Amazon Detective

User Guide

Amazon Detective: User Guide

Table of Contents

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 activities. If you are new to Detective, see What is Detective? and Detective terms and concepts in the *Detective Administration Guide*.

Topics

Investigation phases and starting points

Amazon Detective provides tools to support the overall investigation process. An investigation in Detective can start from a finding or an entity.

Investigation phases

Any 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.

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 continues 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 finding to investigate. Or you might have a concern about unusual activity tied to 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

This is the most common starting point for an investigation process in Detective. GuardDuty uses your log data to uncover suspected instances of malicious or high-risk activity. Detective provides resources that help you dig further into these findings.

Starting with a finding, you can do the following:

- See what entities, such as IP addresses and AWS accounts, are connected to that finding.
- See what other findings might be related to that finding.
- See what activity occurred close in time or location to that finding.

For more information, see *Analyzing finding details* (p. 26).

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. For more information, see *Analyzing entity details* (p. 28).

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.

Investigation flows in Amazon Detective

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

Overview of a typical Detective entity investigation flow

At a high level, the following image shows the process for investigating an entity in Detective.

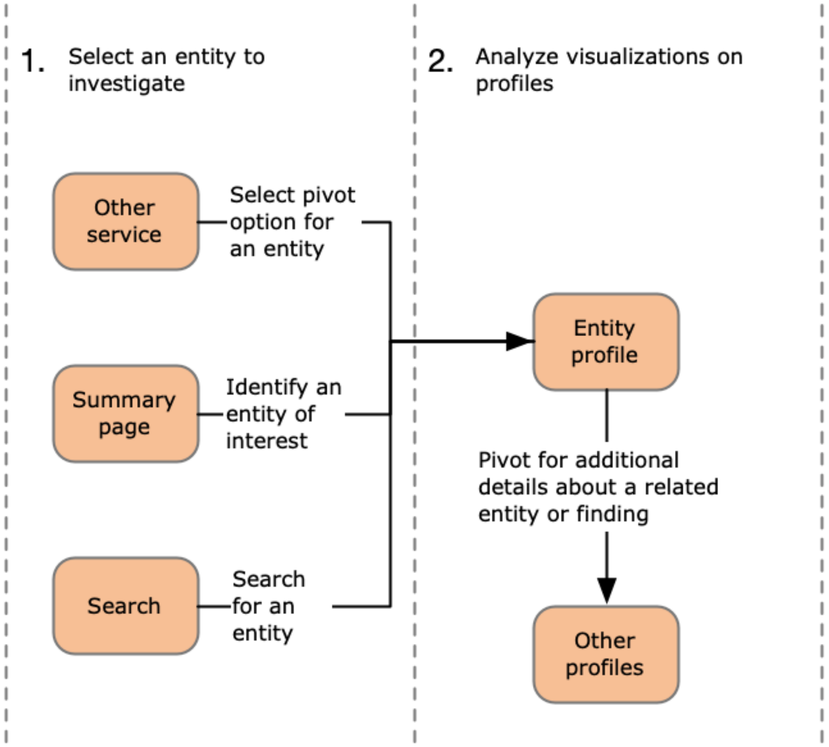
1. Select an entity to investigate
2. Analyze visualizations on profiles
Other service
Select pivot option for an entity
Summary page
Identify an entity of interest
Search
Search for an entity
Entity profile
Pivot for additional details about a related entity or finding
Other profiles

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” (p. 15).

You can use the Detective search function to find and select an entity to investigate. See *Searching for a finding or entity* (p. 20).

You can also use the Detective **Summary** page to identify an entity to investigate. See *Using the Summary page* (p. 22).

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 entity details* (p. 28).

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” (p. 51).

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

Overview of a typical Detective finding investigation flow

At a high level, the following image shows the process for investigating a finding in Detective.

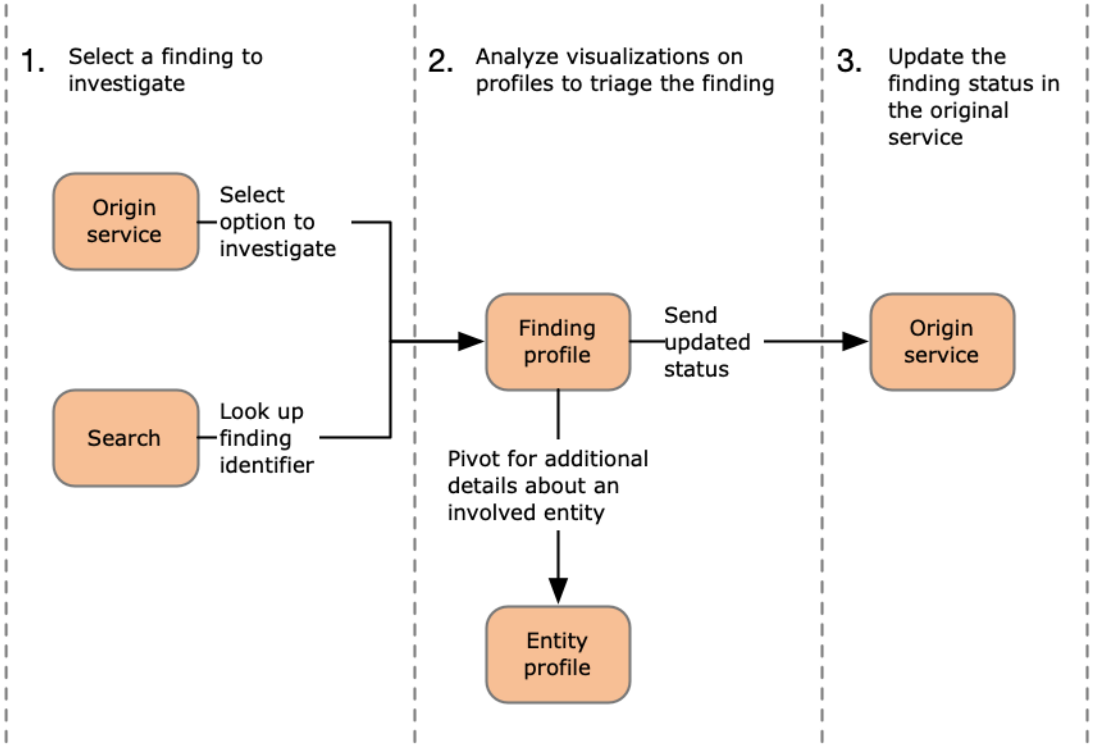
1. Select a finding to investigate
2. Analyze visualizations on profiles to triage the finding
3. Update the finding status in the original service
Origin service
Select option to investigate
Search
Look up finding identifier
Finding profile
Send updated status
Origin service
Pivot for additional details about an involved entity
Entity profile

Step 1: Select a finding to investigate

When you look at a finding in Amazon GuardDuty or AWS Security Hub, you can choose to investigate the finding in Detective. See the section called “Pivoting from another console” (p. 15).

From within Detective, you can use the Detective search function to find and select a finding to triage. See *Searching for a finding or entity* (p. 20).

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

Step 2: Analyze visualizations on profiles

The finding 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.

Most of the visualizations show activity that is related to the entity or entities involved in the finding. You use these visualizations to answer questions that are critical to completing the triage of the finding. See *Analyzing finding details* (p. 26).

To help guide the triage, 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” (p. 51).

From the finding profile, you can pivot to entity profiles to delve deeper into a specific asset that is involved with the finding. See *Analyzing entity details* (p. 28).

Step 3: Update the finding status

Once you determine whether a finding is a true or false positive, you update the finding status in the original service. For GuardDuty findings, Detective provides an option to archive the finding. See *Archiving a GuardDuty finding* (p. 55).

Data in a behavior graph

In Detective, you conduct investigations using data from a Detective behavior graph.

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 specific questions that arise as you conduct investigations

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

For information about how an administrator account manages the member accounts in a behavior graph, see For administrator accounts: Managing the accounts in your behavior graph in *Detective Administration Guide*.

Contents

How Amazon Detective uses source data to populate a behavior graph

To provide the raw material 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 uncovered by Amazon GuardDuty

To learn more about the source data used in a behavior graph, see Source data used in a behavior graph in *Detective Administration Guide*.

How Detective processes source data

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

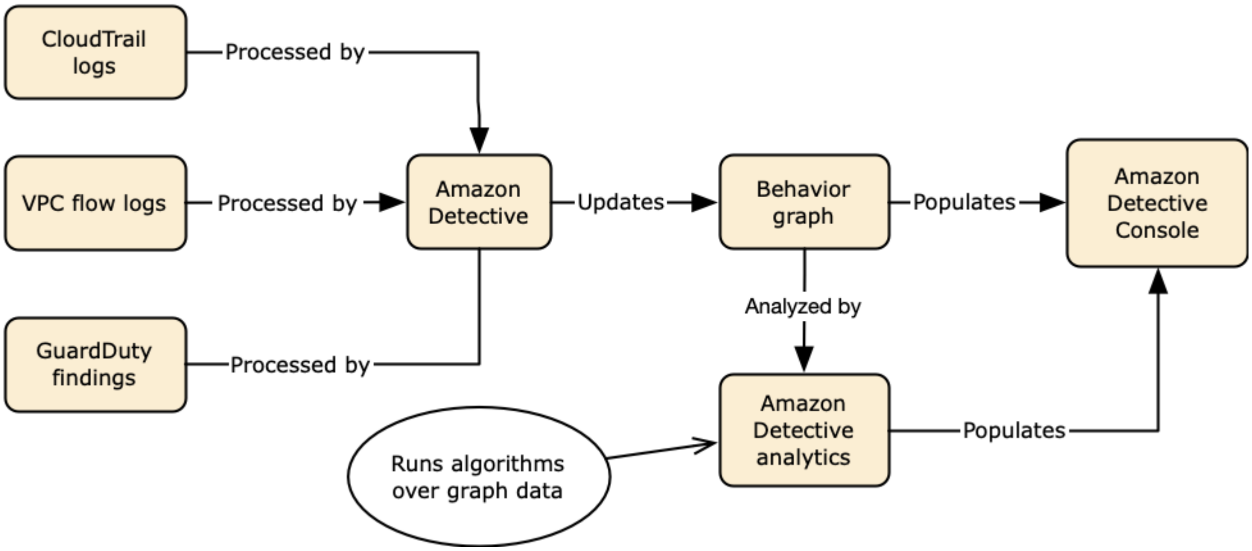

Detective extraction

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 dig deeper into the data to provide insight into activity that is associated with entities.

