



User Guide

Amazon Fraud Detector



Version latest

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

Amazon Fraud Detector: User Guide

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

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

Table of Contents

| | |
|---|-----------|
| What is Amazon Fraud Detector? | 1 |
| Benefits | 1 |
| Core concepts and terms | 3 |
| How Amazon Fraud Detector works | 6 |
| Detecting fraud with Amazon Fraud Detector | 7 |
| Accessing Amazon Fraud Detector | 9 |
| Availability | 9 |
| Interfaces | 9 |
| Pricing | 10 |
| Set up for Amazon Fraud Detector | 11 |
| Sign up for AWS | 11 |
| Sign up for an AWS account | 11 |
| Create a user with administrative access | 12 |
| Set up permissions to access Amazon Fraud Detector interfaces | 13 |
| Set up interfaces to access Amazon Fraud Detector with | 14 |
| Access Amazon Fraud Detector console | 15 |
| Set up AWS CLI | 15 |
| Set up AWS SDK | 15 |
| Get started with Amazon Fraud Detector | 17 |
| Get and upload example dataset | 17 |
| Tutorial: Get started using the Amazon Fraud Detector console | 19 |
| Part A: Build, train, and deploy an Amazon Fraud Detector model | 19 |
| Part B: Generate fraud predictions | 23 |
| Tutorial: Get started using the AWS SDK for Python (Boto3) | 28 |
| Prerequisites | 29 |
| Get started | 29 |
| (Optional) Explore the Amazon Fraud Detector APIs with a Jupyter (iPython) Notebook | 38 |
| Next steps | 38 |
| Event dataset | 40 |
| Event dataset structure | 41 |
| Get event dataset requirements using the Data models explorer | 42 |
| Data models explorer | 42 |
| Gather event data | 43 |
| Dataset validation | 49 |

| | |
|--|-----------|
| Dataset storage | 50 |
| Event type | 51 |
| Create an event type | 51 |
| Create event type in the Amazon Fraud Detector console | 52 |
| Create an event type using the AWS SDK for Python (Boto3) | 53 |
| Delete an event or event type | 53 |
| Event data storage | 56 |
| Store your event data externally with Amazon S3 | 57 |
| Create CSV file | 57 |
| Upload your event data to an Amazon S3 bucket | 60 |
| Store your event data internally with Amazon Fraud Detector | 61 |
| Prepare event data for storage | 61 |
| Store event data using batch import | 63 |
| Store event data using the GetEventPredictions API operation | 76 |
| Store event data using the SendEvent API operation | 77 |
| Get details of a stored event data | 78 |
| View metrics of stored event dataset | 78 |
| Event orchestration | 80 |
| Setting up event orchestration | 81 |
| Enable event orchestration in Amazon Fraud Detector | 82 |
| Enable event orchestration in the Amazon Fraud Detector console | 82 |
| Enable event orchestration using the AWS SDK for Python (Boto3) | 82 |
| Disable event orchestration in Amazon Fraud Detector | 83 |
| Disable event orchestration in the Amazon Fraud Detector console | 83 |
| Disable event orchestration using the AWS SDK for Python (Boto3) | 83 |
| Model | 85 |
| Choose a model type | 85 |
| Online fraud insights | 85 |
| Transaction fraud insights | 87 |
| Account takeover insights | 89 |
| Build a model | 95 |
| Train and deploy a model using the AWS SDK for Python (Boto3) | 95 |
| Model scores | 97 |
| Model performance metrics | 98 |
| Model variable importance | 100 |
| Using model variable importance values | 102 |

| | |
|---|------------|
| Evaluating model variable importance values | 102 |
| Viewing model variable importance ranking | 103 |
| Understanding how the model variable importance value is calculated | 103 |
| Import a SageMaker model | 104 |
| Import a SageMaker model using the AWS SDK for Python (Boto3) | 104 |
| Delete a model or model version | 105 |
| Detector | 108 |
| Create a detector | 108 |
| Create a detector in the Amazon Fraud Detector console | 108 |
| Create a detector using the AWS SDK for Python (Boto3) | 112 |
| Create a detector version | 112 |
| Rule execution mode | 112 |
| Create a detector version using the AWS SDK for Python (Boto3) | 113 |
| Delete a detector, detector version, or rule version | 114 |
| Resources | 116 |
| Variables | 116 |
| Data types | 116 |
| Default value | 117 |
| Variable types | 117 |
| Variable enrichments | 129 |
| Create a variable | 136 |
| Delete a variable | 138 |
| Labels | 139 |
| Create label | 140 |
| Update label | 141 |
| Updating event labels in event data stored in Amazon Fraud Detector | 141 |
| Delete label | 142 |
| Rules | 143 |
| Rule language reference | 143 |
| Create rules | 149 |
| Update rule | 151 |
| Lists | 152 |
| Create a list | 152 |
| Add entries in a list | 154 |
| Assign a variable type to a list | 155 |
| Delete a list | 156 |

| | |
|--|------------|
| Delete entries from a list | 157 |
| Delete all entries from a list | 158 |
| Outcomes | 159 |
| Create an outcome | 159 |
| Delete an outcome | 160 |
| Entity | 161 |
| Create an entity type | 162 |
| Delete an entity type | 163 |
| Manage resources using AWS CloudFormation | 164 |
| Creating Amazon Fraud Detector templates | 164 |
| Managing Amazon Fraud Detector stacks | 164 |
| Understanding Amazon Fraud Detector CloudFormation parameters | 165 |
| Sample AWS CloudFormation template for Amazon Fraud Detector resources | 166 |
| Learn more about AWS CloudFormation | 167 |
| Fraud predictions | 168 |
| Real time prediction | 169 |
| How real time fraud prediction works | 169 |
| Getting real time fraud prediction | 170 |
| Batch predictions | 171 |
| How batch predictions work | 171 |
| Input and output files | 172 |
| Getting batch predictions | 172 |
| Guidance on IAM roles | 173 |
| Get batch fraud predictions using the AWS SDK for Python (Boto3) | 174 |
| Prediction explanations | 175 |
| Viewing prediction explanations | 177 |
| Understanding how prediction explanations are calculated | 178 |
| Security | 180 |
| Data Protection | 180 |
| Encryption at rest | 181 |
| Encryption in transit | 182 |
| Key management | 182 |
| VPC endpoints (AWS PrivateLink) | 184 |
| Opting out | 186 |
| Identity and access management | 187 |
| Audience | 187 |

| | |
|--|------------|
| Authenticating with identities | 188 |
| Managing access using policies | 191 |
| How Amazon Fraud Detector works with IAM | 193 |
| Identity-based policy examples | 197 |
| Confused deputy prevention | 205 |
| Troubleshooting | 208 |
| Monitoring Amazon Fraud Detector | 210 |
| Compliance validation | 211 |
| Resilience | 212 |
| Infrastructure Security | 212 |
| Monitor Amazon Fraud Detector | 214 |
| Monitoring with CloudWatch | 214 |
| Using CloudWatch Metrics for Amazon Fraud Detector. | 215 |
| Amazon Fraud Detector Metrics | 217 |
| Logging Amazon Fraud Detector API Calls with AWS CloudTrail | 221 |
| Amazon Fraud Detector Information in CloudTrail | 221 |
| Understanding Amazon Fraud Detector Log File Entries | 222 |
| Troubleshoot | 224 |
| Troubleshoot training data issues | 224 |
| Unstable fraud rate in the given dataset | 225 |
| Insufficient data | 225 |
| Missing or different EVENT_LABEL values | 227 |
| Missing or incorrect EVENT_TIMESTAMP values | 229 |
| Data not ingested | 230 |
| Insufficient variables | 230 |
| Missing or incorrect variable type | 231 |
| Missing variable values | 231 |
| Insufficient unique variable values | 232 |
| Incorrect variable expression | 233 |
| Insufficient unique entities | 234 |
| Quotas | 235 |
| Amazon Fraud Detector models | 235 |
| Amazon Fraud Detector detectors / variables / outcomes / rules | 235 |
| Amazon Fraud Detector API | 236 |
| Document history | 237 |

What is Amazon Fraud Detector?

Amazon Fraud Detector is a fully managed fraud detection service that automates the detection of potentially fraudulent activities online. These activities include unauthorized transactions and the creation of fake accounts. Amazon Fraud Detector works by using machine learning to analyze your data. It does this in a way that builds off of the seasoned expertise of more than 20 years of fraud detection at Amazon.

You can use Amazon Fraud Detector to build customized fraud-detection models, add decision logic to interpret the model's fraud evaluations, and assign outcomes such as pass or send for review for each possible fraud evaluation. With Amazon Fraud Detector, you don't need machine learning expertise to detect fraudulent activities.

To get started, collect and prepare fraud data that you collected at your organization. Amazon Fraud Detector then uses this data to train, test, and deploy a custom fraud detection model on your behalf. As a part of this process, Amazon Fraud Detector uses machine learning models that have learned patterns of fraud from AWS and Amazon's own fraud expertise to evaluate your fraud data and generate model scores and model performance data. You configure decision logic to interpret the model's score and assign outcomes for how to deal with each fraud evaluation.

Benefits

Amazon Fraud Detector provides the following benefits. These benefits make it possible for you to detect fraud fast without needing to invest the time and resources that are traditionally required to build and maintain a fraud management system.

Automated fraud model creation

Amazon Fraud Detector's fraud detection models are fully automated machine learning models customized to meet your specific business needs. You can use Amazon Fraud Detector models to identify potential fraud in any online transactions such as new account creations, online payments, and guest checkout.

Because fraud models are created through an automated process, you can forgo many of the steps associated with creating and training a model. These steps include data validation and enrichment, feature engineering, algorithm selection, hyperparameter tuning, and model deployment.

To create a fraud detection model using Amazon Fraud Detector, you only upload your company's historical fraud dataset and select the model type. Then, Amazon Fraud Detector automatically

finds the most suitable fraud detection algorithm for your use case and creates the model. You do not need to know coding or have machine learning expertise to create fraud detection models.

Fraud models that evolve and learn

Fraud detection models must constantly evolve to keep up with the changing fraud landscape. Amazon Fraud Detector does this automatically by calculating information including account age, time since last activity, and activity count. The result is that your model learns the difference between trusted customers who frequently make transactions and the continued attempts typical of fraudsters. This helps to maintain the performance of your model longer between retraining sessions.

Fraud model performance visualization

After your model is trained using the data that you provide, Amazon Fraud Detector validates your model performance. It also provides visual tools for you to assess the performance. For each model that you train, you can see the model performance score, the score distribution graph, the confusion matrix, the threshold table, and all of the inputs that you provided ranked by their impact on model performance. Using these performance tools, you can learn how your model is performing and what inputs are driving your model performance. If required, you can tweak your model to improve its overall performance.

Fraud prediction

Amazon Fraud Detector generates fraud prediction for your organization's business activities. Fraud prediction is an evaluation of a business activity for fraud risk. Amazon Fraud Detector generates predictions using the prediction logic with the data that's associated with the activity. You provided this data when you created your fraud detection model. You can get fraud predictions for a single activity in real time or get fraud predictions offline for a set of activities.

Fraud prediction explanation visualization

Amazon Fraud Detector generates prediction explanations as part of the fraud prediction process. Prediction explanations provide insight into how each data element used to train your model has impacted your model's fraud prediction score. Prediction explanations are provided using the visual tools such as tables and graphs. You can use these tools to identify visually how much influence each data element has on the prediction scores. Then, you can use this information to analyze the fraud patterns across your data set and detect bias, if any. Last you can also use the prediction explanations to identify top risk indicators during a manual fraud investigation process. This helps you narrow down the root causes that lead to false positive predictions.

Rule-based actions

After your fraud detection model is trained you can add rules to take actions on the evaluated data, such as accept the data, send data for review, or collect more data. A rule is a condition that tells Amazon Fraud Detector how to interpret data during fraud prediction. For example, you can create a rule that flags suspicious customer accounts to be reviewed. You can set this rule to be initiated if both the detected model score is greater than your predetermined threshold and if the account payment's authorization code (AUTH_CODE) isn't valid.

Core concepts and terms

The following is a list of core concepts and terms that are used in Amazon Fraud Detector:

Event

An event is your organization's business activity that's evaluated for fraud risk. Amazon Fraud Detector generates fraud predictions for events.

Label

A label classifies a single event as fraudulent or legitimate. Labels are used to train machine learning models in Amazon Fraud Detector.

Entity

An entity represents who is performing the event. You provide entity ID as part of your company's fraud data to indicate the specific entity who performed the event.

Event type

An event type defines the structure for an event sent to Amazon Fraud Detector. This includes the data sent as part of the event, the entity performing the event (such as a customer), and the labels that classify the event. Example event types include online payment transactions, account registrations, and authentication.

Entity type

An entity type classifies the entity. Example classifications include customer, merchant, or account.

Event dataset

The event dataset is your company's historical data of a specific business activity or an event. For example, your company's event might be online account registration. Data from a single

event (registration) might include the associated IP address, email address, billing address, and event timestamp. You provide event dataset to Amazon Fraud Detector to create and train fraud detection models.

Model

A model is an output of machine learning algorithms. These algorithms are implemented in code and run on event data you provide.

Model type

The model type defines the algorithms, enrichments, and feature transformations that are used during model training. It also defines the data requirements to train the model. These definitions function to optimize your model for a specific type of fraud. You specify the model type to use when you create your model.

Model training

Model training is the process of using a provided event dataset to create a model that can predict fraudulent events. All steps in the model training process are fully automated. These steps include data validation, data transformation, feature engineering, algorithm selection, and model optimization.

Model score

Model score is the evaluation result of your company's historical fraud data. During the model training process, Amazon Fraud Detector evaluates the dataset for fraudulent activities and generates a score between 0 and 1000. For this score, 0 represents low fraud risk whereas 1000 represents the highest fraud risk. The score itself is directly related to false positive rate (FPR).

Model version

A model version is an output from training a model.

Model deployment

Model deployment is a process for activating a model version and making it available for generating fraud predictions.

Amazon SageMaker model endpoint

In addition to building models using Amazon Fraud Detector, you can optionally use SageMaker-hosted model endpoints in Amazon Fraud Detector evaluations.

For more information about building a model in SageMaker, see [Train a Model with Amazon SageMaker](#).

Detector

A detector contains the detection logic such as the model and rules for a particular event that you want to evaluate for fraud. You create a detector using a model version.

Detector version

A detector can have multiple versions, with each version having a status of Draft, Active, or Inactive. Only one detector version can be in Active status at a time.

Variable

A variable represents a data element associated with an event that you want to use in a fraud prediction. Variables can either be sent with an event as part of a fraud prediction or derived, such as the output of an Amazon Fraud Detector model or Amazon SageMaker.

Rule

A rule is a condition that tells Amazon Fraud Detector how to interpret variable values during a fraud prediction. A rule consists of one or more variables, a logic expression, and one or more outcomes. The variables used in the rule must be part of the event dataset that the detector evaluates. Moreover each detector must have at least one rule associated with it.

Outcome

This is the result, or output, from a fraud prediction. Each rule that is used in a fraud prediction must specify one or more outcomes.

Fraud prediction

Fraud prediction is an evaluation of fraud for either a single event or a set of events. Amazon Fraud Detector generates fraud predictions for a single online event in real time by synchronously providing a model score and an outcome based on the rules. Amazon Fraud Detector generates fraud predictions for a set of events offline. You can use the predictions to perform an offline proof-of-concept, or to retrospectively evaluate fraud risk on an hourly, daily, or weekly basis.

Fraud prediction explanation

Fraud prediction explanations provide insight into how each variable impacted your model's fraud prediction score. It provides information about how each variable influences the risk scores in terms of magnitude (ranging from 0 to 5 with 5 being highest) and direction (driving the score higher or lower).

How Amazon Fraud Detector works

Amazon Fraud Detector builds a machine learning model that is customized to detect potential fraudulent online activities in your business. To get started, you provide your business use case. Depending on your business use case, Amazon Fraud Detector recommends a model type it will use to create a fraud detection model for you. In addition, it also provides insights into the data elements you need to provide as part of your business's historical data. Amazon Fraud Detector uses the historical dataset to automatically create and train a customized model for you.

The automated model training process involves choosing a machine learning algorithm that detects fraud for your specific business use case, validating the data you provided, and performing data manipulations to improve model performance. After training the model, Amazon Fraud Detector generates model scores and other model performance metrics. You can use the score and the performance metrics to evaluate model performance. If needed, you can add or remove data elements from the dataset you provided for training and retrain the model to improve the model score.

After the model is created, trained, and activated, you need to configure decision logic, also known as rules, that tells the model how to interpret the data generated by your business, and assign outcomes for how to deal with interpretation of each activity. The outcomes can represent actions such as, approve or review the activity, or it can represent risk levels of the activity such as high risk, medium risk, and low risk.

A detector is a container that holds your model and the associated rules. You will need to create, test, and deploy the detector to your production environment.

The detector deployed in your production environment provides the fraud detection capability to your business applications. To perform fraud evaluation, the model compares all incoming data from your business activity with your business's historical data and uses its' sophisticated machine learning algorithms with the rules you created to analyze the results and assign outcomes. With Amazon Fraud Detector, you can either evaluate data from a single business activity in real-time or evaluate data from multiple business activities offline.

Let us say you have a business that has online funds transfer as one of its activities. You want to use Amazon Fraud Detector to detect fraudulent requests for funds transfer, in real time. To get started, you will need to first provide Amazon Fraud Detector with data from past fund transfer requests. Amazon Fraud Detector uses this data to create and train a model that is customized to detect fraudulent requests for fund transfers. And then, you create a detector by adding the model

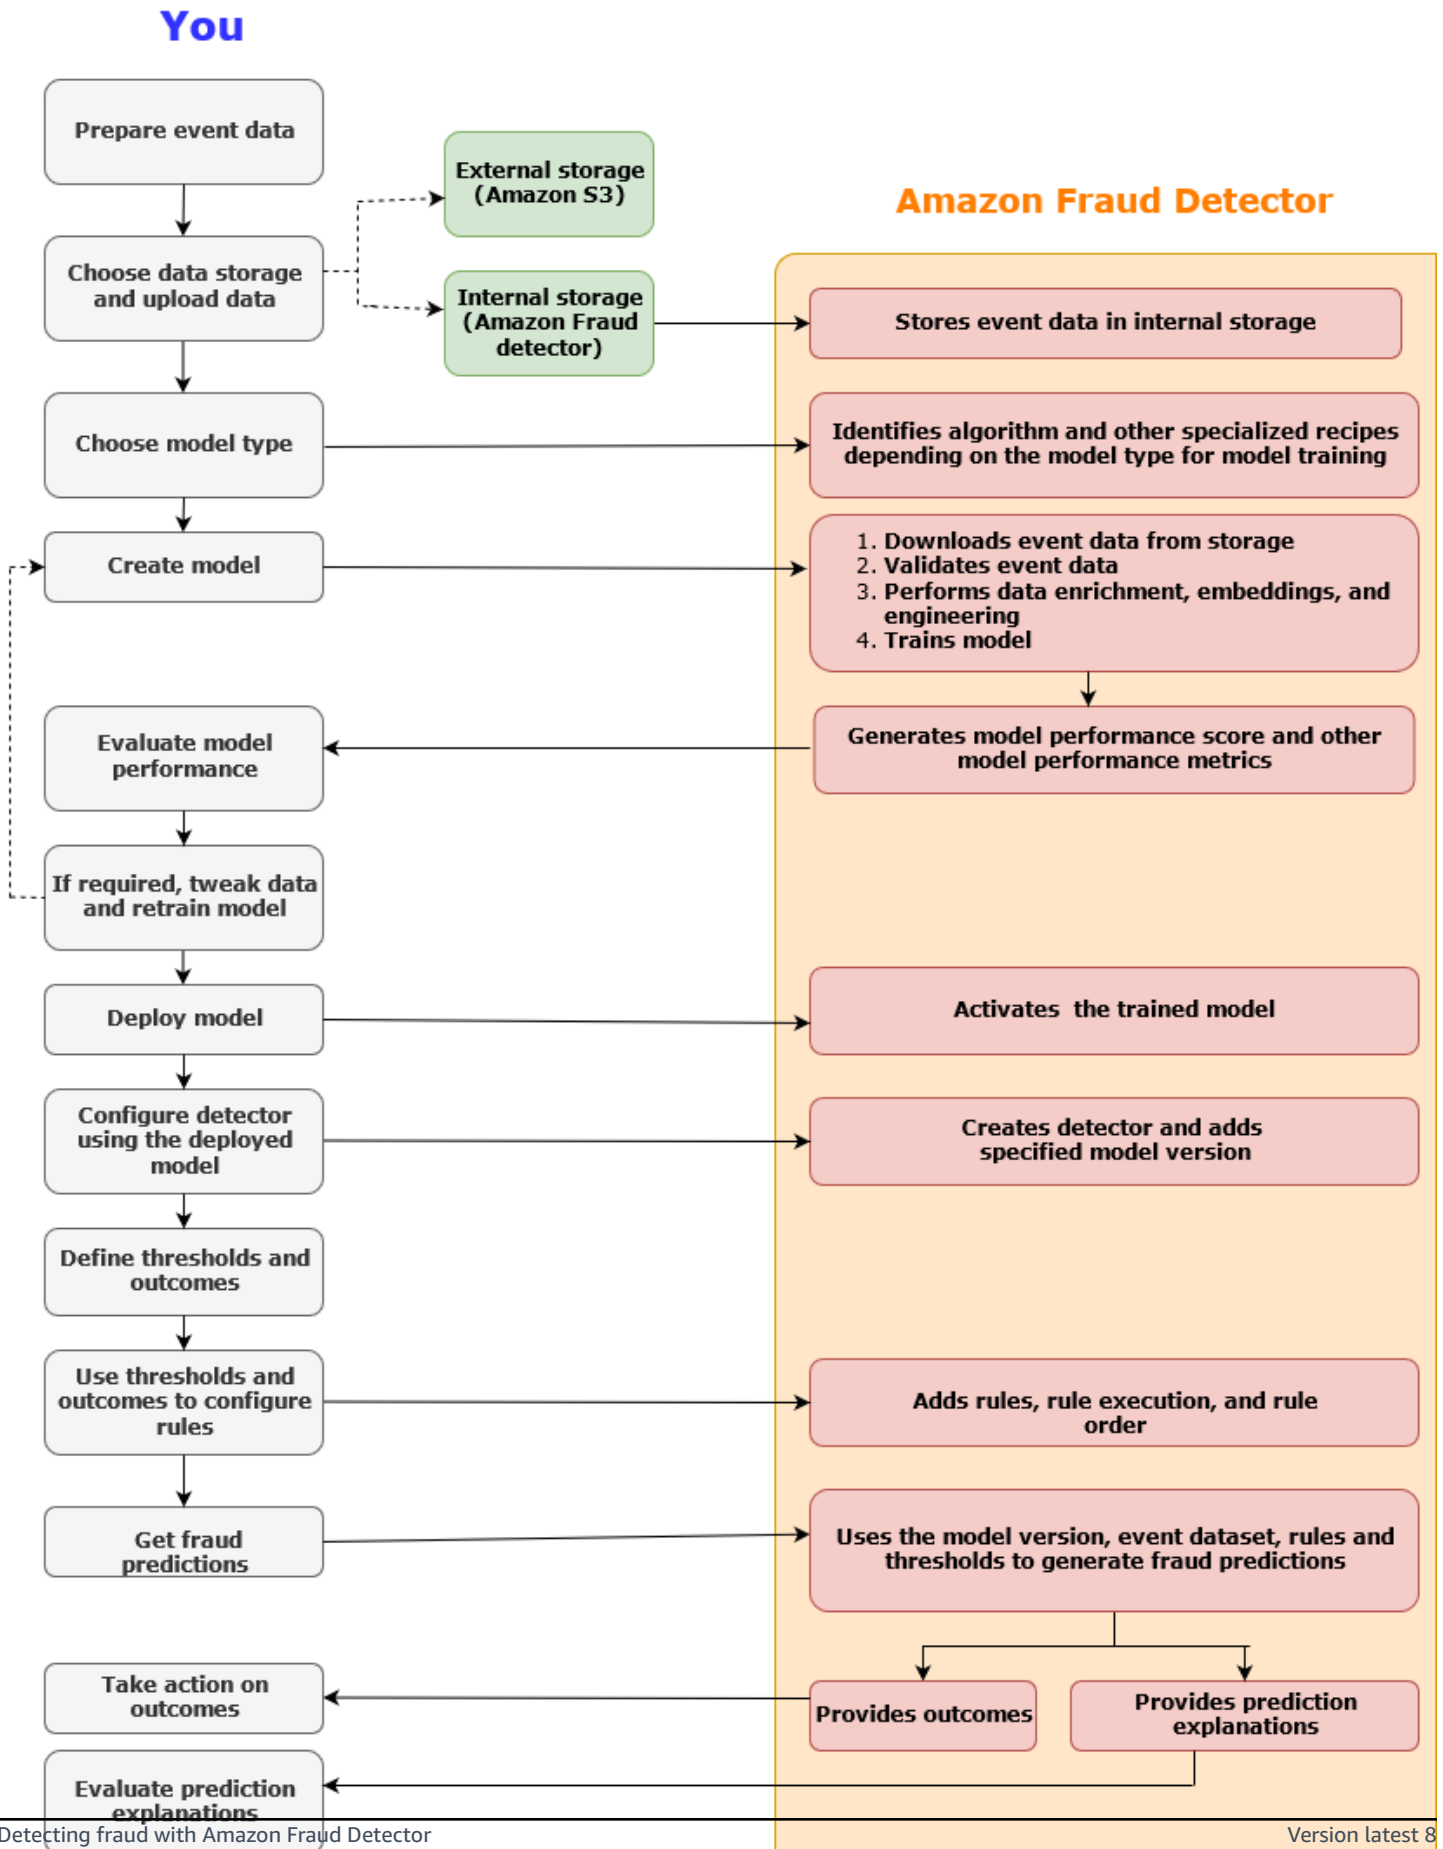
and by configuring rules for your model to interpret the data. An example of a rule for online funds transfer activity can be, if the request for funds transfer is coming from *xyz@example.com* email address, send the request for review. In your business's production environment, when a request for fund transfer comes in, the model analyzes the data that came with the request and uses the rule to assign the outcome. You can then take an action on the request depending on the assigned outcome.

Amazon Fraud Detector uses components such as, training dataset, model, detector, rules, and outcomes to provide your business with a fraud evaluation logic.

For information about the workflow you'll use for detecting fraud using Amazon Fraud Detector, see [Detecting fraud with Amazon Fraud Detector](#)

Detecting fraud with Amazon Fraud Detector

This section describes a typical workflow for detecting fraud with Amazon Fraud Detector. It also summarizes how you can accomplish those tasks. The following diagram provides a high-level view of the workflow for detecting fraud with Amazon Fraud Detector.



Fraud detection is a continuous process. After you deploy your model, make sure to evaluate its performance scores and metrics based on the prediction explanations. By doing so, you can identify top risk indicators, narrow down root causes that lead to false positives, and analyze fraud patterns across your dataset and detect bias, if any exist. To increase the accuracy of the predictions, you can tweak your dataset to include new or revised data. Then, you can retrain your model with the updated dataset. As more data becomes available, you continue retraining your model to increase accuracy.

Accessing Amazon Fraud Detector

Amazon Fraud Detector is available in multiple AWS Regions and can be accessed using AWS interfaces.

Availability

Amazon Fraud Detector is available in the US East (N. Virginia), US East (Ohio), US West (Oregon), Europe (Ireland), Asia Pacific (Singapore), and Asia Pacific (Sydney) AWS Regions.

Interfaces

You can create, train, deploy, test, run, and manage fraud detection models and detectors using any of the following interfaces:

AWS Management Console - Amazon Fraud Detector provides a web-based user interface, the Amazon Fraud Detector console. If you signed up for an AWS account, you can access the Amazon Fraud Detector console. For more information, see [Set up Amazon Fraud Detector](#).

AWS Command Line Interface (AWS CLI) - Provides an interface that you can use to interact with a broad set of AWS services, including Amazon Fraud Detector, using commands in your command-line shell. AWS CLI commands for Amazon Fraud Detector implement functionality that's equivalent to that provided by the Amazon Fraud Detector console.

AWS SDK - Provides language-specific APIs and manage many of the connection details, such as signature calculation, request retry handling, and error handling. For more information, go to [Tools to build AWS](#) page, scroll down to the **SDK** section, and choose plus (+) sign to expand the section.

