastic Kubernetes Service (Amazon EKS) Audit Log version 2.0. — To use Amazon EKS audit logs as a source you must add `ram:ListResources` to the IAM permissions. For more details, see [Add the required IAM permissions to your account](#).

For details on how Security Lake automatically converts logs and events that come from natively-supported AWS services to the OCSF schema, see the [Amazon Security Lake User Guide](#).

After you integrate Detective with Security Lake, Detective begins pulling raw logs from Security Lake related to AWS CloudTrail management events and Amazon VPC Flow Logs. For more details, see [Querying raw logs](#).

Enabling Detective integration with Security Lake

To integrate Detective with Security Lake, you must complete the following steps.

1. [Before you begin](#)

Use an Organizations management account to designate a delegated Security Lake administrator for your organization. Make sure that Security Lake is enabled and verify that Security Lake is collecting logs and events from AWS CloudTrail management events and Amazon Virtual Private Cloud (Amazon VPC) Flow Logs.

In alignment with the Security Reference Architecture, Detective recommends using a Log Archive account and defer from using a Security Tooling account for the Security Lake deployment.

2. [Creating a Security Lake subscriber](#)

To consume logs and events from Amazon Security Lake, you must be a Security Lake subscriber. Follow these steps to grant query access to a Detective account administrator.

3. Adding the required AWS Identity and Access Management (IAM) permissions to your IAM identity.

- Add these permissions to create Detective integration with Security Lake:
 - Attach these AWS Identity and Access Management (IAM) permissions to your IAM identity. For details, see the [Add the required IAM permissions to your account](#) section.
 - Add this IAM policy to the IAM principal that you plan to use to pass the AWS CloudFormation service role. For more details, see the [Add permissions to your IAM principal](#) section.
- If you have already integrated Detective with Security Lake, to use the integration attach these (IAM) permissions to your IAM identity. For details, see the [Add the required IAM permissions to your account](#) section.

4. [Accepting the Resource Share ARN invitation and enable the integration](#)

Use the AWS CloudFormation template to set up the parameters required to create and manage query access for Security Lake subscribers. For the detailed steps to create a stack, see [Create a stack using the AWS CloudFormation template](#). After you finish creating the stack, enable the integration.

For a demonstration of how to integrate Amazon Detective with Amazon Security Lake using the Detective console, watch the following video: [Amazon Detective integration with Amazon Security Lake- How to Setup](#)-->

Before you begin

This topic describes the preliminary steps such as delegating a Security Lake administrator for your organization, enabling Security Lake for your Detective administrator account, and verifying that Security Lake is collecting logs and events.

Security Lake integrates with AWS Organizations to manage log collection across multiple accounts in an organization. To use Security Lake for an organization, your AWS Organizations management account must first designate a delegated Security Lake administrator for your organization. The delegated Security Lake administrator must then enable Security Lake, and enable log and event collection for member accounts in the organization.

Before you integrate Security Lake with Detective, make sure that Security Lake is enabled for the Detective administrator account. You must first configure your data lake settings and set up log collection by enabling Security Lake using the Security Lake console. For the detailed steps on how to enable Security Lake, see [Getting Started](#) in the Amazon Security Lake User Guide.

Also, verify that Security Lake is collecting logs and events from AWS CloudTrail management events and Amazon Virtual Private Cloud (Amazon VPC) Flow Logs. For more details about log collection in Security Lake, see [Collecting data from AWS services](#) in the Amazon Security Lake User Guide.

Step 1: Creating a Security Lake subscriber

This topic explains how to use the Detective console to create a Security Lake subscriber.

To consume logs and events from Amazon Security Lake, you must be a Security Lake subscriber. A Subscriber can query and access the data that Security Lake collects. A subscriber with query access can query AWS Lake Formation tables directly in an Amazon Simple Storage Service (Amazon S3) bucket by using services such as Amazon Athena. To become a subscriber, the Security Lake administrator has to provide you with subscriber access that lets you query the data lake. For information about how the administrator does this, see [Creating a subscriber with query access](#) in the Amazon Security Lake User Guide.

Follow these steps to create a Security Lake subscriber in order to grant query access to a Detective administrator account.

To create a Detective subscriber in Security Lake

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Integrations**.
3. In the Security Lake subscriber pane, note the **Account ID** and **External ID** values.

Ask the Security Lake administrator to use these IDs to:

- To create a Detective subscriber for you in Security Lake.

- To configure the subscriber to have query access.
- To make sure that the Security Lake query subscriber is created with Lake Formation permissions, select **Lake Formation** as the **Data Access Method** in the Security Lake console.

When the Security Lake administrator creates a subscriber for you, Security Lake generates an Amazon Resource Share ARN for you. Ask the administrator to send this ARN to you.

4. Enter the **Resource Share ARN** that is provided by the Security Lake administrator in the **Security Lake subscriber** pane.
5. After you receive the Resource Share ARN from the Security Lake Administrator, enter the ARN in the **Resource Share ARN** box in the **Security Lake subscriber** pane.

Step 2: Adding the required IAM permissions to your account

This topic explains the details of the AWS Identity and Access Management (IAM) permissions policy that you must add to your IAM identity.

To enable Detective integration with Security Lake, you must attach the following AWS Identity and Access Management (IAM) permissions policy to your IAM identity.

Attach the following inline policies to the role. Replace `athena-results-bucket` with your Amazon S3 bucket name if you want to use your own Amazon S3 bucket to store the Athena query results. If you want Detective to automatically generate an Amazon S3 bucket to store the Athena query result, remove the entire `S3ObjectPermissions` from the IAM policy.

If you do not have the required permissions to attach this policy to your IAM identity, contact your AWS administrator. If you have the required permissions but an issue occurs, see [Troubleshoot access denied error messages](#) in the IAM User Guide.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetBucketLocation",
        "s3:ListAllMyBuckets"
      ],
      "Resource": "*"
    }
  ]
}
```

```

    },
    {
      "Sid": "S3ObjectPermissions",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:PutObject"
      ],
      "Resource": [
        "arn:aws:s3:::<athena-results-bucket>",
        "arn:aws:s3:::<athena-results-bucket>/*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "glue:GetDatabases",
        "glue:GetPartitions",
        "glue:GetTable",
        "glue:GetTables"
      ],
      "Resource": [
        "arn:aws:glue:*:<ACCOUNT ID>:database/amazon_security_lake*",
        "arn:aws:glue:*:<ACCOUNT ID>:table/amazon_security_lake*/
amazon_security_lake*",
        "arn:aws:glue:*:<ACCOUNT ID>:catalog"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "athena:BatchGetQueryExecution",
        "athena:GetQueryExecution",
        "athena:GetQueryResults",
        "athena:GetQueryRuntimeStatistics",
        "athena:GetWorkGroup",
        "athena:ListQueryExecutions",
        "athena:StartQueryExecution",
        "athena:StopQueryExecution",
        "lakeformation:GetDataAccess",
        "ram:ListResources"
      ],
      "Resource": "*"
    },
  },

```



```
{
  "Effect": "Allow",
  "Action": [
    "ssm:GetParametersByPath"
  ],
  "Resource": [
    "arn:aws:ssm:*:<ACCOUNT ID>:parameter/Detective/SLI/ResourceShareArn",
    "arn:aws:ssm:*:<ACCOUNT ID>:parameter/Detective/SLI/S3Bucket",
    "arn:aws:ssm:*:<ACCOUNT ID>:parameter/Detective/SLI/TableNames",
    "arn:aws:ssm:*:<ACCOUNT ID>:parameter/Detective/SLI/DatabaseName",
    "arn:aws:ssm:*:<ACCOUNT ID>:parameter/Detective/SLI/StackId"
  ]
},
{
  "Effect": "Allow",
  "Action": [
    "cloudformation:GetTemplateSummary",
    "iam:ListRoles"
  ],
  "Resource": "*"
},
{
  "Effect": "Allow",
  "Action": [
    "organizations:ListDelegatedAdministrators"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "organizations:ServicePrincipal": [
        "securitylake.amazonaws.com"
      ]
    }
  }
}
]
```

Step 3: Accept the Resource Share ARN invitation and enable the integration

This topic explains the steps to accept the Resource Share ARN invitation using a AWS CloudFormation template, which is a required step before you enable Detective integration with Security Lake.

To access raw data logs from Security Lake, you must accept a Resource Share invitation from the Security Lake account that was created by the Security Lake administrator. You also need AWS Lake Formation permissions to set up cross-account table sharing. In addition, you must create an Amazon Simple Storage Service (Amazon S3) bucket that can receive raw query logs.

In this next step, you'll use an AWS CloudFormation template to create a stack for: accepting the Resource Share ARN invitation, create required AWS Glue crawler resources, and grant AWS Lake Formation administrator permissions.

To accept the Resource Share ARN invitation and enable the integration

1. Create a new CloudFormation stack using the CloudFormation template. For more details, see [Creating a stack using the AWS CloudFormation template](#).
2. After you finish creating the stack, choose **Enable integration** to enable Detective integration with Security Lake.

Creating a stack using the AWS CloudFormation template

Detective provides an AWS CloudFormation template, which you can use to set up the parameters required to create and manage query access for Security Lake subscribers.

Step 1: Create an AWS CloudFormation service role

You must create an AWS CloudFormation service role to create a stack using the AWS CloudFormation template. If you do not have the required permissions to create a service role, contact the administrator of the Detective administrator account. For more information about the AWS CloudFormation service role, see [AWS CloudFormation service role](#).

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane of the IAM console, choose **Roles**, and then choose **Create role**.

3. For **Select trusted entity**, choose **AWS service**.
4. Choose **AWS CloudFormation**. Then, choose **Next**.
5. Enter a name for the role. For example, CFN-DetectiveSecurityLakeIntegration.
6. Attach the following inline policies to the role. Replace <Account ID> with your AWS Account ID.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "CloudFormationPermission",
      "Effect": "Allow",
      "Action": [
        "cloudformation:CreateChangeSet"
      ],
      "Resource": [
        "arn:aws:cloudformation:*:aws:transform/*"
      ]
    },
    {
      "Sid": "IamPermissions",
      "Effect": "Allow",
      "Action": [
        "iam:CreateRole",
        "iam>DeleteRole",
        "iam:AttachRolePolicy",
        "iam:DetachRolePolicy",
        "iam:UpdateAssumeRolePolicy",
        "iam:PutRolePolicy",
        "iam>DeleteRolePolicy",
        "iam>CreatePolicy",
        "iam>DeletePolicy",
        "iam:PassRole",
        "iam:GetRole",
        "iam:GetRolePolicy"
      ],
      "Resource": [
        "arn:aws:iam::<ACCOUNT ID>:role/*",
        "arn:aws:iam::<ACCOUNT ID>:policy/*"
      ]
    }
  ],
}
```

```

{
  "Sid": "S3Permissions",
  "Effect": "Allow",
  "Action": [
    "s3:CreateBucket",
    "s3:DeleteBucket*",
    "s3:PutBucket*",
    "s3:GetBucket*",
    "s3:GetObject",
    "s3:PutEncryptionConfiguration",
    "s3:GetEncryptionConfiguration"
  ],
  "Resource": [
    "arn:aws:s3:::*"
  ]
},
{
  "Sid": "LambdaPermissions",
  "Effect": "Allow",
  "Action": [
    "lambda:CreateFunction",
    "lambda:DeleteFunction",
    "lambda:GetFunction",
    "lambda:TagResource",
    "lambda:InvokeFunction"
  ],
  "Resource": [
    "arn:aws:lambda:*:<ACCOUNT ID>:function:*"
  ]
},
{
  "Sid": "CloudwatchPermissions",
  "Effect": "Allow",
  "Action": [
    "logs:CreateLogGroup",
    "logs>DeleteLogGroup",
    "logs:DescribeLogGroups"
  ],
  "Resource": "arn:aws:logs:*:<ACCOUNT ID>:log-group:*"
},
{
  "Sid": "KmsPermission",
  "Effect": "Allow",
  "Action": [

```

```

        "kms:Decrypt"
    ],
    "Resource": "arn:aws:kms:*:<ACCOUNT ID>:key/*"
}
]
}

```

Step 2: Add permissions to your IAM principal.

You'll need the following permissions to create a stack using the CloudFormation service role that you created in the preceding step. Add the following IAM policy to the IAM principal that you plan to use to pass the CloudFormation service role. You will assume this IAM principal to create the stack. If you do not have the required permissions to add the IAM policy, contact the administrator of the Detective administrator account.