For example, one type of Detective analytic analyzes how often activity occurs. For entities that make API calls, the analytic looks for API calls that the entity doesn't normally use. The analytic 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 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.

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 free trial, 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 Amazon 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 Amazon 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?

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?

Supported finding types

Amazon Detective only ingests and provides profiles for the Amazon GuardDuty finding types that are listed below. GuardDuty detects some of these findings from CloudTrail data, and some from VPC flow data.

For a detailed description of each finding type, see the list of finding types in the *Amazon GuardDuty User Guide*

AWS CloudTrail-based findings

These findings are detected using CloudTrail data:

`CredentialAccess:IAMUser/AnomalousBehavior`

An API used to gain access to an AWS environment was invoked in an anomalous way.

`DefenseEvasion:IAMUser/AnomalousBehavior`

An API used to evade defensive measures was invoked in an anomalous way.

`Discovery:IAMUser/AnomalousBehavior`

An API commonly used to discover resources was invoked in an anomalous way.

`Exfiltration:IAMUser/AnomalousBehavior`

An API commonly used to collect data from an AWS environment was invoked in an anomalous way.

`Impact:IAMUser/AnomalousBehavior`

An API commonly used to tamper with data or processes in an AWS environment was invoked in an anomalous way.

`InitialAccess:IAMUser/AnomalousBehavior`

An API commonly used to gain unauthorized access to an AWS environment was invoked in an anomalous way.

`PenTest:IAMUser/KaliLinux`

An API was invoked from a Kali Linux EC2 instance.

`PenTest:IAMUser/ParrotLinux`

An API was invoked from a Parrot Security Linux EC2 instance.

`PenTest:IAMUser/PentooLinux`

An API was invoked from a Pentoo Linux EC2 instance.

`Persistence:IAMUser/AnomalousBehavior`

An API commonly used to maintain unauthorized access to an AWS environment was invoked in an anomalous way.

`Persistence:IAMUser/NetworkPermissions`

A principal invoked an API commonly used to change the network access permissions for security groups, routes, and ACLs in your AWS account.

`Persistence:IAMUser/ResourcePermissions`

A principal invoked an API commonly used to change the security access policies of various resources in your AWS account.

`Persistence:IAMUser/UserPermissions`

A principal invoked an API commonly used to add, modify, or delete IAM users, groups, or policies in your AWS account.

`Policy:IAMUser/RootCredentialUsage`

An API was invoked using root credentials.

`PrivilegeEscalation:IAMUser/AdministrativePermissions`

A principal has attempted to assign a highly permissive policy to themselves.

`PrivilegeEscalation:IAMUser/AnomalousBehavior`

An API commonly used to used to obtain high-level permissions to an AWS environment was invoked in an anomalous way.

`Recon:IAMUser/MaliciousIPCaller`

An API was invoked from a known malicious IP address.

`Recon:IAMUser/MaliciousIPCaller.Custom`

An API was invoked from an IP address on a custom threat list.

`Recon:IAMUser/NetworkPermissions`

A principal invoked an API commonly used to discover the network access permissions of existing security groups, ACLs, and routes in your AWS account.

`Recon:IAMUser/ResourcePermissions`

A principal invoked an API commonly used to discover the permissions associated with various resources in your AWS account.

`Recon:IAMUser/TorIPCaller`

An API was invoked from a Tor exit node IP address

`Recon:IAMUser/UserPermissions`

A principal invoked an API commonly used to discover the users, groups, policies and permissions in your AWS account.

`ResourceConsumption:IAMUser/ComputeResources`

A principal invoked an API commonly used to launch compute resources such as EC2 instances.

`Stealth:IAMUser/CloudTrailLoggingDisabled`

AWS CloudTrail trail was disabled.

`Stealth:IAMUser/LoggingConfigurationModified`

A principal invoked an API commonly used to stop CloudTrail logging, delete existing logs, and otherwise eliminate traces of activity in your AWS account.

`Stealth:IAMUser/PasswordPolicyChange`

Account password policy was weakened.

`UnauthorizedAccess:IAMUser/ConsoleLogin`

An unusual console sign-in by a principal in your AWS account was observed.

`UnauthorizedAccess:IAMUser/ConsoleLoginSuccess.B`

Multiple worldwide successful console sign-ins were observed.

`UnauthorizedAccess:IAMUser/InstanceCredentialExfiltration`

Credentials that were created exclusively for an EC2 instance through an instance launch role are being used from an external IP address.

`UnauthorizedAccess:IAMUser/MaliciousIPCaller`

An API was invoked from a known malicious IP address.

`UnauthorizedAccess:IAMUser/MaliciousIPCaller.Custom`

EC2 instance is communicating outbound with an IP address that is on a custom threat list.

`UnauthorizedAccess:IAMUser/TorIPCaller`

EC2 instance is receiving inbound connections from a Tor exit node.

VPC flow-based findings

These findings are detected using VPC flow data:

`Backdoor:EC2/C&CActivity.B`

An EC2 instance is querying an IP address that is associated with a known command and control server.

`Backdoor:EC2/DenialOfService.Dns`

An EC2 instance is behaving in a manner that may indicate it is being used to perform a Denial of Service (DoS) attack using the DNS protocol.

`Backdoor:EC2/DenialOfService.Tcp`

An EC2 instance is behaving in a manner that may indicate it is being used to perform a Denial of Service (DoS) attack using the TCP protocol.

`Backdoor:EC2/DenialOfService.Udp`

An EC2 instance is behaving in a manner that may indicate it is being used to perform a Denial of Service (DoS) attack using the UDP protocol.

`Backdoor:EC2/DenialOfService.UdpOnTcpPorts`

An EC2 instance is behaving in a manner that may indicate it is being used to perform a Denial of Service (DoS) attack using the UDP protocol on a TCP port.

`Backdoor:EC2/DenialOfService.UnusualProtocol`

An EC2 instance is behaving in a manner that may indicate it is being used to perform a Denial of Service (DoS) attack using an unusual protocol.

`Backdoor:EC2/Spambot`

EC2 instance is exhibiting unusual behavior by communicating with a remote host on port 25.

`Behavior:EC2/NetworkPortUnusual`

EC2 instance is communicating with a remote host on an unusual server port.

`Behavior:EC2/TrafficVolumeUnusual`

EC2 instance is generating unusually large amounts of network traffic to a remote host.

`CryptoCurrency:EC2/BitcoinTool.B`

EC2 instance is querying an IP address that is associated with cryptocurrency-related activity.

`Impact:EC2/PortSweep`

An EC2 instance is probing a port on a large number of IP addresses.

`Impact:EC2/WinRMBruteForce`

An EC2 instance is performing an outbound Windows Remote Management brute force attack.

`Recon:EC2/PortProbeEMRUnprotectedPort`

EC2 instance in an EMR cluster has an unprotected Amazon EMR-related, sensitive port that is being probed by a known malicious host.

`Recon:EC2/PortProbeUnprotectedPort`

EC2 instance has an unprotected port that is being probed by a known malicious host.

`Recon:EC2/Portscan`

EC2 instance is performing outbound port scans to a remote host.

`Trojan:EC2/BlackholeTraffic`

EC2 instance is attempting to communicate with an IP address of a remote host that is a known black hole.

`Trojan:EC2/DropPoint`

An EC2 instance is attempting to communicate with an IP address of a remote host that is known to hold credentials and other stolen data captured by malware.

`UnauthorizedAccess:EC2/MaliciousIPCaller.Custom`

EC2 instance is communicating outbound with an IP address on a custom threat list.

`UnauthorizedAccess:EC2/RDPBruteForce`

EC2 instance has been involved in RDP brute force attacks.

`UnauthorizedAccess:EC2/SSHBruteForce`

EC2 instance has been involved in SSH brute force attacks.

`UnauthorizedAccess:EC2/TorClient`

EC2 instance is making connections to a Tor Guard or an Authority node.

`UnauthorizedAccess:EC2/TorIPCaller`

EC2 instance is receiving inbound connections from a Tor exit node.

`UnauthorizedAccess:EC2/TorRelay`

EC2 instance is making connections to a Tor network as a Tor relay.

Navigating directly to a profile

To navigate directly to a profile 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.

Contents

Pivoting to a profile from Amazon GuardDuty or AWS Security Hub

From the Amazon GuardDuty and AWS Security Hub consoles, you can navigate to Amazon Detective finding profiles. From GuardDuty, you can also navigate to the entity profile for an entity that is related to a finding.

These links can help to streamline the investigation process. When a finding might be a genuine cause for concern, you can quickly use Detective to see the associated resource activity and determine next steps. You can then archive the 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

In Security Hub, the investigation links only work for GuardDuty finding types that Detective supports. See the section called "Supported finding types" (p. 11).

In GuardDuty, the investigation links are available for all GuardDuty findings. GuardDuty presents different investigation options based on whether Detective supports the finding type.

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 always contains related entities, such as IP addresses or EC2 instances associated with the finding.

   If the finding type is supported in Detective, then the list also includes 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**.

   The link is only available for finding types that Detective supports.

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

   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 in *Detective Administration Guide*.

If the pivot does not work, check the following.

- **Does Detective support that finding type?** If the finding type is not one of the types listed in the section called “Supported finding types” (p. 11), then the behavior graph does not contain data for it.
- **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 a profile using a URL

To navigate to a finding or entity profile 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.

Format of a profile URL

The format of the URL is as follows:

https://console.aws.amazon.com/detective/home?
region=*Region*#*type*/*namespace*/*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.

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

namespace

Used to identify the type of finding or entity.

For findings, the namespace identifies the finding provider. For example, for Amazon GuardDuty findings, the namespace is `GuardDuty`.

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

- `AwsAccount`
- `AwsRole`
- `AwsRoleSession`
- `AwsUser`
- `Ec2Instance`
- `FederatedUser`
- `IpAddress`
- `UserAgent`

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>*.

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 *Managing the scope time for profiles* (p. 24).

`scopeStart`

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

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 specify the correct type (`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, does Detective support that finding type?** If the finding type is not one of the types listed in the section called “Supported finding types” (p. 11), then the behavior graph does not contain data for it.
- **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.

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 a finding or entity profile.

Completing the search

To complete the search, you choose the type of entity to search for. Then provide the identifier to search for.

For each entity type, the following identifiers are supported.

- For AWS accounts, the account ID.
- For IP addresses, the address.
- For AWS roles and AWS users, either the principal ID, the name, or the 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.
- For EC2 instances, the instance identifier or the ARN.
- 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 findings, the finding identifier or finding ARN.