AWS CloudFormation - Provides templates that you can use to define your Amazon Fraud Detector resources and properties. For more information, see [Amazon Fraud Detector resource type reference](#) in the AWS CloudFormation User Guide.

Pricing

With Amazon Fraud Detector, you pay only for what you use. There are no minimum fees or upfront commitments. You're charged based on the compute hours that are used to train and host your models, the amount of storage you use, and the quantity of fraud predictions that you make. For more information, see [Amazon Fraud Detector pricing](#).

Set up for Amazon Fraud Detector

To use Amazon Fraud Detector, you first need an Amazon Web Services (AWS) account and then you must set up permissions that give your AWS account access to all interfaces. Later when you start to create your Amazon Fraud Detector resources, you need to grant permissions that allow Amazon Fraud Detector to access your account to perform tasks on your behalf and to access resources that you own.

Complete the following tasks in this section to get set up for using Amazon Fraud Detector:

- Sign up for AWS.
- Set up permissions that allows your AWS account to access Amazon Fraud Detector interfaces.
- Set up interfaces you want to use to access Amazon Fraud Detector.

After you complete these steps, see [Get started with Amazon Fraud Detector](#) to continue getting started with Amazon Fraud Detector.

Sign up for AWS

When you sign up for Amazon Web Services (AWS), your AWS account is automatically signed up for all services on AWS, including Amazon Fraud Detector. You're charged only for the services that you use. If you already have an AWS account, skip to the next task.

Sign up for an AWS account

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

To sign up for an AWS account

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

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

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign

administrative access to a user, and use only the root user to perform [tasks that require root user access](#).

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

Create a user with administrative access

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

Secure your AWS account root user

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

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

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

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

Create a user with administrative access

1. Enable IAM Identity Center.

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

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

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

Sign in as the user with administrative access

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

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

Assign access to additional users

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

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

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

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

Set up permissions to access Amazon Fraud Detector interfaces

To use Amazon Fraud Detector, set up permissions to access the Amazon Fraud Detector console and API operations.

Following security best practice create an AWS Identity and Access Management (IAM) user with access restricted to Amazon Fraud Detector operations and with required permissions. You can add other permissions as needed.

The following policies provide the required permission to use Amazon Fraud Detector:

- `AmazonFraudDetectorFullAccessPolicy`

Allows you to perform the following actions:

- Access all Amazon Fraud Detector resources
 - List and describe all model endpoints in SageMaker
 - List all IAM roles in the account
 - List all Amazon S3 buckets
 - Allow IAM Pass Role to pass a role to Amazon Fraud Detector
- `AmazonS3FullAccess`

Allows full access to Amazon Simple Storage Service. This is required if you need to upload training datasets to Amazon S3.

The following describes how to create an IAM user and assign the needed permissions.

To create a user and assign required permissions

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Users** and then choose **Add user**.
3. For **User name**, enter **AmazonFraudDetectorUser**.
4. Select the **AWS Management Console access** check box, and then configure the user password.
5. (Optional) By default, AWS requires the new user to create a new password when they first sign in. You can clear the check box next to **User must create a new password at next sign-in** to allow the new user to reset their password after they sign in.
6. Choose **Next: Permissions**.
7. Choose **Create group**.
8. For **Group name** enter **AmazonFraudDetectorGroup**.
9. In the policy list, select the check box for **AmazonFraudDetectorFullAccessPolicy** and **AmazonS3FullAccess**. Choose **Create group**.
10. In the list of groups, select the check box for your new group. Choose **Refresh** if you don't see the group in the list.
11. Choose **Next: Tags**.
12. (Optional) Add metadata to the user by attaching tags as key-value pairs. For instructions on how to use tags in IAM, see [Tagging IAM Users and Roles](#).
13. Choose **Next: Review** to see the **User details** and **Permissions summary** for the new user. When you're ready to proceed, choose **Create user**.

Set up interfaces to access Amazon Fraud Detector with

You can access Amazon Fraud Detector using the Amazon Fraud Detector console, AWS CLI, or AWS SDK. Before you can use them, first set up the AWS CLI and AWS SDK.

Access Amazon Fraud Detector console

You can access the Amazon Fraud Detector console and other AWS services through the AWS Management Console. Your AWS account, grants you access to the AWS Management Console.

To access Amazon Fraud Detector console,

1. Go to <https://console.aws.amazon.com> and sign in to your AWS account.
2. Navigate to Amazon Fraud Detector.

With Amazon Fraud Detector console, you can create and manage your models and your fraud detection resources such as Detectors, Variables, Events, Entities, Labels, and Outcomes. You can generate predictions and evaluate the performance and predictions of your model.

Set up AWS CLI

You can use AWS Command Line Interface (AWS CLI) to interact with Amazon Fraud Detector by running commands in your command line shell. With minimal configuration, you can use the AWS CLI to run commands for similar functionality to that provided by the Amazon Fraud Detector console from the command prompt in your terminal.

To set up the AWS CLI

Download and configure the AWS CLI. For instructions, see the following topics in the AWS Command Line Interface User Guide:

- [Getting Set Up with the AWS Command Line Interface](#)
- [Configuring the AWS Command Line Interface](#)

For information about Amazon Fraud Detector commands, see [Available Commands](#)

Set up AWS SDK

You can use the AWS SDKs to write code for creating and managing your fraud detection resources and for getting fraud predictions. The AWS SDKs support Amazon Fraud Detector in [JavaScript](#) and [Python \(Boto3\)](#).

To set up AWS SDK for Python (Boto3)

You can use AWS SDK for Python (Boto3) to create, configure, and manage AWS services. For instructions on how to install Boto, see [AWS SDK for Python \(Boto3\)](#). Make sure that you're using Boto3 SDK version 1.14.29 or higher.

After you install AWS SDK for Python (Boto3), run the following Python example to confirm that your environment is configured correctly. If it's configured correctly, the response contains a list of detectors. If no detectors were created, the list is empty.

```
import boto3
fraudDetector = boto3.client('frauddetector')

response = fraudDetector.get_detectors()
print(response)
```

To set up AWS SDKs for Java

For instructions on how to install and load the AWS SDK for JavaScript, see [Setting up the SDK for JavaScript](#).

Get started with Amazon Fraud Detector

Before you get started, make sure that you have read [Detecting fraud with Amazon Fraud Detector](#) and completed steps in [Set up for Amazon Fraud Detector](#).

Use the hands-on tutorials in this section to learn how you can use Amazon Fraud Detector to build, train, and deploy a fraud detection model. For this tutorial, you assume the role of a fraud analyst using machine learning model to predict whether a new account registration is fraudulent. The model must be trained using data from account registrations. Amazon Fraud Detector provides an example account registration dataset for this tutorial. The example dataset must be uploaded before you get started with the tutorial.

You can get started with Amazon Fraud Detector using one of the following interfaces. Before getting started with the tutorial, make sure that you follow instructions to [Get and upload example dataset](#)

- [Tutorial: Get started using the Amazon Fraud Detector console](#)
- [Tutorial: Get started using the AWS SDK for Python \(Boto3\)](#)

Get and upload example dataset

The example dataset you use in this tutorial provides details of online account registrations. The dataset is in a text file that uses comma-separated value (CSV) in the UTF-8 format. The first row of the CSV dataset file contains the headers. The header row is followed by multiple rows of data. Each of these rows consists of data elements from a single account registration. The data is labeled for your convenience. A column in the dataset identifies whether the account registration is fraudulent.

To get and upload example dataset

1. Go to [Samples](#).

There are two data files that has online account registration data - *registration_data_20K_minimum.csv* and *registration_data_20K_full.csv*. The file *registration_data_20K_minimum* contains only two variables: *ip_address* and *email_address*. The file *registration_data_20K_full* contains other variables. These

variables are for each event and they include *billing_address*, *phone_number*, and *user_agent*. Both data files also contain two mandatory fields:

- **EVENT_TIMESTAMP** – Defines when the event occurred
- **EVENT_LABEL** – Classifies the event as fraudulent or legitimate

You can use either one of the two files for this tutorial. Download the data file you want to use.

2. Create an Amazon Simple Storage Service (Amazon S3) bucket.

In this step, you create an external storage to store the dataset. This external storage is Amazon S3 bucket. For more information about Amazon S3, see [What is Amazon S3?](#)

- Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
 - In **Buckets**, choose **Create bucket**.
 - For **Bucket name**, enter a bucket name. Make sure that you follow the bucket naming rules in the console, and provide a globally unique name. We recommend you use a name that describes the bucket's purpose.
 - For **AWS Region**, choose the AWS Region where you want to create your bucket. The Region that you choose must support Amazon Fraud Detector. To reduce latency, choose the AWS Region that's closest to your geographic location. For a list of Regions that support Amazon Fraud Detector, see the [Region Table](#) in the *Global Infrastructure Guide*.
 - Leave the default settings for **Object Ownership**, **Bucket settings for Block Public Access**, **Bucket Versioning**, and **Tags** for this tutorial.
 - For **Default encryption**, choose **Disable** for this tutorial.
 - Review your bucket configuration, and then choose **Create bucket**.
- ## 3. Upload example data file to Amazon S3 bucket.

Now that you have a bucket, upload one of the example files that you downloaded previously to the Amazon S3 bucket that you just created.

- In the **Buckets**, your bucket name is listed. Choose your bucket.
- Choose **Upload**.
- In **Files and folders**, choose **Add files**.

- d. Choose one of the example data files that you downloaded on your computer, and then choose **Open**.
- e. Leave the default settings for **Destination**, **Permissions**, and **Properties**.
- f. Review configurations, and then choose **Upload**.
- g. The example data file is uploaded to Amazon S3 bucket. Make a note of the bucket location. In the **Objects**, choose the example data file that you just uploaded.
- h. In the **Object overview**, copy the location under **S3 URI**. This is the Amazon S3 location of your example data file. You use it later. You can additionally copy the **Amazon Resource Name (ARN)** of your S3 bucket and save it.

Tutorial: Get started using the Amazon Fraud Detector console

This tutorial consists of two parts. The first part describes how to build, train, and deploy a fraud detection model. The second part covers how to use the model to generate fraud predictions in real time. The model is trained using the example data file that you upload to an S3 bucket. By the end of this tutorial, you complete the following actions:

- Build and train an Amazon Fraud Detector model
- Generate real-time fraud predictions

Important

Before you proceed, make sure that you have followed instructions to [Get and upload example dataset](#)

Part A: Build, train, and deploy an Amazon Fraud Detector model

In part A, you define your business use case, define your event, build a model, train the model, evaluate model's performance, and deploy the model.

Step 1: Choose your business use case

- In this step, you use the **data models explorer** to match your business use case with the fraud detection model types supported by Amazon Fraud Detector. Data models explorer is a tool integrated with the Amazon Fraud Detector console that recommends a model type to use for

creating and training a fraud detection model for your business use case. Data models explorer also provides insights into the mandatory, recommended, and optional data elements you will require to include in your dataset. The dataset will be used to create and train your fraud detection model.

For the purpose of this tutorial, your business use case is new account registrations. After you specify your business use case, the data models explorer will recommend a model type for creating a fraud detection model and will also provide you with a list of data elements you will need to create your dataset. Since you have already uploaded a sample dataset containing data from new account registrations, you do not need to create a new dataset.

- a. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
- b. In the left navigation pane, choose **Data models explorer**.
- c. In the **Data models explorer** page, under **Business use case**, select **New account fraud**.
- d. Amazon Fraud Detector displays the recommended model type to use to create a fraud detection model for the selected business use case. The model type defines the algorithms, enrichments, and transformations Amazon Fraud Detector will use to train your fraud detection model.

Make a note of the recommended model type. You will need this later when you create your model.

- e. The **Data model insights** pane provides insight into the mandatory and recommended data elements required to create and train a fraud detection model.

Take a look at the sample dataset you downloaded and make sure that it has all the mandatory and some recommended data elements listed in the table.

Later when you create a model for your specific business use case, you will use the insights provided to create your dataset.

Step 2: Create event type

- In this step, you define the business activity (event) to evaluate for fraud. Defining the event involves setting the variables that are in your dataset, the entity initiating event, and the labels that classify the event. For this tutorial, you define the account registration event.

- a. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
- b. In the left navigation pane, choose **Events**.
- c. In the **Events type** page, choose **Create**.
- d. Under **Event type details**, enter `sample_registration` as the event type name and, optionally, enter a description of the event.
- e. For **Entity**, choose **Create entity**.
- f. In the **Create entity** page, enter `sample_customer` as the entity type name. Optionally, enter a description of the entity type.
- g. Choose **Create entity**.
- h. Under **Event variables**, for **Choose how to define this event's variables**, choose **Select variables from a training dataset**.
- i. For **IAM role**, choose **Create IAM role**.
- j. In the **Create IAM role** page, enter the name of the S3 bucket that you uploaded your example data to and choose **Create role**.
- k. In **Data location**, enter the path to your example data. This is the S3 URI path that you saved after uploading the example data. The path is similar to this: `S3://your-bucket-name/example dataset filename.csv`.
- l. Choose **Upload**.

Amazon Fraud Detector extracts the headers from your example data file and maps them with a variable type. The mapping is displayed in the console.

- m. Under **Labels - optional**, for **Labels**, choose **Create new labels**.
- n. In **Create label** page, enter `fraud` as the name. This label corresponds to the value that represents the fraudulent account registration in the example dataset.
- o. Choose **Create label**.
- p. Create a second label, then enter `legit` as the name. This label corresponds to the value that represents the legitimate account registration in the example dataset.
- q. Choose **Create event type**.

Step 3: Create model

1. On the **Models** page, choose **Add model**, and then choose **Create model**.

2. For **Step 1 – Define model details**, enter `sample_fraud_detection_model` as the model name. Optionally, add a description of the model.
3. For **Model Type**, choose the **Online Fraud Insights** model.
4. For **Event type**, choose **sample_registration**. This is the event type that you created in Step 1.
5. In **Historical event data**,
 - a. In **Event data source**, choose **Event data stored in S3**.
 - b. For **IAM role**, select the role that you created in Step 1.
 - c. In **Training data location**, enter the S3 URI path to your example data file.
6. Choose **Next**.

Step 4: Train model

1. In **Model inputs**, leave all checkboxes checked. By default, Amazon Fraud Detector uses all variables from your historical event dataset as model inputs.
2. In **Label classification**, for **Fraud labels** choose **fraud** as this label corresponds to the value that represents fraudulent events in the example dataset. For **Legitimate labels**, choose **legit** as this label corresponds to the value that represents legitimate events in the example dataset.
3. For the **Unlabeled events treatment**, keep the default selection **Ignore unlabeled events** for this example dataset.
4. Choose **Next**.
5. After reviewing, choose **Create and train model**. Amazon Fraud Detector creates a model and begins to train a new version of the model.

In **Model versions** the **Status** column indicates the status of model training. Model training that uses the example dataset takes approximately 45 minutes to complete. The status changes to **Ready to deploy** after model training is complete.

Step 5: Review model performance

An important step in using Amazon Fraud Detector is to assess the accuracy of your model using model scores and performance metrics. After model training is complete, Amazon Fraud Detector validates model performance using the 15% of your data that wasn't used to train the model and generates a model performance score and other performance metrics.

1. To view model's performance,

- a. In the left navigation pane of the Amazon Fraud Detector console, choose **Models**.
 - b. In the **Models** page, choose the model that you just trained (**sample_fraud_detection_model**), and then choose **1.0**. This is the version Amazon Fraud Detector created of your model.
2. Look at the **Model performance** overall score and all other metrics that Amazon Fraud Detector generated for this model.

To learn more about the model performance score and performance metrics on this page, see [Model scores](#) and [Model performance metrics](#).

You can expect all your trained Amazon Fraud Detector models to have real-world fraud detection performance metrics that are similar to the performance metrics that you see for the model in this tutorial.

Step 6: Deploy model

After you reviewed the performance metrics of your trained model and are ready to use it generate fraud predictions, you can deploy the model.

1. In left navigation pane of the Amazon Fraud Detector console, choose **Models**.
2. In the **Models** page, choose **sample_fraud_detection_model**, and then choose the specific model version that you want to deploy. For this tutorial, choose **1.0**.
3. On the **Model version** page, choose **Actions** and then choose **Deploy model version**.
4. In the **Model versions**, the **Status** shows the status of the deployment. The status changes to **Active** after the deployment completes. This indicates that the model version is activated and available to generate fraud predictions. Continue with [Part B: Generate fraud predictions](#) to complete steps for generating fraud predictions.

Part B: Generate fraud predictions

Fraud prediction is an evaluation of fraud for a business activity (event). Amazon Fraud Detector uses detectors to generate fraud prediction. A detector contains detection logic, such as models and rules, for a specific event that you want to evaluate for fraud. Detection logic uses rules to tell Amazon Fraud Detector how to interpret the data associated with the model. In this tutorial, you evaluate the account registration event using the account registration example dataset that you uploaded earlier.

In Part A, you created, trained, and deployed your model. In Part B, you build a detector for the `sample_registration` event type, add the deployed model, create rules and a rule execution order, and then create and activate a version of the detector that you use to generate fraud predictions.

Step 1: Build detector

To create detector

1. In the left navigation pane of the Amazon Fraud Detector console, choose **Detectors**.
2. Choose **Create detector**.
3. In the **Define detector details** page, enter `sample_detector` for detector name. Optionally, enter a description for the detector, such as `my sample fraud detector`.
4. For **Event Type**, select **sample_registration**. This is the event that you created in Part A of this tutorial.
5. Choose **Next**.

Step 2: Add model

If you completed Part A of this tutorial, then you likely already have an Amazon Fraud Detector model that's available to add to your detector. If you didn't already create a model, go to Part A and complete the steps to create, train, and deploy a model and then continue with Part B.

1. In the **Add model - optional**, choose **Add Model**.
2. In the **Add model** page, for **Select model**, choose the Amazon Fraud Detector model name that you deployed earlier. For **Select version**, choose the model version of the deployed model.
3. Choose **Add model**.
4. Choose **Next**.

Step 3: Add rules

A rule is a condition that tells Amazon Fraud Detector how to interpret model performance score when evaluating for fraud prediction. For this tutorial, you create three rules: `high_fraud_risk`, `medium_fraud_risk`, and `low_fraud_risk`.

1. In the **Add rules** page, under **Define a rule**, enter `high_fraud_risk` for the rule name and under **Description - optional**, enter **This rule captures events with a high ML model score** as the description for the rule.
2. In **Expression**, enter the following rule expression using the Amazon Fraud Detector simplified rule expression language:

```
$sample_fraud_detection_model_insightscore > 900
```

3. In **Outcomes**, choose **Create a new outcome**. An outcome is the result from a fraud prediction and is returned if the rule matches during an evaluation.
4. In **Create a new outcome**, enter `verify_customer` as the outcome name. Optionally, enter a description.
5. Choose **Save outcome**.
6. Choose **Add rule** to run the rule validation checker and save the rule. After it's created, Amazon Fraud Detector makes the rule available for use in your detector.
7. Choose **Add another rule**, and then choose the **Create rule** tab.
8. Repeat this process twice more to create your `medium_fraud_risk` and `low_fraud_risk` rules using the following rule details:

- `medium_fraud_risk`

Rule name: `medium_fraud_risk`

Outcome: `review`

Expression:

```
$sample_fraud_detection_model_insightscore <= 900 and
```

```
$sample_fraud_detection_model_insightscore > 700
```

- `low_fraud_risk`

Rule name: `low_fraud_risk`

Outcome: `approve`

Expression:

```
$sample_fraud_detection_model_insightscore <= 700
```


These values are examples used for this tutorial. When you create rules for your own detector, use values that are appropriate for your model and your use case,

9. After you have created all three rules, choose **Next**.

For more information about creating and writing rules, see [Rules](#) and [Rule language reference](#).

Step 4: Configure rule execution and rule order

The rule execution mode for the rules that are included in the detector determines if all the rules you define are evaluated, or if rule evaluation stops at the first matched rule. And the rule order determines the order that you want the rule to be run in.

The default rule execution mode is `FIRST_MATCHED`.

First matched

First matched rule execution mode returns the outcomes for the first matching rule based on defined rule order. If you specify `FIRST_MATCHED`, Amazon Fraud Detector evaluates rules sequentially, first to last, stopping at the first matched rule. Amazon Fraud Detector then provides the outcomes for that single rule.

The order that you run rules in can affect the resulting fraud prediction outcome. After you have created your rules, re-order the rules to run them in the desired order by following these steps:

If your `high_fraud_risk` rule isn't already on the top of your rule list, choose **Order**, and then choose **1**. This moves `high_fraud_risk` to the first position.

Repeat this process so that your `medium_fraud_risk` rule is in the second position and your `low_fraud_risk` rule is in the third position.

All matched

All matched rule execution mode returns outcomes for all matched rules, regardless of rule order. If you specify `ALL_MATCHED`, Amazon Fraud Detector evaluates all rules and returns the outcomes for all matched rules.

Select `FIRST_MATCHED` for this tutorial and then choose **Next**.

Step 5: Review and create detector version

A detector version defines the specific models and rules that are used for generating fraud predictions.

1. In the **Review and create** page, review the detector details, models, and rules that you configured. If you need to make any changes, choose **Edit** next to the corresponding section.
2. Choose **Create detector**. After it's created, the first version of your detector appears in the Detector versions table with **Draft** status.

You use the **Draft** version to test your Detector.

Step 6: Test and activate detector version

In the Amazon Fraud Detector console, you can test the logic of your detector using a mock data with the **Run test** feature. For this tutorial, you can use account registration data from the example dataset.

1. Scroll to **Run test** at the bottom of the **Detector version details** page.
2. For **Event metadata**, enter a timestamp of when the event occurred and enter a unique identifier for the entity performing the event. For this tutorial, select a date from the date picker for timestamp, and enter "1234" for the Entity ID.
3. For **Event variable**, enter the variable values that you want to test. For this tutorial, you only need the `ip_address` and `email_address` fields. This is because they are the inputs that are used to train your Amazon Fraud Detector model. You can use the following example values. This assumes that you used the suggested variable names:
 - `ip_address: 205.251.233.178`
 - `email_address: johndoe@exampledomain.com`
4. Choose **Run test**.
5. Amazon Fraud Detector returns the fraud prediction outcome based on the rule execution mode. If the rule execution mode is `FIRST_MATCHED`, the returned outcome corresponds to the first rule that matched. The first rule is the rule with the highest priority. It's matched if it's evaluated as true. If the rule execution mode is `ALL_MATCHED`, the returned outcome corresponds to all rules that matched. That means that they're all evaluated to be true. Amazon Fraud Detector also returns the model score for any models added to your detector.

You can change the inputs and run couple of tests to see different outcomes. You can use the *ip_address* and *email_address* values from your example dataset for the tests and check if the outcomes are as expected.

6. When you're satisfied with how the detector is working, promote it from `Draft` to `Active`. Doing so makes the detector available for use in real-time fraud detection.

On the **Detector version details** page, choose **Actions, Publish, Publish version**. This changes the status of the detector from **Draft** to **Active**.

At this point, your model and the associated detector logic are ready to evaluate online activities for fraud in real time using the Amazon Fraud Detector `GetEventPrediction` API. You can also evaluate events offline using a CSV input file and the `CreateBatchPredictionJob` API. For more information about fraud prediction, see [Fraud predictions](#)

By completing this tutorial, you did the following:

- Uploaded an example event dataset to Amazon S3.
- Created and trained an Amazon Fraud Detector fraud detection model using the example dataset.
- Viewed the model performance score and other performance metrics that Amazon Fraud Detector generated.
- Deployed the fraud detection model.
- Created a detector and added the deployed model.
- Added rules, the rule execution order, and outcomes to the detector.
- Tested the detector by providing different inputs and checking if the rules and rule execution order worked as expected.
- Activated the detector by publishing it.

Tutorial: Get started using the AWS SDK for Python (Boto3)

This tutorial describes how to build and train an Amazon Fraud Detector model and then using this model to generate real-time fraud predictions using the AWS SDK for Python (Boto3). The model is trained using the account registration example data file that you upload to Amazon S3 bucket.

By the end of this tutorial, you complete the following actions:

- Build and train an Amazon Fraud Detector model
- Generate real-time fraud predictions

Prerequisites

The following are prerequisite steps for this tutorial.

- Completed [Set up for Amazon Fraud Detector](#).

If you have already [Set up AWS SDK](#), make sure that you're using Boto3 SDK version 1.14.29 or higher.

- Followed instructions to [Get and upload example dataset](#) file that's required for this tutorial.

Get started

Step 1: Set up and verify your Python environment

Boto is the Amazon Web Services (AWS) SDK for Python. You can use it to create, configure, and manage AWS services. For instructions on how to install Boto3, see [AWS SDK for Python \(Boto3\)](#).

After you install AWS SDK for Python (Boto3), run the following Python example command to confirm that your environment is configured correctly. If your environment is configured correctly, the response contains a list of detectors. If no detectors were created, the list is empty.

```
import boto3
fraudDetector = boto3.client('frauddetector')

response = fraudDetector.get_detectors()
print(response)
```

Step 2: Create variables, entity type, and labels

In this step, you create resources that are used to define model, event, and rules.

Create variable

A variable is a data element from your dataset that you want to use to create event type, model, and rules.

In the following example, the [CreateVariable](#) API is used to create two variables. The variables are `email_address` and `ip_address`. Assign them to the corresponding variable types: `EMAIL_ADDRESS` and `IP_ADDRESS`. These variables are part of the example dataset you uploaded. When you specify the variable type, Amazon Fraud Detector interprets the variable during model training and when getting predictions. Only variables with an associated variable type can be used for model training.

```
import boto3
fraudDetector = boto3.client('frauddetector')

#Create variable email_address
fraudDetector.create_variable(
    name = 'email_address',
    variableType = 'EMAIL_ADDRESS',
    dataSource = 'EVENT',
    dataType = 'STRING',
    defaultValue = '<unknown>'
)

#Create variable ip_address
fraudDetector.create_variable(
    name = 'ip_address',
    variableType = 'IP_ADDRESS',
    dataSource = 'EVENT',
    dataType = 'STRING',
    defaultValue = '<unknown>'
)
```

Create entity type

An entity represents who is performing the event and an entity type classifies the entity. Example classifications include *customer*, *merchant*, or *account*.

In the following example, [PutEntityType](#) API is used to create a `sample_customer` entity type.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_entity_type(
    name = 'sample_customer',
    description = 'sample customer entity type'
```

```
)
```

Create label

A label classifies an event as fraudulent or legitimate and is used to train the fraud detection model. The model learns to classify events using these label values.

In the following example, the [Putlabel](#) API is used to create two labels, fraud and legit.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_label(
    name = 'fraud',
    description = 'label for fraud events'
)

fraudDetector.put_label(
    name = 'legit',
    description = 'label for legitimate events'
)
```

Step 3: Create event type

With Amazon Fraud Detector, you build models that evaluate risks and generate fraud predictions for individual events. An event type defines the structure of an individual event.

In the following example, the [PutEventType](#) API is used to create an event type `sample_registration`. You define the event type by specifying the variables (`email_address`, `ip_address`), entity type (`sample_customer`), and labels (`fraud`, `legit`) that you created in the previous step.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_event_type (
    name = 'sample_registration',
    eventVariables = ['ip_address', 'email_address'],
    labels = ['legit', 'fraud'],
```

```
entityTypes = ['sample_customer'])
```

Step 4: Create, train, and deploy model

Amazon Fraud Detector trains models to learn to detect fraud for a specific event type. In the previous step, you created the event type. In this step, you create and train a model for the event type. The model acts as a container for your model versions. Each time you train a model, a new version is created.

Use following example codes to create and train an Online Fraud Insights model. This model is called `sample_fraud_detection_model`. It's for the event type `sample_registration` using the account registration example dataset that you uploaded to Amazon S3.

For more information about different model types that Amazon Fraud Detector supports, see [Choose a model type](#).

Create a model

In the following example, the [CreateModel](#) API is used to create a model.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.create_model (
    modelId = 'sample_fraud_detection_model',
    eventName = 'sample_registration',
    modelType = 'ONLINE_FRAUD_INSIGHTS')
```

Train a model

In the following example, the [CreateModelVersion](#) API is used to train the model. Specify 'EXTERNAL_EVENTS' for the `trainingDataSource` and the Amazon S3 location where you stored your example dataset and the *RoleArn* of the Amazon S3 bucket for `externalEventsDetail`. For `trainingDataSchema` parameter, specify how Amazon Fraud Detector interprets the example data. More specifically, specify which variables to include and how to classify the event labels.

```
import boto3
fraudDetector = boto3.client('frauddetector')
```

```
fraudDetector.create_model_version (  
    modelId = 'sample_fraud_detection_model',  
    modelType = 'ONLINE_FRAUD_INSIGHTS',  
    trainingDataSource = 'EXTERNAL_EVENTS',  
    trainingDataSchema = {  
        'modelVariables' : ['ip_address', 'email_address'],  
        'labelSchema' : {  
            'labelMapper' : {  
                'FRAUD' : ['fraud'],  
                'LEGIT' : ['legit']  
            }  
        }  
    },  
    externalEventsDetail = {  
        'dataLocation' : 's3://your-S3-bucket-name/your-example-data-  
filename.csv',  
        'dataAccessRoleArn' : 'role_arn'  
    }  
)
```

You can train your model multiple times. Each time that you train a model, a new version is created. After model training is complete, the model version status updates to TRAINING_COMPLETE. You can review the model performance score and other model performance metrics.