Note

In the following policy, `CFN-DetectiveSecurityLakeIntegration` used in this policy refers to the role that you created in the previous [Creating an AWS CloudFormation service role](#) step. Change it to the role name that you entered in the preceding step if it's different.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "PassRole",
      "Effect": "Allow",
      "Action": [
        "iam:GetRole",
        "iam:PassRole"
      ],
      "Resource": "arn:aws:iam::<ACCOUNT ID>:role/CFN-
DetectiveSecurityLakeIntegration"
    },
    {
      "Sid": "RestrictCloudFormationAccess",
      "Effect": "Allow",

```

```

    "Action": [
      "cloudformation:CreateStack",
      "cloudformation>DeleteStack",
      "cloudformation:UpdateStack"
    ],
    "Resource": "arn:aws:cloudformation:*:<ACCOUNT ID>:stack/*",
    "Condition": {
      "StringEquals": {
        "cloudformation:RoleArn": [
          "arn:aws:iam:*:<ACCOUNT ID>:role/CFN-
DetectiveSecurityLakeIntegration"
        ]
      }
    }
  },
  {
    "Sid": "CloudformationDescribeStack",
    "Effect": "Allow",
    "Action": [
      "cloudformation:DescribeStacks",
      "cloudformation:DescribeStackEvents",
      "cloudformation:GetStackPolicy"
    ],
    "Resource": "arn:aws:cloudformation:*:<ACCOUNT ID>:stack/*"
  },
  {
    "Sid": "CloudformationListStacks",
    "Effect": "Allow",
    "Action": [
      "cloudformation:ListStacks"
    ],
    "Resource": "*"
  },
  {
    "Sid": "CloudWatchPermissions",
    "Effect": "Allow",
    "Action": [
      "logs:GetLogEvents"
    ],
    "Resource": "arn:aws:logs:*:<ACCOUNT ID>:log-group:*"
  }
]
}

```

Step 3: Specify custom values in the AWS CloudFormation console

1. Go to the AWS CloudFormation console from Detective.
2. (Optional) Enter a **Stack name**. The stack name is auto-filled. You can change the stack name to a name that does not conflict with existing stack names.
3. Enter the following **Parameters**.
 - **AthenaResultsBucket** – If you don't enter values, this template generates an Amazon S3 bucket. If you want to use your own bucket, enter a bucket name to store the Athena query results. If you use your own bucket, make sure that the bucket is in the same Region as the Resource Share ARN. If you use your own bucket, make sure the LakeFormationPrincipals you choose have permissions to write objects to and read objects from the bucket. For more details about bucket permissions, see [Query results and recent queries](#) in the Amazon Athena User Guide.
 - **DTRegion** – This field is pre-filled. Do not change the values in this field.
 - **LakeFormationPrincipals** – Enter the ARN of the IAM principals (for example, IAM role ARN) that you want to grant access to use the Security Lake integration, separated by commas. These could be your security analysts and security engineers that use Detective.

You can only use the IAM principals that you previously attached the IAM permissions to in step [Step 2: Add the required IAM permissions to your account].

- **ResourceShareARN** – This field is pre-filled. Do not change the values in this field.
4. **Permissions**

IAM role – Select the role that you created in the Creating an AWS CloudFormation Service Role step. Optionally, you can keep it blank if your current IAM role has all the required permissions in the Creating an AWS CloudFormation Service Role step.
 5. Review and check all the **I Acknowledge** boxes and then click the **Create stack** button. For more details, review the following IAM resources that will be created.

- * ResourceShareAcceptorCustomResourceFunction
 - ResourceShareAcceptorLambdaRole
 - ResourceShareAcceptorLogsAccessPolicy
- * SsmParametersCustomResourceFunction
 - SsmParametersLambdaRole
 - SsmParametersLogsAccessPolicy
- * GlueDatabaseCustomResourceFunction

- GlueDatabaseLambdaRole
- GlueDatabaseLogsAccessPolicy
- * GlueTablesCustomResourceFunction
- GlueTablesLambdaRole
- GlueTablesLogsAccessPolicy

Step 4: Add Amazon S3 bucket policy to IAM principals in LakeFormationPrincipals

(Optional) If you let this template generate an AthenaResultsBucket for you, you must attach the following policy to the IAM principals in LakeFormationPrincipals.

```
{
  "Sid": "S3ObjectPermissions",
  "Effect": "Allow",
  "Action": [
    "s3:GetObject",
    "s3:PutObject"
  ],
  "Resource": [
    "arn:aws:s3:::<athena-results-bucket>",
    "arn:aws:s3:::<athena-results-bucket>/*"
  ]
}
```

Replace `athena-results-bucket` with the AthenaResultsBucket name. The AthenaResultsBucket can be found on the AWS CloudFormation console:

1. Open the AWS CloudFormation console at <https://console.aws.amazon.com/cloudformation>.
2. Click on your Stack.
3. Click the **Resources** tab.
4. Search for the logical ID AthenaResultsBucket and copy its physical ID.

Changing the integration configuration

If you want to change any of the parameters that you used to integrate Detective with Security Lake, you can edit them, and then enable the integration again. You can edit the AWS CloudFormation template to re-enable this integration for the following scenarios:

- To update the Security Lake subscription, you can either create a new subscriber, or the Security Lake administrator can update the data source for the existing subscription.
- To specify a different Amazon S3 bucket to store the raw query logs.
- To specify different Lake Formation principals.

When you re-enable Detective integration with Security Lake, you can edit the **Resource Share ARN**, and view the **IAM permissions**. To edit the IAM permissions, you can go to the IAM console from Detective. You can also edit the values you previously entered in the AWS CloudFormation template. You must delete the existing CloudFormation stack and re-create it to re-enable the integration.

To re-enable Detective integration with Security Lake

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Integrations**.
3. You can edit the integration using either of these steps:
 - In the **Security Lake** pane, choose **Edit**.
 - In the **Security Lake** pane, choose **View**. In the view page, choose **Edit**.
4. Enter a new **Resource Share ARN**, to access the data sources in a Region.
5. View the current IAM permissions, and go to the IAM console, if you want to edit the IAM permissions.
6. Edit the values in the CloudFormation template.
 1. Delete the existing stack first, before creating a new stack. If you do not delete the existing stack and you try to create a new stack in the same Region, your request fails. For more details, see [Deleting a CloudFormation stack](#).
 1. Create a new CloudFormation stack. For more details, see [Creating a stack using the AWS CloudFormation template](#).
7. Choose **Enable integration**.

Supported AWS Regions

You can integrate Detective with Security Lake in the following AWS Regions.

Region Name	Region	Endpoint	Protocol;
US East (Ohio)	us-east-2	securitylake.us-east-2.amazonaws.com	HTTPS
US East (N. Virginia)	us-east-1	securitylake.us-east-1.amazonaws.com	HTTPS
US West (N. California)	us-west-1	securitylake.us-west-1.amazonaws.com	HTTPS
US West (Oregon)	us-west-2	securitylake.us-west-2.amazonaws.com	HTTPS
Asia Pacific (Mumbai)	ap-south-1	securitylake.ap-south-1.amazonaws.com	HTTPS
Asia Pacific (Seoul)	ap-northeast-2	securitylake.ap-northeast-2.amazonaws.com	HTTPS
Asia Pacific (Singapore)	ap-southeast-1	securitylake.ap-southeast-1.amazonaws.com	HTTPS
Asia Pacific (Sydney)	ap-southeast-2	securitylake.ap-southeast-2.amazonaws.com	HTTPS
Asia Pacific (Tokyo)	ap-northeast-1	securitylake.ap-northeast-1.amazonaws.com	HTTPS
Canada (Central)	ca-central-1	securitylake.ca-central-1.amazonaws.com	HTTPS
Europe (Frankfurt)	eu-central-1	securitylake.eu-central-1.amazonaws.com	HTTPS
Europe (Ireland)	eu-west-1	securitylake.eu-west-1.amazonaws.com	HTTPS

Region Name	Region	Endpoint	Protocol;
Europe (London)	eu-west-2	securitylake.eu-west-2.amazonaws.com	HTTPS
Europe (Paris)	eu-west-3	securitylake.eu-west-3.amazonaws.com	HTTPS
Europe (Stockholm)	eu-north-1	securitylake.eu-north-1.amazonaws.com	HTTPS
South America (São Paulo)	sa-east-1	securitylake.sa-east-1.amazonaws.com	HTTPS

Querying raw logs in Detective

After you integrate Detective with Security Lake, Detective begins pulling raw logs from Security Lake related to AWS CloudTrail management events and Amazon Virtual Private Cloud (Amazon VPC) Flow Logs.

Note

There are no additional charges to query raw logs in Detective. Usage charges for other AWS Services, including Amazon Athena, still apply at published rates.

AWS CloudTrail management events are available for the following profiles:

- AWS account
- AWS user
- AWS role
- AWS role Session
- Amazon EC2 instance
- Amazon S3 bucket
- IP address
- Kubernetes cluster

- Kubernetes pod
- Kubernetes subject
- IAM role
- IAM role session
- IAM user

Amazon VPC FLOW Logs are available for the following profiles:

- Amazon EC2 instance
- Kubernetes pod

For a demonstration of how to integrate Amazon Detective with Amazon Security Lake using the Detective console, watch the following video: [Amazon Detective integration with Amazon Security Lake- How to Use](#)-->

To query raw logs for an AWS account

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Search** and search for an AWS account.
3. In the **Overall API call volume** section, choose **display details for scope time**.
4. From here, you can start to **Query raw logs**.

Detective > Search > AwsAccount/714603721603

714603721603
AWS account [Info](#)

Scope time [Info](#)
12/21/2023 18:00 UTC > 12/22/2023 18:00 UTC

Activity for time window: 12/21/2023 18:00 UTC - 12/22/2023 18:00 UTC [✎](#)

[Query raw logs](#)

[Observed IP addresses](#) | [API method by service](#) | [Resource](#)

< 1 >

IP address ▾	Successful calls ▾	Failed calls ▾	Location ▾	Actions
▶ [redacted]	6	2	[redacted]	
▶ [redacted]	2	1	-	
▶ [redacted]	1	0	[redacted]	

In the **Raw log preview** table, you can view the logs and events retrieved by querying data from Security Lake. For more details about the raw event logs, you can view the data displayed in Amazon Athena.

Raw log preview: CloudTrail ✕

View raw event logs that were retrieved by querying data from Security Lake. For more details about the raw event logs, you can view the data displayed in Athena.

Raw log preview (500+)							< 1 2 3 4 5 6 7 ... 50 >
date_time ▾	requestor_arn ▾	account_id ▾	region ▾	source_ip ▾	service ▾	apiL	
2023-12-22 09:58:38.000 UTC	[redacted]	[redacted]	us-east-1	[redacted]	s3.amazonaws.com	GetF	
2023-12-22 09:59:49.000 UTC	[redacted]	[redacted]	us-east-1	[redacted]	sts.amazonaws.com	Assu	
2023-12-22 10:00:13.000 UTC	[redacted]	[redacted]	us-east-1	[redacted]	ec2.amazonaws.com	Desc	
2023-12-22 10:00:13.000 UTC	[redacted]	[redacted]	us-east-1	[redacted]	sts.amazonaws.com	Assu	
2023-12-22 10:00:13.000 UTC	[redacted]	[redacted]	us-east-1	[redacted]	iam.amazonaws.com	GetI	
2023-12-22 10:00:13.000 UTC	[redacted]	[redacted]	us-east-1	[redacted]	sts.amazonaws.com	Assu	
2023-12-22 10:00:13.000 UTC	[redacted]	[redacted]	us-east-1	[redacted]	sts.amazonaws.com	GetC	
2023-12-22 10:00:13.000 UTC	[redacted]	[redacted]	us-east-1	[redacted]	autoscaling.amazonaws.com	Desc	
2023-12-22 10:00:14.000 UTC	[redacted]	[redacted]	us-east-1	[redacted]	ec2.amazonaws.com	Desc	
2023-12-22 10:00:14.000 UTC	[redacted]	[redacted]	us-east-1	[redacted]	ec2.amazonaws.com	Desc	

[Close](#)
[Cancel query request](#)
[See results in Athena \[↗\]\(#\)](#)
[Download results](#)

From the Query raw logs table, you can **Cancel query request**, **See results in Amazon Athena**, and **Download results** as a comma-separated values (.csv) file.

If you see logs in Detective, but the query returned no results, it could happen because of the following reasons.

- Raw logs may become available in Detective before showing up in Security Lake log tables. Try again later.
- Logs may be missing from Security Lake. If you waited for an extended period of time, it indicates that logs are missing from Security Lake. Contact your Security Lake administrator to resolve the issue.

Examples


- [Query raw logs for an AWS role](#)
- [Query raw logs for an Amazon EKS cluster](#)
- [Query raw logs for an Amazon EC2 instance](#)

Query raw logs for an AWS role

If you want to understand the activity of an AWS role in a new geolocation, you can do so within the Detective console.

To query raw logs for an AWS role

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. From the Detective **Summary** page **Newly observed geolocations** section, note down the AWS role.
3. In the navigation pane, choose **Search** and search for the AWS role.
4. For the AWS role, expand the resource to display the specific API calls that were issued from that IP address by that resource.
5. Choose the magnifier icon next to the API call that you want to investigate to open the **Raw log preview** table.

Activity for time window: 

Q Query raw logs

Observed IP addresses | **API method by service** | Resource

< 1 >

IP address ▾	Successful calls ▾	Failed calls ▾	Location ▾	Actions
▶ <input type="text"/>	289	284	-	
▶ <input type="text"/>	63	0	<input type="text"/>	
▶ <input type="text"/>	42	0	<input type="text"/>	
▶ <input type="text"/>	21	0	<input type="text"/>	

Query raw logs for an Amazon EKS cluster

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. From the Detective **Summary** page **Container clusters with the most pods created** section, navigate to an Amazon EKS cluster.
3. In the **Amazon EKS cluster details** page, select the **Kubernetes API activity** tab.
4. In the **Overall Kubernetes API activity involving this Amazon EKS cluster** section, choose **display details for scope time**.
5. From here, you can start to **Query raw logs**.

Query raw logs for an Amazon EC2 instance

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Search** and search for an Amazon EC2 instance.
3. In the **Overall VPC Flow volume** section, choose the magnifier icon next to the API call that you want to investigate to open the **Raw log preview** table.
4. From here, you can start to **Query raw logs**.

Activity for time window: 11/21/2023 11:00 (UTC-08:00) - 11/22/2023 11:00 (UTC-08:00) ↗

Toggle overall traffic

Query raw logs

<input type="checkbox"/>	IP address ▾	Local port ▾	Remote port ▾	Inbound traffic ▾	Outbound traffic ▾	Protocol ▾	Directionality ▾	Accept / Reject ▾	Actions
<input type="checkbox"/>		22	-	44.7 kB	57.7 kB	TCP	Inbound	Accept	<input type="checkbox"/>
<input type="checkbox"/>		22	-	240 B	480 B	TCP	Inbound	Accept	<input type="checkbox"/>
<input type="checkbox"/>		22	-	61.1 kB	75 kB	TCP	Inbound	Accept	<input type="checkbox"/>
<input type="checkbox"/>		22	-	59.6 kB	70.8 kB	TCP	Inbound	Accept	<input type="checkbox"/>
<input type="checkbox"/>		22	-	240 B	540 B	TCP	Inbound	Accept	<input type="checkbox"/>

In the **Raw log preview** table, you can view the logs and events retrieved by querying data from Security Lake. For more details about the raw event logs, you can view the data displayed in Amazon Athena.

From the Query raw logs table, you can **Cancel query request**, **See results in Amazon Athena**, and **Download results** as a comma-separated values (.csv) file.

Disabling the integration

If you disable Detective integration with Security Lake, you can no longer query log and event data from Security Lake.

To disable Detective integration with Security Lake

1. Open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, choose **Integrations**.
3. Delete the existing stack. For more details, see [Deleting a CloudFormation stack](#).
4. In the **Disable Security Lake integration** pane, choose **Disable**.

Deleting a CloudFormation stack

If you do not delete the existing stack, new stack creation in the same Region will fail. You can delete a CloudFormation stack by using the CloudFormation console or use the AWS CLI.

To delete the AWS CloudFormation stack (Console)

1. Open the AWS CloudFormation console at <https://console.aws.amazon.com/cloudformation>.

2. On the **Stacks** page in the CloudFormation console, select the stack that you want to delete. The stack must be currently running.
3. In the stack details pane, choose **Delete**.
4. Select **Delete stack** when prompted.

Note

The stack deletion operation can't be stopped once the stack deletion has begun. The stack proceeds to the DELETE_IN_PROGRESS state.

After the stack deletion is complete, the stack will be in the DELETE_COMPLETE state.

Troubleshooting stack deletion errors

If you are seeing a permission error with the message `Failed to delete stack` after clicking the `Delete` button, your IAM role doesn't have CloudFormation permission to delete a stack. Contact your account administrator to delete the stack.

To delete the CloudFormation stack (AWS CLI)

Enter the following command in the AWS CLI interface:

```
aws cloudformation delete-stack --stack-name your-stack-name --role-arn
arn:aws:iam::<ACCOUNT ID>:role/CFN-DetectiveSecurityLakeIntegration
```

CFN-DetectiveSecurityLakeIntegration is the service role that you created in the `Creating an AWS CloudFormation Service Role` step.

Forecasting and monitoring Detective costs

To help you forecast and monitor your costs for using Amazon Detective, Detective calculates and provides estimated usage costs for your account. With this data, you can determine whether to adjust your use of the service or your account quotas. Amazon Detective provides a 30-day free trial for each account in each Region. The free trial for an account starts the first time one of the following actions occurs.

Detective bills each account for the data used in each behavior graph that the account belongs to. Detective charges a tiered flat rate per GB for all data regardless of the source.

Note

Detective pricing is based on uncompressed byte count for usage and cost.

To help you to track your Detective activity, the **Usage** page of the Detective console shows the amount of data ingested and the projected cost.

- For administrator accounts – Shows the data volume and projected cost across the entire behavior graph.
- For member accounts – Shows the data volume and projected cost for their account across the behavior graphs that they contribute to.

Detective also supports AWS CloudTrail logging. CloudTrail captures all API calls for Detective as events. The calls captured include calls from the Detective console and code calls to the Detective API operations. For more details about AWS CloudTrail logging, see [Logging Detective API calls with CloudTrail](#).

In addition to the estimated usage costs that Detective provides, you can review and monitor your actual costs by using AWS Billing and Cost Management. AWS Billing and Cost Management provides features that are designed to help you track and analyze your costs for AWS services, and manage budgets for your account or organization. It also provides features that can help you forecast usage costs based on historical data. To learn more, see the [AWS Billing User Guide](#).

Contents

- [About the free trial for behavior graphs](#)

- [Monitoring usage for a Detective administrator account](#)
- [Monitoring usage for a Detective member account](#)
- [How Amazon Detective calculates projected cost](#)

About the free trial for behavior graphs

Amazon Detective provides a 30-day free trial for each account in each Region. The free trial for an account starts the first time one of the following actions occurs.

- An account enables Detective manually and becomes the administrator account for a behavior graph.
- An account is designated as the Detective administrator account for an organization in AWS Organizations, and has Detective enabled for the first time.
- If the Detective administrator account already had Detective enabled before they were designated, then the account does not start a new 30-day free trial.
- An account accepts an invitation to be a member account in a behavior graph and is enabled as a member account.
- An organization account is enabled as a member account by the Detective administrator account.

The free trial lasts for 30 days from that point. The account is not billed for any data processed during that period. When the trial period ends, Detective begins to bill the account for the data it contributes to behavior graphs. For more information about how you can track your Detective activity, monitor usage and view the projected cost see [Forecasting and monitoring Detective costs](#). For more information on pricing, see [Detective pricing](#)

The same 30-day period is used for all behavior graphs in the Region. For example, an account is enabled as a member account for a behavior graph. This starts the 30-day free trial. After 10 days, the account is enabled for a second behavior graph in the same Region. For the second behavior graph, the account receives 20 days of free data.

The free trial provides multiple benefits:

- Administrator accounts can explore Detective features and functionality to verify its value.
- Administrator and member accounts can monitor the amount of data and the estimated cost before Detective begins to bill them for it. See [the section called "Administrator account usage and cost"](#) and [the section called "Member account usage tracking"](#).

Free trial for optional data sources

Detective also provides a free 30-day trial for optional data sources. This free trial is separate from the free trial provided for the core Detective data sources when Detective is first enabled.

Note

If a customer disables an optional data source package within 7 days of enabling it, Detective does a one-time automatic reset of the free trial for that data source package if it is enabled again.

To enable or disable an optional data source see [Types of optional data sources in Detective](#).

Monitoring usage for a Detective administrator account

Amazon Detective bills each account for the data used in each behavior graph that the account belongs to. Detective charges a tiered flat rate per GB for all data regardless of the source.

For administrator accounts, the **Usage** page of the Detective console allows you to view the volume of data ingested **By data source** or **By account** over the previous 30 days. Administrator accounts also see a projected cost for a typical 30-day period for their account and for the entire behavior graph.

To view Detective usage information

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, under **Settings**, choose **Usage**.
3. Choose a tab to select between viewing usage **By data source** or **By account**.

Volume of data ingested for each account

Ingested volume by member account lists the active accounts in the behavior graph. It does not list member accounts that were removed.

For each account, the ingested volume list provides the following information.

- The AWS account identifier and root user email address.

- The date when the account began to contribute data to the behavior graph.

For the administrator account, this is the date when the account enabled Detective.

For member accounts, this is the date when an account was enabled as a member account after accepting the invitation.

- The volume of ingested data from the account over the previous 30 days. The total includes all source types.
- Whether the account is currently in the free trial period. For accounts that are currently in their free trial period, the list displays the number of days remaining.

If none of the accounts are in the free trial period, then the free trial status column is not displayed.

Projected costs for the behavior graph

This account's projected cost shows a projected cost for 30 days of data for the administrator account. The projected cost is based on the daily average volume for the administrator account.

Important

This amount is a projected cost only. It projects the total cost for the administrator account data for a typical 30-day time period. It is based on the usage from the previous 30 days. See [the section called "How Detective calculates projected cost"](#).

Projected cost for the behavior graph

All accounts' projected cost shows a total projected cost for 30 days of data for the entire behavior graph. The projected cost is based on the daily average volume for each account.

Important

This amount is a projected cost only. It projects the total cost for the behavior graph data for a typical 30-day time period. It is based on the usage from the previous 30 days. The projected cost does not include member accounts that were removed from the behavior graph. See [the section called "How Detective calculates projected cost"](#).

Volume of data ingested by source packages

Select **By source package** to view the volume of data ingested listed by the different source packages enabled in your behavior graph.

All accounts can view this data for their own accounts. An administrator account can see additional panels that list the usage by source package for each member. It does not list member accounts that were removed.

Detective core

Detective core panels show the volume of data ingested from Detective core sources (CloudTrail logs, VPC Flow logs, and GuardDuty findings) for the last 30 days.

EKS audit logs

EKS audit logs panels show the volume of data ingested from EKS audit logs sources for the last 30 days. Panels for this source package are only available if EKS audit logs is enabled for your behavior graph.

Monitoring usage for a Detective member account

Amazon Detective bills each account for the data used in each behavior graph that the account belongs to. Detective charges a tiered flat rate per GB for all data regardless of the source.

For member accounts, the **Usage** page shows the volume of data and projected 30-day cost for that account only.

To view Detective usage information

1. Sign in to the AWS Management Console. Then open the Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, under **Settings**, choose **Usage**.

Ingested volume for each behavior graph

This account's ingested volume lists the behavior graphs that the member account contributes to. It does not include memberships that you resigned, or memberships that the administrator account removed.

For each behavior graph, the list includes the following information.

- The account number of the administrator account
- The volume of ingested data from the member account over the previous 30 days. The total includes all source types.
- The date when the member account was enabled for the behavior graph.

Projected cost across behavior graphs

This account's projected cost shows a projected cost for 30 days of data for the member account across all of the behavior graphs that it contributes to. The projected cost is based on the daily average volume for the member account.

Important

This amount is a projected cost only. It projects the total cost for the administrator account data for a typical 30-day time period. It is based on the usage from the previous 30 days. See [the section called "How Detective calculates projected cost"](#).

How Amazon Detective calculates projected cost

To calculate the projected cost values that it displays on the **Usage** page, Detective does the following.

1. To get the projected cost for an individual account in a behavior graph, Detective does the following.
 - a. Calculates the average volume per day. It adds the data volume across all of the active days and then divides by the number of days that the account has been active.

If the account was enabled more than 30 days ago, then the number of days is 30. If the account was enabled fewer than 30 days ago, then it is the number of days since the acceptance date.

For example, if the account was enabled 12 days ago, then Detective adds the volume ingested for those 12 days and then divides it by 12.

- b. Multiplies the account's daily average by 30. This is the projected 30-day usage for the account.
 - c. Uses its pricing model to calculate the projected 30-day cost for the projected 30-day usage.
 2. To get the total projected cost for a behavior graph, Detective does the following:
 - a. Combines the projected 30-day usage from all of the accounts in the behavior graph.
 - b. Uses its pricing model to calculate the projected 30-day cost for the total projected 30-day usage.
 3. To get the total projected cost for a member account across behavior graphs, Detective does the following:
 - a. Combines the projected 30-day usage across all of the behavior graphs.
 - b. Uses its pricing model to calculate the projected 30-day cost for the total projected 30-day usage.
 4. If you are using a shared Amazon VPC, Detective calculates the projected cost based on monitoring activity. We recommend that you review the projected cost for your investigations specific to your environment.
 - a. If a Detective member account has a shared Amazon VPC and there are other non-Detective accounts using the shared VPC, Detective will monitor all traffic from that VPC. The usage and cost will increase and Detective will provide visualization on all the traffic flow within the VPC.
 - b. If you have an EC2 instance inside a shared Amazon VPC and the shared owner is not a Detective member, Detective will not monitor any traffic from the VPC, and the usage and cost will decrease. If you want to view the traffic flow within the VPC, you must add the Amazon VPC owner as a member of your Detective graph.

Security in Amazon Detective

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely.

Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS compliance programs](#).

To learn about the compliance programs that apply to Amazon Detective, see [AWS Services in Scope by Compliance Program](#).

- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Detective. The following topics show you how to configure Detective to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Detective resources.

Contents

- [Data protection in Amazon Detective](#)
- [Identity and access management for Amazon Detective](#)
- [Compliance validation for Amazon Detective](#)
- [Resilience in Amazon Detective](#)
- [Infrastructure security in Amazon Detective](#)
- [Security best practices for Detective](#)

Data protection in Amazon Detective

The AWS [shared responsibility model](#) applies to data protection in Amazon Detective. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see [Working with CloudTrail trails](#) in the *AWS CloudTrail User Guide*.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-3](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Detective or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Detective encrypts all data that it processes and stores at rest and in transit.

Contents

- [Key management for Amazon Detective](#)

Key management for Amazon Detective

Because Detective does not store any personally identifiable customer data, it uses AWS managed keys.

This type of KMS key can be used across multiple accounts. See the [description of AWS owned keys in the AWS Key Management Service Developer Guide](#).

This type of KMS key rotates automatically every one year (approximately 365 days). See the [description of key rotation in the AWS Key Management Service Developer Guide](#).

Identity and access management for Amazon Detective

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Detective resources. IAM is an AWS service that you can use with no additional charge.

Contents

- [Audience](#)
- [Authenticating With Identities](#)
- [Managing Access Using Policies](#)
- [How Amazon Detective works with IAM](#)
- [Amazon Detective identity-based policy examples](#)
- [AWS managed policies for Amazon Detective](#)
- [Using service-linked roles for Detective](#)
- [Troubleshooting Amazon Detective identity and access](#)

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Detective.

Service user – If you use the Detective service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Detective features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Detective, see [Troubleshooting Amazon Detective identity and access](#).

Service administrator – If you're in charge of Detective resources at your company, you probably have full access to Detective. It's your job to determine which Detective features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Detective, see [How Amazon Detective works with IAM](#).

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Detective. To view example Detective identity-based policies that you can use in IAM, see [Amazon Detective identity-based policy examples](#).

Authenticating With Identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see [Signing AWS API requests](#) in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see [Multi-factor authentication](#) in the *AWS IAM Identity Center User Guide* and [Using multi-factor authentication \(MFA\) in AWS](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

IAM Users and Groups

An [IAM user](#) is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see [Rotate access keys regularly for use cases that require long-term credentials](#) in the *IAM User Guide*.

An [IAM group](#) is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see [When to create an IAM user \(instead of a role\)](#) in the *IAM User Guide*.

IAM Roles

An [IAM role](#) is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in the AWS Management Console by [switching roles](#). You can assume a role by calling an AWS CLI or

AWS API operation or by using a custom URL. For more information about methods for using roles, see [Using IAM roles](#) in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see [Creating a role for a third-party Identity Provider](#) in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the permission set to a role in IAM. For information about permissions sets, see [Permission sets](#) in the *AWS IAM Identity Center User Guide*.
- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.
- **Cross-service access** – Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
- **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).
- **Service role** – A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Creating a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.

- **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see [When to create an IAM role \(instead of a user\)](#) in the *IAM User Guide*.

Managing Access Using Policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-Based Policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Creating IAM policies](#) in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see [Choosing between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-Based Policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are *IAM role trust policies* and *Amazon S3 bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access Control Lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other Policy Types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see [Service control policies](#) in the *AWS Organizations User Guide*.
- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple Policy Types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How Amazon Detective works with IAM

By default, users and roles don't have permission to create or modify Amazon Detective resources. They also can't perform tasks using the AWS Management Console, AWS CLI, or AWS API. A Detective administrator must have AWS Identity and Access Management (IAM) policies that grant IAM users and roles permission to perform specific API operations on the specified resources

they need. The administrator must then attach those policies to the principal that require those permissions.

Detective uses IAM identity-based policies to grant permissions for the following types of users and actions:

- **Administrator accounts** – The administrator account is the owner of a behavior graph, which uses data from their account. Administrator accounts can invite member accounts to contribute their data to the behavior graph. The administrator account can also use the behavior graph for triage and investigation of findings and resources associated with those accounts.

You can set up policies to allow users other than the administrator account to perform different types of tasks. For example, a user from an administrator account might only have permissions to manage member accounts. Another user might only have permissions to use the behavior graph for investigation.

- **Member accounts** – A member account is an account that is invited to contribute data to a behavior graph. A member account responds to an invitation. After accepting an invitation, a member account can remove their account from the behavior graph.

To get a high-level view of how Detective and other AWS services work with IAM, see [Creating policies on the JSON tab](#) in the *IAM User Guide*.

Detective identity-based policies

With IAM identity-based policies, you can specify allowed or denied actions and resources, as well as the conditions under which actions are allowed or denied. Detective supports specific actions, resources, and condition keys.

To learn about all of the elements that you use in a JSON policy, see [IAM JSON Policy Elements Reference](#) in the *IAM User Guide*.

Actions

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The **Action** element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API

operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

Policy statements must include either an `Action` element or a `NotAction` element. The `Action` element lists the actions allowed by the policy. The `NotAction` element lists the actions that are not allowed.

The actions defined for Detective reflect tasks that you can perform using Detective. Policy actions in Detective have the following prefix: `detective:`.

For example, to grant permission to use the `CreateMembers` API operation to invite member accounts to a behavior graph, you include the `detective:CreateMembers` action in their policy.

To specify multiple actions in a single statement, separate them with commas. For example, for a member account, the policy includes the set of actions related to managing an invitation:

```
"Action": [
    "detective:ListInvitations",
    "detective:AcceptInvitation",
    "detective:RejectInvitation",
    "detective:DisassociateMembership"
]
```

You can also use wildcards (*) to specify multiple actions. For example, to manage the data used in their behavior graph, administrator accounts in Detective must be able to perform the following tasks:

- View their list of member accounts (`ListMembers`).
- Get information about selected member accounts (`GetMembers`).
- Invite member accounts to their behavior graph (`CreateMembers`).
- Remove members from their behavior graph (`DeleteMembers`).

Instead of listing these actions separately, you can grant access to all actions that end with the word `Members`. The policy for that could include the following action:

```
"Action": "detective:*Members"
```

To see a list of Detective actions, see [Actions defined by Amazon Detective](#) in the *Service Authorization Reference*.

Resources

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

For more information about the format of ARNs, see [Amazon Resource Names \(ARNs\) and AWS Service Namespaces](#).

For Detective, the only resource type is the behavior graph. The behavior graph resource in Detective has the following ARN:

```
arn:aws:detective:${Region}:${AccountId}:graph:${GraphId}
```

For example, a behavior graph has the following values:

- The Region for the behavior graph is us-east-1.
- The account ID for the administrator account ID is 111122223333.
- The graph ID of the behavior graph is 027c7c4610ea4aacaf0b883093cab899.

To identify this behavior graph in a Resource statement, you would use the following ARN:

```
"Resource": "arn:aws:detective:us-east-1:111122223333:graph:027c7c4610ea4aacaf0b883093cab899"
```

To specify multiple resources in a Resource statement, use commas to separate them.

```
"Resource": [  
    "resource1",  
    "resource2"  
]
```

For example, the same AWS account may be invited to be a member account in more than one behavior graph. In the policy for that member account, the Resource statement would list the behavior graphs they were invited to.

```
"Resource": [  
    "arn:aws:detective:us-  
east-1:111122223333:graph:027c7c4610ea4aacaf0b883093cab899",  
    "arn:aws:detective:us-east-1:444455556666:graph:056d2a9521xi2bb1uw1d164680eby416"  
]
```

Some Detective actions, such as creating a behavior graph, listing behavior graphs, and listing behavior graph invitations, are not performed on a specific behavior graph. For those actions, the Resource statement must use the wildcard (*).

```
"Resource": "*"
```

For administrator account actions, Detective always verifies that the user making the request belongs to the administrator account for the affected behavior graph. For member account actions, Detective always verifies that the user making the request belongs to the member account. Even if an IAM policy grants access to a behavior graph, if the user does not belong to the correct account, the user cannot perform the action.

For all actions that are performed on a specific behavior graph, the IAM policy should include the graph ARN. The graph ARN can be added later. For example, when an account first enables Detective, the initial IAM policy provides access to all Detective actions, using the wildcard for the graph ARN. This allows the user to immediately start to manage member accounts for and conduct investigations in their behavior graph. After the behavior graph is created, you can update the policy to add the graph ARN.

Condition keys

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element (or `Condition block`) lets you specify conditions in which a statement is in effect. The `Condition` element is optional. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple `Condition` elements in a statement, or multiple keys in a single `Condition` element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see [IAM policy elements: variables and tags](#) in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

Detective does not define its own set of condition keys. It does support using global condition keys. To see all AWS global condition keys, see [AWS Global Condition Context Keys](#) in the *IAM User Guide*.

To learn which actions and resources allow you to use a condition key, see [Actions defined by Amazon Detective](#).

Examples

To view examples of Detective identity-based policies, see [Amazon Detective identity-based policy examples](#).

Detective resource-based policies (Not supported)

Detective does not support resource-based policies.

Authorization based on Detective behavior graph tags

Each behavior graph can be assigned tag values. You can use those tag values in condition statements to manage access to the behavior graph.

The condition statement for a tag value uses the following format.

```
{"StringEquals":{"aws:ResourceTag/<tagName>": "<tagValue>"}}
```

For example, use the following code to allow or deny an action when the value of the Department tag is Finance.