  The finding type must be one that Detective supports. See the section called "Supported finding types" (p. 11).

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 are 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 identifier to search for.

   The search is case sensitive.

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 profile for the finding or entity, 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, does Detective support that finding type?** If the finding type is not one of the types listed in the section called “Supported finding types” (p. 11), then the behavior graph does not contain data for it.
- **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.

Using the Summary page to identify an entity of interest

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

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

- 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

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

Newly observed geolocations in the past 24 hours

Newly observed geolocations in the past 24 hours 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.

Roles and users with the most API call volume in the past 24 hours

Roles and users with the most API call volume in the past 24 hours 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.

There is also a timeline of the API call volume for the previous seven 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 pivot to the profile of the account associated with the user or role. To pivot to a profile, choose the user, role, or account identifier.

EC2 instances with the most traffic volume in the past 24 hours

EC2 instances with the most traffic volume in the past 24 hours 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 also see a timeline showing the inbound and outbound traffic over the previous seven days. The timeline can help determine whether the volume of traffic is unusual for that EC2 instance.

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

Approximate value notification

On **Roles and users with the most API call volume in the past 24 hours** and **EC2 instances with the most traffic volume in the past 24 hours**, 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.

Managing the scope time used on finding and entity profiles

The charts, timelines, and other data displayed on finding and entity profiles are all based on the current scope time, which appears at the top right of each profile. 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. All times are displayed in UTC.

Detective analytics use 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.

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 indeed unusual.

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 nor 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. On a finding or entity profile, choose the scope time.
2. 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.

   Remember that in Detective, all times are in UTC.
3. 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.
4. When you're finished editing, to save the changes and update the displayed data, choose **Update scope time**.

Selecting a length of time from the current time

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

To set the scope time length

1. From a finding or entity profile, choose the scope time.
2. 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.

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

Setting the scope time to the finding time window

Each finding has an associated time window, which reflects the first and last times the finding was observed. When you navigate to a finding profile, if the current scope time does not match the finding time window, a warning is displayed.

When you edit the scope time from a finding profile, you can align the scope time to the finding time window.

To align the scope time to the finding time window

1. On a finding profile, choose the scope time.
2. On the **Edit scope time** panel, choose **Align scope time to start and end of current finding**.
3. When you're finished editing, to save the changes and update the displayed data, choose **Update scope time**.

Analyzing finding details

A finding is a possible instance of malicious activity or other risk that was detected by Amazon GuardDuty. Findings are loaded into Amazon Detective so that you can use Detective to investigate them. Detective ingests and provides profiles for a selected set of GuardDuty finding types. See the section called "Supported finding types" (p. 11).

A Detective finding profile is a single page that provides a collection of data visualizations plus supporting guidance. Finding profiles assist with the triage and scoping phases of an investigation.

How to display a finding profile

A finding profile appears when you perform one of the following actions:

- For findings that are loaded into Detective, pivot to Detective from the finding details in GuardDuty or Security Hub.

  See the section called "Pivoting from another console" (p. 15).
- Navigate to the Detective URL for the finding profile.

  See the section called "Navigating using a URL" (p. 16).
- Use the Detective search to look up a finding.

  See *Searching for a finding or entity* (p. 20).
- Choose a link to the finding profile from another entity or finding profile.

Scope time used for the finding profile

When you navigate directly to a finding profile without providing a scope time, the scope time is set to the finding time window. The finding time window reflects the first and last time that the finding activity was observed.

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

All times are in UTC.

If the current scope time does not match the finding time window, a warning is displayed.

For information on setting the scope time, see *Managing the scope time for profiles* (p. 24). When you edit the scope time from a finding profile, you can choose to align the scope time to the finding time window.

Finding title and type

At the top of the profile are the finding title and the finding type. The icon next to the title provides a visual cue to the service that detected the finding.

Profile panels containing finding details and analytics results

A finding 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 tabs and profile panels are tailored to the finding type. The profile panels support the investigation process by providing the critical information that you need to determine how to respond the finding.

The profile panels focus on answering specific questions that an analyst might want to ask when investigating the finding. For example, a finding might involve an AWS role or an IP address. The finding profile panels then highlight the AWS role or IP address activity that contributed to the finding. Each profile panel provides access to guidance on how to use the information. For more information, see the section called “Using profile panel guidance” (p. 51).

For the involved entity, the profile displays a list of all of the findings that the entity was involved in around the scope time. See the section called “Viewing findings for an entity” (p. 51).

For more details about profile panels, the types of data they contain, and available options for interacting with them, see *Viewing and interacting with profile panels* (p. 31).

Analyzing entity details

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" (p. 8).

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.

How to display an entity profile

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

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

  See the section called "Pivoting from another console" (p. 15).
- Go to the Detective URL for the entity profile.

  See the section called "Navigating using a URL" (p. 16).
- Use the Detective search to look up an entity.

  See *Searching for a finding or entity* (p. 20).
- Choose a link to the entity profile from another entity or finding profile.

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 profile, the currently selected scope time remains in place.

All times are displayed in UTC.

For information on setting the scope time, see *Managing the scope time for profiles* (p. 24).

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.

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.

Each profile also contains a panel that lists the findings that the entity was involved in around the scope time. See the section called “Viewing findings for an entity” (p. 51).

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” (p. 51).

For more details about profile panels, the types of data they contain, and available options for interacting with them, see *Viewing and interacting with profile panels* (p. 31).

Navigating in a profile

A finding or 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:

- Finding or entity type
- Finding or entity identifier
- Scope time

When the tabs are no longer visible, the selected tab is added to the links at the top of the page.

To navigate to a different tab, choose the selected tab. Then choose the tab to navigate to.

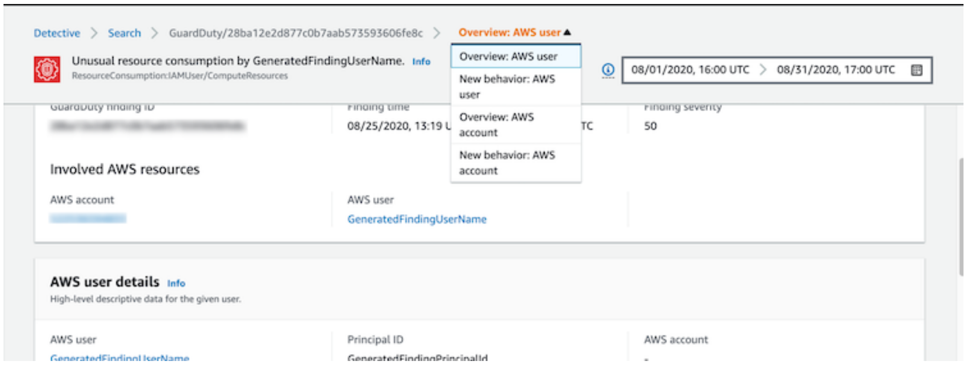
Detective > Search > GuardDuty/28ba12e2d877c0b7aab573593606fe8c > Overview: AWS user ▲
Overview: AWS user
New behavior: AWS user
Overview: AWS account
New behavior: AWS account
Info
08/01/2020, 16:00 UTC > 08/31/2020, 17:00 UTC

Unusual resource consumption by GeneratedFindingUserName.
ResourceConsumption:IAMUser/ComputeResources

GuardDuty finding ID
[redacted]

Finding time
08/25/2020, 13:19 U... TC

Finding severity
50

Involved AWS resources

AWS account
[redacted]

AWS user
GeneratedFindingUserName

AWS user details Info
High-level descriptive data for the given user.

AWS user
GeneratedFindingUserName

Principal ID
GeneratedFindingPrincipalId

AWS account
-

Viewing and interacting with profile panels

Each finding or 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 a finding or 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 findings, the involved entities, and their associated activity. The answers to those questions help lead to a conclusion about whether a finding represents a genuine threat.

For finding profiles, most of the profile panels contain information related to the involved entity. For example, a finding involves the suspected misuse of an Amazon Web Services (AWS) role. For that finding, the finding profile panels analyze how that role was used around the time of the finding activity.

Contents

Profile panel content

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
High-level information about a finding or entity	The simplest type of panel provides some basic information about a finding or entity. Examples of information included on an information panel include the identifier, name, type, and creation date. 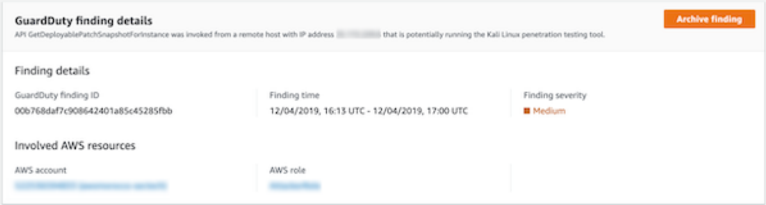 **GuardDuty finding details** API GetDeployablePatchSnapshotForInstance was invoked from a remote host with IP address [redacted] that is potentially running the Kali Linux penetration testing tool. Archive finding Finding details GuardDuty finding ID: 00b768daf7c908642401a85c45285fbb Finding time: 12/04/2019, 16:13 UTC - 12/04/2019, 17:00 UTC Finding severity: ■ Medium Involved AWS resources AWS account: [redacted] AWS role: [redacted]

Panel data type	Description
	Most finding and entity profiles contain an information panel for that finding or entity. Finding profiles can also include an information panel about the involved entity.
General summary of activity over time	Displays a summary of activity for an entity over time. This type of panel provides an overall view of how an entity is behaving during the scope time. 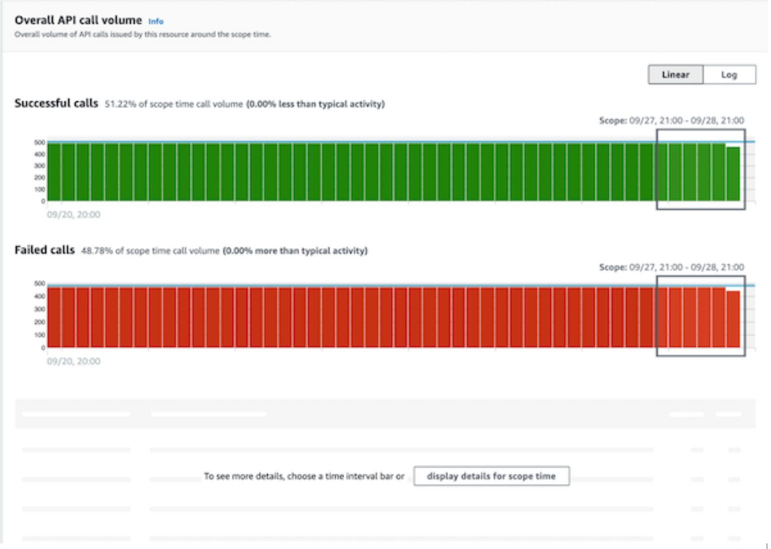 **Overall API call volume** Overall volume of API calls issued by this resource around the scope time. **Successful calls** 51.22% of scope time call volume **(0.00% less than typical activity)** **Scope:** 09/27, 21:00 - 09/28, 21:00 **Failed calls** 48.78% of scope time call volume **(0.00% more than typical activity)** **Scope:** 09/27, 21:00 - 09/28, 21:00 To see more details, choose a time interval bar or Info Linear Log 500 400 300 200 100 0 09/20, 20:00 500 400 300 200 100 0 09/20, 20:00 display details for scope time Here are some examples of summary data provided on Detective profile panels: • Failed and successful API calls • Inbound and outbound VPC volume

<table>
<tr><th>Panel data type</th><th>Description</th></tr>
<tr><td>Summary of activity grouped by values</td><td>

Displays a summary of activity for an entity, grouped by specific values.