Review model performance

An important step in using Amazon Fraud Detector is to assess the accuracy of your model using model scores and performance metrics. After model training is complete, Amazon Fraud Detector validates model performance using the 15% of your data that wasn't used to train the model. It generates a model performance score and other performance metrics.

Use the [DescribeModelVersions](#) API to review model performance. Look at the **Model performance** overall score and all other metrics generated by Amazon Fraud Detector for this model.

To learn more about the model performance score and performance metrics, see [Model scores](#) and [Model performance metrics](#).

You can expect all your trained Amazon Fraud Detector models to have real-world fraud detection performance metrics, which are similar to the metrics in this tutorial.

Deploy a model

After you reviewed the performance metrics of your trained model, deploy the model and make it available to Amazon Fraud Detector to generate fraud predictions. To deploy the trained model, use the [UpdateModelVersionStatus](#) API. In the following example, it's used to update the model version status to ACTIVE.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.update_model_version_status (
    modelId = 'sample_fraud_detection_model',
    modelType = 'ONLINE_FRAUD_INSIGHTS',
    modelVersionNumber = '1.00',
    status = 'ACTIVE'
)
```

Step 5: Create detector, outcomes, rules, and detector version

A detector contains the detection logic, such as the models and rules. This logic is for a particular event that you want to evaluate for fraud. A rule is a condition that you specify to tell Amazon Fraud Detector how to interpret variable values during prediction. And outcome is the result of a fraud prediction. A detector can have multiple versions with each version having a status of *DRAFT*, *ACTIVE*, or *INACTIVE*. A detector version must have at least one rule that's associated with it.

Use the following example codes to create detector, rules, outcome, and to publish the detector.

Create a detector

In the following example, the [PutDetector](#) API is used to create a `sample_detector` detector for `sample_registration` event type.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_detector (
    detectorId = 'sample_detector',
    eventName = 'sample_registration'
)
```

Create outcomes

Outcomes are created for each possible fraud prediction result. In the following example, the [PutOutcome](#) API is used to create three outcomes - `verify_customer`, `review`, and `approve`. These outcomes are later assigned to rules.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_outcome(
    name = 'verify_customer',
    description = 'this outcome initiates a verification workflow'
)

fraudDetector.put_outcome(
    name = 'review',
    description = 'this outcome sidelines event for review'
)

fraudDetector.put_outcome(
    name = 'approve',
    description = 'this outcome approves the event'
)
```

Create rules

Rule consists of one or more variables from your dataset, a logic expression, and one or more outcomes.

In the following example, the [CreateRule](#) API is used to create three different rules: `high_risk`, `medium_risk`, and `low_risk`. Create rule expressions to compare the model performance score `sample_fraud_detection_model_insightscore` value against various thresholds. This is to determine the level of risk for an event and assign outcome that was defined in the previous step.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.create_rule(
    ruleId = 'high_fraud_risk',
    detectorId = 'sample_detector',
    expression = '$sample_fraud_detection_model_insightscore > 900',
    language = 'DETECTORPL',
```

```
        outcomes = ['verify_customer']
    )

    fraudDetector.create_rule(
        ruleId = 'medium_fraud_risk',
        detectorId = 'sample_detector',
        expression = '$sample_fraud_detection_model_insightscore <= 900 and
    $sample_fraud_detection_model_insightscore > 700',
        language = 'DETECTORPL',
        outcomes = ['review']
    )

    fraudDetector.create_rule(
        ruleId = 'low_fraud_risk',
        detectorId = 'sample_detector',
        expression = '$sample_fraud_detection_model_insightscore <= 700',
        language = 'DETECTORPL',
        outcomes = ['approve']
    )
```

Create a detector version

A detector version defines model and rules that are used to get fraud prediction.

In the following example, the [CreateDetectorVersion](#) API is used to create a detector version. It does this by providing model version details, rules, and a rule execution mode `FIRST_MATCHED`. A rule execution mode specifies the sequence for evaluating rules. The rule execution mode `FIRST_MATCHED` specifies that the rules are evaluated sequentially, first to last, stopping at first matched rule.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.create_detector_version(
    detectorId = 'sample_detector',
    rules = [{
        'detectorId' : 'sample_detector',
        'ruleId' : 'high_fraud_risk',
        'ruleVersion' : '1'
    }],
    {
```

```
        'detectorId' : 'sample_detector',
        'ruleId' : 'medium_fraud_risk',
        'ruleVersion' : '1'
    },
    {
        'detectorId' : 'sample_detector',
        'ruleId' : 'low_fraud_risk',
        'ruleVersion' : '1'
    }
],
modelVersions = [{
    'modelId' : 'sample_fraud_detection_model',
    'modelType': 'ONLINE_FRAUD_INSIGHTS',
    'modelVersionNumber' : '1.00'
} ],
ruleExecutionMode = 'FIRST_MATCHED'
)
```

Step 6: Generate fraud predictions

The last step of this tutorial uses the detector `sample_detector` created in the previous step to generate fraud predictions for `sample_registration` event type in real time. The detector evaluates the example data that's uploaded to Amazon S3. The response includes model performance scores as well as any outcomes that are associated to the matched rules.

In the following example, the [GetEventPrediction](#) API is used to provide data from a single account registration with each request. For this tutorial, take data (`email_address` and `ip_address`) from the account registration example data file. Each line (row) after the top header line represents data from a single account registration event.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.get_event_prediction(
    detectorId = 'sample_detector',
    eventId = '802454d3-f7d8-482d-97e8-c4b6db9a0428',
    eventName = 'sample_registration',
    eventTimestamp = '2020-07-13T23:18:21Z',
    entities = [{'entityType': 'sample_customer', 'entityId': '12345'}],
    eventVariables = {
```

```
'email_address': 'johndoe@exampldomain.com',  
'ip_address': '1.2.3.4'  
}  
)
```

After you completed this tutorial, you did the following:

- Uploaded an example event dataset to Amazon S3.
- Created variables, entities, and labels that are used to create and train a model.
- Created and trained a model using the example dataset.
- Viewed the model performance score and other performance metrics that Amazon Fraud Detector generated.
- Deployed the fraud detection model.
- Created a detector and added the deployed model.
- Added rules, the rule execution order, and outcomes to the detector.
- Created detector version.
- Tested the detector by providing different inputs and checking if the rules and rule execution order worked as expected.

(Optional) Explore the Amazon Fraud Detector APIs with a Jupyter (iPython) Notebook

For more examples for how to use the Amazon Fraud Detector APIs, see [aws-fraud-detector-samples GitHub repository](#). The topics that the notebooks cover include both building models and detectors using the Amazon Fraud Detector APIs and making batch fraud prediction requests using the `GetEventPrediction` API.

Next steps

Now that you created a model and a detector, you can take a deeper dive and start to create models and detectors and generate fraud predictions.

The following sections in the Amazon Fraud Detector User Guide describe how your business or organization can use Amazon Fraud Detector to detect fraud.

- Prepare and create your event dataset for training your model.
- Create event type
- Create model
- Create detector
- Get fraud predictions
- Manage your Amazon Fraud Detector resources (specifically, Variables, Entities, Outcomes, and Labels)
- Configure Amazon Fraud Detector to meet your security and compliance objectives
- Monitor Amazon Fraud Detector and log Amazon Fraud Detector API calls
- Troubleshoot issues with Amazon Fraud Detector

Event dataset

An event dataset is the historical fraud data for your company. You provide this data to Amazon Fraud Detector to create fraud detection models.

Amazon Fraud Detector uses machine learning models for generating fraud predictions. Each model is trained using a model type. The model type specifies the algorithms and transformations that are used for training the model. Model training is the process of using a dataset that you provide to create a model that can predict fraudulent events. For more information, see [How Amazon Fraud Detector works](#)

The dataset used for creating fraud detection model provides details of an event. An event is a business activity that is evaluated for fraud risk. For example, an account registration can be an event. The data associated with the account registration event can be event dataset. Amazon Fraud Detector uses this dataset to evaluate account registration fraud.

Before you provide your dataset to Amazon Fraud Detector for creating a model, make sure to define your goal for creating the model. You also need to determine how you want to use the model and define your metrics for evaluating if the model is performing based on your specific requirements.

For example, your goals for creating a fraud detection model that evaluates account registration fraud can be the following:

- To auto-approve legitimate registrations.
- To capture fraudulent registrations for later investigation.

After you determined your goal, the next step is to decide how you want to use the model. Some examples for using fraud detection model to evaluate registration fraud are the following:

- For real-time fraud detection for each account registration.
- For offline evaluation of all account registrations every hour.

Some examples of metrics that can be used to measure the performance of the model include the following:

- Performs consistently better than the current baseline in production.

- Captures X% fraud registrations with Y% false positives rate.
- Accepts up to 5% of auto-approved registrations that are fraudulent.

Event dataset structure

Amazon Fraud Detector requires that you provide your event dataset in a text file using comma-separated value (CSV) in the UTF-8 format. The first line of your CSV dataset file must contain file headers. The file header consists of event metadata and event variables that describe each data element that's associated with the event. The header is followed by event data. Each line consists of data elements from a single event.

- **Event metadata**- provides information about the event. For example, `EVENT_TIMESTAMP` is an event metadata that specifies the time event occurred. Depending on your business use case and the model type used for creating and training your fraud detection model, Amazon Fraud Detector requires you to provide specific event metadata. When specifying event metadata in your CSV file header, use the same event metadata name as specified by Amazon Fraud Detector and use upper case letters only.
- **Event variable**- represents the data elements that are specific to your event which you want to use for creating and training your fraud detection model. Depending on your business use case and the model type used for creating and training a fraud detection model, Amazon Fraud Detector might require or recommend that you to provide specific event variables. You can also optionally provide other event variables from your event that you want to include in training the model. Some examples of event variables for an online registration event can be email address, ip address, and phone number. When specifying event variable name in your CSV file header, use any variable name of your choice and use lower case letters only.
- **Event data**- represents the data collected from the actual event. In your CSV file, each row following the file header consists of data elements from a single event. For example, in an online registration event data file, each row contains data from a single registration. Each data element in the row must match with the corresponding event metadata or the event variable.

The following is an example of a CSV file containing data from an account registration event. The header row contains both event metadata in uppercase and event variables in lowercase followed by the event data. Each row in the dataset contains data elements associated with single account registration with each data element corresponding with the header.

| Event metadata | | | Event variables | | | | |
|-------------------------------|----------|-----------------------|-----------------|---------------|----------------|-----------------|---------------|
| EVENT_TIMESTAMP | EVENT_ID | EVENT_LABEL | email_address | phone_number | billing_street | billing_state | ip_address |
| 2020-12-06T03:13:34Z, R12345, | fraud, | regular1@example.com, | 110-345-0990, | mayhem ave, | OH, | 112.136.132.151 | ← Header |
| 2020-11-13T12:47:00Z, P56890, | legit, | premium1@example.com, | 112-890-4532, | howie lane, | KY, | 192.169.234.143 | ← Event data |
| 2021-02-19T22:52:43Z, R10001, | legit, | regular2@example.net, | 078-777-5555, | lankhurst dr, | HI, | 185.112.224.79 | Event dataset |
| 2020-11-29T00:16:09Z, R56099, | fraud, | regular3@example.edu, | 777-213-0033, | noland ave, | IL, | 68.73.183.186 | |
| 2021-01-16T07:30:03Z, P08954, | legit, | premium2@example.net, | 444-040-8344, | oakwood apt, | MA, | 117.65.246.206 | |

Get event dataset requirements using the Data models explorer

The model type you choose to create your model defines the requirements for your dataset. Amazon Fraud Detector uses the dataset you provide to create and train your fraud detection model. Before Amazon Fraud Detector starts to create your model, it checks if the dataset meets the size, format, and other requirements. If the dataset does not meet the requirements, the model creation and training fails. You can use the **data models explorer** to identify a model type to use for your business use case and to gain insights into the dataset requirements for the identified model type.


Data models explorer

The **data models explorer** is a tool in Amazon Fraud Detector console that aligns your business use case with the model type supported by Amazon Fraud Detector. The data models explorer also provides insights into the data elements required by Amazon Fraud Detector to create your fraud detection model. Before you start to prepare your event dataset, use the data models explorer to figure out the model type Amazon Fraud Detector recommends for your business use and also to see a list of mandatory, recommended, and optional data elements you will need to create your dataset.

To use data models explorer,

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Data models explorer**.
3. In the **Data models explorer** page, under **Business use case**, select the business use case you want to evaluate for fraud risk.
4. Amazon Fraud Detector displays the recommended model type that matches with your business use case. The model type defines the algorithms, enrichments, and transformations Amazon Fraud Detector will use to train your fraud detection model.

Make a note of the recommended model type. You will need this later when you create your model.

 **Note**

If you do not find your business use case, use the **reach us** link in the description to provide us the details of your business use case. We will recommend the model type to use for creating a fraud detection model for your business use case.

5. The **Data model insights** pane provides insight into the mandatory, recommended, and optional data elements required to create and train a fraud detection model for your business use case. Use the information in the insights pane to gather your event data and to create your data set.

Gather event data

Gathering your event data is an important step in creating your model. This is because the performance of your model in predicting fraud is dependent on the quality of your dataset. As you start to gather your event data, keep in mind the list of data elements that the **Data models explorer** provided for you to create your dataset. You will need to gather all the mandatory (event metadata) data and decide what recommended and optional data elements (event variables) to include based on your goals for creating the model. It's also important to decide the format of each event variable you intend to include and the total size of your dataset.

Event dataset quality

To gather high quality dataset for your model, we recommend the following:

- **Collect mature data-** Using the most recent data helps to identify the most recent fraud pattern. However, to detect fraud use cases, allow the data to mature. The maturity period is dependent on your business, and can take anywhere from two weeks to three months. For example, if your event includes credit card transaction, then the maturity of the data might be determined by the chargeback period of the credit card or time taken by an investigator to make determination.

Ensure that the dataset used to train the model have had sufficient time to mature as per your business.

- **Make sure the data distribution doesn't drift significantly-** Amazon Fraud Detector model training process samples and partitions your dataset based on `EVENT_TIMESTAMP`. For example, if your dataset consists of fraud events pulled from last 6 months, but only the last month of legitimate events are included, the data distribution is considered to be drifting and unstable. An unstable dataset might lead to biases in model performance evaluation. If you find the data distribution to be drifting significantly, consider balancing your dataset by collecting data similar to the current data distribution.
- **Make sure the dataset is representative of the use case where the model is implemented/ tested-** Otherwise, the estimated performance could be biased. Let us say that you are using a model to automatically decline all in-door applicants, but your model is trained with a dataset that has historical data/labels which were previously approved. Then, your model's evaluation might be inaccurate because the evaluation is based on the dataset that does not have representation from declined applicants.

Event data format

Amazon Fraud Detector transforms most of your data to the required format as part of its model training process. However, there are some standard formats you can easily use for providing your data that can help avoid issues later when Amazon Fraud Detector validates your dataset. The following table provides guidance on the formats for providing the recommended event metadata.

Note

When you create your CSV file, make sure to enter event metadata name as listed below, in uppercase letters.

| Metadata name | Format | Required |
|-----------------------|--|---------------------------|
| <code>EVENT_ID</code> | <p>If provided, it must meet the following requirements:</p> <ul style="list-style-type: none"> • It is unique for that event. • It represents information that's meaningful to your business. | Depends on the model type |

| Metadata name | Format | Required |
|---------------|--|----------|
| | <ul style="list-style-type: none">• It follows the regular expression pattern (for example, <code>^[0-9a-z_-]+\$.)</code>• In addition to the above requirements, we recommend that you do not append a timestamp to the <code>EVENT_ID</code>. Doing so might cause issues when you update the event. This because you must provide the exact same <code>EVENT_ID</code> if you do this. | |

| Metadata name | Format | Required |
|-----------------|---|----------|
| EVENT_TIMESTAMP | <ul style="list-style-type: none"> • It must be specified in one of the following formats: <ul style="list-style-type: none"> • %yyyy-%mm-%ddT%hh:%mm:%ssZ (ISO 8601 standard in UTC only with no milliseconds) <p>Example: 2019-11-30T13:01:01Z</p> • %yyyy/%mm/%dd %hh:%mm:%ss (AM/PM) <p>Examples: 2019/11/30 1:01:01 PM, or 2019/11/30 13:01:01</p> <ul style="list-style-type: none"> • %mm/%dd/%yyyy %hh:%mm:%ss <p>Examples: 11/30/2019 1:01:01 PM, 11/30/2019 13:01:01</p> <ul style="list-style-type: none"> • %mm/%dd/%yy %hh:%mm:%ss <p>Examples: 11/30/19 1:01:01 PM, 11/30/19 13:01:01</p> <ul style="list-style-type: none"> • Amazon Fraud Detector makes the following assumptions when parsing date/timestamp formats for event timestamps: <ul style="list-style-type: none"> • If you are using the ISO 8601 standard, it must | Yes |

| Metadata name | Format | Required |
|---------------|--|----------|
| | <p>be an exact match of the preceding specification</p> <ul style="list-style-type: none">• If you are using one of the other formats, there is additional flexibility:<ul style="list-style-type: none">• For months and days, you can provide single or double digits. For example, 1/12/2019 is a valid date.• You do not need to include hh:mm:ss if you do not have them (that is, you can simply provide a date). You can also provide a subset of just the hour and minutes (for example, hh:mm). Just providing hour is not supported. Milliseconds are also not supported.• If you provide AM/PM labels, a 12-hour clock is assumed. If there is no AM/PM information, a 24-hour clock is assumed.• You can use "/" or "-" as delimiters for the date elements. ":" is assumed for the timestamp elements. | |

| Metadata name | Format | Required |
|-----------------|---|---|
| ENTITY_ID | <ul style="list-style-type: none"> It must follow the regular expression pattern: $^{\wedge}[0-9A-Za-z_.\@+-]^{\dagger}$ $\\$.$ If the entity id isn't available at the time of evaluation, specify the entity id as <i>unknown</i>. | Depends on the model type |
| ENTITY_TYPE | You can use any string | Depends on the model type |
| EVENT_LABEL | You can use any labels, such as "fraud", "legit", "1", or "0". | Required if LABEL_TIMESTAMP is included |
| LABEL_TIMESTAMP | It must follow the timestamp format. | Required if EVENT_LABEL is included |

For information about event variables, see [Variables](#).

Important

If you are creating Account Takeover Insights (ATI) model, see [Preparing data](#) for details on preparing and selecting data.

Null or missing values

The EVENT_TIMESTAMP and EVENT_LABEL variables must not contain any null or missing values. You can have null or missing values for other variables. However, we recommend that you only use a small number nulls for those variables. If Amazon Fraud Detector determines that there are too many null or missing values for an event variables, it will automatically omit variable from your model.

Minimum variables

When you create your model, the dataset must include at least two event variables in addition to the required event metadata. The two event variables must pass the validation check.

Event dataset size

Required

Your dataset must meet the following basic requirements for successful model training.

- Data from at least 100 events.
- Dataset must include at least 50 events (rows) classified as fraudulent.

Recommended

We recommend that your dataset includes the following for successful model training and a good model performance.

- Include a minimum of three weeks of historic data, but at best six months of data.
- Include a minimum of 10K total events data.
- Include at least 400 events (rows) classified as fraudulent and 400 events (rows) classified as legitimate.
- Include more than 100 unique entities, if your model type requires ENTITY_ID.

Dataset validation

Before Amazon Fraud Detector starts to create your model, it checks if the variables included in the dataset for training the model meets the size, format, and other requirements. If the dataset doesn't pass the validation, model isn't created. You must first fix the variables that didn't pass the validation before you create the model. Amazon Fraud Detector provides you with a *Data profiler* which you can use to help you identify and fix issues with your dataset before you start to train your model

Data profiler

Amazon Fraud Detector provides an open-source tool for profiling and preparing your data for model training. This automated data profiler helps you avoid common data preparation errors and identify potential issues like mis-mapped variable types that would negatively impact model performance. The profiler generates an intuitive and comprehensive report of your dataset,

including variable statistics, label distribution, categorical and numeric analysis, and variable and label correlations. It provides guidance on variable types as well as an option to transform the dataset into a format that Amazon Fraud Detector requires.

Using data profiler

The automated data profiler is built with an AWS CloudFormation stack, which you can easily launch with a few clicks. All codes are available on [Github](#). For information on how to use data profiler, follow directions in our blog [Train models faster with an automated data profiler for Amazon Fraud Detector](#)

Common event dataset errors

The following are some of the common issues Amazon Fraud Detector comes across when validating an event dataset. After you run the data profiler, use this list to check your dataset for errors before creating your model.

- CSV file isn't in the UTF-8 format.
- The number of events in the dataset is less than 100.
- The number of events identified as fraud or legitimate is less than 50.
- The number of unique entities associated to a fraud event is less than 100.
- More than 0.1% of values in EVENT_TIMESTAMP contains nulls or values other than the supported date/timestamp formats.
- More than 1% of the values in EVENT_LABEL contains nulls or values other than those defined in the event type.
- Less than two variables are available for model training.

Dataset storage

After you gathered your dataset, you store your dataset internally using Amazon Fraud Detector or externally with Amazon Simple Storage Service (Amazon S3). We recommend that you choose where to store your dataset based on the model you use for generating fraud predictions. For more information on model types, see [Choose a model type](#). For more information on storing your dataset, see [Event data storage](#).

Event type

With Amazon Fraud Detector you generate fraud predictions for events. An event type defines the structure for an individual event sent to Amazon Fraud Detector. Once defined, you can build models and detectors that evaluate the risk for specific event types.

The structure of an event includes the following:

- **Entity Type:** Classifies who is performing the event. During prediction, specify the entity type and entity Id to define who performed the event.
- **Variables:** Defines what variables can be sent as part of the event. Variables are used by models and rules to evaluate fraud risk. Once added, variables cannot be removed from an event type.
- **Labels:** Classifies an event as fraudulent or legitimate. Used during model training. Once added, labels cannot be removed from an event type.

Create an event type

Before you create your fraud detection model, you must first create an event type. Creating an event type involves defining your business activity (event) to evaluate for fraud. Defining the event involves identifying the event variables in your dataset to include for fraud evaluation, specifying the entity initiating event, and the labels that classify the event.

Prerequisites for creating an event type

Before you start to create your event type, make sure that you have completed the following:

- Used the [Data models explorer](#) tool to gain insights into the data elements required by Amazon Fraud Detector to create your fraud detection model.
- Used the insights you got from the Data Models Explorer to create your event dataset and have uploaded your dataset to Amazon S3 bucket.
- Created [Variables](#), [Entity](#), and [Labels](#) you want Amazon Fraud Detector to use for creating a fraud detection model for this event. Make sure that the variables, entity type, and labels you created are included in your event dataset.


You can create your event type in the Amazon Fraud Detector console, using the API, using the AWS CLI, or using the AWS SDK.

Create event type in the Amazon Fraud Detector console

To create an event type,

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Events**.
3. In the **Events type** page, choose **Create**.
4. Under **Event type details**,
 - a. In the **Name**, enter the name of your event.
 - b. In the **Description**, optionally, enter a description.
 - c. In the **Entity**, select the entity type you created for your event.
5. Under **Event variables**,
 - In the **Choose how to define this event's variables**,
 - If you have already created your event variables for this event, select **Select variables from your variable list** and in the **Variables**, select the variables you created for this event.
 - If you have not created variables for this event, select **Select variables from a training dataset**,
 - In the **IAM role** select the IAM Role you want Amazon Fraud Detector to use to access the Amazon S3 bucket that contains your dataset
 - In the **Data location** enter the path to your dataset location. Use the S3 URI path that is similar to this: `S3://your-bucket-name/example dataset filename.csv`.
 - Choose **Upload**.
 - Under **Variables**, all the event variable names that Amazon Fraud Detector has extracted from your dataset file is displayed.
 - If you want the variable to be included for detecting fraud, in the **Variable type**, select the variable type. Choose **Remove** to remove the variables from being included for fraud detection. Repeat this step for each variable in the list.
6. Under **Labels (optional)**, in the **Labels**, select the labels you created for this event. Make sure to select one label each for fraudulent and legitimate events.

7. If you want to setup automatic downstream processing for this event, under **Event orchestration with Amazon EventBridge - optional**, turn on **Enable event orchestration with Amazon EventBridge**. For more information about event orchestration, see [Event orchestration](#).

 **Note**

You can also enable event orchestration later after creating your event type.

8. Choose **Create event type**.

Create an event type using the AWS SDK for Python (Boto3)

The following example shows a sample request for the PutEventType API. The example assumes you have created the variables `ip_address` and `email_address`, the labels `legit` and `fraud`, and the entity type `sample_customer`. For information about how to create these resources, see [Resources](#).

 **Note**

You must first create variables, entity types, and labels prior to adding them to the event type.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_event_type (
    name = 'sample_registration',
    eventVariables = ['ip_address', 'email_address'],
    labels = ['legit', 'fraud'],
    entityTypees = ['sample_customer'])
```

Delete an event or event type

When you delete an event, Amazon Fraud Detector permanently deletes that event and the data associated with the event is no longer stored in Amazon Fraud Detector.

To delete an event that Amazon Fraud Detector has evaluated through GetEventPrediction API

1. Sign in to the AWS Management Console and open the Amazon Fraud Detector console at <https://console.aws.amazon.com/frauddetector>.
2. In the left navigation pane of the console, choose **Search past predictions**.
3. Choose the event that you want to delete.
4. Choose **Actions**, and then choose **Delete event**.
5. Enter **delete**, and then choose **Delete event**.

Note

This deletes all records that are associated to that Event ID, including any event data sent to the `SendEvent` operation and any prediction data generated through the `GetEventPrediction` operation.

To delete an event that is stored in Amazon Fraud Detector but has not been evaluated (that is, it was stored via the `SendEvent` operation), you must make a `DeleteEvent` request and specify the Event ID and Event Type ID. If you want to delete both the event and any prediction history associated with the event, set the value of the `deleteAuditHistory` parameter to `"true"`. With `deleteAuditHistory` parameter set to `"true"`, the event data is available through search for up to 30 seconds after the delete operation is completed.

To delete all events associated with an event type

1. In the left navigation pane of the console, choose **Event types**
2. Choose the event type for which you want all events deleted.
3. Navigate to **Stored events** tab and choose **Delete stored events**

Depending on the number of stored events for the event type, it might take some time to delete all stored events. For example, a 1 GB dataset (approximately 1-2 million events for the average customer) takes around 2 hours to delete. During this time, new events that you send to Amazon Fraud Detector of this event type are not stored, but you can continue to generate fraud predictions via the `GetEventPrediction` operation.

To delete an event type

You cannot delete an event type that is used in a detector or a model, or has associated stored events. Before deleting an event type, you must delete all events that are associated with that event type.

When you delete an event type, Amazon Fraud Detector permanently deletes that event type and the data is no longer stored in Amazon Fraud Detector.

1. In the left navigation pane of the Amazon Fraud Detector console, choose **Resources**, and then choose **Events**.
2. Choose the event type that you want to delete.
3. Choose **Actions**, and then choose **Delete event type**.
4. Enter the event type name, and then choose **Delete event type**.

Event data storage

After you've gathered your dataset, you store your dataset internally using Amazon Fraud Detector or externally with Amazon Simple Storage Service (Amazon S3). We recommend that you choose where to store your dataset based on the model you use for generating fraud predictions. The following is a detailed breakdown of these two storage options.

- **Internal storage-** Your dataset is stored with Amazon Fraud Detector. All event data associated with an event is stored together. You can upload the event dataset that's stored with Amazon Fraud Detector at any time. You can either stream events one at a time to an Amazon Fraud Detector API, or import large datasets (up to 1GB) using the batch import feature. When you train a model using the dataset stored with Amazon Fraud Detector, you can specify a time range to limit the size of your dataset.
- **External storage-** Your dataset is stored in an external data source other than Amazon Fraud Detector. Currently, Amazon Fraud Detector supports using Amazon Simple Storage Service (Amazon S3) for this purpose. If your model is on a file that's uploaded to Amazon S3, that file can't be more than 5GB of uncompressed data. If it's more than that, make sure to shorten the time range of your dataset.

The following table provides details about the model type and the data source it supports.

| Model type | Compatible training data source |
|----------------------------|------------------------------------|
| Online Fraud Insights | External storage, Internal storage |
| Transaction Fraud Insights | Internal storage |
| Account Takeover Insights | Internal storage |

For information on storing your dataset externally with Amazon Simple Storage Service, see [Store your event data externally with Amazon S3](#). For information on storing your dataset internally with Amazon Fraud Detector see [Store your event data internally with Amazon Fraud Detector](#).

Store your event data externally with Amazon S3

If you are training an Online Fraud Insights model, you can choose to store your event data externally with Amazon S3. To store your event data in Amazon S3 you must first create a text file in CSV format, add your event data, and then upload the CSV file to an Amazon S3 bucket.

Note

The **Transaction Fraud Insights** and **Account Takeover Insights** model types do not support datasets stored externally with Amazon S3

Create CSV file

Amazon Fraud Detector requires that the first row of your CSV file contain column headers. The column headers in your CSV file must map to the variables that are defined in the event type. For an example dataset, see [Get and upload example dataset](#)

The Online Fraud Insights model requires a training dataset that has at least 2 variables and up to 100 variables. In addition to the event variables, the training dataset must contain the following headers:

- **EVENT_TIMESTAMP** - Defines when the event occurred
- **EVENT_LABEL** - Classifies the event as fraudulent or legitimate. The values in the column must correspond to the values defined in the event type.

The following sample CSV data represents historical registration events from an online merchant:

```
EVENT_TIMESTAMP,EVENT_LABEL,ip_address,email_address
4/10/2019 11:05,fraud,209.146.137.48,fake_burtonlinda@example.net
12/20/2018 20:04,legit,203.0.112.189,fake_davidbutler@example.org
3/14/2019 10:56,legit,169.255.33.54,fake_shelby76@example.net
1/3/2019 8:38,legit,192.119.44.26,fake_curtis40@example.com
9/25/2019 3:12,legit,192.169.85.29,fake_rmiranda@example.org
```

Note

The CSV data file can contain double quotes and commas as part of your data.

A simplified version of the corresponding event type is represented below. The event variables correspond to the headers in the CSV file and the values in `EVENT_LABEL` correspond to the values in the labels list.