```
{"StringEquals":{"aws:ResourceTag/Department": "Finance"}}
```

For examples of policies that use resource tag values, see [the section called “Administrator account: Restricting access based on tag values”](#).

Detective IAM Roles

An [IAM role](#) is an entity within your AWS account that has specific permissions.

Using temporary credentials with Detective

You can use temporary credentials to sign in with federation, assume an IAM role, or to assume a cross-account role. You obtain temporary security credentials by calling AWS STS API operations such as [AssumeRole](#) or [GetFederationToken](#).

Detective supports using temporary credentials.

Service-linked roles

[Service-linked roles](#) allow AWS services to access resources in other services to complete an action on your behalf. Service-linked roles appear in your IAM account and are owned by the service. An IAM administrator can view but not edit the permissions for service-linked roles.

For details about creating or managing Detective service-linked roles, see [the section called “Using service-linked roles”](#).

Service roles (Not supported)

This feature allows a service to assume a [service role](#) on your behalf. This role allows the service to access resources in other services to complete an action on your behalf. Service roles appear in your IAM account and are owned by the account. This means that an IAM administrator can change the permissions for this role. However, doing so might break the functionality of the service.

Detective does not support service roles.

Amazon Detective identity-based policy examples

By default, IAM users and roles don't have permission to create or modify Detective resources. They also can't perform tasks using the AWS Management Console, AWS CLI, or AWS API.

An IAM administrator must create IAM policies that grant users and roles permission to perform specific API operations on the specified resources they need. The administrator then attaches those policies to the IAM users or groups that require those permissions.

To learn how to create an IAM identity-based policy using these example JSON policy documents, see [Creating Policies on the JSON Tab](#) in the *IAM User Guide*.

Topics

- [Policy best practices](#)
- [Using the Detective console](#)
- [Allowing users to view their own permissions](#)
- [Administrator account: Managing the member accounts in a behavior graph](#)
- [Administrator account: Using a behavior graph for investigation](#)
- [Member account: Managing behavior graph invitations and memberships](#)
- [Administrator account: Restricting access based on tag values](#)

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Detective resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see [AWS managed policies](#) or [AWS managed policies for job functions](#) in the *IAM User Guide*.
- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on

specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.

- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see [IAM Access Analyzer policy validation](#) in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see [Configuring MFA-protected API access](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

Using the Detective console

To use the Amazon Detective console, the user or role must have access to the relevant actions, which match corresponding actions in the API.

To enable Detective and become an administrator account for a behavior graph, the user or role must be granted permission for the `CreateGraph` action.

To use the Detective console to perform any administrator account actions, the user or role must be granted permission for the `ListGraphs` action. This grants permission to retrieve the behavior graphs their account is an administrator account for. They also must be granted permission to perform specific administrator account actions.

The most basic administrator account actions are to view a list of member accounts in a behavior graph, and to use the behavior graph for investigation.

- To view the list of member accounts in a behavior graph, the principal must be granted permission for the `ListMembers` action.
- To conduct investigation in a behavior graph, the principal must be granted permission for the `SearchGraph` action.

To use the Detective console to perform any member account actions, the user or role must be granted permission for the `ListInvitations` action. This grants permission to view behavior graph invitations. They can then be granted permission for specific member account actions.

Allowing users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
      "Effect": "Allow",
      "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsWithUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
      ],
      "Resource": ["arn:aws:iam::*:user/${aws:username}"]
    },
    {
      "Sid": "NavigateInConsole",
      "Effect": "Allow",
      "Action": [
        "iam:GetGroupPolicy",
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedGroupPolicies",
        "iam:ListGroupPolicies",
        "iam:ListPolicyVersions",
        "iam:ListPolicies",

```

```

        "iam:ListUsers"
      ],
      "Resource": "*"
    }
  ]
}

```

Administrator account: Managing the member accounts in a behavior graph

This example policy is intended for administrator account users who are only responsible for managing the member accounts used in the behavior graph. The policy also allows the user to view the usage information and deactivate Detective. The policy does not grant permission to use the behavior graph for investigation.

```

{"Version":"2012-10-17",
  "Statement":[
    {
      "Effect":"Allow",
      "Action":
["detective:ListMembers","detective:CreateMembers","detective:DeleteMembers","detective:DeleteG
      "Resource":"arn:aws:detective:us-
east-1:111122223333:graph:027c7c4610ea4aacaf0b883093cab899"
    },
    {
      "Effect":"Allow",
      "Action":["detective:CreateGraph","detective:ListGraphs"],
      "Resource":"*"
    }
  ]
}

```

Administrator account: Using a behavior graph for investigation

This example policy is intended for administrator account users who use the behavior graph for investigation only. They cannot view or edit the list of member accounts in the behavior graph.

```

{"Version":"2012-10-17",
  "Statement":[
    {
      "Effect":"Allow",
      "Action":["detective:SearchGraph"],

```

```

    "Resource": "arn:aws:detective:us-
east-1:111122223333:graph:027c7c4610ea4aacaf0b883093cab899"
  },
  {
    "Effect": "Allow",
    "Action": ["detective:ListGraphs"],
    "Resource": "*"
  }
]
}

```

Member account: Managing behavior graph invitations and memberships

This example policy is intended for users belonging to a member account. In the example, the member account belongs to two behavior graphs. The policy grants permission to respond to invitations and remove the member account from the behavior graph.

```

{"Version": "2012-10-17",
 "Statement": [
  {
    "Effect": "Allow",
    "Action":
["detective:AcceptInvitation", "detective:RejectInvitation", "detective:DisassociateMembership"],
    "Resource": [
      "arn:aws:detective:us-
east-1:111122223333:graph:027c7c4610ea4aacaf0b883093cab899",
      "arn:aws:detective:us-
east-1:444455556666:graph:056d2a9521xi2bb1uw1d164680eby416"
    ]
  },
  {
    "Effect": "Allow",
    "Action": ["detective:ListInvitations"],
    "Resource": "*"
  }
]
}

```

Administrator account: Restricting access based on tag values

The following policy allows the user to use a behavior graph for investigation if the SecurityDomain tag of the behavior graph matches the SecurityDomain tag of the user.

```
{
  "Version": "2012-10-17",
  "Statement": [ {
    "Effect": "Allow",
    "Action": ["detective:SearchGraph"],
    "Resource": "arn:aws:detective:*:*:graph:*",
    "Condition": {
      "StringEquals": {
        "aws:ResourceTag/SecurityDomain": "aws:PrincipalTag/SecurityDomain"
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": ["detective:ListGraphs"],
    "Resource": "*"
  } ]
}
```

The following policy prevents the users from using a behavior graph for investigation if the value of the `SecurityDomain` tag for the behavior graph is `Finance`.

```
{
  "Version": "2012-10-17",
  "Statement": [ {
    "Effect": "Deny",
    "Action": ["detective:SearchGraph"],
    "Resource": "arn:aws:detective:*:*:graph:*",
    "Condition": {
      "StringEquals": {"aws:ResourceTag/SecurityDomain": "Finance"}
    }
  } ]
}
```

AWS managed policies for Amazon Detective

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you

reduce permissions further by defining [customer managed policies](#) that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see [AWS managed policies](#) in the *IAM User Guide*.

AWS managed policy: AmazonDetectiveFullAccess

You can attach the AmazonDetectiveFullAccess policy to your IAM identities.

This policy grants administrative permissions that allow a principal full access to all Amazon Detective actions. You can attach this policy to a principal before they enable Detective for their account. It must also be attached to the role that is used to run the Detective Python scripts to create and manage a behavior graph.

Principals with these permissions can manage member accounts, add tags to their behavior graph, and use Detective for investigation. They can also archive GuardDuty findings. The policy provides permissions that the Detective console needs to display account names for accounts that are in AWS Organizations.

Permissions details

This policy includes the following permissions:

- `detective` – Allows principals full access to all Detective actions.
- `organizations` – Allows principals to retrieve from AWS Organizations information about the accounts in an organization. If an account belongs to an organization, these permissions allow the Detective console to display account names in addition to account numbers.
- `guardduty` – Allows principals to get and archive GuardDuty findings from within Detective.
- `securityhub` – Allows principals to get Security Hub findings from within Detective.

```
{
```

```
"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Allow",
    "Action": [
      "detective:*",
      "organizations:DescribeOrganization",
      "organizations:ListAccounts"
    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "guardduty:ArchiveFindings"
    ],
    "Resource": "arn:aws:guardduty:*:*:detector/*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "guardduty:GetFindings",
      "guardduty:ListDetectors"
    ],
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "securityHub:GetFindings"
    ],
    "Resource": "*"
  }
]
```

AWS managed policy: AmazonDetectiveMemberAccess

You can attach the AmazonDetectiveMemberAccess policy to your IAM entities.

This policy provides member access to Amazon Detective and scoped access to the console.

With this policy, you can:

- View invitations to Detective graph membership and accept or reject those invitations.
- View how your activity in Detective contributes to the cost of using this service on the **Usage** page.
- Resign from your membership in a graph.

This policy grants read-only permissions that allow scoped access to the Detective console.

Permissions details

This policy includes the following permissions:

- `detective` – Allows member access to Detective.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "detective:AcceptInvitation",
        "detective:BatchGetMembershipDatasources",
        "detective:DisassociateMembership",
        "detective:GetFreeTrialEligibility",
        "detective:GetPricingInformation",
        "detective:GetUsageInformation",
        "detective:ListInvitations",
        "detective:RejectInvitation"
      ],
      "Resource": "*"
    }
  ]
}
```


AWS managed policy: AmazonDetectiveInvestigatorAccess

You can attach the AmazonDetectiveInvestigatorAccess policy to your IAM entities.

This policy provides investigator access to the Detective service and scoped access to the Detective console UI dependencies. This policy grants permissions to enable Detective investigations in Detective for IAM users and IAM roles. You can investigate to identify indicators of compromise such as findings using an investigation report, which provides analysis and insights about security indicators. The report is ranked by severity, which is determined using Detective's behavioral analysis and machine learning. You can use the report to prioritize remediation of resources.

Permissions details

This policy includes the following permissions:

- `detective` – Allows principals investigator access to Detective actions, to enable Detective investigations, and to enable finding groups summary.
- `guardduty` – Allows principals to get and archive GuardDuty findings from within Detective.
- `securityhub` – Allows principals to get Security Hub findings from within Detective.
- `organizations` – Allows principals to retrieve information about the accounts in an organization from AWS Organizations. If an account belongs to an organization, then these permissions allow the Detective console to display account names in addition to account numbers.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DetectivePermissions",
      "Effect": "Allow",
      "Action": [
        "detective:BatchGetGraphMemberDatasources",
        "detective:BatchGetMembershipDatasources",
        "detective:DescribeOrganizationConfiguration",
        "detective:GetFreeTrialEligibility",
        "detective:GetGraphIngestState",
```

```

    "detective:GetMembers",
    "detective:GetPricingInformation",
    "detective:GetUsageInformation",
    "detective:ListDatasourcePackages",
    "detective:ListGraphs",
    "detective:ListHighDegreeEntities",
    "detective:ListInvitations",
    "detective:ListMembers",
    "detective:ListOrganizationAdminAccount",
    "detective:ListTagsForResource",
    "detective:SearchGraph",
    "detective:StartInvestigation",
    "detective:GetInvestigation",
    "detective:ListInvestigations",
    "detective:UpdateInvestigationState",
    "detective:ListIndicators",
    "detective:InvokeAssistant"
  ],
  "Resource": "*"
},
{
  "Sid": "OrganizationsPermissions",
  "Effect": "Allow",
  "Action": [
    "organizations:DescribeOrganization",
    "organizations:ListAccounts"
  ],
  "Resource": "*"
},
{
  "Sid": "GuardDutyPermissions",
  "Effect": "Allow",
  "Action": [
    "guardduty:ArchiveFindings",
    "guardduty:GetFindings",
    "guardduty:ListDetectors"
  ],
  "Resource": "*"
},
{
  "Sid": "SecurityHubPermissions",
  "Effect": "Allow",
  "Action": [
    "securityHub:GetFindings"
  ]
}

```

```
    ],  
    "Resource": "*"    
  }  
]  
}
```

AWS managed policy: AmazonDetectiveOrganizationsAccess

You can attach the AmazonDetectiveOrganizationsAccess policy to your IAM entities.

This policy grants permission to enable and manage Amazon Detective within an organization. You can enable Detective across the organization and determine the delegated administrator account for Detective.