You can see this type of profile panel on the profile for an Amazon EC2 instance. The profile panel shows the average volume of Amazon VPC flow data to and from an Amazon EC2 instance for common ports that are associated with specific types of services.

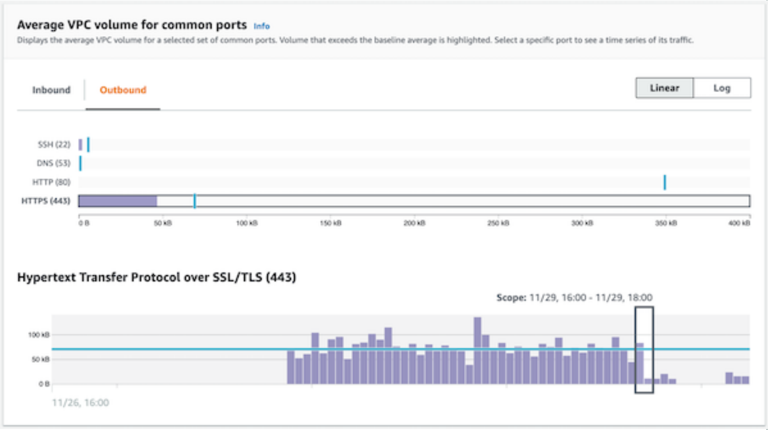

Average VPC volume for common ports

Displays the average VPC volume for a selected set of common ports. Volume that exceeds the baseline average is highlighted. Select a specific port to see a time series of its traffic.

Hypertext Transfer Protocol over SSL/TLS (443)

Scope: 11/29, 16:00 - 11/29, 18:00

</td></tr>
<tr><td>Activity that only started during the scope time</td><td>

During an investigation, it is valuable to see what activity only began to occur at the same time as the finding activity.

For example, are there API calls, geographic locations, or user agents that were not seen before?

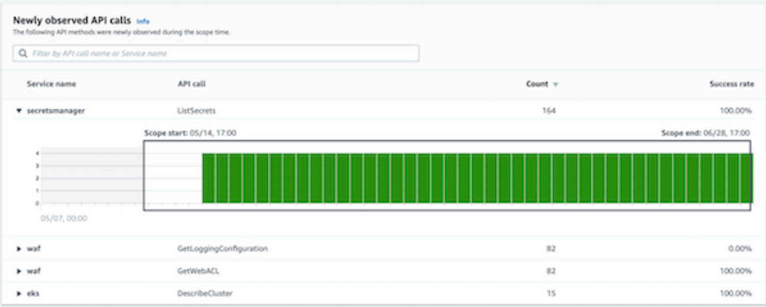

Newly observed API calls

The following API methods were newly observed during the scope time.

<table>
<tr><th>Service name</th><th>API call</th><th>Count</th><th>Success rate</th></tr>
<tr><td>secretsmanager</td><td>ListSecrets</td><td>164</td><td>100.00%</td></tr>
<tr><td>waf</td><td>GetLoggingConfiguration</td><td>82</td><td>0.00%</td></tr>
<tr><td>waf</td><td>GetWebACL</td><td>82</td><td>100.00%</td></tr>
<tr><td>eks</td><td>DescribeCluster</td><td>15</td><td>100.00%</td></tr>
</table>

Scope start: 05/14, 17:00

Scope end: 06/28, 17:00

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. 7).

</td></tr>
</table>

Panel data type	Description
Activity that changed significantly during the scope time	Similar to the new activity panels, profile panels can also display activity that changed significantly during the scope time. 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. 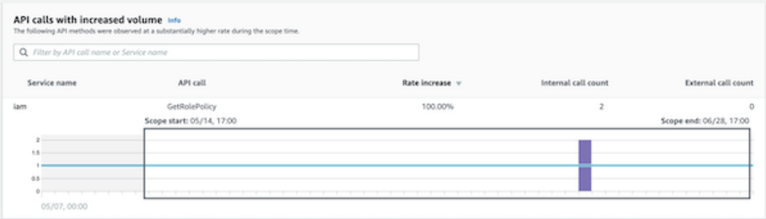 **API calls with increased volume** Info The following API methods were observed at a substantially higher rate during the scope time. Filter by API call name or Service name Service name · API call · Rate increase · Internal call count · External call count iam · GetRolePolicy · 100.00% · 2 · 0 Scope start: 05/14, 17:00 Scope end: 06/28, 17:00 2 1.5 1 0.5 0 05/07, 00:00 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. 7).

Types of profile panel visualizations

Profile panel content can take one of the following forms:

Visualization type	Description
Key-value pairs	The simplest type of visualization is a set of key-value pairs. A finding or entity information panel is the most common example of a key-value pair panel. 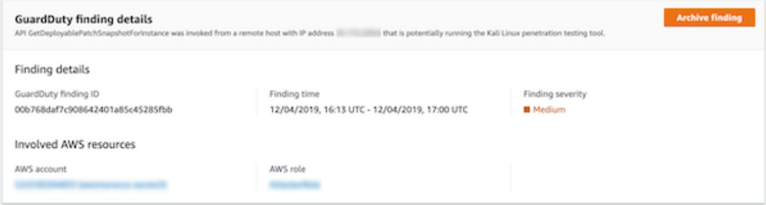 **GuardDuty finding details** API GetDeployablePatchSnapshotForInstance was invoked from a remote host with IP address [redacted] that is potentially running the Kali Linux penetration testing tool. Finding details GuardDuty finding ID: 00b768daf7c908642401a85c45285fbb Finding time: 12/04/2019, 16:13 UTC - 12/04/2019, 17:00 UTC Finding severity: Medium Involved AWS resources AWS account: [redacted] AWS role: [redacted] Key-value pairs can also be used to add additional information to other types of panels. From a key-value pair panel, if a value is an identifier of a finding or entity, then you can pivot to its profile.
Table	A table is a simple multiple-column list of items.

<table>
<tr><th>Visualization type</th><th>Description</th></tr>
<tr><td></td><td>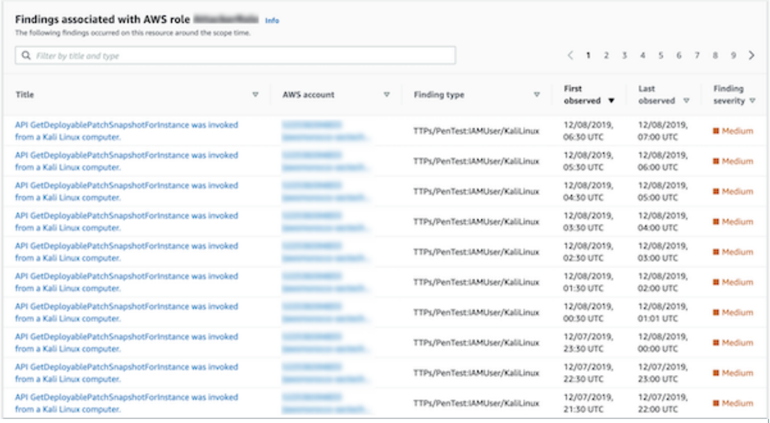
Findings associated with AWS role [redacted] Info
The following findings occurred on this resource around the scope time.
Filter by title and type
&lt; 1 2 3 4 5 6 7 8 9 &gt;
<table>
<tr><th>Title ▿</th><th>AWS account ▿</th><th>Finding type ▿</th><th>First observed ▼</th><th>Last observed ▿</th><th>Finding severity ▿</th></tr>
<tr><td>API GetDeployablePatchSnapshotForInstance was invoked from a Kali Linux computer.</td><td>[redacted]</td><td>TTPs/PenTest:IAMUser/KaliLinux</td><td>12/08/2019, 06:30 UTC</td><td>12/08/2019, 07:00 UTC</td><td>■ Medium</td></tr>
<tr><td>API GetDeployablePatchSnapshotForInstance was invoked from a Kali Linux computer.</td><td>[redacted]</td><td>TTPs/PenTest:IAMUser/KaliLinux</td><td>12/08/2019, 05:30 UTC</td><td>12/08/2019, 06:00 UTC</td><td>■ Medium</td></tr>
<tr><td>API GetDeployablePatchSnapshotForInstance was invoked from a Kali Linux computer.</td><td>[redacted]</td><td>TTPs/PenTest:IAMUser/KaliLinux</td><td>12/08/2019, 04:30 UTC</td><td>12/08/2019, 05:00 UTC</td><td>■ Medium</td></tr>
<tr><td>API GetDeployablePatchSnapshotForInstance was invoked from a Kali Linux computer.</td><td>[redacted]</td><td>TTPs/PenTest:IAMUser/KaliLinux</td><td>12/08/2019, 03:30 UTC</td><td>12/08/2019, 04:00 UTC</td><td>■ Medium</td></tr>
<tr><td>API GetDeployablePatchSnapshotForInstance was invoked from a Kali Linux computer.</td><td>[redacted]</td><td>TTPs/PenTest:IAMUser/KaliLinux</td><td>12/08/2019, 02:30 UTC</td><td>12/08/2019, 03:00 UTC</td><td>■ Medium</td></tr>
<tr><td>API GetDeployablePatchSnapshotForInstance was invoked from a Kali Linux computer.</td><td>[redacted]</td><td>TTPs/PenTest:IAMUser/KaliLinux</td><td>12/08/2019, 01:30 UTC</td><td>12/08/2019, 02:00 UTC</td><td>■ Medium</td></tr>
<tr><td>API GetDeployablePatchSnapshotForInstance was invoked from a Kali Linux computer.</td><td>[redacted]</td><td>TTPs/PenTest:IAMUser/KaliLinux</td><td>12/08/2019, 00:30 UTC</td><td>12/08/2019, 01:01 UTC</td><td>■ Medium</td></tr>
<tr><td>API GetDeployablePatchSnapshotForInstance was invoked from a Kali Linux computer.</td><td>[redacted]</td><td>TTPs/PenTest:IAMUser/KaliLinux</td><td>12/07/2019, 23:30 UTC</td><td>12/08/2019, 00:00 UTC</td><td>■ Medium</td></tr>
<tr><td>API GetDeployablePatchSnapshotForInstance was invoked from a Kali Linux computer.</td><td>[redacted]</td><td>TTPs/PenTest:IAMUser/KaliLinux</td><td>12/07/2019, 22:30 UTC</td><td>12/07/2019, 23:00 UTC</td><td>■ Medium</td></tr>
<tr><td>API GetDeployablePatchSnapshotForInstance was invoked from a Kali Linux computer.</td><td>[redacted]</td><td>TTPs/PenTest:IAMUser/KaliLinux</td><td>12/07/2019, 21:30 UTC</td><td>12/07/2019, 22:00 UTC</td><td>■ Medium</td></tr>
</table>

You can sort, filter, and page through the table.