```
(  
  name = 'sample_registration',  
  eventVariables = ['ip_address', 'email_address'],  
  labels = ['legit', 'fraud'],  
  entityType = ['sample_customer']  
)
```

Event Timestamp formats

Ensure that your event timestamp is in the required format. As part of the model build process, the Online Fraud Insights model type orders your data based on the event timestamp, and splits your data for training and testing purposes. To get a fair estimate of performance, the model first trains on the training dataset, and then tests this model on the test dataset.

Amazon Fraud Detector supports the following date/timestamp formats for the values in `EVENT_TIMESTAMP` during model training:

- `%yyyy-%mm-%ddT%hh:%mm:%ssZ` (ISO 8601 standard in UTC only with no milliseconds)

Example: 2019-11-30T13:01:01Z

- `%yyyy/%mm/%dd %hh:%mm:%ss (AM/PM)`

Examples: 2019/11/30 1:01:01 PM, or 2019/11/30 13:01:01

- `%mm/%dd/%yyyy %hh:%mm:%ss`

Examples: 11/30/2019 1:01:01 PM, 11/30/2019 13:01:01

- `%mm/%dd/%yy %hh:%mm:%ss`

Examples: 11/30/19 1:01:01 PM, 11/30/19 13:01:01

Amazon Fraud Detector makes the following assumptions when parsing date/timestamp formats for event timestamps:

- If you are using the ISO 8601 standard, it must be an exact match of the preceding specification
- If you are using one of the other formats, there is additional flexibility:

- For months and days, you can provide single or double digits. For example, 1/12/2019 is a valid date.
- You do not need to include hh:mm:ss if you do not have them (that is, you can simply provide a date). You can also provide a subset of just the hour and minutes (for example, hh:mm). Just providing hour is not supported. Milliseconds are also not supported.
- If you provide AM/PM labels, a 12-hour clock is assumed. If there is no AM/PM information, a 24-hour clock is assumed.
- You can use "/" or "-" as delimiters for the date elements. ":" is assumed for the timestamp elements.

Sampling your dataset across time

We recommend that you provide examples of fraud and legitimate samples from the same time range. For example, if you provide fraud events from the past 6 months, you should also provide legitimate events that evenly span the same time period. If your dataset contains a highly uneven distribution of fraud and legitimate events, you might receive the following error: *"The fraud distribution across time is unacceptably fluctuant. Cannot split dataset properly."* Typically, the easiest fix for this error is to ensure that the fraud events and legitimate events are sampled evenly across the same timeframe. You also might need to remove data if you experienced a large spike in fraud within a short time period.

If you cannot generate enough data to create an evenly distributed dataset, one approach is to randomize the `EVENT_TIMESTAMP` of your events such that they are evenly distributed. However, this often results in performance metrics being unrealistic because Amazon Fraud Detector uses `EVENT_TIMESTAMP` to evaluate models on the appropriate subset of events in your dataset.

Null and missing values

Amazon Fraud Detector handles null and missing values. However, the percentage of nulls for variables should be limited. `EVENT_TIMESTAMP` and `EVENT_LABEL` columns should not contain any missing values.

File validation

Amazon Fraud Detector will fail to train a model if any of the following conditions are triggered:

- If the CSV is unable to be parsed
- If the datatype for a column is incorrect

Upload your event data to an Amazon S3 bucket

After you create a CSV file with your event data, upload the file to your Amazon S3 bucket.

To upload to an Amazon S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.

2. Choose **Create bucket**.

The **Create bucket** wizard opens.

3. In **Bucket name**, enter a DNS-compliant name for your bucket.

The bucket name must:

- Be unique across all of Amazon S3.
- Be between 3 and 63 characters long.
- Not contain uppercase characters.
- Start with a lowercase letter or number.

After you create the bucket, you can't change its name. For information about naming buckets, see [Bucket naming rules](#) in the *Amazon Simple Storage Service User Guide*.

Important

Avoid including sensitive information, such as account numbers, in the bucket name. The bucket name is visible in the URLs that point to the objects in the bucket.

4. In **Region**, choose the AWS Region where you want the bucket to reside. You must select the same Region in which you are using Amazon Fraud Detector, that is US East (N. Virginia), US East (Ohio), US West (Oregon), Europe (Ireland), Asia Pacific (Singapore) or Asia Pacific (Sydney).
5. In **Bucket settings for Block Public Access**, choose the Block Public Access settings that you want to apply to the bucket.

We recommend that you leave all settings enabled. For more information about blocking public access, see [Blocking public access to your Amazon S3 storage](#) in the *Amazon Simple Storage Service User Guide*.

6. Choose **Create bucket**.
7. Upload training data file to your Amazon S3 bucket. Note the Amazon S3 location path for your training file (for example, `s3://bucketname/object.csv`).

Store your event data internally with Amazon Fraud Detector

You can choose to store event data in Amazon Fraud Detector and use the stored data later to train your models. By storing event data in Amazon Fraud Detector, you can train models that use auto-computed variables to improve performance, simplify model retraining, and update fraud labels to close the machine learning feedback loop. Events are stored at the Event Type resource level, so all events of the same event type are stored together in a single event type dataset. As part of defining an event type, you can optionally specify whether to store events for that event type by toggling the *Event Ingestion* setting in the Amazon Fraud Detector console.

You can either store single events or import large number of event datasets in Amazon Fraud Detector. Single events can be streamed using the [GetEventPrediction](#) API or the [SendEvent](#) API. Large datasets can be quickly and easily imported to Amazon Fraud Detector using the batch import feature in the Amazon Fraud Detector console or using the [CreateBatchImportJob](#) API.

You can use the Amazon Fraud Detector console at any time to check the number of events already stored for each event type.

Prepare event data for storage

Event data that is stored internally with Amazon Fraud Detector is stored at the Event Type resource level. So, all event data that are from the same event are stored in a single Event Type. The stored events can later be used to train a new model or re-train an existing model. When training a model using the stored event data, you can optionally specify a time range of events to limit the size of your training dataset.

Each time you store your data in Amazon Fraud Detector, using the Amazon Fraud Detector console, the [SendEvent](#) API, or the [CreateBatchImportJob](#) API, Amazon Fraud Detector validates your data before storing. If your data fails validation, the event data is not stored.

Prerequisites for storing data internally with Amazon Fraud Detector

- To ensure that your event data passes validation and the dataset gets stored successfully, make sure you have used the insights provided by [Data models explorer](#) to prepare your dataset.

- Created an event type for the event data you want to store with Amazon Fraud Detector. If you haven't, follow instructions to [Create an event type](#).

Smart Data Validation

When you upload your dataset in Amazon Fraud Detector console for batch import, Amazon Fraud Detector uses Smart Data Validation (SDV) to validate your dataset before importing your data. SDV scans the uploaded data file and identifies issues such as missing data, and incorrect format or data types. In addition to validating your dataset, SDV also provides a validation report that lists all issues that were identified and suggests actions to fix issues that are most impactful. Some of the issues identified by SDV might be critical and must be addressed before Amazon Fraud Detector can successfully import your dataset. For more information, see [Smart Data Validation report](#).

The SDV validates your dataset at the file level and at the data (row) level. At the file level, SDV scans your data file and identifies issues such as inadequate permissions to access the file, incorrect file size, file format, and headers (event metadata and event variables). At the data level, SDV scans each event data (row) and identifies issues such as incorrect data format, data length, timestamp format, and null values.

Smart Data Validation is currently available in the Amazon Fraud Detector console only and the validation is turned on by default. If you don't want Amazon Fraud Detector to use the Smart Data Validation before importing your dataset, turn off the validation in the Amazon Fraud Detector console when uploading your dataset.

Validating stored data when using APIs or AWS SDK

When uploading events via the `SendEvent`, `GetEventPrediction`, or `CreateBatchImportJob` API operation, Amazon Fraud Detector validates the following:

- The `EventIngestion` setting for that event type is `ENABLED`.
- Event timestamps cannot be updated. An event with a repeated event ID and different `EVENT_TIMESTAMP` will be treated as an error.
- Variable names and values match their expected format. For more information, see [Create a variable](#)
- Required variables are populated with a value.
- All event timestamps are not older than 18 months and are not in the future.

Store event data using batch import

With the batch import feature, you can quickly and easily upload large historical event datasets in Amazon Fraud Detector using the console, the API, or the AWS SDK. To use batch import, create an input file in CSV format that contains all your event data, upload the CSV file onto Amazon S3 bucket, and start an *Import* job. Amazon Fraud Detector first validates the data based on the event type, and then automatically imports the entire dataset. After the data is imported, it's ready to be used for training new models or for re-training existing models.

Input and output files

The input CSV file must contain headers that match the variables defined in the associated event type plus four mandatory variables. See [Prepare event data for storage](#) for more information. The maximum size of the input data file is 20 Gigabytes (GB), or about 50 million events. The number of events will vary by your event size. If the import job was successful, the output file is empty. If the import was unsuccessful, the output file contains the error logs.

Create a CSV file

Amazon Fraud Detector imports data only from files that are in the comma-separated values (CSV) format. The first row of your CSV file must contain column headers that exactly match the variables defined in the associated event type plus four mandatory variables: `EVENT_ID`, `EVENT_TIMESTAMP`, `ENTITY_ID`, and `ENTITY_TYPE`. You can also optionally include `EVENT_LABEL` and `LABEL_TIMESTAMP` (`LABEL_TIMESTAMP` is required if `EVENT_LABEL` is included).

Define mandatory variables

Mandatory variables are considered as event metadata and they must be specified in uppercase. Event metadata are automatically included for model training. The following table lists the mandatory variables, description of each variable, and required format for the variable.

| Name | Description | Requirements |
|-----------------------|---|--|
| <code>EVENT_ID</code> | An identifier for the event. For example, if your event is an online transaction, the <code>EVENT_ID</code> might be the transaction reference number | <ul style="list-style-type: none">The <code>EVENT_ID</code> is required for batch import jobs.It must be unique for that event. |

| Name | Description | Requirements |
|------|-------------------------------------|---|
| | that was provided to your customer. | <ul style="list-style-type: none">• It should represent information that's meaningful to your business.• It must satisfy the regular expression pattern (for example, <code>^[0-9a-z_-]+\$.)</code>• We don't recommend that you append a timestamp to the <code>EVENT_ID</code>. Doing so might cause issues when you update the event. This because you must provide the exact same <code>EVENT_ID</code> if you do this. |

| Name | Description | Requirements |
|-----------------|--|---|
| EVENT_TIMESTAMP | The timestamp of when the event occurred. The timestamp must be in ISO 8601 standard in UTC. | <ul style="list-style-type: none"> • The EVENT_TIMESTAMP is required for batch import jobs. • It must be specified in one of the following formats: <ul style="list-style-type: none"> • %yyyy-%mm-%ddT%hh:%mm:%ssZ (ISO 8601 standard in UTC only with no milliseconds) <p>Example: 2019-11-30T13:01:01Z</p> • %yyyy/%mm/%dd %hh:%mm:%ss (AM/PM) <p>Examples: 2019/11/30 1:01:01 PM, or 2019/11/30 13:01:01</p> <ul style="list-style-type: none"> • %mm/%dd/%yyyy %hh:%mm:%ss <p>Examples: 11/30/2019 1:01:01 PM, 11/30/2019 13:01:01</p> <ul style="list-style-type: none"> • %mm/%dd/%yy %hh:%mm:%ss <p>Examples: 11/30/19 1:01:01 PM, 11/30/19 13:01:01</p> • Amazon Fraud Detector makes the following assumptions when parsing date/timestamp formats for event timestamps: |

| Name | Description | Requirements |
|------|-------------|--|
| | | <ul style="list-style-type: none">• If you are using the ISO 8601 standard, it must be an exact match of the preceding specification• If you are using one of the other formats, there is additional flexibility:<ul style="list-style-type: none">• For months and days, you can provide single or double digits. For example, 1/12/2019 is a valid date.• You do not need to include hh:mm:ss if you do not have them (that is, you can simply provide a date). You can also provide a subset of just the hour and minutes (for example, hh:mm). Just providing hour is not supported. Milliseconds are also not supported.• If you provide AM/PM labels, a 12-hour clock is assumed. If there is no AM/PM information, a 24-hour clock is assumed.• You can use "/" or "-" as delimiters for the date elements. |

| Name | Description | Requirements |
|-----------------|--|---|
| | | ":" is assumed for the timestamp elements. |
| ENTITY_ID | An identifier for the entity performing the event. | <ul style="list-style-type: none"> ENTITY_ID is required for batch import jobs It must follow the regular expression pattern: $^[\text{0-9A-Za-z_}.\text{@+-}]^{\\$}$ If the entity id isn't available at the time of evaluation, specify the entity id as <i>unknown</i>. |
| ENTITY_TYPE | The entity that performs the event, such as a merchant or a customer | ENTITY_TYPE is required for batch import jobs |
| EVENT_LABEL | Classifies the event as <code>fraudulent</code> or <code>legitimate</code> | EVENT_LABEL is required if LABEL_TIMESTAMP is included |
| LABEL_TIMESTAMP | The timestamp when the event label was last populated or updated | <ul style="list-style-type: none"> LABEL_TIMESTAMP is required if EVENT_LABEL is included. It must follow the timestamp format. |

Upload CSV file to Amazon S3 for batch import

After you create a CSV file with your data, upload the file to your Amazon Simple Storage Service (Amazon S3) bucket.

To upload event data to an Amazon S3 bucket

1. Sign in to the AWS Management Console and open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
2. Choose **Create bucket**.

The **Create bucket** wizard opens.

3. In **Bucket name**, enter a DNS-compliant name for your bucket.

The bucket name must:

- Be unique across all of Amazon S3.
- Be between 3 and 63 characters long.
- Not contain uppercase characters.
- Start with a lowercase letter or number.

After you create the bucket, you can't change its name. For information about naming buckets, see [Bucket naming rules](#) in the *Amazon Simple Storage Service User Guide*.

Important

Avoid including sensitive information, such as account numbers, in the bucket name. The bucket name is visible in the URLs that point to the objects in the bucket.

4. In **Region**, choose the AWS Region where you want the bucket to reside. You must select the same Region in which you are using Amazon Fraud Detector, that is US East (N. Virginia), US East (Ohio), US West (Oregon), Europe (Ireland), Asia Pacific (Singapore) or Asia Pacific (Sydney).
5. In **Bucket settings for Block Public Access**, choose the Block Public Access settings that you want to apply to the bucket.

We recommend that you leave all settings enabled. For more information about blocking public access, see [Blocking public access to your Amazon S3 storage](#) in the *Amazon Simple Storage Service User Guide*.

6. Choose **Create bucket**.
7. Upload training data file to your Amazon S3 bucket. Note the Amazon S3 location path for your training file (for example, `s3://bucketname/object.csv`).

Batch import event data in Amazon Fraud Detector console

You can easily import large number of your event datasets in Amazon Fraud Detector console, using the `CreateBatchImportJob` API or using AWS SDK. Before you proceed, make sure that you have followed instructions to prepare your dataset as a CSV file. Make sure that you also uploaded the CSV file to an Amazon S3 bucket.

Using Amazon Fraud Detector console

To batch import event data in console

1. Open the AWS Console and sign in to your account, and navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Events**.
3. Choose your event type.
4. Select **Stored events** tab.
5. In the **Stored events details** pane, make sure that the **Event ingestion** is **ON**.
6. In the **Import events data** pane, choose **New Import**.
7. In the **New events import** page, provide the following information:
 - [Recommended] Leave **Enable Smart Data Validation for this dataset - new** set to the default setting.
 - For **IAM role for data**, select the IAM role that you created for the Amazon S3 bucket that holds the CSV file you are planning to import.
 - For **Input data location**, enter the S3 location where you have your CSV file.
 - If you want to specify a separate location to store your import results, click **Separate data location for inputs and results** button and provide a valid Amazon S3 bucket location.
8. Choose **Start**.
9. The **Status** column in **Import events data** pane displays the status of your validation and import job. The banner at the top provides high level description of the status as your dataset first goes through validation and then the import.

Important

Make sure that the IAM role you selected has read permissions to your input Amazon S3 bucket and write permissions to your output Amazon S3 bucket.

10. Follow the guidance provided to [Monitor the progress of dataset validation and import job](#).

Monitor the progress of dataset validation and import job

If you are using the Amazon Fraud Detector console to perform a batch import job, by default, Amazon Fraud Detector validates your dataset before import. You can monitor the progress and status of validation and import jobs in the **New events import** page of the Amazon Fraud Detector console. A banner at the top of the page provides a brief description of the validation findings and the status of the import job. Depending on the validation findings and the status of your import job you might be required to take actions to ensure successful validation and import of your dataset.

The following table provides details of the actions you must take depending on the outcome of validation and import operations.

| Banner message | Status | What it means | What should I do |
|---|------------------------|--|--|
| Data validation has started | Validation in progress | SDV has started validating your dataset | Wait for the status to change |
| Data validation cannot proceed due to errors in your dataset. Fix errors in your data file and start a new import job. See the validation report for more information | Validation failed | SDV identified issues in your data file. These issues must be addressed for successful import of your dataset. | In the Import events data pane, select the Job Id and view the validation report. Follow the Recommendations in the report to address all the errors listed. For more information, see Using the validation report . |

| Banner message | Status | What it means | What should I do |
|---|--------------------|---|---|
| Data import has started. Validation completed successfully | Import in progress | Your dataset passed the validation. AFD has started to import your dataset | Wait for the status to change |
| Validation completed with warnings. Data import has started | Import in progress | Some of the data in your dataset failed validation. However, the data that passed validation meets the minimum data size requirements for import. | Monitor the message in the banner and wait for the status to change |

| Banner message | Status | What it means | What should I do |
|--|--|---|---|
| Your data was partially imported. Some of the data failed validation and did not get imported. See validation report for more information. | Imported. The status shows a warning icon. | Some of the data in your data file that failed validation did not get imported. The rest of the data that passed validation was imported. | In the Import events data pane, select the Job Id and view the validation report. Follow the Recommendations in the Data level warnings table to address the listed warnings. You need not address all the warnings. However, make sure that your dataset has more than 50% of data that passes validation for a successful import. After you have addressed the warnings, start a new import job. For more information, see Using the validation report . |
| Data import failed due to a processing error. Start a new data import job | Import failed | The import failed due to a transient run-time error | Start a new import job |
| Data was imported successfully | Imported | Both validation and import completed successfully | Select the Job Id of your import job to view details and then proceed with model training |

Note

We recommend waiting 10 minutes after the dataset has imported successfully into Amazon Fraud Detector to ensure that they are fully ingested by the system.

Smart Data Validation report

The Smart Data Validation creates a validation report after validation is complete. The validation report provides details of all the issues that the SDV has identified in your dataset, with suggested actions to fix the most impactful issues. You can use the validation report to determine what the issues are, where the issues are located in the dataset, the severity of the issues, and how to fix them. The validation report is created even when the validation completes successfully. In this case, you can view the report to see if there are any issues listed and if there are, decide if you want to fix any of those.

Note

The current version of SDV scans your dataset for issues that might cause the batch import to fail. If validation and batch import succeed, your dataset can still have issues that might cause model training to fail. We recommend that you view your validation report even if validation and import were successful, and address any issues listed in the report for successful model training. After you have addressed the issues, create a new batch import job.

Accessing the validation report

You can access the validation report any time after the validation completes using one of the following options:

1. After the validation completes and while the import job is in progress, in the top banner, choose **View validation report**.
2. After the import job completes, in the **Import events data** pane, choose the Job ID of the import job that just completed.

Using the validation report

The validation report page of your import job provides the details of this import job, a list of critical errors if any are found, a list of warnings about specific events (rows) in your dataset if found, and a brief summary of your dataset that includes information such as values that are not valid, and missing values for each variable.

- **Import job details**

Provides details of the import job. If your import job has failed or your dataset was partially imported, choose **Go to results file** to view the error logs of the events that failed to import.

- **Critical errors**


Provides details of the most impactful issues in your dataset identified by SDV. All the issues listed in this pane are critical and you must address them before you proceed with import. If you try to import your dataset without addressing the critical issues, your import job might fail.

To address the critical issues, follow the recommendations provided for each warning. After you have addressed all the issues listed in the Critical errors pane, create a new batch import job.

- **Data level warnings**

Provides a summary of the warnings for specific events (rows) in your dataset. If the Data level warnings pane is populated, some of the events in your dataset failed validation and were not imported.

For each warning, the **Description** column displays the number of events that has the issue. And the **Sample event IDs** provides a partial list of sample event IDs you can use as a starting point to locate the rest of the events that have the issue. Use the **Recommendation** provided for the warning to fix the issue. Also use the error logs from your output file for additional information about the issue. The error logs are generated for all the events that failed batch import. To access error logs, in the **Import job details** pane, choose **Go to results file**.

 **Note**

If more than 50% of the events (rows) in your dataset failed validation, the import job also fails. In this case, you must fix the data before you start a new import job.

- **Dataset summary**

Provides a summary of the validation report of your dataset. If the Number of warnings column shows more than 0 warnings, decide if you need to fix those warning. If the **Number of warnings** column shows 0s, continue to train your model.

Batch import event data using the AWS SDK for Python (Boto3)

The following example shows a sample request for [CreateBatchImportJob](#) API. A batch import job must include a **jobID**, **inputPath**, **outputPath**, **eventName** and **iamRoleArn**. The jobID can't contain the same ID of a past job, unless the job exists in CREATE_FAILED state. The inputPath and outputPath must be valid S3 paths. You can opt out of specifying the file name in the outputPath, however, you will still need to provide a valid S3 bucket location. The eventName and iamRoleArn must exist. The IAM role must grant read permissions to input Amazon S3 bucket and write permissions to output Amazon S3 bucket.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.create_batch_import_job (
    jobId = 'sample_batch_import',
    inputPath = 's3://bucket_name/input_file_name.csv',
    outputPath = 's3://bucket_name/',
    eventName = 'sample_registration',
    iamRoleArn: 'arn:aws:iam:*****:role/service-role/AmazonFraudDetector-
DataAccessRole-*****'
)
```

Cancel batch import job

You can cancel an in-progress batch import job at any time in the Amazon Fraud Detector console, using the `CancelBatchImportJob` API, or AWS SDK.

To cancel a batch import job in console,

1. Open the AWS Console and sign in to your account, and navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Events**.
3. Choose your event type.
4. Select **Stored events** tab.

5. In the **Import events data** pane, choose the job Id of an in-progress import job you want to cancel.
6. In the event job page, click **Actions** and select **Cancel events import**.
7. Choose **Stop events import** to cancel the batch import job.

Canceling batch import job using the AWS SDK for Python (Boto3)

The following example shows a sample request for the `CancelBatchImportJob` API. The cancel import job must include the job ID of an in-progress batch import job.

```
import boto3
fraudDetector = boto3.client('frauddetector')
fraudDetector.cancel_batch_import_job (
    jobId = 'sample_batch'
)
```

Store event data using the `GetEventPredictions` API operation

By default, all events sent to the `GetEventPrediction` API for evaluation are stored in Amazon Fraud Detector. This means that Amazon Fraud Detector will automatically store event data when you generate a prediction and use that data to update calculated variables in near-real time. You can disable data storage by navigating to the event type in the Amazon Fraud Detector console and setting **Event ingestion** OFF or updating the `EventIngestion` value to `DISABLED` using the `PutEventType` API operation. For more information about the `GetEventPrediction` API operation, see [Fraud predictions](#).

Important

We highly recommend that once you enable *Event ingestion* for an Event type, keep it enabled. Disabling the Event ingestion for the same Event type and then generating predictions might result in inconsistent behavior.

Store event data using the SendEvent API operation

You can use the SendEvent API operation to store events in Amazon Fraud Detector without generating fraud predictions for those events. For example, you can use the SendEvent operation to upload a historical dataset, which you can later use to train a model.

Event Timestamp formats for SendEvent API

When storing event data using SendEvent API, you must ensure that your event timestamp is in the required format. Amazon Fraud Detector supports the following date/timestamp formats:

- %yyyy-%mm-%ddT%hh:%mm:%ssZ (ISO 8601 standard in UTC only with no milliseconds)

Example: 2019-11-30T13:01:01Z

- %yyyy/%mm/%dd %hh:%mm:%ss (AM/PM)

Examples: 2019/11/30 1:01:01 PM, or 2019/11/30 13:01:01

- %mm/%dd/%yyyy %hh:%mm:%ss

Examples: 11/30/2019 1:01:01 PM, 11/30/2019 13:01:01

- %mm/%dd/%yy %hh:%mm:%ss

Examples: 11/30/19 1:01:01 PM, 11/30/19 13:01:01

Amazon Fraud Detector makes the following assumptions when parsing date/timestamp formats for event timestamps:

- If you are using the ISO 8601 standard, it must be an exact match of the preceding specification
- If you are using one of the other formats, there is additional flexibility:
 - For months and days, you can provide single or double digits. For example, 1/12/2019 is a valid date.
 - You do not need to include hh:mm:ss if you do not have them (that is, you can simply provide a date). You can also provide a subset of just the hour and minutes (for example, hh:mm). Just providing hour is not supported. Milliseconds are also not supported.
 - If you provide AM/PM labels, a 12-hour clock is assumed. If there is no AM/PM information, a 24-hour clock is assumed.
 - You can use "/" or "-" as delimiters for the date elements. ":" is assumed for the timestamp elements.