Permissions details

This policy includes the following permissions:

- `detective` – Allows principals access to Detective actions.
- `iam` – Specifies that a service linked role is created when Detective calls `EnableOrganizationAdminAccount`.
- `organizations` – Allows principals to retrieve information about the accounts in an organization from AWS Organizations. If an account belongs to an organization, then these permissions allow the Detective console to display account names in addition to account numbers. Enables the integration of an AWS service, allows register and deregister of the specified member account as a Delegated administrator, and allows principals to retrieve Delegated administrator accounts in other security services like Amazon Detective, Amazon GuardDuty, Amazon Macie, and AWS Security Hub.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "detective:DisableOrganizationAdminAccount",  
        "detective:EnableOrganizationAdminAccount",  
      ]  
    }  
  ]  
}
```

```
    "detective:ListOrganizationAdminAccount"
  ],
  "Resource": "*"
},
{
  "Effect": "Allow",
  "Action": [
    "iam:CreateServiceLinkedRole"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "iam:AWSServiceName": "detective.amazonaws.com"
    }
  }
},
{
  "Effect": "Allow",
  "Action": [
    "organizations:EnableAWSServiceAccess",
    "organizations:RegisterDelegatedAdministrator",
    "organizations:DeregisterDelegatedAdministrator"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "organizations:ServicePrincipal": [
        "detective.amazonaws.com"
      ]
    }
  }
},
{
  "Effect": "Allow",
  "Action": [
    "organizations:DescribeAccount",
    "organizations:DescribeOrganization",
    "organizations:ListAccounts"
  ],
  "Resource": "*"
},
{
  "Effect": "Allow",
  "Action": [
```



```

        "organizations:ListAccounts"
      ],
      "Resource": "*"
    }
  ]
}

```

Detective updates to AWS managed policies

View details about updates to AWS managed policies for Detective since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Document history page](#).

Change	Description	Date
AmazonDetectiveInvestigatorAccess – Updates to existing policies	<p>Added Detective investigations and finding groups summary actions to the AmazonDetectiveInvestigatorAccess policy.</p> <p>These actions allow starting, retrieving, and updating Detective investigations; and obtaining a summary of finding groups from within Detective.</p>	November 26, 2023
AmazonDetectiveFullAccess and AmazonDetectiveInvestigatorAccess – Updates to existing policies	<p>Detective added Security Hub GetFindings actions to the AmazonDetectiveFullAccess and AmazonDetectiveInvestigatorAccess policies.</p> <p>These actions allow getting Security Hub findings from within Detective.</p>	May 16, 2023

Change	Description	Date
AmazonDetectiveOrganizationsAccess – New policy	<p>Detective added AmazonDetectiveOrganizationsAccess policy.</p> <p>This policy grants permission to enable and manage Detective within an organization</p>	March 02, 2023
AmazonDetectiveMemberAccess – New policy	<p>Detective added the AmazonDetectiveMemberAccess policy.</p> <p>This policy provides member access to Detective and scoped access to the console UI dependencies.</p>	January 17, 2023
AmazonDetectiveFullAccess – Updates to an existing policy	<p>Detective added GuardDutyGetFindings actions to the AmazonDetectiveFullAccess policy.</p> <p>These actions allow getting GuardDuty findings from within Detective.</p>	January 17, 2023
AmazonDetectiveInvestigatorAccess – New policy	<p>Detective added the AmazonDetectiveInvestigatorAccess policy.</p> <p>This policy allows the principal to conduct investigations in Detective.</p>	January 17, 2023
AmazonDetectiveServiceLinkedRole – New policy	<p>Detective added a new policy for its service-linked role.</p> <p>The policy allows the service-linked role to retrieve information about the accounts in an organization.</p>	December 16, 2021

Change	Description	Date
Detective started to track changes	Detective started to track changes for its AWS managed policies.	May 10, 2021

Using service-linked roles for Detective

Amazon Detective uses AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique type of IAM role that is linked directly to Detective. Service-linked roles are predefined by Detective and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up Detective easier because you do not have to manually add the necessary permissions. Detective defines the permissions of its service-linked roles, and unless defined otherwise, only Detective can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting their related resources. This protects your Detective resources because you cannot inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see [AWS services that work with IAM](#) and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Service-linked role permissions for Detective

Detective uses the service-linked role named **AWSServiceRoleForDetective** – Allows Detective to access AWS Organizations information on your behalf.

The AWSServiceRoleForDetective service-linked role trusts the following services to assume the role:

- `detective.amazonaws.com`

The AWSServiceRoleForDetective service-linked role uses the managed policy [AmazonDetectiveServiceLinkedRolePolicy](#).

For details about updates to the `AmazonDetectiveServiceLinkedRolePolicy` policy see [Amazon Detective updates to AWS managed policies](#). For automatic alerts about changes to this policy, subscribe to the RSS feed on the [Detective document history](#) page.

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see [Service-linked role permissions](#) in the *IAM User Guide*.

Creating a service-linked role for Detective

You do not need to manually create a service-linked role. When you designate the Detective administrator account for an organization in the AWS Management Console, the AWS CLI, or the AWS API, Detective creates the service-linked role for you.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you designate the Detective administrator account for an organization, Detective creates the service-linked role for you again.

Editing a service-linked role for Detective

Detective does not allow you to edit the `AWSServiceRoleForDetective` service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see [Editing a service-linked role](#) in the *IAM User Guide*.

Deleting a service-linked role for Detective

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that is not actively monitored or maintained. However, you must clean up the resources for your service-linked role before you can manually delete it.

Note

If the Detective service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes and then try the operation again.

To delete Detective resources used by the AWSServiceRoleForDetective

1. Remove the Detective administrator account. See [the section called “Designating the Detective administrator account”](#).
2. Repeat the process in each Region where you designated the Detective administrator account.

To manually delete the service-linked role using IAM

Use the IAM console, the AWS CLI, or the AWS API to delete the AWSServiceRoleForDetective service-linked role. For more information, see [Deleting a Service-Linked Role](#) in the *IAM User Guide*.

Supported Regions for Detective service-linked roles

Detective supports using service-linked roles in all of the Regions where the service is available. For more information, see [AWS Regions and Endpoints](#).

Troubleshooting Amazon Detective identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Detective and IAM. If you encounter access denied issues or similar difficulties when working with AWS Identity and Access Management(IAM), consult the [Troubleshooting IAM](#) topics in the *IAM User Guide*.

I am not authorized to perform an action in Detective

If the AWS Management Console tells you that you're not authorized to perform an action, then you must contact your administrator for assistance. Your administrator is the person that provided you with your user name and password.

The following example error occurs when the mateojackson IAM user tries to use the console to accept an invitation to become a member account for a behavior graph, but does not have `detective:AcceptInvitation` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to
perform: detective:AcceptInvitation on resource: arn:aws:detective:us-
east-1:444455556666:graph:567856785678
```

In this case, Mateo asks his administrator to update his policies to allow him to access the `arn:aws:detective:us-east-1:444455556666:graph:567856785678` resource using the `detective:AcceptInvitation` action.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Detective.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Detective. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Detective resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Detective supports these features, see [How Amazon Detective works with IAM](#).
- To learn how to provide access to your resources across AWS accounts that you own, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see [Providing access to AWS accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.

- To learn the difference between using roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Compliance validation for Amazon Detective

Amazon Detective is in Scope of the AWS assurance program. For more information, see [Health Information Trust Alliance Common Security Framework \(HITRUST\) CSF](#).

For a list of AWS services in scope of specific compliance programs, see [AWS Services in Scope by Compliance Program](#). For general information, see [AWS Compliance Programs](#).

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

AWS provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying security and compliance focused baseline environments on AWS.
- [Evaluating resources with rules](#) in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [AWS Security Hub](#) – This AWS service provides a comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices.

Resilience in Amazon Detective

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).

In addition to the AWS global infrastructure, Detective makes use of the resiliency built into Amazon DynamoDB and Amazon Simple Storage Service (Amazon S3). For more information, see

[resiliency and disaster recovery in Amazon DynamoDB](#) and [Resilience in Amazon Simple Storage Service](#).

The Detective architecture is also resilient to the failure of a single Availability Zone. This resilience is built into Detective, and does not require any configuration.

Infrastructure security in Amazon Detective

As a managed service, Amazon Detective; is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see [AWS Cloud Security](#). To design your AWS environment using the best practices for infrastructure security, see [Infrastructure Protection](#) in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access Detective; through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

Security best practices for Detective

Detective provides a number of security features to consider as you develop and implement your own security policies. The following best practices are general guidelines and don't represent a complete security solution. Because these best practices might not be appropriate or sufficient for your environment, treat them as helpful considerations rather than prescriptions.

For Detective, the security best practices are associated with managing the accounts in a behavior graph.

Best practices for Detective administrator accounts

When inviting member accounts to your Detective behavior graph, only invite accounts that you oversee.

Limit access to the behavior graph. Users with the [AmazonDetectiveFullAccess](#) policy can grant access to all Detective actions. Principals with these permissions can manage member accounts, add tags to their behavior graph, and use Detective for investigation. When a user has access to a behavior graph, they can see all of the findings for the member accounts. Such findings might expose sensitive security information.

Best practices for member accounts

When you receive an invitation to a behavior graph, make sure to validate the source of the invitation.

Check the AWS account identifier of the administrator account that sent the invitation. Verify that you know who the account belongs to, and that the inviting account has a legitimate reason to monitor your security data.

Logging Amazon Detective API calls with AWS CloudTrail

Detective is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Detective. CloudTrail captures all API calls for Detective as events. The calls captured include calls from the Detective console and code calls to the Detective API operations.

- If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Detective.
- If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**.

Using the information collected by CloudTrail, you can determine the following:

- The request that was made to Detective
- The IP address from which the request was made
- Who made the request
- When it was made
- Additional details about the request

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

Detective information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in Detective, that activity is recorded in a CloudTrail event, along with other AWS service events, in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing Events with CloudTrail Event History](#).

For an ongoing record of events in your AWS account, including events for Detective, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket.

By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. You also can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs.

For more information, see the following:

- [Overview for Creating a Trail](#)
- [CloudTrail Supported Services and Integrations](#)
- [Configuring Amazon SNS Notifications for CloudTrail](#)
- [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#)

CloudTrail logs all Detective operations, which are documented in the [Detective API Reference](#).

For example, calls to the `CreateMembers`, `AcceptInvitation`, and `DeleteMembers` operations generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials
- Whether the request was made with temporary security credentials for a role or a federated user
- Whether the request was made by another AWS service


For more information, see the [CloudTrail userIdentity Element](#).

Understanding Detective log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries.

An event represents a single request from any source. Events include information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so the entries don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `AcceptInvitation` action.




```

{
  "EventId": "f2545ee3-170f-4340-8af4-a983c669ce37",
  "Username": "JaneRoe",
  "EventTime": 1571956406.0,
  "CloudTrailEvent": "{ \"eventVersion\": \"1.05\", \"userIdentity\":
{ \"type\": \"AssumedRole\", \"principalId\": \"AR0AJZARKEP6WKJ5JHSUS:JaneRoe\", \"arn
\": \"arn:aws:sts::111122223333:assumed-role/1A4R5SKSPGG9V/JaneRoe\", \"accountId
\": \"111122223333\", \"accessKeyId\": \"AKIAIOSFODNN7EXAMPLE\", \"sessionContext\":
{ \"attributes\": { \"mfaAuthenticated\": \"false\", \"creationDate\": \"2019-10-24T21:54:56Z
\"}, \"sessionIssuer\": { \"type\": \"Role\", \"principalId\": \"AR0AJZARKEP6WKJ5JHSUS
\", \"arn\": \"arn:aws:iam::111122223333:role/1A4R5SKSPGG9V\", \"accountId\":
\"111122223333\", \"userName\": \"JaneRoe\" } } }, \"eventTime\": \"2019-10-24T22:33:26Z
\", \"eventSource\": \"detective.amazonaws.com\", \"eventName\": \"AcceptInvitation
\", \"awsRegion\": \"us-east-2\", \"sourceIPAddress\": \"192.0.2.123\", \"userAgent
\": \"aws /3 aws-sdk-java/1.11.648 Linux/4.14.133-97.112.amzn2.x86_64 OpenJDK_64-
Bit_Server_VM/25.201-b09 java/1.8.0_201 vendor/Oracle_Corporation exec-env/
AWS_Lambda_java8\", \"errorCode\": \"ValidationException\", \"requestParameters\":
{ \"masterAccount\": \"111111111111\" }, \"responseElements\": { \"message\": \"Invalid
request body\" }, \"requestID\": \"8437ff99-5ec4-4b1a-8353-173be984301f\", \"eventID\":
\"f2545ee3-170f-4340-8af4-a983c669ce37\", \"readOnly\": false, \"eventType\": \"AwsApiCall
\", \"recipientAccountId\": \"111122223333\" }",
  "EventName": "AcceptInvitation",
  "EventSource": "detective.amazonaws.com",
  "Resources": []
},

```

Amazon Detective Regions and quotas

When using Amazon Detective, be aware of these quotas.

Detective Regions and endpoints

To see the list of AWS Regions where Detective is available, see [Detective service endpoints](#).

Detective quotas

Detective has the following quotas, which cannot be configured.

Resource	Quota	Comments
Number of member accounts	1,200	The number of member accounts that an administrator account can add to a behavior graph.
Behavior graph data volume – volume warning	9 TB per day	If the behavior graph data volume is larger than 9 TB per day, then Detective displays a warning that the behavior graph is nearing the maximum allowed volume.
Behavior graph data volume – no new accounts	10 TB per day	If the behavior graph data volume is larger than 10 TB per day, then you cannot add new member accounts to the behavior graph.
Behavior graph data volume – stop data ingest into the behavior graph	15 TB per day	<p>If the behavior graph data volume is larger than 15 TB per day, then Detective stops ingesting data into the behavior graph.</p> <p>The 15 TB per day reflects both normal data volume and spikes in the data volume.</p>

Resource	Quota	Comments
		To re-enable the data ingest, you must contact AWS Support.