You can change the number of entries to display per page. See the section called “Setting the number of table rows per page” (p. 38).

If a value in the table is an identifier of a finding or entity, then you can pivot to its profile.</td></tr>
<tr><td>Timeline</td><td>A timeline visualization shows an aggregated value for each time interval across time.
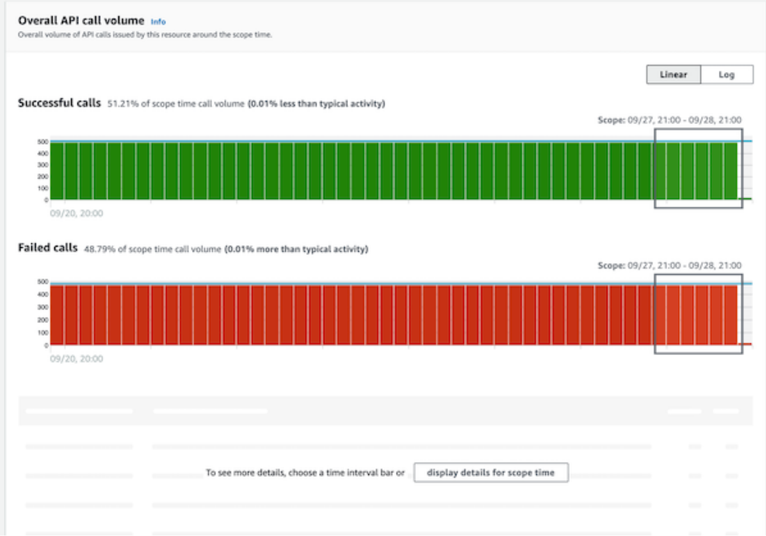
Overall API call volume Info
Overall volume of API calls issued by this resource around the scope time.
Successful calls 51.21% of scope time call volume (0.01% less than typical activity)
Scope: 09/27, 21:00 - 09/28, 21:00
Failed calls 48.79% of scope time call volume (0.01% more than typical activity)
Scope: 09/27, 21:00 - 09/28, 21:00
To see more details, choose a time interval bar or display details for scope time

Linear
Log
500
400
300
200
100
0
09/20, 20:00
500
400
300
200
100
0
09/20, 20:00

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.

Pause on a time interval to display a summary of the data for that time interval.</td></tr>
</table>

Visualization type	Description
Expandable table	An expandable table combines tables and timelines. 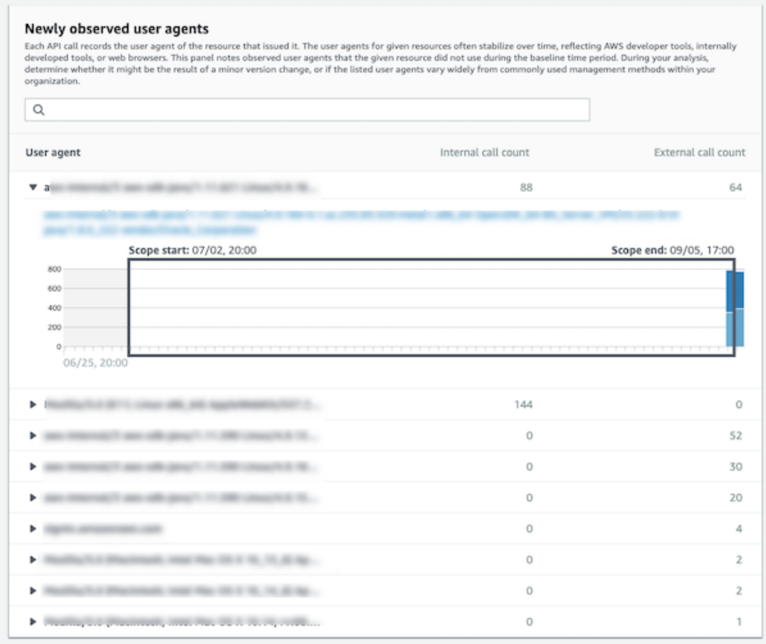 **Newly observed user agents** Each API call records the user agent of the resource that issued it. The user agents for given resources often stabilize over time, reflecting AWS developer tools, internally developed tools, or web browsers. This panel notes observed user agents that the given resource did not use during the baseline time period. During your analysis, determine whether it might be the result of a minor version change, or if the listed user agents vary widely from commonly used management methods within your organization. User agent — Internal call count — External call count ▼ a[redacted] — 88 — 64 [redacted] Scope start: 07/02, 20:00 Scope end: 09/05, 17:00 800 600 400 200 0 06/25, 20:00 ▶ [redacted] — 144 — 0 ▶ [redacted] — 0 — 52 ▶ [redacted] — 0 — 30 ▶ [redacted] — 0 — 20 ▶ [redacted] — 0 — 4 ▶ [redacted] — 0 — 2 ▶ [redacted] — 0 — 2 ▶ [redacted] — 0 — 1 The visualization starts as a table. You can sort, filter, and page through the table. You can change the number of entries to display per page. See the section called “Setting the number of table rows per page” (p. 38). You can then expand each row to show a timeline visualization specific to that row.
Bar chart	A bar chart shows values based on groupings. Depending on the chart, you might be able to choose a bar to display a timeline of related activity. 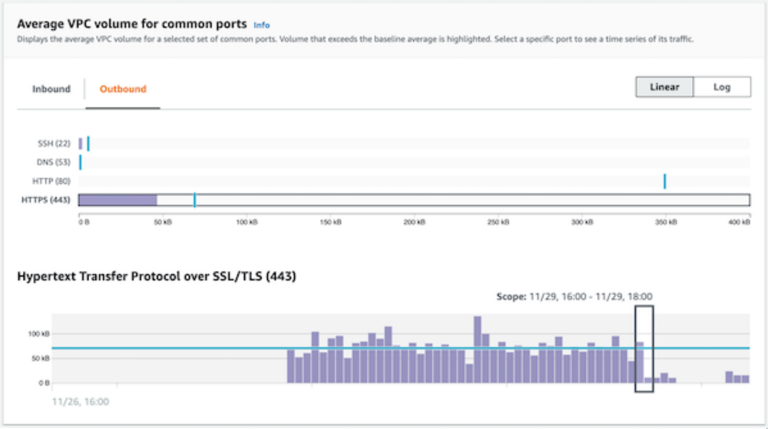 **Average VPC volume for common ports** Info Displays the average VPC volume for a selected set of common ports. Volume that exceeds the baseline average is highlighted. Select a specific port to see a time series of its traffic. Inbound Outbound Linear Log SSH (22) DNS (53) HTTP (80) HTTPS (443) 0 B 50 kB 100 kB 150 kB 200 kB 250 kB 300 kB 350 kB 400 kB **Hypertext Transfer Protocol over SSL/TLS (443)** Scope: 11/29, 16:00 - 11/29, 18:00 100 kB 50 kB 0 B 11/26, 16:00

<table>
<tr><th>Visualization type</th><th>Description</th></tr>
<tr><td>Geolocation chart</td><td>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.
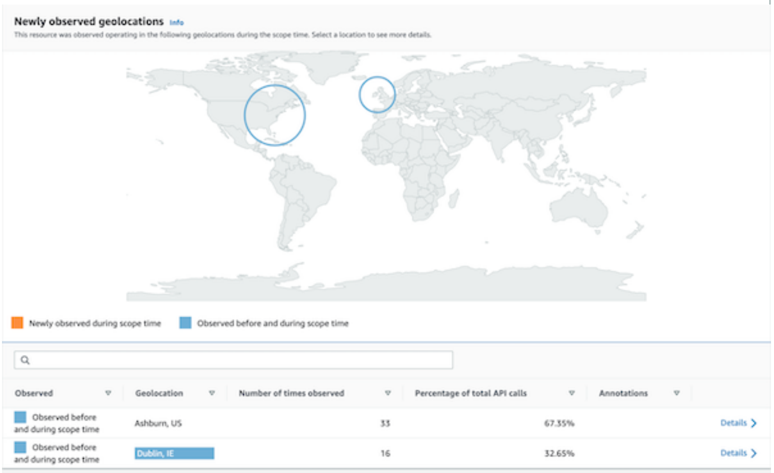
Newly observed geolocations Info
This resource was observed operating in the following geolocations during the scope time. Select a location to see more details.
🟧 Newly observed during scope time 🟦 Observed before and during scope time
<table><tr><th>Observed</th><th>Geolocation</th><th>Number of times observed</th><th>Percentage of total API calls</th><th>Annotations</th><th></th></tr><tr><td>🟦 Observed before and during scope time</td><td>Ashburn, US</td><td>33</td><td>67.35%</td><td></td><td>Details ›</td></tr><tr><td>🟦 Observed before and during scope time</td><td>Dublin, IE</td><td>16</td><td>32.65%</td><td></td><td>Details ›</td></tr></table>
Note that when processing incoming geographic data, Detective rounds the latitude and longitude values to a single decimal point.</td></tr>
</table>

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 *Managing the scope time for profiles* (p. 24).