The following is an example `SendEvent` API call.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.send_event(
    eventId          = '802454d3-f7d8-482d-97e8-c4b6db9a0428',
    eventName       = 'sample_registration',
    eventTimestamp  = '2020-07-13T23:18:21Z',
    eventVariables = {
        'email_address' : 'johndoe@examplomain.com',
        'ip_address'   : '1.2.3.4'},
    assignedLabel   = 'legit',
    labelTimestamp  = '2020-07-13T23:18:21Z',
    entities        = [{'entityType':'sample_customer', 'entityId':'12345'}],
)
```

Get details of a stored event data

After you store event data in Amazon Fraud Detector, you can check the latest data that was stored for an event using the [GetEvent](#) API. The following example code checks the latest data stored for the `sample_registration` event.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.get_event(
    eventId          = '802454d3-f7d8-482d-97e8-c4b6db9a0428',
    eventName       = 'sample_registration'
)
```

View metrics of stored event dataset

For each event type, you can view metrics such as, number of stored events, total size of your stored events, and timestamps of the earliest and the latest stored events, in the Amazon Fraud Detector console.

To view stored event metrics of an event type,

1. Open the AWS Console and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Events**.
3. Choose your event type.
4. Select **Stored events** tab.
5. The **Stored events details** pane displays the metrics. These metrics are automatically updated once per day.
6. Optionally click **Refresh event metrics** to manually update your metrics.

Note

If you have just imported your data, we recommend waiting 5 - 10 minutes after you have finished importing data to refresh and view metrics.

Event orchestration

Event orchestration makes it easy for you to send events to AWS services for downstream processing, using [Amazon EventBridge](#). Amazon Fraud Detector provides you with simple rules you can use to automate processing of events after fraud detection. With event orchestration, you can automate downstream event processes such as, sending events to dashboards to get insights from event data, generating notifications based on the fraud detection outcomes, and updating events with a label based on the learning from fraud detection.


Event orchestration provides easy access to services in the AWS environment, through Amazon EventBridge. You can configure Amazon EventBridge to either send events directly to AWS services or indirectly using [API destinations](#). The AWS services you use to orchestrate your downstream processes are also called *targets*. Some of the targets you can use to orchestrate downstream processing are as follows:

- For monitoring and analytics — [Amazon QuickSight](#), [Amazon CloudWatch](#)
- For storage — [Amazon S3](#), [Amazon RDS](#), [Amazon DynamoDB](#)
- For sending notifications — [Amazon SNS](#), [Amazon SES](#)
- For custom processing — [AWS Lambda](#), [AWS Step Functions](#)

For more information on the orchestration targets supported by Amazon EventBridge, see [Amazon EventBridge targets](#).

The following diagram provides a high-level view of how event orchestration works.

Event Orchestration



Setting up event orchestration

Setting up event orchestration for your events requires you to set up processes in your target service, configure Amazon EventBridge to receive and send event data, and create rules in Amazon EventBridge that specifies the conditions for starting the downstream processes. Complete the following steps to set up event orchestration:

To set up event orchestration

1. Go to [Amazon EventBridge User Guide](#) and learn how to use Amazon EventBridge. Make sure to learn how to create [Rules](#) in Amazon EventBridge for your use case.
2. Follow instructions to [Enable event orchestration in Amazon Fraud Detector](#).

Note

The event orchestration for your event is *disabled* by default.

3. Set up your target service to receive and process the event data. For example, if your downstream process involves sending notifications and you want to use Amazon SNS, go to Amazon SNS console, create an SNS topic, and then subscribe an endpoint to the topic.

4. Follow instructions to [Create Amazon EventBridge rules](#).

Important

When building the event pattern in Amazon EventBridge, make sure to provide `aws.frauddetector` for the *source* field and `Event Prediction Result Returned` for the *detail-type* field.

Enable event orchestration in Amazon Fraud Detector

You can enable event orchestration for an event either when you are creating your event type or after you have created your event type. Event orchestration can be enabled in the Amazon Fraud Detector console, using the `put-event-type` command, using the `PutEventType` API, or using the AWS SDK for Python (Boto3).

Enable event orchestration in the Amazon Fraud Detector console

This example enables event orchestration for an event type that has already been created. If you are creating a new event type and want to enable orchestration, follow instructions to [Create an event type](#).

To enable event orchestration

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Events**.
3. In the **Events type** page, choose your event type.
4. Turn on **Enable event orchestration with Amazon EventBridge**.
5. Continue with step 3 instructions for [Setting up event orchestration](#).

Enable event orchestration using the AWS SDK for Python (Boto3)

The following example shows a sample request for updating an event type `sample_registration` to enable event orchestration. The example uses the `PutEventType` API and assumes you have created the variables `ip_address` and `email_address`, the labels

legit and fraud, and the entity type `sample_customer`. For information on how to create these resources, see [Resources](#).

```
import boto3
fraudDetector = boto3.client('frauddetector')
fraud_detector.put_event_type(
    name = 'sample_registration',
    eventVariables = ['ip_address', 'email_address'],
    eventOrchestration = {'eventBridgeEnabled': True},
    labels = ['legit', 'fraud'],
    entityTypes = ['sample_customer'])
```

Disable event orchestration in Amazon Fraud Detector

You can disable event orchestration for an event anytime in the Amazon Fraud Detector console, using the `put-event-type` command, using the `PutEventType` API, or using the AWS SDK for Python (Boto3).

Disable event orchestration in the Amazon Fraud Detector console

To disable event orchestration

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Events**.
3. In the **Events type** page, choose your event type.
4. Turn off **Enable event orchestration with Amazon EventBridge**.

Disable event orchestration using the AWS SDK for Python (Boto3)

The following example shows a sample request for updating an event type `sample_registration` to disable event orchestration using the `PutEventType` API.

```
import boto3
fraudDetector = boto3.client('frauddetector')
fraud_detector.put_event_type(
    name = 'sample_registration',
    eventVariables = ['ip_address', 'email_address'],
```

```
eventOrchestration = {'eventBridgeEnabled': False},  
entityTypes = ['sample_customer'])
```

Model

Amazon Fraud Detector uses machine learning models for generating fraud predictions. Each model is trained using a *model type*. The model type specifies the algorithms and transformations used for training the model. Model training is the process of using a dataset that you provide to create a model that can predict fraudulent events.

To create a model, you must first choose a model type and then prepare and provide data that will be used to train the model.

Choose a model type

The following model types are available in Amazon Fraud Detector. Choose a model type that works for your use case.

- **Online Fraud Insights**

The *Online Fraud Insights* model type is optimized to detect fraud when little historical data is available about the entity being evaluated, for example, a new customer registering online for a new account.

- **Transaction Fraud Insights**

The *Transaction Fraud Insights* model type is best suited for detecting fraud use cases where the entity that is being evaluated might have a history of interactions that the model can analyze to improve prediction accuracy (for example, an existing customer with history of past purchases).

- **Account Takeover Insights**

The *Account Takeover Insights* model type detects if an account was compromised by phishing or another type of attack. The login data of a compromised account, such as the browser and device used at login, is different from the historical login data that's associated with the account.

Online fraud insights

Online Fraud Insights is a supervised machine learning model, which means that it uses historical examples of fraudulent and legitimate transactions to train the model. The Online Fraud Insights model can detect fraud based on little historical data. The model's inputs are flexible, so you

can adapt it to detect a variety of fraud risks including fake reviews, promotion abuse, and guest checkout fraud.

The Online Fraud Insights model uses an ensemble of machine learning algorithms for data enrichment, transformation, and fraud classification. As part of the model training process, Online Fraud Insights enriches raw data elements like IP address and BIN number with third-party data such as the geolocation of the IP address or the issuing bank for a credit card. In addition to third-party data, Online Fraud Insights uses deep learning algorithms that take into account fraud patterns that have been seen at Amazon and AWS. These fraud patterns become input features to your model using a gradient tree boosting algorithm.

To increase performance, Online Fraud Insights optimizes the hyper parameters of the gradient tree boosting algorithm via a Bayesian optimization process. It sequentially trains dozens of different models with varying model parameters (such as number of trees, depth of trees, and number of samples per leaf). It also uses different optimization strategies like upweighting the minority fraud population to take care of very low fraud rates.

Selecting data source

When training an Online Fraud Insights model, you can choose to train the model on event data that is either stored externally (outside of Amazon Fraud Detector) or stored within Amazon Fraud Detector. The external storage Amazon Fraud Detector currently supports is Amazon Simple Storage Service (Amazon S3). If you are using external storage, your event dataset must be uploaded as a comma-separated values (CSV) format to an Amazon S3 bucket. These data storage options are referred to within the model training configuration as `EXTERNAL_EVENTS` (for external storage) and `INGESTED_EVENTS` (for internal storage). For more information about the available data sources and how to store data in them, see [Event data storage](#).

Preparing data

Regardless of where you choose to store your event data (Amazon S3 or Amazon Fraud Detector), the requirements for Online Fraud Insights model type are the same.

Your dataset must contain the column header `EVENT_LABEL`. This variable classifies an event as fraudulent or legitimate. When using a CSV file (external storage), you must include `EVENT_LABEL` for each event in the file. For internal storage, the `EVENT_LABEL` field is optional but all events must be labeled to be included within a training dataset. When configuring your model training, you can choose whether to ignore unlabeled events, assume a legitimate label for unlabeled events, or assume a fraudulent label for all unlabeled events.

Selecting data

See [Gather event data](#) for information on selecting data for training your Online Fraud Insights model.

The Online Fraud Insights training process samples and partitions historic data based on `EVENT_TIMESTAMP`. There is no need to manually sample the data and doing so may negatively impact your model results.

Event variables

The Online Fraud Insights model requires at least two variables, apart from the required event metadata, that has passed [data validation](#) for model training and allows up to 100 variables per model. Generally, the more variables you provide, the better the model can differentiate between fraud and legitimate events. While the Online Fraud Insights model can support dozens of variables, including custom variables, we recommend including IP address and email address because these variables are typically most effective at identifying the entity being evaluated.

Validating data

As part of the training process, Online Fraud Insights will validate the dataset for data quality issues that may impact model training. After validating the data, Amazon Fraud Detector will take appropriate action to build the best possible model. This includes issuing warnings for potential data quality issues, automatically removing variables that have data quality issues, or issuing an error and stopping the model training process. For more information, see [dataset validation](#).

Transaction fraud insights

The Transaction Fraud Insights model type is designed to detect online, or card-not-present, transaction fraud. Transaction Fraud Insights is a supervised machine learning model, which means that it uses historical examples of fraudulent and legitimate transactions to train the model.

The Transaction Fraud Insights model uses an ensemble of machine learning algorithms for data enrichment, transformation, and fraud classification. It leverages a feature engineering engine to create entity-level and event-level aggregates. As part of the model training process, Transaction Fraud Insights enriches raw data elements like IP address and BIN number with third-party data such as the geolocation of the IP address or the issuing bank for a credit card. In addition to third-party data, Transaction Fraud Insights uses deep learning algorithms that take into account fraud

patterns that have been seen at Amazon and AWS. These fraud patterns become input features to your model using a gradient tree boosting algorithm.

To increase performance, Transaction Fraud Insights optimizes the hyper parameters of the gradient tree boosting algorithm via a Bayesian optimization process, sequentially training dozens of different models with varying model parameters (such as number of trees, depth of trees, number of samples per leaf) as well as different optimization strategies like upweighting the minority fraud population to take care of very low fraud rates.

As part of the model training process, the Transaction Fraud model's feature engineering engine calculates values for each unique entity within your training dataset to help improve fraud predictions. For example, during the training process, Amazon Fraud Detector computes and stores the last time an entity made a purchase and dynamically updates this value each time you call the `GetEventPrediction` or `SendEvent` API. During a fraud prediction, the event variables are combined with other entity and event metadata to predict whether the transaction is fraudulent.

Selecting data source

Transaction Fraud Insights models are trained on dataset stored internally with Amazon Fraud Detector (`INGESTED_EVENTS`) only. This allows Amazon Fraud Detector to continuously update calculated values about the entities you are evaluating. For more information about the available data sources, see [Event data storage](#)

Preparing data

Before you train a Transaction Fraud Insights model, ensure that your data file contains all headers as mentioned in [Prepare event dataset](#). The Transaction Fraud Insights model compares new entities that are received with the examples of fraudulent and legitimate entities in the dataset, so it is helpful to provide many examples for each entity.

Amazon Fraud Detector automatically transforms the stored event dataset into the correct format for training. After the model has completed training, you can review the performance metrics and determine whether you should add entities to your training dataset.

Selecting data

By default, Transaction Fraud Insights trains on your entire stored dataset for the Event Type that you select. You can optionally set a time range to reduce the events that are used to train your model. When setting a time range, ensure that the records that are used to train the model have had sufficient time to mature. That is, enough time has passed to ensure legitimate and fraud

records have been correctly identified. For example, for chargeback fraud, it often takes 60 days or more to correctly identify fraudulent events. For the best model performance, ensure that all records in your training dataset are mature.

There is no need to select a time range that represents an ideal fraud rate. Amazon Fraud Detector automatically samples your data to achieve balance between fraud rates, time range, and entity counts.

Amazon Fraud Detector returns a validation error during model training if you select a time range for which there are not enough events to successfully train a model. For stored datasets, the `EVENT_LABEL` field is optional, but events must be labeled to be included in your training dataset. When configuring your model training, you can choose whether to ignore unlabeled events, assume a legitimate label for unlabeled events, or assume a fraudulent label for unlabeled events.

Event variables

The event type used to train the model must contain at least 2 variables, apart from required event metadata, that has passed [data validation](#) and can contain up to 100 variables. Generally, the more variables you provide, the better the model can differentiate between fraud and legitimate events. Although the Transaction Fraud Insight model can support dozens of variables, including custom variables, we recommend that you include IP address, email address, payment instrument type, order price, and card BIN.

Validating data

As part of the training process, Transaction Fraud Insights validates the training dataset for data quality issues that might impact model training. After validating the data, Amazon Fraud Detector takes appropriate action to build the best possible model. This includes issuing warnings for potential data quality issues, automatically removing variables that have data quality issues, or issuing an error and stopping the model training process. For more information, see [Dataset validation](#).

Amazon Fraud Detector will issue a warning but continue training a model if the number of unique entities is less than 1,500 because this can impact the quality of the training data. If you receive a warning, review the [performance metric](#).

Account takeover insights

The Account Takeover Insights (ATI) model type identifies fraudulent online activity by detecting if accounts were compromised through malicious takeovers, phishing, or from credentials being

stolen. Account Takeover Insights is a machine learning model that uses login events from your online business to train the model.

You can embed a trained Account Takeover Insights model within your real time login flow to detect if an account is compromised. The model assesses a variety of authentication and login types. They include web application logins, API-based authentications, and single-sign-on (SSO). To use the Account Takeover Insights model, call the [GetEventPrediction](#) API after a valid login credentials is presented. The API generates a score that quantifies the risk of the account being compromised. Amazon Fraud Detector uses the score and the rules that you defined to return one or more outcomes for the login events. The outcomes are ones that you configured. Based on the outcomes you receive, you can take appropriate actions for each login. That is, you can either approve or challenge the credentials presented for the login. For example, you can challenge the credentials by asking for an account PIN as an additional verification.

You can also use the Account Takeover Insights model to evaluate account logins asynchronously and take actions on high-risk accounts. For example, a high-risk account can be added to investigation queue for a human reviewer to determine if further action needs to be taken, such as suspend the account.

The Account Takeover Insights model is trained using a dataset that contains the historical login events of your business. You provide this data. You can optionally label the accounts as legitimate or fraudulent. However, this isn't required to train the model. The Account Takeover Insights model detects anomalies based on the history of successful logins of an account. It also learns how to detect anomalies in a user's behavior that suggest increased risk of an event of malicious account takeover. For example, a user that typically logs in from the same set of devices and IP addresses. A fraudster typically logs in from a different device and geolocation. This technique produces a risk score of an activity being anomalous, which typically is a primary characteristic of malicious account takeovers.

Before training an Account Takeover Insights model, Amazon Fraud Detector uses a combination of machine learning techniques to perform data enrichment, data aggregation, and data transformation . Then, during the training process, Amazon Fraud Detector enriches the raw data elements that you provide. Examples of raw data elements include IP address and user agent. Amazon Fraud Detector uses these elements to create additional inputs that describe the login data. These inputs include the device, browser, and geolocation inputs. Amazon Fraud Detector also uses the login data that you provide to continuously compute aggregated variables that describe the past user behavior. Examples of user behavior include the number of times that the user signed in from a specific IP address. Using these additional enrichments and aggregates,

Amazon Fraud Detector can generate strong model performance from a small set of inputs from your login events.

The Account Takeover Insights model detects instances where a legitimate account is accessed by a bad actor, regardless of whether the bad actor is human or a robot. The model produces a single score that indicates the relative risk of account compromise. Accounts that might have been compromised are flagged as high-risk accounts. You can process high-risk accounts by one of two ways. Either, you can enforce an additional identity verification. Or, you can send the account to a queue for manual investigation.

Selecting data source

Account Takeover Insights models are trained on a dataset that's stored internally, in Amazon Fraud Detector. To store your login events data with Amazon Fraud Detector, create a CSV file with login events of users. For each event, include login data such as the event timestamp, user ID, IP address, user agent, and whether the login data is valid. After creating the CSV file, first upload the file to Amazon Fraud Detector, and then use import feature to store the data. You can then train your model using the stored data. For more information on storing your event dataset with Amazon Fraud Detector see [Store your event data internally with Amazon Fraud Detector](#)

Preparing data

Amazon Fraud Detector requires that you provide your user account login data in a comma-separated values (CSV) file that's encoded in the UTF-8 format. The first line of your CSV file must contain a file header. The file header consists of event metadata and event variables that describe each data element. Event data follows the header. Each line in the event data consists of data from a single login event.

For the Accounts Takeover Insights model, you must provide the following event metadata and event variables in the header line of your CSV file.

Event metadata

We recommend that you provide the following metadata in your CSV file header. The event metadata must be in uppercase letters.

- **EVENT_ID** - A unique identifier for the login event.
- **ENTITY_TYPE** - The entity that performs the login event, such as a merchant or a customer.
- **ENTITY_ID** - An identifier for the entity performing the login event.

- **EVENT_TIMESTAMP** - The timestamp when the login event occurred. The timestamp must be in ISO 8601 standard in UTC.
- **EVENT_LABEL** (recommended) - A label that classifies the event as fraudulent or legitimate. You can use any labels, such as "fraud", "legit", "1", or "0".

Note

- Event metadata must be in uppercase letters. It's case sensitive.
- Labels aren't required for login events. However, we recommend that you include **EVENT_LABEL** metadata and provide labels for your login events. It's fine if the labels are incomplete or sporadic. If you provide labels, Amazon Fraud Detector will use them to automatically calculate an Account Takeover Discovery Rate and display it in model performance chart and table.

Event variables

For Accounts Takeover Insights model, there are both required (mandatory) variables that you must provide and optional variables. When you create your variables, make sure to assign the variable to the right variable type. As part of the model training process, Amazon Fraud Detector uses the variable type that's associated with the variable to perform variable enrichment and feature engineering.

Note

Event variable names must be in lowercase letters. They're case sensitive.

Mandatory variables

The following variables are required for training an Accounts Takeover Insights model.

| Category | Variable type | Description |
|------------|---------------|--|
| IP address | IP_ADDRESS | The IP address used in the login event |

| Category | Variable type | Description |
|--------------------|---------------|---|
| Browser and device | USERAGENT | The browser, device, and OS used in the login event |
| Valid credentials | VALIDCRED | Indicates if the credentials that were used for login are valid |

Optional variables

The following variables are optional for training an Accounts Takeover Insights model.

| Category | Type | Description |
|--------------------|-----------------|---|
| Browser and device | FINGERPRINT | The unique identifier for a browser or device fingerprint |
| Session Id | SESSION_ID | The identifier for an authentication session |
| Label | EVENT_LABEL | A label that classifies the event as fraudulent or legitimate. You can use any labels, such as "fraud", "legit", "1", or "0". |
| Timestamp | LABEL_TIMESTAMP | The timestamp when the label was last updated. This is required if EVENT_LABEL is provided. |

Note

- You can provide any variable names for both mandatory variables optional variables. It's important that each mandatory and optional variable is assigned to the right variable type.

- You can provide additional variables. However, Amazon Fraud Detector won't include these variables for training an Accounts Takeover Insights model.

Selecting data

Gathering data is an important step to creating your Account Takeover Insights model. As you start to gather your login data, consider the following requirements and recommendations:

Required

- Provide at least 1,500 user account examples, each with at least two associated login events.
- Your dataset must cover at least 30 days of login events. You can later specify the specific time range of the events to use to train the model.

Recommended

- Your dataset includes examples of unsuccessful login events. You can optionally label these unsuccessful logins as "fraudulent" or "legitimate."
- Prepare historic data with login events spanning more than six months and include 100K entities.

If you don't have a dataset that already meets the minimum requirements, consider streaming event data to Amazon Fraud Detector by calling the [SendEvent](#) API operation.

Validating data

Before creating your Account Takeover Insights model, Amazon Fraud Detector checks if the metadata and variables you included in your dataset for training the model meet size and format requirements. For more information, see [Dataset validation](#). It also checks for other requirements. If the dataset doesn't pass validation, model isn't created. For the model to be successfully created, make sure to fix the data that didn't pass the validation before you train again.

Common dataset errors

When validating a dataset for training an Account Takeover Insights model, Amazon Fraud Detector scans for these and other issues and throws an error if it encounters one or more of the issues.

- CSV file isn't in the UTF-8 format.

- The CSV file header doesn't contain at least one of the following metadata: `EVENT_ID`, `ENTITY_ID`, or `EVENT_TIMESTAMP`.
- The CSV file header doesn't contain at least one variable of the following variable types: `IP_ADDRESS`, `USERAGENT`, or `VALIDCRED`.
- There's more than one variable that's associated with the same variable type.
- More than 0.1% of values in `EVENT_TIMESTAMP` contains nulls or values other than the supported date and timestamp formats.
- The number of days between the first and last event is fewer than 30 days.
- More than 10% of variables of the `IP_ADDRESS` variable type are either invalid or null.
- More than 50% of variables of the `USERAGENT` variable type contain nulls.
- All of the variables of the `VALIDCRED` variable type are set to `false`.

Build a model

Amazon Fraud Detector models learn to detect fraud for a specific event type. In Amazon Fraud Detector, you first create a model, which acts as a container for your model versions. Each time you train a model, a new version is created. For details on how to create and train a model using the AWS Console see [Step 3: Create model](#).

Each model has a corresponding model score variable. Amazon Fraud Detector creates this variable on your behalf when you create a model. You can use this variable in your rule expressions to interpret your model scores during a fraud evaluation.

Train and deploy a model using the AWS SDK for Python (Boto3)

A model version is created by calling the `CreateModel` and `CreateModelVersion` operations. `CreateModel` initiates the model, which acts as a container for your model versions. `CreateModelVersion` starts the training process, which results in a specific version of the model. A new version of the solution is created each time you call `CreateModelVersion`.

The following example shows a sample request for the `CreateModel` API. This example creates *Online Fraud Insights* model type and assumes you have created an event type `sample_registration`. For additional details about creating an event type, see [Create an event type](#).

```
import boto3
fraudDetector = boto3.client('frauddetector')
```

```
fraudDetector.create_model (
    modelId = 'sample_fraud_detection_model',
    eventTypeName = 'sample_registration',
    modelType = 'ONLINE_FRAUD_INSIGHTS')
```

Train your first version using the [CreateModelVersion](#) API. For the `TrainingDataSource` and `ExternalEventsDetail` specify the source and Amazon S3 location of the training data set. For the `TrainingDataSchema` specify how Amazon Fraud Detector should interpret the training data, specifically which event variables to include and how to classify the event labels. By default, Amazon Fraud Detector ignores the unlabeled events. This example code uses `AUTO` for `unlabeledEventsTreatment` to specify that Amazon Fraud Detector decides how to use the unlabeled events.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.create_model_version (
    modelId = 'sample_fraud_detection_model',
    modelType = 'ONLINE_FRAUD_INSIGHTS',
    trainingDataSource = 'EXTERNAL_EVENTS',
    trainingDataSchema = {
        'modelVariables' : ['ip_address', 'email_address'],
        'labelSchema' : {
            'labelMapper' : {
                'FRAUD' : ['fraud'],
                'LEGIT' : ['legit']
            }
        }
        unlabeledEventsTreatment = 'AUTO'
    }
},
externalEventsDetail = {
    'dataLocation' : 's3://bucket/file.csv',
    'dataAccessRoleArn' : 'role_arn'
}
)
```

A successful request will result in a new model version with status `TRAINING_IN_PROGRESS`. At any point during the training, you can cancel the training by calling `UpdateModelVersionStatus` and updating the status to `TRAINING_CANCELLED`. Once training is complete, the model version status will update to `TRAINING_COMPLETE`. You can review model performance using the Amazon

Fraud Detector console or by calling `DescribeModelVersions`. For more information on how to interpret model scores and performance, see [Model scores](#) and [Model performance metrics](#).

After reviewing the model performance, activate the model to make it available to use by Detectors in real-time fraud predictions. Amazon Fraud Detector will deploy the model in multiple availability zones for redundancy with auto-scaling turned on to ensure the model scales with the number of fraud predictions you are making. To activate the model, call the `UpdateModelVersionStatus` API and update the status to `ACTIVE`.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.update_model_version_status (
    modelId = 'sample_fraud_detection_model',
    modelType = 'ONLINE_FRAUD_INSIGHTS',
    modelVersionNumber = '1.00',
    status = 'ACTIVE'
)
```

Model scores

Amazon Fraud Detector generates model scores differently for different model types.

For **Account Takeover Insights (ATI)** models, Amazon Fraud Detector uses only aggregated value (a value calculated by combining a set of raw variables) to generate the model score. A score of -1 is generated for the first event of a new entity, indicating an *unknown risk*. This is because for a new entity, the values used for calculating the aggregate will be zero or null. Account Takeover Insights (ATI) model generates model scores between 0 and 1000 for all subsequent events for the same entity and for existing entities, where 0 indicates *low fraud risk* and 1000 indicates *high fraud risk*. For ATI models, the model scores are directly related to the challenge rate (CR). For example, a score of 500 corresponds to an estimated 5% challenge rate whereas a score of 900 corresponds to an estimated 0.1% challenge rate.

For **Online Fraud Insights (OFI)** and **Transaction Fraud Insights (TFI)** models, Amazon Fraud Detector uses both aggregated value (a value calculated by combining a set of raw variables) and raw value (the value provided for the variable) to generate the model scores. The model scores can be between 0 and 1000, where 0 indicates *low fraud risk* and 1000 indicates *high fraud risk*. For the OFI and TFI models, the model scores are directly related to the false positive rate (FPR). For example, a score of 600 corresponds to an estimated 10% false positive rate whereas a score

of 900 corresponds to an estimated 2% false positive rate. The following table provides details of how certain model scores correlate to estimated false positive rates.

| Model score | Estimated FPR |
|-------------|---------------|
| 975 | 0.50% |
| 950 | 1% |
| 900 | 2% |
| 860 | 3% |
| 775 | 5% |
| 700 | 7% |
| 600 | 10% |

Model performance metrics

After model training is complete, Amazon Fraud Detector validates model performance using 15% of your data that was not used to train the model. You can expect your trained Amazon Fraud Detector model to have real-world fraud detection performance that is similar to the validation performance metrics.

As a business, you must balance between detecting more fraud, and adding more friction to legitimate customers. To assist in choosing the right balance, Amazon Fraud Detector provides the following tools to assess model performance:

- **Score distribution chart** – A histogram of model score distributions assumes an example population of 100,000 events. The left Y axis represents the legitimate events and the right Y axis represents the fraud events. You can select a specific model threshold by clicking on the chart area. This will update the corresponding views in the confusion matrix and ROC chart.
- **Confusion matrix** – Summarizes the model accuracy for a given score threshold by comparing model predictions versus actual results. Amazon Fraud Detector assumes an example population of 100,000 events. The distribution of fraud and legitimate events simulates the fraud rate in your businesses.

- **True positives** – The model predicts fraud and the event is actually fraud.
- **False positives** – The model predicts fraud but the event is actually legitimate.
- **True negatives** – The model predicts legitimate and the event is actually legitimate.
- **False negatives** – The model predicts legitimate but the event is actually fraud.
- **True positive rate (TPR)** – Percentage of total fraud the model detects. Also known as capture rate.
- **False positive rate (FPR)** – Percentage of total legitimate events that are incorrectly predicted as fraud.
- **Receiver Operator Curve (ROC)** – Plots the true positive rate as a function of false positive rate over all possible model score thresholds. View this chart by choosing **Advanced Metrics**.
- **Area under the curve (AUC)** – Summarizes TPR and FPR across all possible model score thresholds. A model with no predictive power has an AUC of 0.5, whereas a perfect model has a score of 1.0.
- **Uncertainty range** – It shows the range of AUC expected from the model. Larger range (difference in upper and lower bound of AUC > 0.1) means higher model uncertainty. If the uncertainty range is large (>0.1), consider providing more labeled events and retrain the model.