Internet Explorer 11 not supported

You cannot use Detective with Internet Explorer 11.

Managing tags for a behavior graph

A tag is an optional label that you can define and assign to AWS resources, including certain types of Detective resources. Tags can help you identify, categorize, and manage resources in different ways, such as by purpose, owner, environment, or other criteria. For example, you can use tags to apply policies, allocate costs, distinguish between versions of resources, or identify resources that support certain compliance requirements or workflows.

You can assign tags to your behavior graph. You can then use the tag values in IAM policies to manage access to behavior graph functions in Detective. See [the section called “Authorization based on Detective behavior graph tags”](#).

You also can use tags as a tool for cost reporting. For example, to track costs associated with security, you could assign the same tag to your Detective behavior graph, AWS Security Hub hub resource, and Amazon GuardDuty detectors. In AWS Cost Explorer, you could then search for that tag to see a consolidated view of the costs across those resources.

Viewing the tags for a behavior graph

You manage the tags for your behavior graph from the **General** page.

Console

To view the list of tags assigned to the behavior graph

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the navigation pane, under **Settings**, choose **General**.

Detective API, AWS CLI

You can use the Detective API or the AWS Command Line Interface to get the list of tags for your behavior graph.

To get the list of tags for a behavior graph (Detective API, AWS CLI)

- **Detective API:** Use the [ListTagsForResource](#) operation. You must provide the ARN of your behavior graph.
- **AWS CLI:** At the command line, run the `list-tags-for-resource` command.

```
aws detective list-tags-for-resource --resource-arn <behavior graph ARN>
```

Example

```
aws detective list-tags-for-resource --resource-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```

Adding tags to a behavior graph

Console

From the tag list on the **General** page, you can add tag values to the behavior graph.

To add a tag to your behavior graph

1. Choose **Add new tag**.
2. For **Key**, enter the name of the tag.
3. For **Value**, enter the value of the tag.

Detective API, AWS CLI

You can use the Detective API or the AWS CLI to add tag values to your behavior graph.

To add tags to a behavior graph (Detective API, AWS CLI)

- **Detective API:** Use the [TagResource](#) operation. You provide the behavior graph ARN and the tag values to add.
- **AWS CLI:** At the command line, run the `tag-resource` command.

```
aws-detective tag-resource --aws detective tag-resource --resource-arn <behavior graph ARN> --tags '{"TagName":"TagValue"}
```

Example

```
aws detective tag-resource --resource-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234 --tags '{"Department":"Finance"}
```

Removing tags from a behavior graph

Console

To remove a tag from the list on the **General** page, choose the **Remove** option for that tag.

Detective API, AWS CLI

You can use the Detective API or the AWS CLI to remove tag values from your behavior graph.

To remove tags from a behavior graph (Detective API, AWS CLI)

- **Detective API:** Use the [UntagResource](#) operation. You provide the behavior graph ARN, and the names of the tags to remove.
- **AWS CLI:** At the command line, run the `untag-resource` command.

```
aws detective untag-resource --resource-arn <behavior graph ARN> --tag-keys  
"TagName"
```

Example

```
aws detective untag-resource --resource-arn arn:aws:detective:us-  
east-1:111122223333:graph:123412341234 --tag-keys "Department"
```

Disabling Amazon Detective

The administrator account for a behavior graph can disable Amazon Detective from the Detective console, the Detective API, or AWS Command Line Interface. When you disable Detective, the behavior graph and its associated Detective data are deleted.

Once a behavior graph is deleted, it cannot be restored.

Contents

- [Disabling Detective \(Console\)](#)
- [Disabling Detective \(Detective API, AWS CLI\)](#)
- [Disabling Detective across Regions \(Python script on GitHub\)](#)

Disabling Detective (Console)

You can disable Amazon Detective from the AWS Management Console.

To disable Amazon Detective (console)

1. Open the Amazon Detective console at <https://console.aws.amazon.com/detective/>.
2. In the Detective navigation pane, under **Settings**, choose **General**.
3. On the **General** page, under **Disable Amazon Detective**, choose **Disable Amazon Detective**.
4. When prompted to confirm, type **disable**.
5. Choose **Disable Amazon Detective**.

Disabling Detective (Detective API, AWS CLI)

You can disable Amazon Detective from the Detective API or the AWS Command Line Interface. To get the ARN of your behavior graph to use in the request, use the [ListGraphs](#) operation.

To disable Detective (Detective API, AWS CLI)

- **Detective API:** Use the [DeleteGraph](#) operation. You must provide the graph ARN.
- **AWS CLI:** At the command line, run the [delete-graph](#) command.

```
aws detective delete-graph --graph-arn <graph ARN>
```

Example:

```
aws detective delete-graph --graph-arn arn:aws:detective:us-east-1:111122223333:graph:123412341234
```

Disabling Detective across Regions (Python script on GitHub)

Detective provides an open-source script in GitHub that allows you to disable Detective for an administrator account across a specified list of Regions.

For information on how to configure and use the GitHub scripts, see [the section called “Amazon Detective Python scripts”](#).

Document history for Detective User Guide

The following table describes the important changes to the documentation since the last release of Detective. For notification about updates to this documentation, you can subscribe to an RSS feed.

- **Latest documentation update:** August 27, 2024

Change	Description	Date
Added support for Amazon GuardDuty findings	Detective now provides support for the following GuardDuty Runtime Monitoring finding types . <ul style="list-style-type: none">• Execution:Runtime/SuspiciousShell• PrivilegeEscalation:Runtime/ElevationToRoot	August 27, 2024
Added support for Amazon GuardDuty findings	Detective now provides support for GuardDuty Malware protection for S3 . This helps you scan newly uploaded objects to Amazon S3 buckets for potential malware and suspicious uploads, and take action to isolate them before they are ingested into downstream processes.	July 9, 2024
Updated functionality	Detective added a new Radial layout to the finding group Visualization panel , to provide	June 26, 2024

improved visualization for easier data interpretation.

[New Security Lake source versions](#)

In addition to source version 1 (OCSF 1.0.0-rc.2), Detective now ingests data from source version 2 (OCSF 1.1.0) for the [Security Lake sources](#) that are supported by Detective.

May 15, 2024

[New Security Lake log source](#)

You can use the Detective integration with Security Lake to collect logs and events from [Amazon EKS Audit Logs](#).

May 15, 2024

[Documentation update](#)

The content from the Amazon Detective Administration Guide is now consolidated into the Amazon Detective User Guide. Amazon Detective Administration Guide will reach its end of standard support on May 08, 2024.

April 15, 2024

[Added support for Amazon GuardDuty findings](#)

Detective now provides support for the following [GuardDuty Runtime Monitoring finding types](#).

April 5, 2024

- Execution:Runtime/MaliciousFileExecuted
- Execution:Runtime/SuspiciousTool
- DefenseEvasion:Runtime/PtraceAntiDebugging
- Execution:Runtime/SuspiciousCommand
- DefenseEvasion:Runtime/SuspiciousCommand

[Removed the Amazon GuardDuty membership requirement](#)

You are no longer required to be a GuardDuty customer to enable Amazon Detective . The requirement to have GuardDuty enabled in your account for 48 hours before enabling Detective has been removed.

February 2, 2024

[Added support for Amazon GuardDuty findings](#)

Detective extends support for [GuardDuty EC2 Runtime Monitoring](#) finding types to ECS and EC2 resources.

January 30, 2024

[Updated functionality](#)

You can now run a Detective investigation from the Investigations page for a specific resource that you want to investigate. Detective recommends resources based on its activity in findings and finding groups. [Detective Investigations](#) lets you investigate IAM users and IAM roles with indicators of compromise, which can help you determine if a resource is involved in a security incident.

January 16, 2024

[Updated functionality](#)

You can now run a Detective investigation from the Investigations page on a recommended resource. Detective recommends resources based on its activity in findings and finding groups. [Detective Investigations](#) lets you investigate IAM users and IAM roles with indicators of compromise, which can help you determine if a resource is involved in a security incident.

December 26, 2023

[Changes in how Detective reads the flow traffic for shared VPCs](#)

If you are using a shared Amazon VPC, you may see changes in the traffic monitored by Detective. We recommend that you review the changes in [Activity details for overall VPC flow volume](#) to understand the potential effects on your coverage, and review [how Detective calculates projected cost](#) to understand how that can impact your service costs.

December 20, 2023

[Regional availability](#)

Added Europe (Stockholm), Europe (Paris), and Canada (Central) Regions to the list of AWS Regions where [Detective integration with Security Lake](#) is available.

December 8, 2023

[New feature](#)

[Detective investigations](#) lets you investigate IAM users and IAM roles with indicators of compromise, which can help you determine if a resource is involved in a security incident.

November 26, 2023

New feature	By default, Detective automatically generates finding group summaries for finding groups, powered by generative artificial intelligence (generative AI). Finding group summary, rapidly analyzes relationships between findings and affected resources, and then summarizes potential threats in natural language.	November 26, 2023
New feature	Detective integration with Security Lake lets you can query and retrieve the raw log data stored by Security Lake. Using this integration, you can collect logs and events from CloudTrail management events and Amazon Virtual Private Cloud (Amazon VPC) Flow Logs.	November 26, 2023
Added managed policy information to the security chapter	Added Detective investigations and finding groups summary actions to the AmazonDetectiveInvestigatorAccess policy.	November 26, 2023
Viewing a finding overview	If a finding is correlated to a larger activity, Detective now notifies you to navigate to that finding group.	September 18, 2023
Amazon Detective endpoints and quotas	Detective is now available in the Israel (Tel Aviv) Region.	August 25, 2023

Enhanced finding groups visualization	Detective finding groups visualization now includes finding groups with aggregated findings making it more efficient to analyze related evidences, entities, and findings.	August 8, 2023
Enhanced finding groups	Finding groups now include vulnerability findings from Amazon Inspector.	June 13, 2023
Added support for Amazon GuardDuty Lambda Protection	Detective now provides support for GuardDuty Lambda Protection.	May 26, 2023
Added AWS security findings as a new optional data source package.	Detective now provides AWS security findings as an optional data source package. This optional data source package allows Detective to ingest data from Security Hub and adds that data to your behavior graph.	May 16, 2023
Added support for Amazon GuardDuty EKS Runtime Monitoring finding types	Detective now provides support for GuardDuty EKS Runtime Monitoring finding types.	May 3, 2023
Added support for Amazon GuardDuty RDS Protection finding types	Detective now provides support for GuardDuty RDS Protection finding types.	April 20, 2023

[Added support for additional Amazon GuardDuty finding types](#)

Detective now provides profiles for the following additional GuardDuty finding types: DefenseEvasion: EC2UnusualDNSResolver DefenseEvasion: EvasionEC2UnusualDoHActivity DefenseEvasion: DefenseEvasionEC2UnusualDoTActivity

April 12, 2023

[Added new console panels in the Detective console to help users select the appropriate AWS managed policy for their specific use case.](#)

Detective offers managed policies to securely choose the permissions that you need.

April 3, 2023

[Displaying the VPC flow traffic for EKS clusters](#)

Added new section for Amazon Virtual Private Cloud (Amazon VPC) flow traffic with Amazon Elastic Kubernetes Service (Amazon EKS) clusters.

March 2, 2023

[Finding group now includes a dynamic visual representation of Detective's behavior graph](#)

Detective finding group now includes a dynamic visual representation of Detective's behavior graph to emphasize the relationship between entities and findings within the finding group.

February 28, 2023

Export data from Detective Summary page and search results page. The data is exported in comma-separated values (CSV) format.	Detective now provides the option to export data to your browser from the Detective console.	February 7, 2023
Added overall VPC flow volume for EKS Amazon EKS workloads	Detective now adds visual summaries and analytics about your Amazon Virtual Private Cloud (VPC) flow logs from your Amazon Elastic Kubernetes Service Amazon EKS workloads.	January 19, 2023
Added managed policy information to the security chapter	Detective now supports GuardDuty get findings actions through the AmazonDetectiveFullAccess policy. The security chapter now provides details about the following new managed policies for Detective: AmazonDetectiveMemberAccess and AmazonDetectiveInvestigatorAccess.	January 17, 2023
Added data retention	With Detective, you can access up to a year of historical event data.	December 20, 2022
Added the option to adjust scope time on the summary page.	Detective now provides the option to adjust the scope time so view the activity for any 24-hour time frame in the previous 365 days.	October 5, 2022

Searching for a finding or entity	Detective now provides case insensitive search.	October 3, 2022
Added the ability to set scope timestamp	Detective now provides a way to configure the scope timestamp format preference. This preference will be applied to all timestamps in Detective.	October 3, 2022
Added terms related to finding groups	Detective now supports finding groups that connect related findings together in a single display to help you investigate potential malicious activity in your environment. From a finding group profile, you can pivot to entity profiles and finding overviews related to that group.	August 3, 2022
Added new profiles associated with Amazon EKS audit logs	Detective now provides profiles to allow you to investigate activity associated with the following container-related entities: Amazon EKS clusters, container images, Kubernetes pods, and Kubernetes subjects.	July 26, 2022

[Added a new optional data source](#)

Detective now supports EKS audit logs as an optional data source package. An administrator account can enable this new data source for their existing behavior graph. Graphs created after this date will have this data source enabled by default. Administrators can disable this data source manually at any time.

July 26, 2022

[New service-linked role and managed policy for Detective](#)

Detective now has a service-linked role, `AWSServiceRoleForDetective`. The service-linked role is used to access Organizations data on your behalf. The role uses a new `AmazonDetectiveServiceLinkedRolePolicy` managed policy.

December 16, 2021

[Added integration with AWS Organizations](#)

Detective is now integrated with Organizations. The organization management account designates a Detective administrator account for the organization. The Detective administrator account can view all of the accounts in the organization, and enable those accounts as member accounts in the organization behavior graph.

December 16, 2021

[Replaced finding profiles with finding overviews](#)

Finding profiles contained visualizations that analyzed activity for the involved resource. The new finding overview contains finding details ingested from GuardDuty, and a list of involved entities. From the finding overview, you can pivot to the profiles for related entities.

September 20, 2021

[Removed the limit on supported GuardDuty finding types](#)

Detective is no longer limited to a selected set of GuardDuty finding types. Detective automatically collects finding details for all finding types, and provides access to the entity profiles for the related entities.

September 20, 2021

[Link to finding details from the associated findings profile panel](#)

On an entity profile, when you choose a finding in the associated findings list, the finding details are displayed in the panel to the right. The scope time is set to the finding time window.

September 20, 2021

[Added S3 buckets to the available entity types in Detective](#)

Detective now provides profiles for S3 buckets. The S3 bucket profiles provide details about the principals that interacted with the S3 bucket and the API operations that they performed on the S3 bucket.

September 20, 2021

[New option to generate Detective URLs in Splunk](#)

The Splunk Trumpet project allows you to send AWS content to Splunk. The project now allows you to add Detective URLs to navigate to profiles for GuardDuty findings.

September 8, 2021

[Replaced AKIDs in the activity details for accounts and roles](#)

On account profiles, the activity details for **Overall API call volume** now show users or roles instead of access key identifiers (AKIDs). On role profiles, the activity details for **Overall API call volume** now show role sessions instead of AKIDs. For activity that occurred before this change, the caller is listed as **Unknown resource**.

July 14, 2021

[Added the calling service to information about API calls](#)

On the Detective console, information about API calls now includes the service that issued the call. Added a **Service** column to the lists on the **Overall API call volume**, **Newly observed API calls**, and **API calls with increased volume**. On the activity details for **Overall API call volume** and **Newly observed geolocations**, API methods are grouped under the services that issued them. For activity that occurred before this change, the API methods are grouped under **Unknown service**.

July 14, 2021

[New Resource interaction tab for users, roles, and role sessions](#)

The **Resource interaction** tab for users, roles, and role sessions contains information about role assumption activity that involved those entities. For role sessions, this is a new tab. For users and roles, this is an existing tab with new content.

June 29, 2021

[Updated values for behavior graph data volume quotas](#)

Increased the data volume quotas for behavior graphs. At 3.24 TB per day, Detective issues a warning. At 3.6 TB per day, no new accounts can be added. At 4.5 TB per day, Detective stops ingesting data into the behavior graph.

June 10, 2021

[Added tag values to the Python script options](#)

When you use the Detective Python script `enableDetective.py` to enable Detective, you can now assign tag values to the behavior graph.

May 19, 2021

[Added automatic enabling of member accounts that pass the data volume check](#)

When member accounts accept an invitation, their status is **Accepted (Not enabled)** until Detective verifies that their data will not cause the behavior graph data volume to exceed the quota. If the data volume is not a problem, Detective automatically changes the status to **Accepted (Enabled)**. Note that existing member accounts that are currently **Accepted (Not enabled)** cannot be enabled automatically.

May 12, 2021

[Added managed policy information to the security chapter](#)

A new section in the security chapter provides details about managed policies for Detective. Detective currently provides a single managed policy, `AmazonDetectiveFullAccess`.

May 10, 2021

[Changed the data volume values in the member accounts list](#)

On the account management page, the member accounts list now displays the daily data volume for each member account. Previously the list displayed the volume as a percentage of the total allowed volume.

April 29, 2021

[Revised options for managing member accounts](#)

Replaced the **Manage accounts** menu with an **Actions** menu. Combined the options for adding individual accounts and adding accounts from a .csv file. Moved **Enable accounts** from **Manage accounts** to a separate option next to **Actions**.

April 5, 2021

[Added behavior graph tags and authorization based on tags](#)

When you enable Detective, you can add tags to the behavior graph. You can manage tags for a behavior graph from the **General** page. Detective also supports authorization based on tag values.

March 31, 2021

[Added support for additional Amazon GuardDuty finding types](#)

Detective now provides profiles for the following additional GuardDuty finding types: CredentialAccess:IAMUser/AnomalousBehavior , DefenseEvasion:IAMUser/AnomalousBehavior , Discovery:IAMUser/AnomalousBehavior , Exfiltration:IAMUser/AnomalousBehavior , Impact:IAMUser/AnomalousBehavior , InitialAccess:IAMUser/AnomalousBehavior , Persistence:IAMUser/AnomalousBehavior , PrivilegeEscalation:IAMUser/AnomalousBehavior

March 29, 2021

[Added differences for AWS GovCloud \(US\) Regions](#)

Detective is now available in the AWS GovCloud (US) Regions. In AWS GovCloud (US-East) and AWS GovCloud (US-West), Detective does not send invitation emails to member accounts. Detective also does not automatically remove member accounts that are shut down in AWS.

March 24, 2021

[Added tabs to filter the member account list based on the member account status](#)

The list of member accounts now displays tabs that you can use to filter the list based on the member account status. You can view all member accounts, those that have a status of **Accepted (Enabled)**, or those that have a status other than **Accepted (Enabled)**.

March 16, 2021

[Added support for additional Amazon GuardDuty finding types](#)

Detective now provides profiles for the following additional GuardDuty finding types: Backdoor: EC2/C&CActivity.B , Impact:EC2/PortSweep , Impact:EC2/WinRMBruteForce , and Privilege Escalation:IAMUser/AdministrativePermissions

March 4, 2021

[Added option to Python script to suppress invitation emails](#)

The Detective enableDetective.py script now provides a --disable_email option. When you include that option, Detective does not send invitation emails to the member accounts.

February 26, 2021

[Changed "master account" to "administrator account"](#)

The term "master account" is changed to "administrator account." The term is also changed in the Detective console and API.

February 25, 2021

[Changed "master account" to "administrator account"](#)

The term "master account" is changed to "administrator account." The term is also changed in the Detective console and API.

February 25, 2021

[Added activity details for the profile panel VPC flow volume to and from the finding's IP address](#)

The profile panel **VPC flow volume to and from the finding's IP address** now allows you to display activity details. The activity details are available only if the finding is associated with a single IP address. The activity details show the volume for each combination of ports, protocol, and direction.

February 25, 2021

[Added API option to not send invitation emails to member accounts](#)

When using the Detective API to add member accounts, administrator accounts can choose to not send invitation emails to member accounts.

February 25, 2021

[New activity details for the Overall API call volume profile panel on IP address profiles](#)

You can now display activity details for IP addresses from the **Overall API call volume** profile panel. The activity details show the number of successful and failed calls for each resource that issued the call from the IP address.

February 23, 2021

[New Overall VPC flow volume profile panel on IP address profiles](#)

The IP address profile now contains the **Overall VPC flow volume** profile panel. The profile panel shows the volume of VPC flow traffic to and from the IP address. You can display activity details to show the volume for each EC2 instance that the IP address communicated with.

January 21, 2021

[Added the Detective Summary page](#)

The Detective **Summary** page contains visualizations to guide analysts to entities of interest based on geolocation, numbers of API calls, and Amazon EC2 traffic volume.

January 21, 2021

[Updated the option to pivot from Amazon GuardDuty to Detective](#)

In GuardDuty, the **Investigate in Detective** option is moved from the **Actions** menu to the finding details panel. It displays a list of related entities. If the finding type is supported, the list also includes the finding. You can then choose to navigate to either an entity profile or a finding profile.

January 15, 2021

[Added option to set the activity details window to the default scope time](#)

On the activity details for **Overall API call volume** and **Overall VPC flow volume**, you can set the time window for the activity details to the default scope time for the profile.

January 15, 2021

Added handling of high-volume time intervals for entities	Added a new notice to indicate when an entity has one or more high-volume time intervals. A new High-volume entities page displays all of the high-volume intervals for the current scope time.	December 18, 2020
Member account quota increased to 1,200	Master accounts can now invite up to 1,200 member accounts to their behavior graph. Previously the quota was 1,000.	December 11, 2020
Added values for behavior graph data volume quotas	Updated the information about behavior graph data volume quotas to add the specific quota values.	December 11, 2020
Added time range selection for activity details on the Overall API call volume profile panel	On the Overall API flow volume panel, you can now display activity details for any selected time range. The panel initially displays an option to display the activity details for the scope time.	September 29, 2020
Added time interval selection for activity details on the Overall VPC flow volume profile panel	On the Overall VPC flow volume panel, you can display activity details for a single time interval from the chart. To display the details for time interval, choose the time interval.	September 25, 2020

[New role session and federated user entities](#)

Detective now allows you to explore and investigate federated authentication. You can see what resources have assumed each role, and when those authentications occurred.

September 17, 2020

[Updates to scope time management](#)

Removed the option to lock or unlock the scope time. It is always locked. On a finding profile, a warning is displayed if the scope time is different from the finding time window.

September 4, 2020

[Profile header remains visible as you scroll through a profile](#)

On profiles, the type, identifier, and scope time remain visible as you scroll through the profile panels on a tab. When the tabs are not visible, you can use the tab drop down list in the breadcrumbs to navigate to a different tab.

September 4, 2020

[Search always displays search results](#)

When you conduct a search, it now displays the results on the **Search** page. From the results, you can pivot to a finding or entity profile.

August 27, 2020

[Added to the allowed criteria for searches](#)

The allowed criteria for searches has expanded. You can search for AWS users and AWS roles by name. You can use the ARN to search for findings, AWS roles, AWS users, and EC2 instances.

August 27, 2020

[Links to other consoles from profile panels](#)

On the **EC2 instance details** profile panel, the EC2 instance identifier is linked to the Amazon EC2 console. On the **User details**, and **Role details** profile panels, the user name and role name are linked to the IAM console.

August 14, 2020

[Activity details for VPC flow data](#)

The **Overall VPC flow volume** profile panel now provides access to activity details. The activity details show the traffic flow between IP addresses and an EC2 instance during a selected time period.

July 23, 2020

[Member accounts can now see their usage and projected cost](#)

Member accounts can now view their own usage information. For member accounts, the **Usage** page shows the amount of data ingested into each behavior graph that they contribute to. Member accounts can also see their projected 30-day cost.

May 26, 2020

[Free trial is now per account instead of per behavior graph](#)

Each account Amazon Detective now receives a separate free trial within each Region. The free trial starts either when the account enables Detective, or the first time the account is enabled as a member account.

May 26, 2020

[New open source Python scripts on GitHub](#)

The new [amazon-detective-multiaccount-scripts](#) repository on GitHub provides open source Python scripts that you can use to manage behavior graphs across Regions. You can enable Detective, add member accounts, remove member accounts, and disable Detective.

January 21, 2020

[Introducing Amazon Detective](#)

Detective uses machine learning and purpose-built visualizations to help you analyze and investigate security issues across your Amazon Web Services (AWS) workloads.

December 2, 2019