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 that date, 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 number of rows per page for profile panel tables

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

To set your preference for the number of entries per page

1. In the Detective navigation pane, under **Settings**, choose **Preferences**.
2. On the **Preferences** page, under **Table preferences**, choose your preferred option.
3. Choose **Save changes**.

Pivoting from a profile panel to another console

For EC2 instances, AWS users, and AWS 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 or finding profile

When a profile panel contains an identifier of a different finding or entity, it is usually a link to that finding or entity profile. The exceptions are the links to the Amazon EC2 and IAM consoles on the EC2 instance, AWS user, and AWS role profiles. See the section called "Pivoting to another console" (p. 38).

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.

When you pivot to a finding profile, if the scope time does not match the finding time window, a warning is displayed.

For information on setting the scope time, see *Managing the scope time for profiles* (p. 24). When you edit the scope time from a finding profile, you can choose to align the scope time to the finding time window.

Exploring activity details on a profile panel

After you assess that a finding might be a true positive, the next step is to dig deeper into the pattern of activity for the related resources.

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**, except for the profile panel on the federated user profile
- **Overall VPC flow volume**
- **VPC flow volume to and from the finding IP address**, for findings that are associated with a single IP address

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 data was sent to or from each IP address?

Topics

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)

For users, roles, accounts, and role sessions, 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 called them. 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 also 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. 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.

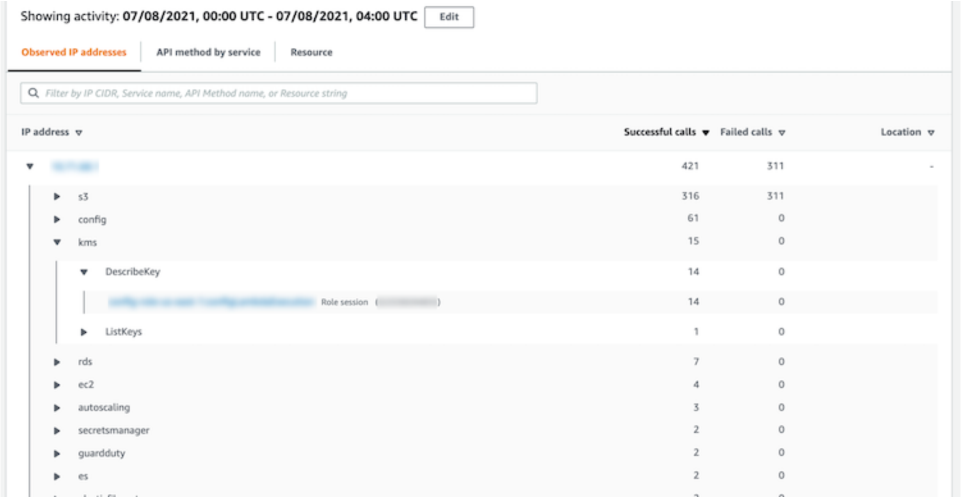

Showing activity: **07/08/2021, 00:00 UTC - 07/08/2021, 04:00 UTC**

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
▼ [redacted]	421	311	-
▶ s3	316	311	
▶ config	61	0	
▼ kms	15	0	
▼ DescribeKey	14	0	
[redacted] Role session ([redacted])	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. 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.

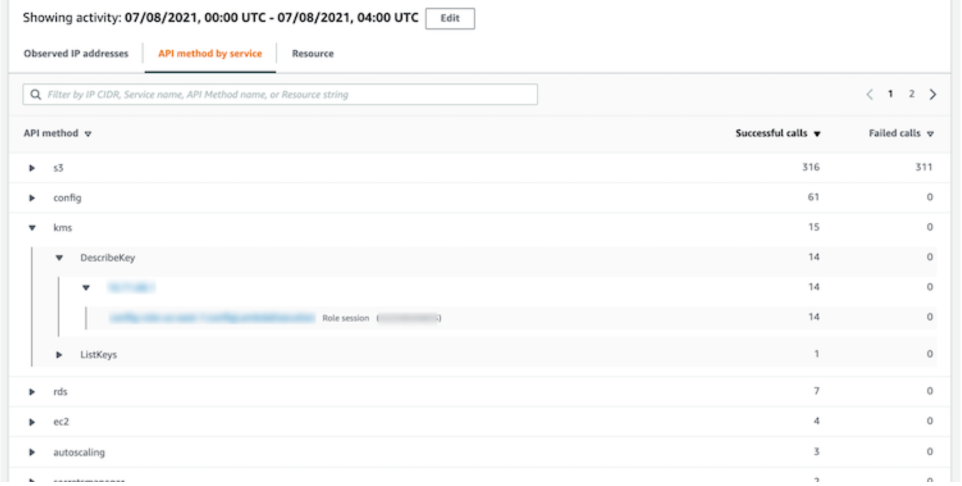

Showing activity: **07/08/2021, 00:00 UTC - 07/08/2021, 04:00 UTC**

Edit
Observed IP addresses
API method by service
Resource
Filter by IP CIDR, Service name, API Method name, or Resource string
< 1 2 >

API method	Successful calls	Failed calls
▶ s3	316	311
▶ config	61	0
▼ kms	15	0
▼ DescribeKey	14	0
▼ [redacted]	14	0
[redacted] Role session ([redacted])	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. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

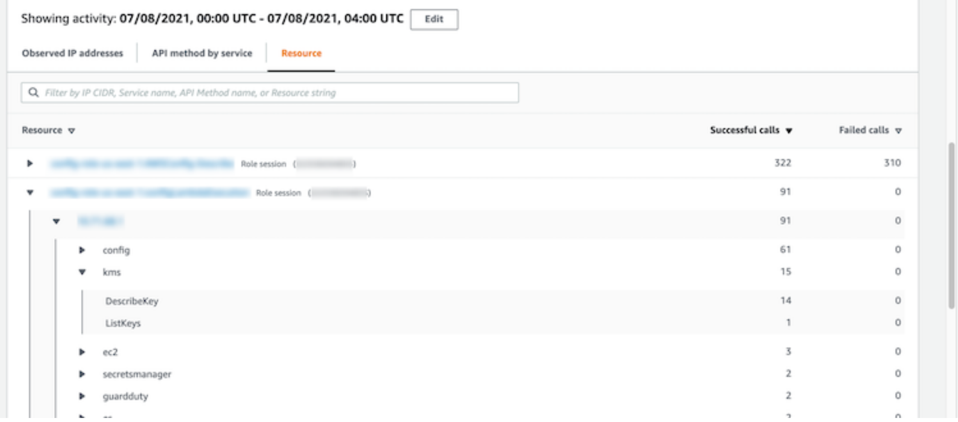

Showing activity: **07/08/2021, 00:00 UTC - 07/08/2021, 04:00 UTC** Edit

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
▶ [redacted] Role session ([redacted])	322	310
▼ [redacted] Role session ([redacted])	91	0
▼ [redacted]	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**.

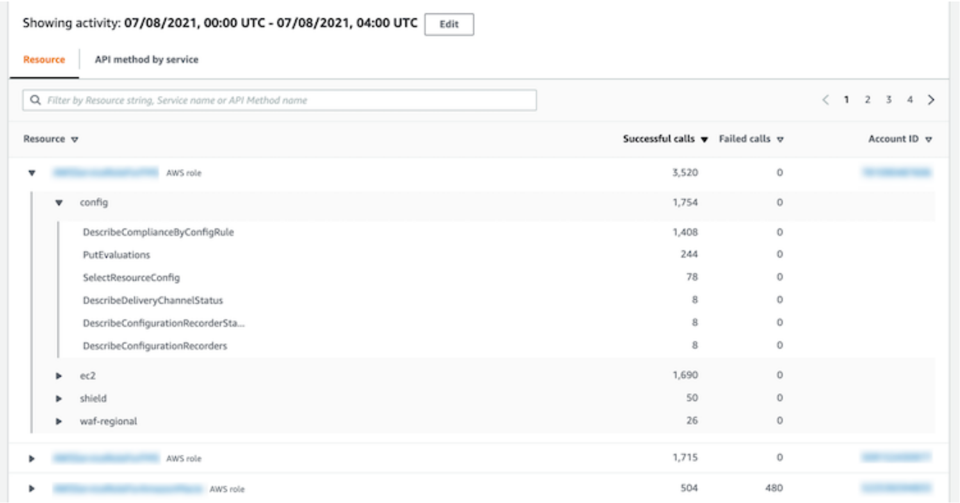

Showing activity: **07/08/2021, 00:00 UTC - 07/08/2021, 04:00 UTC**

Resource	Successful calls	Failed calls	Account ID
[redacted] AWS role	3,520	0	[redacted]
config	1,754	0	
DescribeComplianceByConfigRule	1,408	0	
PutEvaluations	244	0	
SelectResourceConfig	78	0	
DescribeDeliveryChannelStatus	8	0	
DescribeConfigurationRecorderSta...	8	0	
DescribeConfigurationRecorders	8	0	
ec2	1,690	0	
shield	50	0	
waf-regional	26	0	
[redacted] AWS role	1,715	0	[redacted]
[redacted] AWS role	504	480	[redacted]

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.

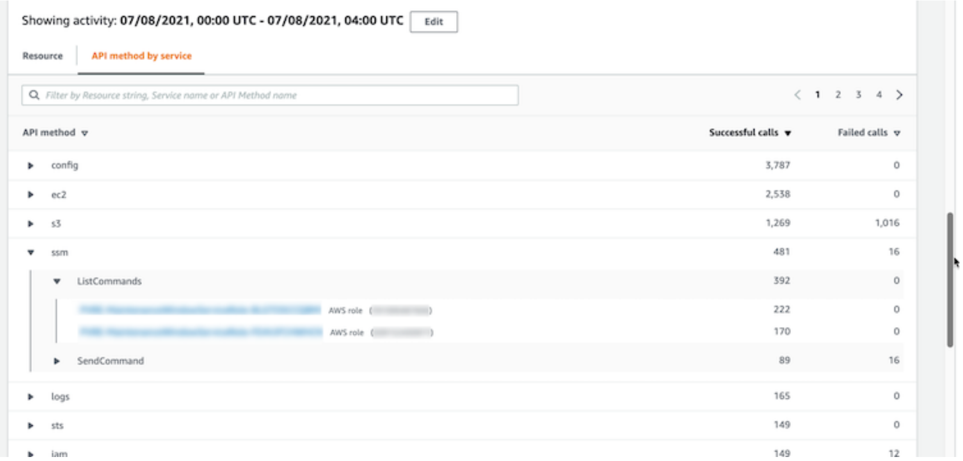

Showing activity: **07/08/2021, 00:00 UTC - 07/08/2021, 04:00 UTC**

API method	Successful calls	Failed calls
config	3,787	0
ec2	2,538	0
s3	1,269	1,016
ssm	481	16
ListCommands	392	0
[redacted] AWS role ([redacted])	222	0
[redacted] AWS role ([redacted])	170	0
SendCommand	89	16
logs	165	0
sts	149	0
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.