To use the model performance metrics


1. Start with the **Score distribution** chart to review the distribution of model scores for your fraud and legitimate events. Ideally, you will have a clear separation between the fraud and legitimate events. This indicates the model can accurately identify which events are fraudulent and which are legitimate. Select a model threshold by clicking on the chart area. You can see how adjusting the model score threshold impacts your true positive and false positive rates.

Note

The score distribution chart plots the fraud and legitimate events on two different Y axis. The left Y axis represents the legitimate events and the right Y axis represents the fraud events.

2. Review the **Confusion matrix**. Depending on your selected model score threshold, you can see the simulated impact based on a sample of 100,000 events. The distribution of fraud and legitimate events simulates the fraud rate in your businesses. Use this information to find the right balance between true positive rate and false positive rate.

3. For additional details, choose **Advanced Metrics**. Use the ROC chart to understand the relationship between true positive rate and false positive rate for any model score threshold. The ROC curve can help you fine-tune the tradeoff between true positive rate and false positive rate.

 **Note**

You can also review metrics in table form by choosing **Table**.

The table view also shows the metric **Precision**. **Precision** is the percentage of fraud events correctly predicted as fraudulent as compared to all events predicted as fraudulent.

4. Use the performance metrics to determine the optimal model thresholds for your businesses based on your goals and fraud-detection use case. For example, if you plan to use the model to classify new account registrations as either high, medium, or low risk, you need to identify two threshold scores so you can draft three rule conditions as follows:
 - Scores $> X$ are high risk
 - Scores $< X$ but $> Y$ are medium risk
 - Scores $< Y$ are low risk

Model variable importance

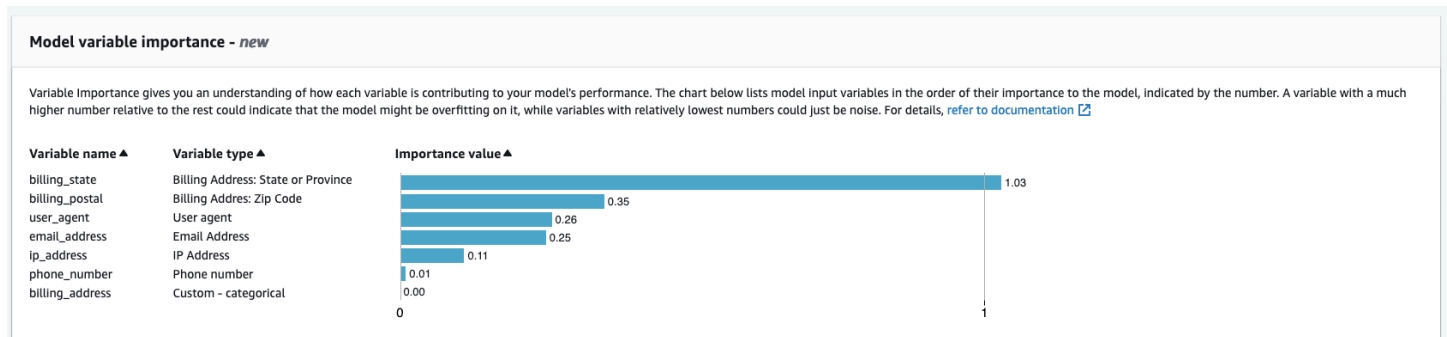
Model variable importance is a feature of Amazon Fraud Detector that ranks model variables within a model version. Each model variable is provided a value based on its relative importance to the overall performance of your model. The model variable with the highest value is more important to the model than the other model variables in the dataset for that model version, and is listed at the top by default. Likewise, the model variable with the lowest value is listed at the bottom by default and is least important compared to the other model variables. Using model variable importance values, you can gain insight into what inputs are driving your model's performance.

You can view model variable importance values for your trained model version in the Amazon Fraud Detector console or by using the [DescribeModelVersion](#) API.

Model variable importance provides the following set of values for each [Variable](#) used to train the [Model Version](#).

- **Variable Type:** Type of variable (for example, IP address or Email). For more information, see [Variable types](#). For Account Takeover Insights (ATI) models, Amazon Fraud Detector provides variable importance value for both raw and aggregate variable type. Raw variable types are assigned to the variables that you provide. Aggregate variable type is assigned to a set of raw variables that Amazon Fraud Detector has combined to calculate an aggregated importance value.
- **Variable Name:** Name of the event variable that was used to train the model version (for example, `ip_address`, `email_address`, `are_credentials_valid`). For aggregated variable type, the names of all variables that were used to calculate the aggregated variable importance value are listed.
- **Variable Importance Value:** A number that represents the relative importance of the raw or aggregated variable to the model's performance. Typical range: 0–10

In the Amazon Fraud Detector console, the model variable importance values are displayed as follows for either an Online Fraud Insights (OFI) or an Transaction Fraud Insights (TFI) model. An Account Takeover Insight (ATI) model will provide aggregated variable importance values in addition to the raw variable's importance values. The visual chart makes it easy to see the relative importance between variables with the vertical dotted line providing reference to the importance value of the highest ranked variable.



Amazon Fraud Detector generates variable importance values for every Fraud Detector model version at no additional cost.

⚠ Important

Model versions that were created before *July 9, 2021* do not have variable importance values. You must train a new version of your model to generate the model variable importance values.

Using model variable importance values

You can use model variable importance values to gain insight into what is driving performance of your model up or down and which of variables contribute the most. And then tweak your model to improve overall performance.

More specifically, to improve your model performance, examine the variable importance values against your domain knowledge and debug issues in the training data. For example, if Account Id was used as an input to the model and it is listed at the top, take a look at its variable importance value. If the variable importance value is significantly higher than the rest of the values, then your model might be overfitting on a specific fraud pattern (for example, all the fraud events are from the same Account Id). However, it might also be the case that there is a label leakage if the variable depends on the fraud labels. Depending on the outcome of your analysis based on your domain knowledge, you might want to remove the variable and train with a more diverse dataset, or keep the model as it is.

Similarly, take a look at the variables ranked last. If the variable importance value is significantly lower than the rest of the values, then this model variable might not have any importance in training your model. You could consider removing the variable to train a simpler model version. If your model has few variables, such as only two variables, Amazon Fraud Detector still provides the variable importance values and rank the variables. However, the insights in this case will be limited.

Important

1. If you notice variables missing in the **Model variable importance** chart, it might be due to one of the following reasons. Consider modifying the variable in your dataset and retrain your model.
 - The count of unique values for the variable in the training dataset is lower than 100.
 - Greater than 0.9 of values for the variable are missing from the training data-set.
2. You need to train a new model version every time that you want to adjust your model's input variables.

Evaluating model variable importance values

We recommend that you consider the following when you evaluate model variable importance values:

- Variable importance values must always be evaluated in combination with the domain knowledge.
- Examine variable importance value of a variable relative to the variable importance value of the other variables within the model version. Do not consider variable importance value for a single variable independently.
- Compare variable importance values of the variables within the same model version. Do not compare variable importance values of the same variables across model versions because the variable importance value of a variable in a model version might differ from the value of the same variable in a different model version. If you use the same variables and dataset to train different model versions, this does not necessarily generate the same variable importance values.

Viewing model variable importance ranking

After model training is complete, you can view model variable importance ranking of your trained model version in the Amazon Fraud Detector console or by using the [DescribeModelVersion](#) API.

To view the model variable importance ranking using console,

1. Open the AWS Console and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Models**.
3. Choose your model and then your model version.
4. Make sure that the **Overview** tab is selected.
5. Scroll down to view the **Model variable importance** pane.

Understanding how the model variable importance value is calculated

Upon completion of each model version training, Amazon Fraud Detector automatically generates model variable importance values and model's performance metrics. For this, Amazon Fraud Detector uses SHapley Additive exPlanations ([SHAP](#)). SHAP is essentially the average expected contribution of a model variable after all possible combinations of all model variables have been considered.

SHAP first assigns contribution of each model variable for prediction of an event. Then, it aggregates these predictions to create a ranking of the variables at the model level. To assign contributions of each model variable for a prediction, SHAP considers differences in model outputs

among all possible variable combinations. By including all possibilities of including or removing specific set of variables to generate a model output, SHAP can accurately access the importance of each model variable. This is particularly important when the model variables are highly correlated with one another.

ML models, in most cases, do not allow you to remove variables. You can instead replace a removed or missing variable in the model with the corresponding variable values from one or more baselines (for example, non-fraud events). Choosing proper baseline instances can be difficult, but Amazon Fraud Detector makes this easy by setting this baseline as the population average for you.

Import a SageMaker model

You can optionally import SageMaker-hosted models to Amazon Fraud Detector. Similar to models, SageMaker models can be added to detectors and generate fraud predictions using the `GetEventPrediction` API. As part of the `GetEventPrediction` request, Amazon Fraud Detector will invoke your SageMaker endpoint and pass the results to your rules.

You can configure Amazon Fraud Detector to use the event variables sent as part of the `GetEventPrediction` request. If you choose to use event variables, you must provide an input template. Amazon Fraud Detector will use this template to transform your event variables into the required input payload to invoke the SageMaker endpoint. Alternatively, you can configure your SageMaker model to use a `byteBuffer` that is sent as part of the `GetEventPrediction` request.

Amazon Fraud Detector supports importing SageMaker algorithms that use JSON or CSV input formats and JSON or CSV output formats. Examples of supported SageMaker algorithms include XGBoost, Linear Learner, and Random Cut Forest.

Import a SageMaker model using the AWS SDK for Python (Boto3)

To import a SageMaker model, use the `PutExternalModel` API. The following example assumes the SageMaker endpoint `sagemaker-transaction-model` has been deployed, is `InService` status, and uses the XGBoost algorithm.

The input configuration specifies that will use the event variables to construct the model input (`useEventVariables` is set to `TRUE`). The input format is `TEXT_CSV`, given XGBoost requires a CSV input. The `csvInputTemplate` specifies how to construct the CSV input from the variables sent as part of the `GetEventPrediction` request. This example assumes you have created the variables `order_amt`, `prev_amt`, `hist_amt` and `payment_type`.

The output configuration specifies the response format of the SageMaker model, and maps the appropriate CSV index to the Amazon Fraud Detector variable `sagemaker_output_score`. Once configured, you can use the output variable in rules.

Note

The output from a SageMaker model must be mapped to a variable with source `EXTERNAL_MODEL_SCORE`. You cannot create these variables in the console using **Variables**. You must instead create them when you configure your model import.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_external_model (
    modelSource = 'SAGEMAKER',
    modelEndpoint = 'sagemaker-transaction-model',
    invokeModelEndpointRoleArn = 'your_SagemakerExecutionRole_arn',
    inputConfiguration = {
        'useEventVariables' : True,
        'eventName' : 'sample_transaction',
        'format' : 'TEXT_CSV',
        'csvInputTemplate' : '{{order_amt}}, {{prev_amt}}, {{hist_amt}}, {{payment_type}}'
    },
    outputConfiguration = {
        'format' : 'TEXT_CSV',
        'csvIndexToVariableMap' : {
            '0' : 'sagemaker_output_score'
        }
    },
    modelEndpointStatus = 'ASSOCIATED'
)
```

Delete a model or model version

You can delete models and model versions in Amazon Fraud Detector, provided that they are not associated with a detector version. When you delete a model, Amazon Fraud Detector permanently deletes that model and the data is no longer stored in Amazon Fraud Detector.

You can also remove Amazon SageMaker models if they are not associated with a detector version. Removing a SageMaker model disconnects it from Amazon Fraud Detector, but the model remains available in SageMaker.

To delete a model version

You can only delete model versions that are in the Ready to deploy status. To change a model version from ACTIVE to Ready to deploy status, undeploy the model version.

1. Sign in to the AWS Management Console and open the Amazon Fraud Detector console at <https://console.aws.amazon.com/frauddetector>.
2. In the left navigation pane of the Amazon Fraud Detector console, choose **Models**.
3. Choose the model that contains the model version you want to delete.
4. Choose the model version that you want to delete.
5. Choose **Actions**, and then choose **Delete**.
6. Enter the model version name, and then choose **Delete model version**.

To undeploy a model version

You can't undeploy a model version that is in use by any detector version (ACTIVE, INACTIVE, DRAFT). Therefore, to undeploy a model version that is in use by a detector version, first remove the model version from the detector version.

1. In the left navigation pane of the Amazon Fraud Detector console, choose **Models**.
2. Choose the model that contains the model version you want to undeploy.
3. Choose the model version that you want to delete.
4. Choose **Actions**, and then choose **Undeploy model version**.

To delete a model

Before deleting a model, you must first delete all model versions and are associated with the model.

1. In the left navigation pane of the Amazon Fraud Detector console, choose **Models**.
2. Choose the model that you want to delete.
3. Choose **Actions**, and then choose **Delete**.

4. Enter the model name, and then choose **Delete model**.

To remove an Amazon SageMaker model

1. In the left navigation pane of the Amazon Fraud Detector console, choose **Models**.
2. Choose the SageMaker model that you want to remove.
3. Choose **Actions**, and then choose **Remove model**.
4. Enter the model name and then choose **Remove SageMaker model**.

Detector

A detector is a container that contains fraud detection logic, such as the models and rules, for one specific business event that you want to evaluate for fraud. You first create a detector by specifying the event that you have already defined and optionally add a model version that is already created and trained by Amazon Fraud Detector for the event.

You then add rules and rule execution order to a detector to create a version of the detector. A detector version defines the rules and optionally a model that will be run as part of the request for generating fraud predictions. You can add any of the rules defined within a detector to the detector version. You can also add any model trained on the evaluated event type to the detector version. A detector can have multiple versions, with each version having different rules and rule execution order to meet multiple use cases.

Each detector version must have a status of DRAFT, ACTIVE, or INACTIVE. Only one detector version can be in ACTIVE status at a time. Amazon Fraud Detector uses the detector version with ACTIVE status to generate fraud predictions.

Create a detector

You create a detector by specifying the event type that you have already defined. You can optionally add a model that is already trained and deployed by Amazon Fraud Detector. If you add a model, you can use the model score generated by Amazon Fraud Detector in your rule expression when creating a rule (for example, `$model score < 90`).

You can create a detector in the Amazon Fraud Detector console, using the [PutDetector](#) API, using the [put-detector](#) command, or using the AWS SDK. If you are using API, command, or SDK for creating a detector, after you've created the detector follow instructions to [Create a detector version](#).

Create a detector in the Amazon Fraud Detector console

This example assumes that you've created an event type and also have created and deployed a model version you want to use for fraud prediction.

Step 1: Build detector

1. In the left navigation pane of the Amazon Fraud Detector console, choose **Detectors**.

2. Choose **Create detector**.
3. In the **Define detector details** page, enter `sample_detector` for detector name. Optionally, enter a description for the detector, such as `my sample fraud detector`.
4. For **Event Type**, select the event type you have created for fraud prediction.
5. Choose **Next**.

Step 2: Add a deployed model version

1. Note that this is an optional step. You do not need to add a model to your detector. To skip this step, choose **Next**.
2. In the **Add model - optional**, choose **Add Model**.
3. In the **Add model** page, for **Select model**, choose the Amazon Fraud Detector model name that you deployed earlier. For **Select version**, choose the model version of the deployed model.
4. Choose **Add model**.
5. Choose **Next**.

Step 3: Add rules

A rule is a condition that tells Amazon Fraud Detector how to interpret variable values when evaluating for fraud prediction. This example will create three rules using the model scores as variable values: `high_fraud_risk`, `medium_fraud_risk`, and `low_fraud_risk`. To create your own rules, rule expressions, rule execution order, and outcomes, use values that are appropriate for your model and your use case.

1. In the **Add rules** page, under **Define a rule**, enter `high_fraud_risk` for the rule name and under **Description - optional**, enter **This rule captures events with a high ML model score** as the description for the rule.
2. In **Expression**, enter the following rule expression using the Amazon Fraud Detector simplified rule expression language:

```
$sample_fraud_detection_model_insightscore > 900
```
3. In **Outcomes**, choose **Create a new outcome**. An outcome is the result from a fraud prediction and is returned if the rule matches during an evaluation.

4. In **Create a new outcome**, enter `verify_customer` as the outcome name. Optionally, enter a description.
5. Choose **Save outcome**.
6. Choose **Add rule** to run the rule validation checker and save the rule. After it's created, Amazon Fraud Detector makes the rule available for use in your detector.
7. Choose **Add another rule**, and then choose the **Create rule** tab.
8. Repeat this process twice more to create your `medium_fraud_risk` and `low_fraud_risk` rules using the following rule details:

- `medium_fraud_risk`

Rule name: `medium_fraud_risk`

Outcome: `review`

Expression:

```
$sample_fraud_detection_model_insightscore <= 900 and
```

```
$sample_fraud_detection_model_insightscore > 700
```

- `low_fraud_risk`

Rule name: `low_fraud_risk`

Outcome: `approve`

Expression:

```
$sample_fraud_detection_model_insightscore <= 700
```

9. After you have created all the rules for your use case, choose **Next**.

For more information about creating and writing rules, see [Rules](#) and [Rule language reference](#).

Step 4: Configure rule execution and rule order

The rule execution mode for the rules that are included in the detector determines if all the rules you define are evaluated, or if rule evaluation stops at the first matched rule. And the rule order determines the order that you want the rule to be run in.

The default rule execution mode is `FIRST_MATCHED`.

First matched

First matched rule execution mode returns the outcomes for the first matching rule based on defined rule order. If you specify `FIRST_MATCHED`, Amazon Fraud Detector evaluates rules sequentially, first to last, stopping at the first matched rule. Amazon Fraud Detector then provides the outcomes for that single rule.

The order that you run rules in can affect the resulting fraud prediction outcome. After you have created your rules, re-order the rules to run them in the desired order by following these steps:

If your `high_fraud_risk` rule isn't already on the top of your rule list, choose **Order**, and then choose **1**. This moves `high_fraud_risk` to the first position.

Repeat this process so that your `medium_fraud_risk` rule is in the second position and your `low_fraud_risk` rule is in the third position.

All matched

All matched rule execution mode returns outcomes for all matched rules, regardless of rule order. If you specify `ALL_MATCHED`, Amazon Fraud Detector evaluates all rules and returns the outcomes for all matched rules.

Select `FIRST_MATCHED` for this tutorial and then choose **Next**.

Step 5: Review and create detector version

A detector version defines the specific models and rules that are used for generating fraud predictions.

1. In the **Review and create** page, review the detector details, models, and rules that you configured. If you need to make any changes, choose **Edit** next to the corresponding section.
2. Choose **Create detector**. After it's created, the first version of your detector appears in the Detector versions table with `Draft` status.

You use the **Draft** version to test your Detector.

Create a detector using the AWS SDK for Python (Boto3)

The following example shows a sample request for the `PutDetector` API. A detector acts as a container for your detector versions. The `PutDetector` API specifies what event type the detector will evaluate. The following example assumes you have created an event type `sample_registration`.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_detector (
    detectorId = 'sample_detector',
    eventName = 'sample_registration'
)
```

Create a detector version

A detector version defines the rules, rule execution order, and optionally a model version, that will be used as part of the request for generating fraud predictions. You can add any of the rules defined within a detector to the detector version. You can also add any model trained on the evaluated event type.

Each detector version has a status of `DRAFT`, `ACTIVE`, or `INACTIVE`. Only one detector version can be in `ACTIVE` status at a time. During the `GetEventPrediction` request, Amazon Fraud Detector will use the `ACTIVE` detector if no `DetectorVersion` is specified.

Rule execution mode

Amazon Fraud Detector supports two different rule execution modes: `FIRST_MATCHED` and `ALL_MATCHED`.

- If the rule execution mode is `FIRST_MATCHED`, Amazon Fraud Detector evaluates rules sequentially, first to last, stopping at the first matched rule. Amazon Fraud Detector then provides the outcomes for that single rule. If a rule evaluates to false (not matched), the next rule in the list is evaluated.
- If the rule execution mode is `ALL_MATCHED`, then all rules in an evaluation are executed in parallel, regardless of their order. Amazon Fraud Detector executes all rules and returns the defined outcomes for every matched rule.

Create a detector version using the AWS SDK for Python (Boto3)

The following example shows a sample request for the `CreateDetectorVersion` API. The rule execution mode is set to `FIRST_MATCHED`, therefore Amazon Fraud Detector will evaluate rules sequentially, first to last, stopping at the first matched rule. Amazon Fraud Detector then provides the outcomes for that single rule during the `GetEventPrediction` response.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.create_detector_version(
    detectorId = 'sample_detector',
    rules = [{
        'detectorId' : 'sample_detector',
        'ruleId' : 'high_fraud_risk',
        'ruleVersion' : '1'
    },
    {
        'detectorId' : 'sample_detector',
        'ruleId' : 'medium_fraud_risk',
        'ruleVersion' : '1'
    },
    {
        'detectorId' : 'sample_detector',
        'ruleId' : 'low_fraud_risk',
        'ruleVersion' : '1'
    }
    ],
    modelVersions = [{
        'modelId' : 'sample_fraud_detection_model',
        'modelType': 'ONLINE_FRAUD_INSIGHTS',
        'modelVersionNumber' : '1.00'
    }],
    ruleExecutionMode = 'FIRST_MATCHED'
)
```

To update the status of a detector version, use the `UpdateDetectorVersionStatus` API. The following example updates the detector version status from `DRAFT` to `ACTIVE`. During a `GetEventPrediction` request, if a detector ID is not specified, Amazon Fraud Detector will use the `ACTIVE` version of the detector.

```
import boto3
```



```
fraudDetector = boto3.client('frauddetector')