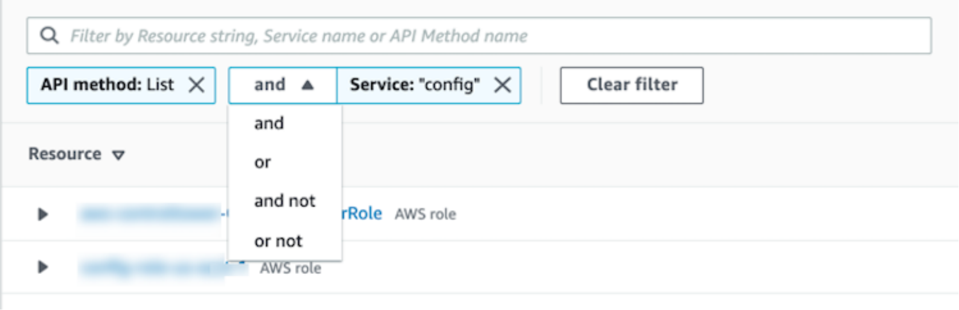

5. To remove a filter, choose the **x** icon in the top-right corner.
6. To clear all of the filters, choose **Clear filter**.

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.

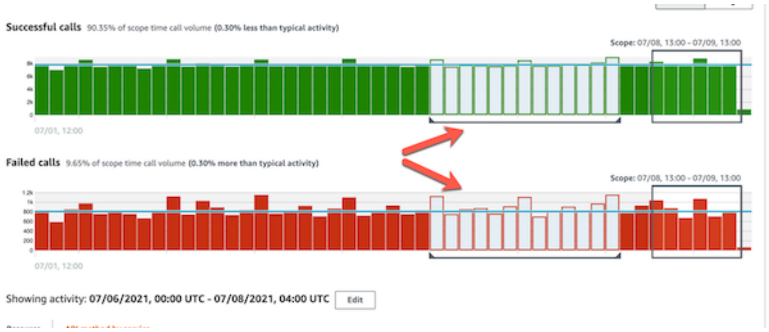
Successful calls 90.35% of scope time call volume (0.30% less than typical activity)
Scope: 07/08, 13:00 - 07/09, 13:00
8k
6k
4k
2k
0
07/01, 12:00
Failed calls 9.65% of scope time call volume (0.30% more than typical activity)
Scope: 07/08, 13:00 - 07/09, 13:00
1.2k
1k
800
600
400
200
0
07/01, 12:00
Showing activity: 07/06/2021, 00:00 UTC - 07/08/2021, 04:00 UTC
Edit

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.

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**.

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**.

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, pause on 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. If Detective cannot determine the service that issued a call, the call is listed under **Unknown service**.

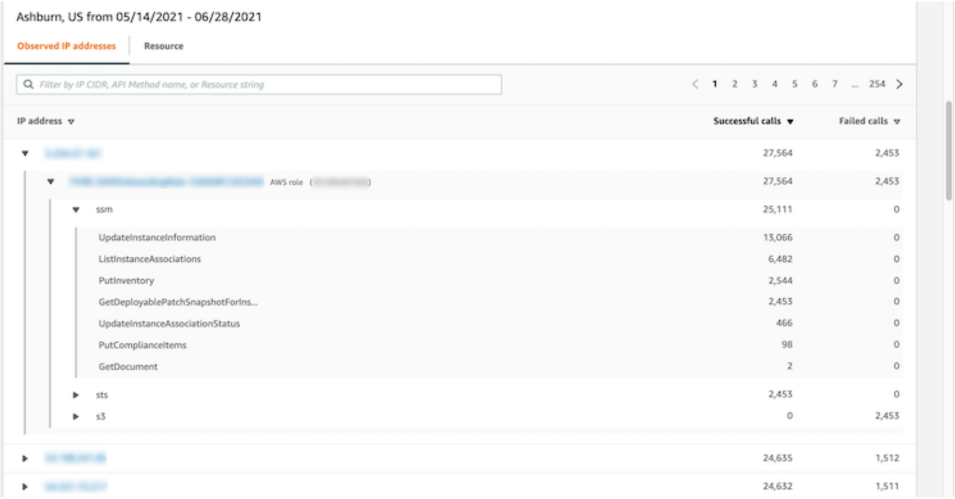

Ashburn, US from 05/14/2021 - 06/28/2021

IP address	Successful calls	Failed calls
[redacted]	27,564	2,453
[redacted] AWS role ([redacted])	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
[redacted]	24,635	1,512
[redacted]	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. 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.

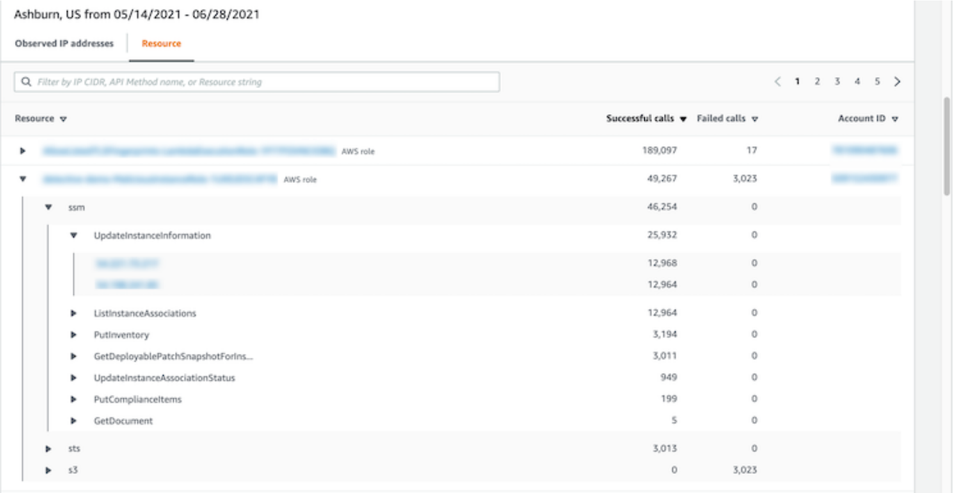

Ashburn, US from 05/14/2021 - 06/28/2021

Resource	Successful calls	Failed calls	Account ID
[redacted] AWS role	189,097	17	[redacted]
[redacted] AWS role	49,267	3,023	[redacted]
ssm	46,254	0	
UpdateInstanceInformation	25,932	0	
[redacted]	12,968	0	
[redacted]	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.

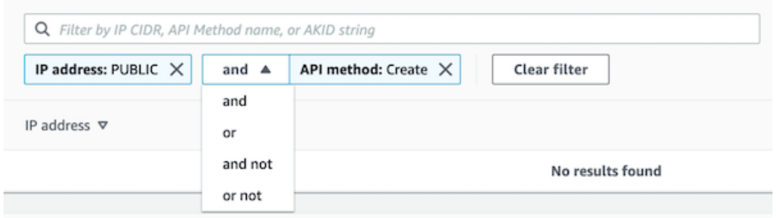

5. To remove a filter, choose the **x** icon in the top-right corner.
6. To clear all of the filters, choose **Clear filter**.

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 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.

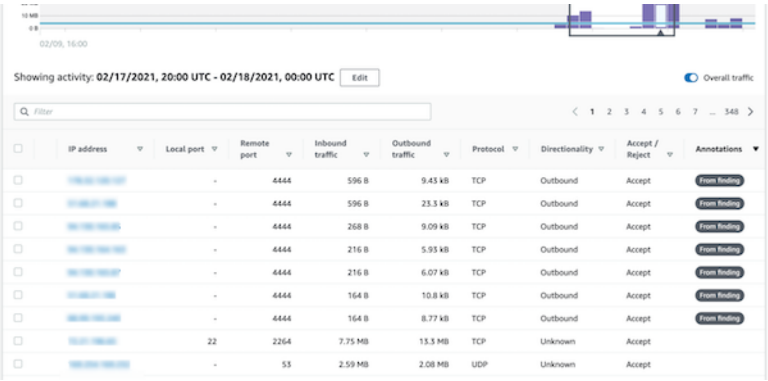

Showing activity: **02/17/2021, 20:00 UTC - 02/18/2021, 00:00 UTC** Edit

Overall traffic

Filter

< 1 2 3 4 5 6 7 ... 348 >

□	IP address	Local port	Remote port	Inbound traffic	Outbound traffic	Protocol	Directionality	Accept / Reject	Annotations
□	[redacted]	-	4444	596 B	9.43 kB	TCP	Outbound	Accept	From finding
□	[redacted]	-	4444	596 B	23.3 kB	TCP	Outbound	Accept	From finding
□	[redacted]	-	4444	268 B	9.09 kB	TCP	Outbound	Accept	From finding
□	[redacted]	-	4444	216 B	5.93 kB	TCP	Outbound	Accept	From finding
□	[redacted]	-	4444	216 B	6.07 kB	TCP	Outbound	Accept	From finding
□	[redacted]	-	4444	164 B	10.8 kB	TCP	Outbound	Accept	From finding
□	[redacted]	-	4444	164 B	8.77 kB	TCP	Outbound	Accept	From finding
□	[redacted]	22	2264	7.75 MB	13.3 MB	TCP	Unknown	Accept	
□	[redacted]	-	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.

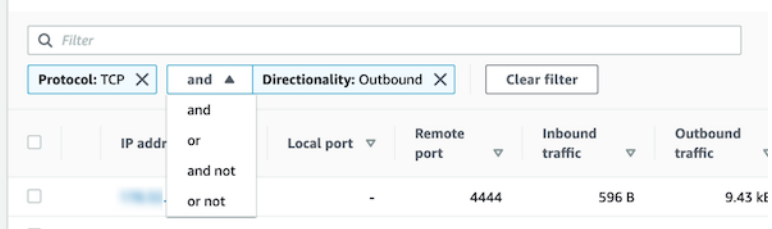

□	IP addr	Local port	Remote port	Inbound traffic	Outbound traffic
□	[redacted]	-	4444	596 B	9.43 kE

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.

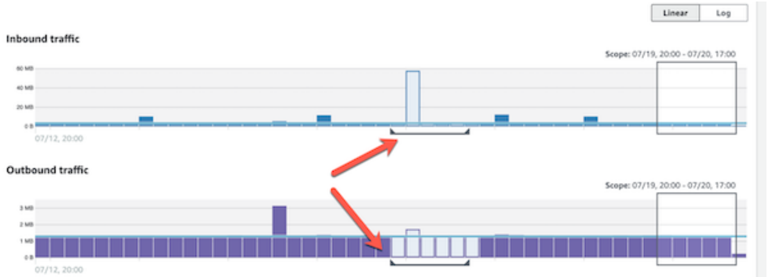

Displaying the volume of traffic for selected rows

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.

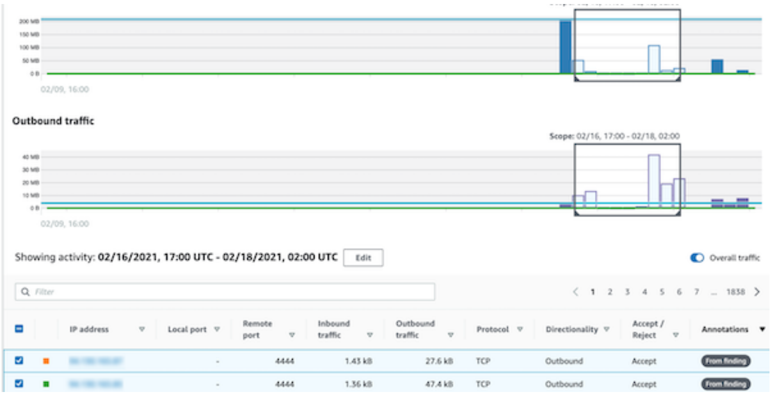

Showing activity: **02/16/2021, 17:00 UTC - 02/18/2021, 02:00 UTC** Edit

● Overall traffic

Filter

< 1 2 3 4 5 6 7 ... 1838 >

☐		IP address	Local port	Remote port	Inbound traffic	Outbound traffic	Protocol	Directionality	Accept / Reject	Annotations
☑	🟧	[redacted]	-	4444	1.43 kB	27.6 kB	TCP	Outbound	Accept	From finding
☑	🟩	[redacted]	-	4444	1.36 kB	47.4 kB	TCP	Outbound	Accept	From finding

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**.

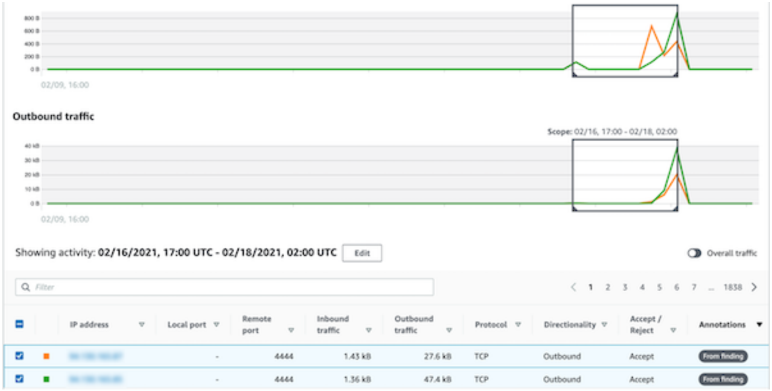

Showing activity: **02/16/2021, 17:00 UTC - 02/18/2021, 02:00 UTC**

☐		IP address	Local port	Remote port	Inbound traffic	Outbound traffic	Protocol	Directionality	Accept / Reject	Annotations
☑	🟧	[redacted]	-	4444	1.43 kB	27.6 kB	TCP	Outbound	Accept	From finding
☑	🟩	[redacted]	-	4444	1.36 kB	47.4 kB	TCP	Outbound	Accept	From finding

Activity details for VPC flow volume to and from the finding's IP address

For a finding that involves a single IP address and an EC2 instance, the activity details for **VPC flow volume to and from the finding's IP address** show the interactions between the finding EC2 instance and the finding IP address 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.

If the finding is associated with multiple IP addresses, then the profile panel does not provide activity details.

Content of the activity details

The content reflects the activity during the selected time range.

The activity details contain an entry for each unique combination of 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.

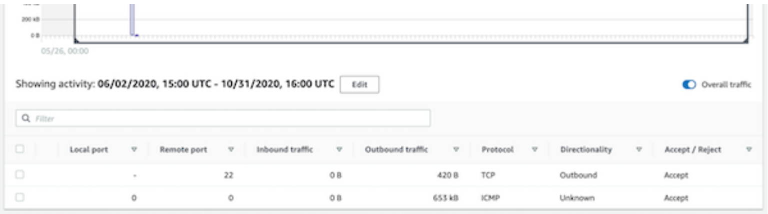

Showing activity: **06/02/2020, 15:00 UTC - 10/31/2020, 16:00 UTC**

☐	Local port	Remote port	Inbound traffic	Outbound traffic	Protocol	Directionality	Accept / Reject
☐	-	22	0 B	420 B	TCP	Outbound	Accept
☐	0	0	0 B	653 kB	ICMP	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 by the inbound traffic.

Filtering the activity details

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

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

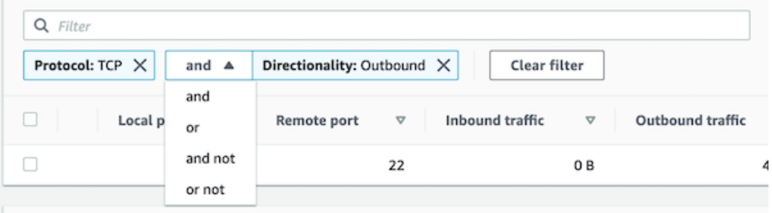

□	Local p	Remote port	Inbound traffic	Outbound traffic
□		22	0 B	4

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.
3. To set the time window to the default scope time for the profile, choose **Set to default scope time**.
4. 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.

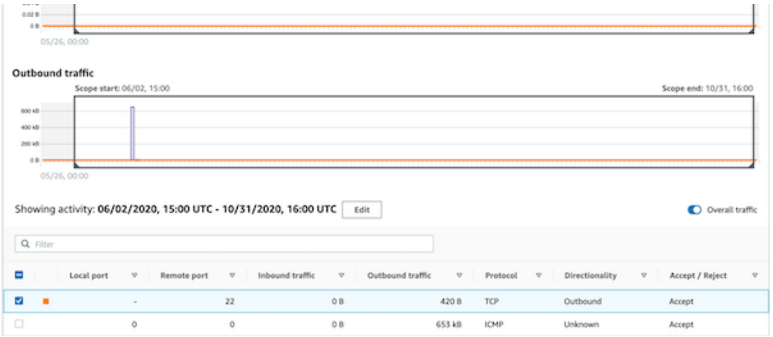

Showing activity: **06/02/2020, 15:00 UTC - 10/31/2020, 16:00 UTC**

☐		Local port	Remote port	Inbound traffic	Outbound traffic	Protocol	Directionality	Accept / Reject
☑	🟧	-	22	0 B	420 B	TCP	Outbound	Accept
☐		0	0	0 B	653 kB	ICMP	Unknown	Accept

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**.

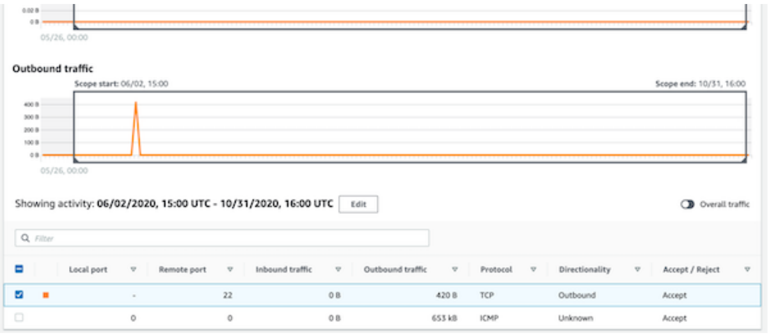

Showing activity: **06/02/2020, 15:00 UTC - 10/31/2020, 16:00 UTC** Edit

☐		Local port	Remote port	Inbound traffic	Outbound traffic	Protocol	Directionality	Accept / Reject
☑	🟧	-	22	0 B	420 B	TCP	Outbound	Accept
☐		0	0	0 B	653 kB	ICMP	Unknown	Accept

Viewing a list of findings that involve an entity

One indication that an entity has been compromised is its involvement in multiple findings.

Each finding and entity profile contains an associated findings profile panel. This profile panel is only populated if you have Amazon GuardDuty enabled. It also only includes findings that are supported by Detective. For the list of supported finding types, see the section called “Supported finding types” (p. 11).

For an entity profile, the panel lists findings that involved that entity.

For a finding profile, the panel lists findings for that finding's involved entity. For example, if a finding involves an AWS role, the associated findings contains a list of findings that involve that role. The list includes the finding from the profile.

The listed findings were observed during the current scope time plus additional peripheral time. This is the same time window that is displayed on timeline profile panels.

For each finding, the list includes the following information:

- The finding title, which is also a link to the finding profile
- 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

Using profile panel guidance during an investigation

Each profile panel is designed to provide answers to specific questions that arise as you investigate a finding 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

Viewing details for high-volume entities

In the behavior graph (p. 6), 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 Detective console.
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.

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. In the Detective console, on the finding profile, in the finding details panel, choose **Archive finding**.
2. When prompted to confirm, choose **Archive**.

Document history for Detective User Guide

The following table provides a history of the updates to this guide.

Change	Description	Date
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
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
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

Change	Description	Date
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 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` - `PrivilegeEscalation:IAMUser/AdministrativePermissions`	March 4, 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
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

Change	Description	Date
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 EC2 traffic volume.	January 21, 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.	January 21, 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 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
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

Change	Description	Date
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
Added time range selection for activity details on the Overall API call volume profile panel	On the **Overall API flow volume** profile 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 a 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 dropdown list in the breadcrumbs to navigate to a different tab.	September 4, 2020

Change	Description	Date
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
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
Links to other consoles on 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
New activity details for **Overvall VPC flow volume** profile panel	From the **Overall VPC flow volume** profile panel, you can now display activity details. The details show a list of interactions between the EC2 instance and IP addresses.	July 23, 2020
Amazon Detective general availability release	Detective is now generally available.	March 31, 2020
Introducing Amazon Detective (preview)	Detective uses machine learning and purpose-built visualizations to help you analyze and investigate security issues across your Amazon Web Services (AWS) workloads. Detective is currently in preview.	December 3, 2019