fraudDetector.update_detector_version_status(
    detectorId = 'sample_detector',
    detectorVersionId = '1',
    status = 'ACTIVE'
)
```

Delete a detector, detector version, or rule version

Before deleting a detector in Amazon Fraud Detector, you must first delete all detector versions and rule versions that are associated with the detector.

When you delete a detector, detector version, or rule version, Amazon Fraud Detector permanently deletes that resource and the data is no longer stored in Amazon Fraud Detector.

To delete a detector version

You can only delete detector versions that are in DRAFT or INACTIVE status.

1. Sign in to the AWS Management Console and open the Amazon Fraud Detector console at <https://console.aws.amazon.com/frauddetector>.
2. In the left navigation pane of the Amazon Fraud Detector console, choose **Detectors**.
3. Choose the detector that contains the detector version you want to delete.
4. Choose the detector version that you want to delete.
5. Choose **Actions**, and then choose **Delete**.
6. Enter **delete**, and then choose **Delete detector**.

To delete a rule version

You can delete a rule version only if it is not used by any ACTIVE or INACTIVE detector versions. If necessary, before deleting a rule version, first move the ACTIVE detector version to INACTIVE, then delete the INACTIVE detector version.

1. In the left navigation pane of the Amazon Fraud Detector console, choose **Detectors**.
2. Choose the detector that contains the rule version you want to delete.
3. Choose the **Associated rules** tab, and choose the rule that you want to delete.

4. Choose the rule version that you want to delete.
5. Choose **Actions**, and then choose **Delete rule version**.
6. Enter **delete**, and then choose **Delete version**.

To delete a detector

Before deleting a detector, you must first delete all detector versions and rule versions that are associated with the detector.

1. In the left navigation pane of the Amazon Fraud Detector console, choose **Detectors**.
2. Choose the detector that you want to delete.
3. Choose **Actions**, and then choose **Delete detector**.
4. Enter **delete**, and then choose **Delete detector**.

Resources

Models, rules, and detectors use resources such as variables, outcomes, labels, lists, and entities to evaluate events for fraud risk. This section provides information about creating and managing the resources.

Topics

- [Variables](#)
- [Labels](#)
- [Rules](#)
- [Lists](#)
- [Outcomes](#)
- [Entity](#)
- [Manage Amazon Fraud Detector resources using AWS CloudFormation](#)

Variables

Variables represent data elements that you want to use in a fraud prediction. These variables can be taken from the event dataset that you prepared for training your model, from your Amazon Fraud Detector model's risk score outputs, or from Amazon SageMaker models. For more information about variables taken from the event dataset, see [Get event dataset requirements using the Data models explorer](#).

The variables you want to use in your fraud prediction must first be created and then added to the event when creating your event type. Each variable you create must be assigned a datatype, a default value, and optionally a variable type. Amazon Fraud Detector enriches some of the variables that you provide such as IP addresses, bank identification numbers (BINs), and phone numbers, to create additional inputs and boost performance for the models that use these variables.

Data types

Variables must have a data type for the data element that the variable represents and can optionally be assigned one of the predefined [Variable types](#). For variables that are assigned to a variable type, the data type is pre-selected. Possible data types include the following types :

| Data type | Description | Default value | Example values |
|-----------|--|---------------|----------------------|
| String | Any combination of letters, whole numbers, or both | <empty> | abc, 123, 1D3B |
| Integer | Positive or negative whole numbers | 0 | 1, -1 |
| Boolean | True or False | False | True, False |
| DateTime | Date and time specified in the ISO 8601 standard UTC format only | <empty> | 2019-11-30T13:01:01Z |
| Float | Numbers with decimal points | 0.0 | 4.01, 0.10 |

Default value

Variables must have a default value. When Amazon Fraud Detector generates fraud predictions, this default value is used to run a rule or model if Amazon Fraud Detector doesn't receive a value for a variable. Default values you provide must match the selected data type. In the AWS Console, Amazon Fraud Detector assigns the default value of `0` for integers, `false` for Booleans, `0.0` for floats, and `(empty)` for strings. You can set a custom default value for any of these data types.

Variable types

When you create a variable, you can optionally assign the variable to a variable type. Variable type represents the common data elements that are used to train models and to generate fraud predictions. Only variables with an associated variable type can be used for model training. As part of the model training process, Amazon Fraud Detector uses the variable type associated with the variable to perform variable enrichments, feature engineering, and risk scoring.

Amazon Fraud Detector has pre-defined the following variable types that can be used to assign to your variables.

| Category | Variable type | Description | Data type | Example |
|----------|---------------|--|-----------|---|
| Session | IPADDRESS | The IP address that's collected during the event | String | 192.0.2.0 Note: Amazon Fraud Detector enriches this data. For more information, see Geolocation enrichment |
| | USERAGENT | The user agent that's collected during the event | String | Mozilla 5.0 (Windows NT 10.0, Win64, x64;rv:68 .0) Gecko 20100101 |
| | FINGERPRINT | The unique identifier for a | String | sadfow987 u234 |

| Ca | Variable type | Description | Data type | Ex |
|----|-----------------------|---|-----------|----------------|
| | | device used for the event | | |
| | SESSION_ID | The session ID for the event's active session | String | sid123456789 |
| | ARE_CREDENTIALS_VALID | Indicates if the credentials used for event login are valid | Boolean | True |
| Us | EMAIL_ADDRESS | The email address that's collected during the event | String | abc@domain.com |

| Category | Variable type | Description | Data type | Example |
|----------|---------------|---|-----------|---|
| | PHONE_NUMBER | The phone number collected during the event | String | +1 555-0100 Note: Amazon Fraud Detector enriches this data. For more information, see Phone number enrichment |
| Billing | BILLING_NAME | The name that's associated with the billing address | String | John Doe |

| Category | Variable type | Description | Data type | Example |
|----------|--------------------|---|-----------|---|
| | BILLING_PHONE | The phone number that's associated with the billing address | String | +1 555-0100 Note: Amazon Fraud Detector enriches this data. For more information, see Phone number enrichment |
| | BILLING_ADDRESS_L1 | The first line of the billing address | String | Any street |
| | BILLING_ADDRESS_L2 | The second line of the billing address | String | Any unit 123 |
| | BILLING_CITY | The city that's in the billing address | String | Any City |

| Ca | Variable type | Description | Data type | Ex |
|----|-----------------|---|-----------|---|
| | BILLING_STATE | The state or province that's in the billing address | String | Any state or province |
| | BILLING_COUNTRY | The country that's in the billing address | String | Any country Note: Amazon Fraud Detector enriches this data. For more information, see Geolocation enrichment |

| Category | Variable type | Description | Data type | Example |
|----------|---------------|--|-----------|---|
| | BILLING_ZIP | The postal code that's in the billing address | String | 01234 Note: Amazon Fraud Detector enriches this data. For more information, see Geolocation enrichment |
| Shipping | SHIPPING_NAME | The name that's associated with the shipping address | String | John Doe |

| Ca | Variable type | Description | Data type | Ex |
|----|---------------------|--|-----------|---|
| | SHIPPING_PHONE | The phone number that's associated with the shipping address | Strir | +1 555-0100 Note: Amazon Fraud Detector enriches this data. For more information, see Phone number enrichment |
| | SHIPPING_ADDRESS_L1 | The first line of the shipping address | Strir | 123 Any Street |
| | SHIPPING_ADDRESS_L2 | The second line of the shipping address | Strir | Unit 123 |
| | SHIPPING_CITY | The city that's in the shipping address | Strir | Any City |

| Ca | Variable type | Description | Data type | Ex |
|----|------------------|--|-----------|--|
| | SHIPPING_STATE | The state or province that's in the shipping address | String | Any State |
| | SHIPPING_COUNTRY | The country that's in the shipping address | String | Any Country Note: Amazon Fraud Detector enriches this data. For more information, see Geolocation enrichment |

| Category | Variable type | Description | Data type | Example |
|----------|---------------|---|-----------|---|
| Payment | SHIPPING_ZIP | The postal code that's in the shipping address | String | 01234 Note: Amazon Fraud Detector enriches this data. For more information, see Geolocation enrichment |
| | ORDER_ID | The unique identifier for the transaction | String | LUX60 |
| | PRICE | The total order price | String | 560.00 |
| | CURRENCY_CODE | The ISO 4217 currency code | String | USD |
| | PAYMENT_TYPE | The payment method that's used for payment during the event | String | Credit card |

| Category | Variable type | Description | Data type | Example |
|----------|------------------|---|-----------|-----------------------------------|
| | AUTH_CODE | The alphanumeric code that's sent by a credit card issuer or issuing bank | String | 0000 |
| | AVS | The address verification system (AVS) response code from the card processor | String | Y |
| Product | PRODUCT_CATEGORY | The product category of order item | String | Kitchen |
| Customer | NUMERIC | Any variable that can be represented as a real number | Float | 1.224 |
| | CATEGORICAL | Any variable that describes categories, segments, or groups | String | Large |
| | FREE_FORM_TEXT | Any free form text that's captured as part of the event (for example, a customer review or comment) | String | Example of a free form text input |

Assigning variable to a variable type

If you are planning to use a variable for training your model, it is important that you choose a right variable type to assign to the variable. Incorrect variable type assignment can negatively impact your model performance. It can also become very difficult for you change the assignment later, especially if multiple models and events have used the variable.

You can assign your variable any one of the pre-defined variable types or one of the custom variable types – FREE_FORM_TEXT, CATEGORICAL, or NUMERIC.

Important notes for assigning variables to the right variable types

1. If the variable matches one of predefined variable types, use it. Make sure the variable type corresponds to the variable. For example, if you assign an *ip_address* variable to EMAIL_ADDRESS variable type, the *ip_address* variable will not get enriched with enrichments such as ASN, ISP, geo-location, and risk score. For more information, see [Variable enrichments](#).
2. If the variable doesn't match any of predefined variable types, follow the recommendations listed below to assign one of the custom variable types.
3. Assign CATEGORICAL variable type to variables that typically do not have natural ordering and can be put into categories, segments, or groups. The dataset you are using to train your model might have ID variables such as, *merchant_id*, *campaign_id*, or *policy_id*. These variables represent groups (for example, all customers with same *policy_id* represent a group). Variables that have the following data must be assigned CATEGORICAL variable type -
 - Variables that contain data such as *customer_ID*, *segment_ID*, *color_ID*, *department_code*, or *product_ID*.
 - Variables that contain Boolean data with true, false, or null values.
 - Variables that can be put into groups or categories such as company name, product category, card type, or referral medium.

Note

ENTITY_ID is a reserved variable type used by Amazon Fraud Detector to assign to ENTITY_ID variable. The ENTITY_ID variable is the ID of the entity initiating the action you want to evaluate. If you are creating a Transaction Fraud Insight (TFI) model type, you are required to provide ENTITY_ID variable. You will need to decide which variable in your data uniquely identifies the entity initiating the action and pass it on as ENTITY_ID variable. Assign CATEGORICAL variable type to all the other IDs in your dataset, if they

are present and if you are using them for model training. Examples of other IDs that are not an entity in your dataset can be *merchant_ID*, *policy_ID*, and *campaign_ID*.

4. Assign `FREE_FORM_TEXT` variable type to variables that contain a block of text. Examples of `FREE_FORM_TEXT` variable types are – *user reviews*, *comments*, *dates*, and *referral codes*. The `FREE_FORM_TEXT` data contains multiple tokens separated by a delimiter. The delimiters can be any character other than alpha-numeric and underscore symbol. For example, user reviews and comments can be separated by “space” delimiter, dates and referral codes can use hyphens as delimiters to separate out prefix, suffix, and intermediate parts. Amazon Fraud Detector uses the delimiters to extract data from `FREE_FORM_TEXT` variables.
5. Assign `NUMERIC` variable type to variables that are real numbers and have inherent ordering. Examples of `NUMERIC` variables include *day_of_the_week*, *incident_severity*, *customer_rating*. Although, you can assign `CATEGORICAL` variable type to these variables, we strongly recommend to assign all real number variables with inherent order to `NUMERIC` variable type.

Variable enrichments

Amazon Fraud Detector enriches some of the raw data elements that you provide such as IP addresses, bank identification numbers (BINs), and phone numbers, to create additional inputs and boost performance for the models that use these data elements. The enrichment helps identify potentially suspicious situations and help the models to capture more fraud.

Phone number enrichment

Amazon Fraud Detector enriches phone number data with additional information that relates to geolocation, the original carrier, and the validity of the phone number. Phone number enrichment is automatically enabled for all the models that are trained on or after *December 13, 2021* and have a phone number that includes a country code (+xxx). If you have included phone number variable in your model and have trained it before *December 13, 2021*, retrain your model so it can take advantage of this enrichment.

We highly recommend that you use the following format for phone number variables to ensure that your data is enriched successfully.

| Variable | Format | Description |
|------------------|------------------------------------|--|
| PHONE_NUM BER | The E.164 standard | Make sure to include country code (+xxx) |

| Variable | Format | Description |
|----------------------------------|------------------------------------|---|
| | | with the phone number. |
| BILLING_PHONE and SHIPPING_PHONE | The E.164 standard | Make sure to include country code (+xxx) with the phone number. |

Geolocation enrichment

Starting on *February 8, 2022* Amazon Fraud Detector calculates the physical distance between the IP_ADDRESS, BILLING_ZIP, and SHIPPING_ZIP values that you provide for an event. The calculated distances are used as inputs to your fraud detection model.

To enable geolocation enrichment, your event data must include at least two of the three variables: IP_ADDRESS, BILLING_ZIP, or SHIPPING_ZIP. In addition, each BILLING_ZIP and SHIPPING_ZIP value must have a valid BILLING_COUNTRY code and SHIPPING_COUNTRY code respectively. If you have a model that was trained before *February 8, 2022* and it includes these variables, you must retrain the model to enable the geolocation enrichment.

If Amazon Fraud Detector can't determine the location that's associated with the IP_ADDRESS, BILLING_ZIP, or SHIPPING_ZIP values for an event due to the data being not valid, a special placeholder value is used instead. For example, suppose that an event has valid IP_ADDRESS and BILLING_ZIP values, but SHIPPING_ZIP value isn't valid. In this case, enrichment is done only for IP_ADDRESS→BILLING_ZIP. The enrichment isn't done for IP_ADDRESS→SHIPPING_ZIP and BILLING_ZIP→SHIPPING_ZIP. Instead, the placeholder values are used in their place. No matter if geolocation enrichment is enabled for your model or not, the performance of your model doesn't change.

You can opt out of geolocation enrichment by mapping your BILLING_ZIP and SHIPPING_ZIP variables to the CUSTOM_CATEGORICAL variable type. Changing the variable type doesn't affect your model's performance.

Geolocation variable format

We highly recommend that you use the following format for geolocation variables to ensure that your location data is enriched successfully.

| Variable | Format | Description |
|--------------------------------------|--|---|
| IP_ADDRESS | IPv4 address | For example - 1.1.1.1 |
| BILLING_ZIP and SHIPPING_ZIP | The ISO 3166-1 alpha-2 postal code for the specified country | For more information, see the Country and territory codes section in this topic. |
| BILLING_COUNTRY and SHIPPING_COUNTRY | The ISO 3166-1 alpha-2 two-letter standard country code | For more information, see the Country and territory codes section in this topic. Amazon Fraud Detector tries to match all the common variations of a country's name to their ISO 3166-1 two-letter standard country code. However, we cannot guarantee they will be matched correctly. |

Country and territory codes

The following table provides a complete list of the countries and territories that are supported by Amazon Fraud Detector for geolocation enrichment. Each country and territory has an assigned country code (specifically, the ISO 3166-1 alpha-2 two-letter country code) and a postal code.

Postal code format

- 9 - number
- a - letter
- [X] - X is optional. For example, Guernsey "GY9[9] 9aa" means both "GY9 9aa" and "GY99 9aa" are valid. Use one format.

- [X/XX] - either X or XX can be used. For example, Bermuda "aa[aa/99]" means both "aa aa" and "aa 99" are valid. Use either one of these formats, but *do not* use both.
- Some countries have fixed prefix. For example, the postal code for Andorra is AD999. This means the country code must start with letters AD followed by three numbers.

| Code | Name | Postal code |
|------|----------------------|-------------|
| AD | Andorra | AD999 |
| AR | Netherlands Antilles | 9999 |
| AT | Austria | 9999 |
| AU | Australia | 9999 |
| AZ | Azerbaijan | AZ 9999 |
| BD | Bangladesh | 9999 |
| BE | Belgium | 9999 |
| BG | Bulgaria | 9999 |
| BM | Bermuda | aa[aa/99] |
| BY | Belarus | 999999 |
| CA | Canada | a9a 9a9 |
| CH | Switzerland | 9999 |
| CL | Chile | 9999999 |
| CO | Colombia | 999999 |
| CR | Costa Rica | 99999 |
| CY | Cyprus | 9999 |
| CZ | Czechia | 999 99 |

| Code | Name | Postal code |
|------|------------------------------------|-------------------------|
| DE | Germany | 99999 |
| DK | Denmark | 9999 |
| DO | Dominican Republic | 99999 |
| DZ | Algeria | 99999 |
| EE | Estonia | 99999 |
| ES | Spain | 99999 |
| FI | Finland | 99999 |
| FM | Federated States of Micronesi a | 99999 |
| FO | Faroe Islands | 999 |
| FR | France | 99999 |
| GB | United Kingdom | a[a]9[a/9] 9aa |
| GG | Guernsey | GY9[9] 9aa |
| GL | Greenland | 9999 |
| GP | Guadeloupe | 99999 |
| GT | Guatemala | 99999 |
| GU | Guam | 99999 |
| HR | Croatia | 99999 |
| HU | Hungary | 9999 |
| IE | Ireland | a99[a/9][a/9][a/9][a/9] |
| IM | Isle of Man | IM9[9]9aa |

| Code | Name | Postal code |
|------|-----------------------|-------------|
| IN | India | 999999 |
| IS | Iceland | 999 |
| IT | Italy | 99999 |
| JE | Jersey | JE9[9]9aa |
| JP | Japan | 999-9999 |
| KR | Republic of Korea | 99999 |
| LI | Liechtenstein | 9999 |
| LK | Sri Lanka | 99999 |
| LT | Lithuania | 99999 |
| LU | Luxembourg | L-9999 |
| LV | Latvia | LV-9999 |
| MC | Monaco | 99999 |
| MD | Republic of Moldova | 9999 |
| MH | Marshall Islands | 99999 |
| MK | North Macedonia | 9999 |
| MP | North Mariana Islands | 99999 |
| MQ | Martinique | 99999 |
| MT | Malta | aaa 9999 |
| MX | Mexico | 99999 |
| MY | Malaysia | 99999 |

| Code | Name | Postal code |
|------|--------------------|-------------|
| NL | Netherlands | 9999 aa |
| NO | Norway | 9999 |
| NZ | New Zealand | 9999 |
| PH | Philippines | 9999 |
| PK | Pakistan | 99999 |
| PL | Poland | 99-999 |
| PR | Puerto Rico | 99999 |
| PT | Portugal | 9999-999 |
| PW | Palau | 99999 |
| RE | Reunion | 99999 |
| RO | Romania | 999999 |
| RU | Russian Federation | 999999 |
| SE | Sweden | 999 99 |
| SG | Singapore | 999999 |
| SI | Slovenia | 9999 |
| SK | Slovakia | 999 99 |
| SM | San Marino | 99999 |
| TH | Thailand | 99999 |
| TR | Turkey | 99999 |
| UA | Ukraine | 99999 |

| Code | Name | Postal code |
|------|--------------------|-------------|
| US | United States | 99999 |
| UY | Uruguay | 99999 |
| VI | Virgin Islands, US | 99999 |
| WF | Wallis and Futuna | 99999 |
| YT | Mayotte | 99999 |
| ZA | South Africa | 9999 |

Useragent enrichment

If you create the Account Takeover Insights (ATI) model, you must provide a variable of the useragent variable type in your dataset. This variable contains the browser, device, and OS data of a login event. Amazon Fraud Detector enriches the useragent data with additional information such as `user_agent_family`, `OS_family`, and `device_family`.

Create a variable

You can create variables in the Amazon Fraud Detector console, using the [create-variable](#) command, using the [CreateVariable](#), or using the AWS SDK for Python (Boto3)

Create a variable using the Amazon Fraud Detector console

This example creates two variables, `email_address` and `ip_address`, and assigns them to the corresponding variable types (`EMAIL_ADDRESS` and `IP_ADDRESS`). These variables are used as examples. If you are creating variables to use for your model training, use the variables from your dataset that are appropriate for your use case. Make sure to read about [Variable types](#) and [Variable enrichments](#) before you create your variables.

To create a variable,

1. Open the [AWS Management Console](#) and sign in to your account.
2. Navigate to Amazon Fraud Detector, choose **Variables** in the left navigation, then choose **Create**.

3. In the **New variable** page, enter `email_address` as the variable name. Optionally, enter a description of the variable.
4. In the **Variable type**, choose **Email Address**.
5. Amazon Fraud Detector automatically selects the data type for this variable type because this variable type is predefined. If your variable is not automatically assigned a variable type, select a variable type from the list. For more information, see [Variable types](#).
6. If you want to provide a default value for your variable, select **Define a custom default value** and enter a default value for your variable. Skip this step if you are following this example.
7. Choose **Create**.
8. In the **email_address** overview page, confirm the details of the variable you just created.

If you need to update, choose **Edit** and provide the updates. Choose **Save changes**.
9. Repeat the process to create another variable `ip_address` and choose **IP Address** for the variable type.
10. The **Variables** page shows the newly created variables.

Important

We recommend that you create as many variables as you want from your dataset. You can decide later when creating your event type which variables you want to include for training your model to detect fraud and to generate fraud detections.

Create a variable using the AWS SDK for Python (Boto3)

The following example shows requests for the [CreateVariable](#) API. The example creates two variables, `email_address` and `ip_address`, and assigns them to the corresponding variable types (`EMAIL_ADDRESS` and `IP_ADDRESS`).

These variables are used as examples. If you are creating variables to use for your model training, use the variables from your dataset that are appropriate for your use case. Make sure to read about [Variable types](#) and [Variable enrichments](#) before you create your variables.

Be sure to specify a variable source. It helps to identify where the variable value is derived. If the variable source is **EVENT**, the variable value is sent as part of the [GetEventPrediction](#)

request. If the variable value is `MODEL_SCORE`, it's populated by an Amazon Fraud Detector. If `EXTERNAL_MODEL_SCORE`, the variable value is populated by an imported SageMaker model.

```
import boto3
fraudDetector = boto3.client('frauddetector')

#Create variable email_address
fraudDetector.create_variable(
    name = 'email_address',
    variableType = 'EMAIL_ADDRESS',
    dataSource = 'EVENT',
    dataType = 'STRING',
    defaultValue = '<unknown>'
)

#Create variable ip_address
fraudDetector.create_variable(
    name = 'ip_address',
    variableType = 'IP_ADDRESS',
    dataSource = 'EVENT',
    dataType = 'STRING',
    defaultValue = '<unknown>'
)
```

Delete a variable

When you delete a variable, Amazon Fraud Detector permanently deletes that variable and the data is no longer stored in Amazon Fraud Detector.

You can't delete variables that are included in an event type in Amazon Fraud Detector. You will have to first delete the event type the variable is associated with and then delete the variable.

You can't delete Amazon Fraud Detector model output variables and SageMaker model output variables manually. Amazon Fraud Detector automatically deletes model output variables when you delete the model.

You can delete variable in Amazon Fraud Detector console, using the [delete-variable](#) CLI command, using the [DeleteVariable](#) API, or using the AWS SDK for Python (Boto3)

Delete variable using the console

To delete a variable,

1. Sign in to the AWS Management Console and open the Amazon Fraud Detector console at <https://console.aws.amazon.com/frauddetector>.
2. In the left navigation pane of the Amazon Fraud Detector console, choose **Resources**, then choose **Variables**.
3. Choose the variable that you want to delete.
4. Choose **Actions**, and then choose **Delete**.
5. Enter the variable name, and then choose **Delete variable**.

Delete variable using the AWS SDK for Python (Boto3)

The following code sample deletes a variable *customer_name* using the [DeleteVariable](#) API.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.delete_variable (

name = 'customer_name'

)
```

Labels

A label classifies an event as fraudulent or legitimate. Labels are associated with event types and used to train machine learning models in Amazon Fraud Detector. If you are planning to train either an Online Fraud Insights (OFI) or a Transaction Fraud Insights (TFI) model, a minimum of 400 events in your training dataset must be classified as either *fraudulent* or *legitimate*. You can use any labels such as *fraud*, *legit*, *1*, or *0* for classifying events in your training dataset. After the training is complete, the trained model evaluates events for fraud and uses these values to classify events as fraudulent or legitimate.

You will have to first create the labels with the values used in your training dataset and then associate the labels with the event type that is used to build and train your fraud detection model.

Create label

You can create labels in the Amazon Fraud Detector console, using the [put-label](#) command, using the [PutLabel](#) API, or using the AWS SDK for Python (Boto3).

Create a label using the Amazon Fraud Detector console

To create labels,

1. Open the [AWS Management Console](#) and sign in to your account.
2. Navigate to Amazon Fraud Detector, choose **Labels** in the left navigation, then choose **Create**.
3. In the **Create label** page, enter your label name for fraudulent event as the label name. The label name must correspond to the label that represents fraudulent activity in your training dataset. Optionally, enter a description of the label.
4. Choose **Create label**.
5. Create a second label and enter a label name for legitimate event. Make sure the label name corresponds to the value that represents the legitimate activity in your training dataset.

Create a label using the AWS SDK for Python (Boto3)

The following AWS SDK for Python (Boto3) example code creates two labels (fraud, legit) using the [PutLabel](#) API. After creating the labels, you can add them to an event type to classify specific events.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_label(
    name = 'fraud',
    description = 'label for fraud events'
)

fraudDetector.put_label(
    name = 'legit',
    description = 'label for legitimate events'
)
```

Update label

If your event dataset is stored with Amazon Fraud Detector, you might need to add or update labels for the stored events, such as when you perform an offline fraud investigation for an event and want to close the machine learning feed back loop.

You can add or update labels for stored events using the [update-event-label](#) command, using the [UpdateEventLabel](#) API, or using the AWS SDK for Python (Boto3)

The following AWS SDK for Python (Boto3) example code adds a label *fraud* that is associated with the event type *registration* using the `UpdateEventLabel` API.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.update_event_label(
    eventId          = '802454d3-f7d8-482d-97e8-c4b6db9a0428',
    eventTypeName    = 'registration',
    assignedLabel    = 'fraud',
    labelTimestamp   = '2020-07-13T23:18:21Z'
)
```

Updating event labels in event data stored in Amazon Fraud Detector

You might need to add or update fraud labels for events that are already stored in Amazon Fraud Detector, such as when you perform an offline fraud investigation for an event and want to close the machine learning feed back loop. To update the label for an event that is already stored in Amazon Fraud Detector, use the `UpdateEventLabel` API operation. The following shows an example `UpdateEventLabel` API call.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.update_event_label(
    eventId          = '802454d3-f7d8-482d-97e8-c4b6db9a0428',
    eventTypeName    = 'sample_registration',
    assignedLabel    = 'fraud',
```

```
        labelTimestamp = '2020-07-13T23:18:21Z'  
    )
```

Delete label

When you delete a label, Amazon Fraud Detector permanently deletes that label and the data is no longer stored in Amazon Fraud Detector.

You cannot delete a label that is included in an event type in Amazon Fraud Detector. And you also cannot delete a label that is assigned to an event ID. You must first delete the relevant event ID.

You can delete labels in Amazon Fraud Detector console, using the [delete-label](#) command, using the [DeleteLabel](#) API, or using the AWS SDK for Python (Boto3)

Delete label using the console

To delete a label

1. Sign in to the AWS Management Console and open the Amazon Fraud Detector console at <https://console.aws.amazon.com/frauddetector>.
2. In the left navigation pane of the Amazon Fraud Detector console, choose **Resources**, then choose **Labels**.
3. Choose the label that you want to delete.
4. Choose **Actions**, and then choose **Delete**.
5. Enter the label name, and then choose **Delete label**.

Delete a label using the AWS SDK for Python (Boto3)

The following AWS SDK for Python (Boto3) example code deletes a label *legit* using the [DeleteLabel](#) API.

```
import boto3  
fraudDetector = boto3.client('frauddetector')  
  
fraudDetector.delete_event_label (  
    name = 'legit'
```

)

Rules

A rule is a condition that tells Amazon Fraud Detector how to interpret variable values during a fraud prediction. A rule is part of a detector logic and it consists of the following elements:

- **Variable or List** – Variable represents a data element in your event dataset that you want to use in a fraud prediction. A list is a set of input data elements for a variable in your event dataset. Variables used in a rule must be predefined in the evaluated event type and lists used in a rule must be associated with a variable type. For more information, see [Variables](#) and [Lists](#).
- **Expression** – An expression in a rule captures your business logic. If you are using variable in your rule, a simple rule expression is constructed using a variable, a comparison operator such as `>`, `<`, `<=`, `>=`, `==`, and a value. If you are using a list, rule expression is constructed as list entry, `in`, and the list name. For more information, see [Rule language reference](#). You can combine multiple expressions together using `and` and `or`. All expressions must evaluate to a Boolean value (true or false) and be less than 4,000 characters in length. If-else type conditions are not supported.
- **Outcome** – An outcome is a response returned by Amazon Fraud Detector when a rule is matched. The outcome indicates the result of a fraud prediction. You can create outcomes for each possible fraud prediction and add them to a rule. For more information, see [Outcomes](#).

A detector must have at least one associated rule. A rule can have up to 3 lists, and a detector can have up to 30 lists. You create rule as part of the detector creation process. You can also create and associate new rules with an existing detector.

Rule language reference

The following section outlines the expression (that is, rule writing) capabilities in Amazon Fraud Detector.

Using variables

You can use any variable defined in the evaluated event type as part of your expression. Use the dollar sign to indicate a variable:

```
$example_variable < 100
```

Using lists

You can use any list that is associated with a variable type and is populated with entries as part of your rule expression. Use the dollar sign to indicate a list entry value:

```
$example_list_variable in @list_name
```

Comparison, membership, and identity operators

Amazon Fraud Detector includes the following comparison operators: >, >=, <, <=, !=, ==, in, not in

The following are examples:

Example: <

```
$variable < 100
```

Example: in, not in

```
$variable in [5, 10, 25, 100]
```

Example: !=

```
$variable != "US"
```

Example: ==

```
$variable == 1000
```

Operator Tables

| Operator | Amazon Fraud Detector Operator |
|--------------|--------------------------------|
| Equal to | == |
| Not equal to | != |
| Greater than | > |

| Operator | Amazon Fraud Detector Operator |
|------------------------|--------------------------------|
| Less than | < |
| Great than or equal to | >= |
| Less than or equal to | <= |
| In | in |
| And | and |
| Or | or |
| Not | ! |

Basic math

You can use basic math operators in your expression (for example, +, -, *, /). A typical use case is when you need to combine variables during your evaluation.

In the rule below, we are adding the variable `$variable_1` with `$variable_2`, and checking whether the total is less than 10.

```
$variable_1 + $variable_2 < 10
```

Basic Math Table Data

| Operator | Amazon Fraud Detector Operator |
|----------|--------------------------------|
| Plus | + |
| Minus | - |
| Multiply | * |
| Divide | / |
| Modulo | % |

Regular Expression (regex)

You can use regex to search for specific patterns as part of your expression. This is particularly useful if you are looking to match a specific string or numerical value for one of your variables. Amazon Fraud Detector only supports match when working with regular expressions (for example, it returns True/False depending on whether the provided string is matched by the regular expression). Amazon Fraud Detector's regular expression support is based on `.matches()` in java (using the RE2J Regular Expression library). There are several helpful websites on the internet that are useful for testing different regular expression patterns.

In the first example below, we first transform the variable `email` to lowercase. We then check whether the pattern `@gmail.com` is in the `email` variable. Notice the second period is escaped so that we can explicitly check for the string `.com`.

```
regex_match(".*@gmail\.com", lowercase($email))
```

In the second example, we check whether the variable `phone_number` contains the country code `+1` to determine if the phone number is from the US. The plus symbol is escaped so that we can explicitly check for the string `+1`.

```
regex_match(".*\+1", $phone_number)
```

Regex Table

| Operator | Amazon Fraud Detector Example |
|---|--|
| Match any string that starts with | <code>regex_match("^mystring", \$variable)</code> |
| Match entire string exactly | <code>regex_match("mystring", \$variable)</code> |
| Match any character except new line | <code>regex_match(".", \$variable)</code> |
| Match any number of characters except new line prior 'mystring' | <code>regex_match(".*mystring", \$variable)</code> |
| Escape special characters | <code>\</code> |

Checking for missing values

Sometimes it is beneficial to check whether the value is missing. In Amazon Fraud Detector this is represented by null. You can do this by using the following syntax:

```
$variable != null
```

Similarly, if you wanted to check whether a value is not present, you could do the following:

```
$variable == null
```

Multiple conditions

You can combine multiple expressions together using `and` and `or`. Amazon Fraud Detector stops in an OR expression when a single true value is found, and it stops in an AND when a single false value is found.

In the example below, we are checking for two conditions using the `and` condition. In the first statement, we are checking whether variable 1 is less than 100. In the second we check whether variable 2 is not the US.

Given the rule uses an `and`, both must be TRUE for the entire condition to evaluate to TRUE.

```
$variable_1 < 100 and $variable_2 != "US"
```

You can use parenthesis to group Boolean operations, as shown following:

```
$variable_1 < 100 and $variable_2 != "US" or ($variable_1 * 100.0 > $variable_3)
```

Other expression types

DateTime functions

| Function | Description | Example |
|-----------------------------------|--|---|
| <code>getcurrentdatetime()</code> | Gives the current time of the rule execution in ISO8601 UTC format. You can use getepochm | <code>getcurrentdatetime() == "2023-03-28T18:34:02Z"</code> |

| Function | Description | Example |
|--|--|---|
| | milliseconds(getcurrentdatetime()) to perform additional operations | |
| isbefore(DateTime1, DateTime2) | Returns a boolean(True/False) if the caller DateTime1 is before DateTime2 | isbefore(getcurrentdatetime(), "2019-11-30T01:01:01Z") == "False" isbefore(getcurrentdatetime(), "2050-11-30T01:05:01Z") == "True" |
| isafter(D ateTime1, DateTime2) | Returns a boolean(True/False) if the caller DateTime1 is after DateTime2 | isafter(getcurrentdatetime(), "2019-11-30T01:01:01Z") == "True" isafter(getcurrentdatetime(), "2050-11-30T01:05:01Z") == "False" |
| getepochm illiseconds(DateTi me) | Takes a DateTime and returns that DateTime in epoch milliseconds. Useful for performing mathematical operations on the date | getepochmilliseconds("2019- 11-30T01:01:01Z") == 1575032461 |

String Operators

| Operator | Example |
|-------------------------------|-----------------------|
| Transform string to uppercase | uppercase(\$variable) |
| Transform string to lowercase | lowercase(\$variable) |

Other

| Operator | Comment |
|---------------|--------------|
| Add a comment | # my comment |

Create rules

You can create rules in Amazon Fraud Detector console, using the [create-rule](#) command, using the [CreateRule](#) API, or using the AWS SDK for Python (Boto3).

Each rule must contain a single expression that captures your business logic. All expressions must evaluate to a Boolean value (true or false) and be less than 4,000 characters in length. If-else type conditions are not supported. All variables used in the expression must be predefined in the evaluated event type. Similarly, all lists used in the expression must be predefined, associated with a variable type, and be populated with entries.

The following example creates a rule `high_risk` for an existing detector `payments_detector`. The rule associates an expression and an outcome `verify_customer` with the rule.

Prerequisites

To follow the steps mentioned below, make sure that you complete the following before you proceed with creating rules:

- [Create a detector](#)
- [Create an outcome](#)

If you are creating a detector, rule, and outcome for your use case, replace the example detector name, rule name, rule expression and outcome name with the names and expressions relevant to your use case.

Create a new rule in the Amazon Fraud Detector console

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Detectors** and select the detector you created for your use case, example `payments_detector`.
3. In the `payments_detector` page, choose **Associated rules** tab and then choose **Create rule**.
4. In the **New rule** page, enter the following:
 - a. In the **Name**, enter a name for the rule, example `high_risk`
 - b. In the **Description - optional**, optionally enter a rule description, example, **This rule captures events with a high ML model score**

- c. In the **Expression**, enter a rule expression for your use case using the **Expression quick reference guide**. Example `$sample_fraud_detection_model_insightscore >900`
 - d. In the **Outcomes**, choose the outcome you created for your use case, example **verify_customer**. An outcome is the result from a fraud prediction and is returned if the rule matches during an evaluation.
5. Choose **Save rule**

You created a new rule for your detector. This is the version 1 of the rule which Amazon Fraud Detector automatically makes it available for the detector to use.

Create a rule using the AWS SDK for Python (Boto3)

The following example code uses [CreateRule](#) API to create a rule `high_risk` for an existing detector `payments_detector`. The example code also adds a rule expression and an outcome `verify_customer` to the rule.

Prerequisites

To use the example code, make sure that you have complete the following before you proceed with creating rules:

- [Create a detector](#)
- [Create an outcome](#)

If you are creating a detector, rule, and outcome for your use case, replace the example detector name, rule name, rule expression and outcome name with names and expression relevant to your use case.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.create_rule(
    ruleId = 'high_risk',
    detectorId = 'payments_detector',
    expression = '$sample_fraud_detection_model_insightscore > 900',
    language = 'DETECTORPL',
    outcomes = ['verify_customer']
)
```

You have created the version 1 of the rule which Amazon Fraud Detector automatically makes it available for the detector to use.

Update rule

You can update a rule anytime by adding or updating the rule description, updating the rule expression, or adding or removing the outcome for the rule. When you update a rule a new rule version is created.

You can update a rule in the Amazon Fraud Detector console, using the [update-rule-version](#) command, using the [UpdateRuleVersion](#) API, or using the using the AWS SDK.

After you have updated the rule, make sure to update your detector version to use the new rule version.

Update rule in the Amazon Fraud Detector console

To update a rule,

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Detectors**.
3. In the **Detectors** pane, select the detector that is associated with the rule you want to update.
4. In your detector page, choose **Associated rules** tab and select the rule you want to update.
5. In your rule page, choose **Actions** and select **Create version**.
6. Note that the version has changed. Enter updated description, expression, or outcome.
7. Choose **Save new version**

Update rule using the AWS SDK for Python (Boto3)

The following example code uses the [UpdateRuleVersion](#) API to update the threshold for the rule `high_risk` from 900 to 950. This rule is associated with the detector `payments_detector`.

```
fraudDetector.update_rule_version(  
rule = {  
    'detectorId' : 'payments_detector',  
    'ruleId' : 'high_risk',  
    'ruleVersion' : '1'
```

```
},  
expression = '$sample_fraud_detection_model_insightscore > 950',  
language = 'DETECTORPL',  
outcomes = ['verify_customer']  
)
```

Lists

A list is a set of input data for a variable in your event dataset. You use the input data in a rule that's associated with your detector. A rule is a condition that tells Amazon Fraud Detector how to interpret input data during a fraud prediction. For example, you can create a list of IP addresses and then create a rule to deny access if a specific IP address is in the list. Rules that use lists are expressed in the `$ip_address_value` in `@list_name` format.

With Amazon Fraud Detector, you can manage a list by adding or removing data without needing to update an associated rule. A rule associated with your list automatically incorporates newly added or removed data.

A list can contain up to 100,000 unique entries and each entry can be up to 320 characters long. Every list you use in a rule is, by default, associated with Amazon Fraud Detector's [Variable types](#) `FREE_FORM_TEXT`. You can assign a variable type to your list at any time. You can use up to 3 lists in a rule.

You can create a list, add entries to the list, delete a list, or delete one or more entries in the list, or assign a variable type to your list in the Amazon Fraud Detector console, using the API, using the AWS CLI, or using the AWS SDK.

Create a list

You can create a list containing input data (entries) of a variable in your event dataset and use the list in rule expression. The entries in the list can be managed dynamically without updating the rule that is using the list.

To create a list, you must first specify a name and then optionally associate the list with a [Variable types](#) supported by Amazon Fraud Detector. By default, Amazon Fraud Detector assumes the list to be of `FREE_FORM_TEXT` variable type.

You can create a list in the Amazon Fraud Detector console, using the API, using the AWS CLI, or using the AWS SDK.

Create list using the Amazon Fraud Detector console

To create a list

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Lists**.
3. Under **Lists** details
 - a. In the **List name**, enter a name for your list.
 - b. In the **Description**, optionally, enter a description.
 - c. (Optional) In the **Variable type**, select a variable type for your list.

Important

If your list contains IP addresses, make sure to select **IP_ADDRESS** as the variable type. If you don't select a variable type, Amazon Fraud Detector assumes the list to be of **FREE_FORM_TEXT** variable type.

4. In the **Add list data**, add list entries, one entry in each line. You can also copy and paste entries from a spreadsheet.

Note

Make sure that the entries aren't separated using a comma and are unique in the list. If two identical entries are entered, only one will be added.

5. Choose **Create**.

Create list using the AWS SDK for Python (Boto3)

You create a list by specifying a list name. You can optionally provide a description, associate a variable type, or add entries to your list when you are creating a list. Or, you can update the list later on by adding entries or a description. You can assign a variable type to the list later if you haven't assigned it when at the time of list creation. Variable type of a list cannot be changed after it is assigned.

⚠ Important

If your list contains IP addresses, make sure to assign **IP_ADDRESS** as the variable type. If you don't assign a variable type, Amazon Fraud Detector assumes the list to be of **FREE_FORM_TEXT** variable type.

The following example uses [CreateList](#) API operation to create an `allow_email_ids` list by providing a description, a variable type, and by adding four list entries.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.create_list (
    name = 'allow_email_ids',
    description = 'legitimate email_ids'
    variableType = 'EMAIL_ADDRESS',
    elements = ['emailId_1', 'emailId_2', 'emailId_3', 'emailId_4']
)
```

Add entries in a list

After you created your list you can add or append entries in your list at any time. When you add or append entries in your list, you don't need to update the rule the list is associated with. The rule automatically incorporates the newly added entries.

Your list can contain up to 100,000 unique entries and each entry can be up to 320 characters.

You can add entries in the Amazon Fraud Detector console, using the API, using the AWS CLI, or using the AWS SDK.

Add entries in a list using the Amazon Fraud Detector console

To add one or more entries in a list

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Lists**.

3. In the **Lists** page, select the list that you want to add entries to.
4. In your list details page, select **List data** tab and choose **Add data**.
5. In the **Add list data** box, add one entry on each line or copy and paste entries from a spreadsheet. Make sure to not use comma to separate entries.
6. Choose **Add**.

Add entries in a list using the AWS SDK for Python (Boto3)

The following example uses the [UpdateList](#) API operation to add two new entries in the `allow_email_ids` list. Make sure that the entries you are adding are unique in the list.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.update_list (
    name = 'allow_email_ids',
    updateMode = 'APPEND'
    elements = ['emailId_11', 'emailId_12']
```

Assign a variable type to a list

Every list you use in a rule must be associated with an Amazon Fraud Detector's [Variable types](#) variable type. By default, Amazon Fraud Detector assumes the list to be of `FREE_FORM_TEXT` variable type. It is important to note that a list that consists of IP addresses must be associated with `IP_ADDRESS` variable type.

You can associate your list with a variable type either at the time of list creation or anytime later. If you already associated your list with a variable type and want to change it later, you must create a new list. You can't change the variable type of a list.

You can assign a variable type in the Amazon Fraud Detector console, using the API, using the AWS CLI, or using the AWS SDK.

Assign variable type to a list using the Amazon Fraud Detector console

To assign a variable type to a list

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Lists**.
3. In the **Lists** page, select the list that you want to assign a variable type.
4. In your list details page, choose **Actions** and select **Edit list**.
5. In the **Edit list** box, select the variable type for your list.
6. Choose **Save**.

Assign variable type to a list using the AWS SDK for Python (Boto3)

The following example uses the [UpdateList](#) API operation to assign a variable type to `allow_ip_address` list.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.update_list (
    name = 'allow_ip_address',
    variableType = 'IP_ADDRESS'
)
```

Delete a list

You can delete a list that isn't used in any rule. When you delete a list, Amazon Fraud Detector permanently deletes that list and all entries in the list.

You can delete a list in the Amazon Fraud Detector console, using the API, using the AWS CLI or the AWS SDK.

Delete list using the Amazon Fraud Detector console

To delete a list

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Lists**
3. In the **Lists** page, select the list you want to delete.
4. In your list details page, choose **Actions** and select **Delete list**.
5. Choose **Delete list**.

Delete list using the AWS SDK for Python (Boto3)

The following example uses the [DeleteList](#) API operation to delete `allow_email_ids`.

```
import boto3

fraudDetector = boto3.client('frauddetector')

fraudDetector.delete_list(
    name = 'allow_email_ids'
)
```

Delete entries from a list

You can delete one or more entries from your lists at any time. When you delete entries in your list you don't need to update the rule the list is associated with. The rule automatically incorporates the updated list.

You can delete entries from a list in the Amazon Fraud Detector console, using the API, using the AWS CLI or the AWS SDK.

Delete entries from a list using the Amazon Fraud Detector console

To delete one or more entries from a list

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Lists**

3. In the **Lists** page, select the list that contains entries you want to delete.
4. In your list details page, select **List data** tab and select entries you want to delete.
5. Choose **Delete** and choose **Delete** again to confirm.

Delete entries from a list using the AWS SDK for Python (Boto3)

In the following example the [UpdateList](#) API operation deletes entries from allow_email_ids list.

```
import boto3

fraudDetector = boto3.client('frauddetector')

fraudDetector.update_list(
    name = 'allow_email_ids',
    updateMode = 'REMOVE',
    elements = ['emailId_4', 'emailId_12']
)
```

Delete all entries from a list

You can delete all entries in your list, if the list isn't being used in a rule. You can delete all the entries that are in the list and later add entries in the same list.

You can delete entries from a list in the Amazon Fraud Detector console, using the API, using the AWS CLI or the AWS SDK.

Delete all entries from a list using the Amazon Fraud Detector console

To delete all entries from a list

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Lists**
3. In the **Lists** page, select the list that contains entries you want to delete.
4. In your list details page, select **List data** tab and choose **Delete all**.
5. In the **Delete all** box, type delete all to confirm and then choose **Delete all list data**.

Delete all entries from a list using the AWS SDK for Python (Boto3)

In the following example the [UpdateList](#) API operation deletes all entries from `allow_email_ids` list.

```
import boto3

fraudDetector = boto3.client('frauddetector')

fraudDetector.update_list(
    name = 'allow_email_ids',
    updateMode = 'REPLACE',
    elements = []
)
```

Outcomes

An outcome is the result of a fraud prediction. You can create an outcome for each possible fraud prediction result. For example, you might want outcomes to represent risk levels (`high_risk`, `medium_risk`, and `low_risk`) or actions (`approve`, `review`). After an outcome is created, you can add one or more outcomes to a rule. As part of the [GetEventPrediction](#) response, Amazon Fraud Detector returns the defined outcomes for any matched rule.

Create an outcome

You can create outcomes in the Amazon Fraud Detector console, using the [put-outcome](#) command, using the [PutOutcome](#) API, or using the AWS SDK for Python (Boto3).

Create an outcome using the Amazon Fraud Detector console

To create one or more outcomes,

1. Open the [AWS Management Console](#) and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Outcomes**.
3. In the **Outcomes** page, choose **Create**.
4. In your **New outcome** page, enter the following:
 - a. In the **Outcome name**, enter a name for your outcome.

- b. In the **Outcome description**, optionally, enter a description.
5. Choose **Save outcome**.
6. Repeat steps 2 to 5 for creating additional outcomes.

Create an outcome using the AWS SDK for Python (Boto3)

The following example uses the PutOutcome API to create three outcomes. They are `verify_customer`, `review`, and `approve`. After the outcomes are created, you can assign them to rules.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_outcome(
    name = 'verify_customer',
    description = 'this outcome initiates a verification workflow'
)

fraudDetector.put_outcome(
    name = 'review',
    description = 'this outcome sidelines event for review'
)

fraudDetector.put_outcome(
    name = 'approve',
    description = 'this outcome approves the event'
)
```

Delete an outcome

You cannot delete an outcome that is used in a rule version.

When you delete an outcome, Amazon Fraud Detector permanently deletes that outcome and the data is no longer stored in Amazon Fraud Detector.

You can delete an outcome in the Amazon Fraud Detector console, using the [delete-outcome](#) command, using the [DeleteOutcome](#) API, or using the AWS SDK for Python (Boto3)

Delete an outcome in the Amazon Fraud Detector console

To delete an outcome

1. Sign in to the AWS Management Console and open the Amazon Fraud Detector console at <https://console.aws.amazon.com/frauddetector>.
2. In the left navigation pane of the Amazon Fraud Detector console, choose **Resources**, and then choose **Outcomes**.
3. Choose the outcome that you want to delete.
4. Choose **Actions**, and then choose **Delete**.
5. Enter the outcome name, and then choose **Delete outcome**.

Delete an outcome using the AWS SDK for Python (Boto3)

The following example uses the [DeleteOutcome](#) API to delete the `verify_customer` outcome. After the outcome is deleted, you can no longer assign it to a rule.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.delete_outcome(
    name = 'verify_customer'
)
```

Entity

An entity represents a person or thing that's performing the event. An entity type classifies the entity. Example classifications include customer, merchant, user, or account. You provide the entity type (`ENTITY_TYPE`) and an entity identifier (`ENTITY_ID`) as part of your event dataset to indicate the specific entity that performed the event.

Amazon Fraud Detector uses the entity type when generating fraud prediction for an event to indicate who performed the event. The entity type you want to use in your fraud predictions must first be created in Amazon Fraud Detector and then added to the event when creating your event type.

Create an entity type

You can create an entity type in the Amazon Fraud Detector console, using the [put-entity-type](#) command, using the [PutEntityType](#) API, or using the AWS SDK for Python (Boto3). The examples below creates an entity type `customer` in the Amazon Fraud Detector console and using the SDK for Python (Boto3). If you are creating an entity type to associate with an event type for training a fraud detection model, use the entity type from your event dataset that is appropriate for your use case.

Create an entity type using the Amazon Fraud Detector console

To create an entity type,

1. Open the [AWS Management Console](#) and sign in to your account.
2. Navigate to Amazon Fraud Detector, choose **Entities** in the left navigation, then choose **Create**.
3. In the **Create entity** page, enter **customer** as the entity type name. Optionally, enter a description of the entity.
4. Choose **Create entity**.

Create an entity type using the AWS SDK for Python (Boto3)

The following AWS SDK for Python (Boto3) code example uses the `PutEntityType` API to create an entity type `customer`. If you are creating an entity type to associate with an event type for training a fraud detection model, use the entity from your event dataset that is appropriate for your use case.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.put_entity_type(
    name = 'customer',
    description = 'customer'
)
```

Delete an entity type

In Amazon Fraud Detector, you cannot delete an entity type that is included in an event type. You will have to first delete the event type the entity is associated with and then delete the entity type.

When you delete an entity type, Amazon Fraud Detector permanently deletes that entity type and the data is no longer stored in Amazon Fraud Detector.

An entity type can be deleted in Amazon Fraud Detector console, using the [delete-entity-type](#) command, using the [DeleteEntityType](#) API, or using the AWS SDK for Python (Boto3)

Delete an entity type in Amazon Fraud Detector console

To delete an entity type,

1. Sign in to the AWS Management Console and open the Amazon Fraud Detector console at <https://console.aws.amazon.com/frauddetector>.
2. In the left navigation pane of the Amazon Fraud Detector console, choose **Resources**, then choose **Entities**.
3. Choose the entity type that you want to delete.
4. Choose **Actions**, and then choose **Delete**.
5. Enter the entity type name, and then choose **Delete entity type**.

Delete entity type using the AWS SDK for Python (Boto3)

The following AWS SDK for Python (Boto3) example code deletes the entity type *customer* using the [DeleteEntityType](#) API.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.delete_entity_type (

name = 'customer'

)
```

Manage Amazon Fraud Detector resources using AWS CloudFormation

Amazon Fraud Detector is integrated with AWS CloudFormation, a service that helps you to model and set up your Amazon Fraud Detector resources so that you can spend less time creating and managing your resources and infrastructure. You create a template that describes all the Amazon Fraud Detector resources that you want (such as Detector, Variables, EntityType, EventType, Outcome, and Label), and AWS CloudFormation provisions and configures those resources for you. You can reuse the template to provision and configure the resources consistently and repeatedly in multiple AWS accounts and Regions.

There is no additional charge for using AWS CloudFormation.

Creating Amazon Fraud Detector templates

To provision and configure resources for Amazon Fraud Detector and related services, you must understand [AWS CloudFormation templates](#). Templates are formatted text files in JSON or YAML. These templates describe the resources that you want to provision in your AWS CloudFormation stacks. If you're unfamiliar with JSON or YAML, you can use AWS CloudFormation Designer to help you get started with AWS CloudFormation templates. For more information, see [What is AWS CloudFormation Designer?](#) in the *AWS CloudFormation User Guide*.

You can also create, update, and delete your Amazon Fraud Detector resources using AWS CloudFormation templates. For more information, including examples of JSON and YAML templates for your resources, see the [Amazon Fraud Detector resource type reference](#) in the *AWS CloudFormation User Guide*.

If you are already using CloudFormation, there is no need to manage additional IAM policies or CloudTrail logging.

Managing Amazon Fraud Detector stacks

You can create, update, and delete your Amazon Fraud Detector stacks through the CloudFormation console or through the AWS CLI.

To create a stack, you must have a template that describes what resources AWS CloudFormation will include in your stack. You can also bring Amazon Fraud Detector resources that you have already created into CloudFormation management by [importing them](#) into a new or existing stack.

For detailed instructions for managing your stacks, see the *AWS CloudFormation User Guide* to learn how to [create](#), [update](#), and [delete](#) stacks.

Organizing your Amazon Fraud Detector stacks

The way you organize your AWS CloudFormation stacks is entirely up to you. It is generally a best practice is to organize stacks by lifecycle and ownership. This means grouping resources by how frequently they change or by teams that are responsible for updating them.

You can choose to organize your stacks by creating a stack for each detector and its detection logic (for example, rules, variables, etc.). If you are using other services, you should consider whether you want to stack together Amazon Fraud Detector resources with resources from other services. For example, you could create a stack that includes Kinesis resources that help gather data and Amazon Fraud Detector resources that process the data. This can be an effective way to ensure that all of your fraud team's products are working together.

Understanding Amazon Fraud Detector CloudFormation parameters

In addition to the standard parameters that are available in all CloudFormation templates, Amazon Fraud Detector introduces two additional parameters that will help you manage deployment behavior. If you do not include one or both of these parameters, CloudFormation will use the default value shown below.

| Parameter | Values | Default Value |
|-----------------------|---|---------------|
| DetectorVersionStatus | <p>ACTIVE: Set the new/updated detector version to Active status</p> <p>DRAFT: Set the new/updated detector version to Draft status</p> | DRAFT |
| Inline | <p>TRUE: Allow CloudFormation to create/update/delete the resource when creating/updating/deleting the stack.</p> <p>FALSE: Allow CloudFormation to validate that the object exists but not make any changes to the object.</p> | TRUE |

Sample AWS CloudFormation template for Amazon Fraud Detector resources

The following is a sample AWS CloudFormation YAML template for managing a detector and associated detector versions.

```
# Simple Detector resource containing inline Rule, EventType, Variable, EntityType and
Label resource definitions
Resources:
  TestDetectorLogicalId:
    Type: AWS::FraudDetector::Detector
    Properties:
      DetectorId: "sample_cfn_created_detector"
      DetectorVersionStatus: "DRAFT"
      Description: "A detector defined and created in a CloudFormation stack!"

    Rules:
      - RuleId: "over_threshold_investigate"
        Description: "Automatically sends transactions of $10000 or more to an
investigation queue"
        DetectorId: "sample_cfn_created_detector"
        Expression: "$amount >= 10000"
        Language: "DETECTORPL"
        Outcomes:
          - Name: "investigate"
            Inline: true
      - RuleId: "under_threshold_approve"
        Description: "Automatically approves transactions of less than $10000"
        DetectorId: "sample_cfn_created_detector"
        Expression: "$amount <10000"
        Language: "DETECTORPL"
        Outcomes:
          - Name: "approve"
            Inline: true
    EventType:
      Inline: "true"
      Name: "online_transaction"
      EventVariables:
        - Name: "amount"
          DataSource: 'EVENT'
          DataType: 'FLOAT'
          DefaultValue: '0'
          VariableType: "PRICE"
```

```
    Inline: 'true'  
EntityTypes:  
  - Name: "customer"  
    Inline: 'true'  
Labels:  
  - Name: "legitimate"  
    Inline: 'true'  
  - Name: "fraudulent"  
    Inline: 'true'
```

Learn more about AWS CloudFormation

To learn more about AWS CloudFormation, see the following resources:

- [AWS CloudFormation](#)
- [AWS CloudFormation User Guide](#)
- [AWS CloudFormation API Reference](#)
- [AWS CloudFormation Command Line Interface User Guide](#)

Fraud predictions

You can use Amazon Fraud Detector to get fraud predictions for single event in real time or get fraud predictions offline for a set of events. To generate fraud predictions for either a single event or a set of events, you will need to provide Amazon Fraud Detector with the following information:

- Fraud prediction logic
- Event metadata

Fraud detection logic

Fraud prediction logic uses one or more rules to evaluate data associated with an event and then provides result and a fraud prediction score. You create your fraud prediction logic using the following components:

- Event types - Defines structure of the event
- Models - Defines algorithm and data requirements for predicting fraud
- Variables - Represents a data element associated with the event
- Rules - Tells Amazon Fraud Detector how to interpret the variable values during fraud prediction
- Outcomes - Results generated from a fraud prediction
- Detector version - Contains fraud prediction logic for a particular event

For more information about the components used to create fraud detection logic, see [Amazon Fraud Detector concepts](#). Before you start generating fraud predictions, make sure you have created and published the detector version that contains your fraud prediction logic. You can create and publish detector version using the Fraud Detector Console or API. For instructions on using the console, see [Get started \(console\)](#). For instructions on using the API, see [Create a detector version](#).

Event metadata

Event metadata provides details of the event being evaluated. Each event you want to evaluate must include value for each variable in the event type associated with your detector version. In addition, your event metadata must include the following:

- **EVENT_ID** – An identifier for the event. For example, if your event is an online transaction the **EVENT_ID** might be the transaction reference number provided to your customer.

Important notes about **EVENT_ID**

- Must be unique for that event
- Should represent information that is meaningful to your business
- Must satisfy the regular expression pattern : `^[0-9a-z_-]+$`.
- Must be saved. **EVENT_ID** is the reference for the event and is used to perform operations on the event such as deleting the event.
- Appending timestamp to the **EVENT_ID** is not recommended as it might cause issues when later you want to update the event, since you will need to provide the exact same **EVENT_ID**.
- **ENTITY_TYPE** – The entity that performs the event, such as a merchant or a customer.
- **ENTITY_ID** - An identifier for the entity performing the event. The **ENTITY_ID** must satisfy the following regular expression pattern: `^[0-9a-z_-]+$`. If the **ENTITY_ID** is not available at the time of evaluation, pass the string `unknown`.
- **EVENT_TIMESTAMP** - The timestamp when the event occurred. The timestamp must be in ISO 8601 standard in UTC.

Real time prediction

You can evaluate online activities for fraud in real time by calling `GetEventPrediction` API. You provide information about a single event in each request and synchronously receive a model score and an outcome based on the fraud prediction logic associated with the specified detector.

How real time fraud prediction works

The `GetEventPrediction` API uses a specified detector version to evaluate the event metadata provided for the event. During the evaluation, Amazon Fraud Detector first generates model scores for models that are added to the detector version, then passes the results to the rules for evaluation. The rules are executed as specified by the rule execution mode (see [Create a detector version](#)). As part of the response, Amazon Fraud Detector provides model scores as well as any outcomes associated to the matched rules.

Getting real time fraud prediction

To get real time fraud predictions, make sure you have created and published a detector that contains your fraud prediction model and rules, or simply a ruleset.

You can get fraud prediction for an event in real time by calling the [GetEventPrediction](#) API operation using the AWS Command Line Interface (AWS CLI) or one of the Amazon Fraud Detector SDKs.

To use the API, supply information of a single event with each request. As part of the request you must specify `detectorId` that Amazon Fraud Detector will use to evaluate the event. You can optionally specify a `detectorVersionId`. If a `detectorVersionId` is not specified, Amazon Fraud Detector will use the `ACTIVE` version of the detector.

You can optionally send data to invoke an SageMaker model by passing the data in the field `externalModelEndpointBlobs`.

Get a fraud prediction using the AWS SDK for Python (Boto3)

To generate a fraud prediction, call the `GetEventPrediction` API. The example below assumes you have completed [Part B: Generate fraud predictions](#). As part of the response, you will receive a model score as well as any matched rules and corresponding outcomes. You can find additional examples of `GetEventPrediction` requests on the [aws-fraud-detector-samples GitHub repository](#).

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.get_event_prediction(
    detectorId = 'sample_detector',
    eventId = '802454d3-f7d8-482d-97e8-c4b6db9a0428',
    eventName = 'sample_registration',
    eventTimestamp = '2020-07-13T23:18:21Z',
    entities = [{'entityType': 'sample_customer', 'entityId': '12345'}],
    eventVariables = {
        'email_address' : 'johndoe@exampldomain.com',
        'ip_address' : '1.2.3.4'
    }
)
```

Batch predictions

You can use a *batch predictions* job in Amazon Fraud Detector to get predictions for a set of events that do not require real-time scoring. For example, you could create a batch predictions job to perform an offline proof-of-concept, or to retrospectively evaluate the risk of events on an hourly, daily, or weekly basis.

You can create a batch prediction job using the [Amazon Fraud Detector console](#), or by calling the [CreateBatchPredictionJob](#) API operation using the AWS Command Line Interface (AWS CLI) or one of the Amazon Fraud Detector SDKs.

Topics

- [How batch predictions work](#)
- [Input and output files](#)
- [Getting batch predictions](#)
- [Guidance on IAM roles](#)
- [Get batch fraud predictions using the AWS SDK for Python \(Boto3\)](#)

How batch predictions work

The `CreateBatchPredictionJob` API operation uses a specified detector version to make predictions based on data provided in an input CSV file that is located in an Amazon S3 bucket. The API then returns the resulting CSV file to an S3 bucket.

Batch prediction jobs calculate model scores and prediction outcomes in the same way as the `GetEventPrediction` operation. Similar to `GetEventPrediction`, to create a batch predictions job, you first create an event type, optionally train a model, and then create a detector version that evaluates the events in your batch job.

The pricing for event risk scores evaluated by batch prediction jobs is the same as the pricing for scores created by the `GetEventPrediction` API. For details, see [Amazon Fraud Detector pricing](#).

You can only run one batch prediction job at a time.

Input and output files

The input CSV file should contain headers that match the event type that is associated with the selected detector version. The maximum size of the input data file is 1GB. The number of events will vary by your event size.

Amazon Fraud Detector creates the output file in the same bucket as the input file, unless you specify a separate location for the output data. The output file contains the original data from the input file and the following appended columns:

- **MODEL_SCORES** — Details the model scores for the event from each model associated with the selected detector version.
- **OUTCOMES** — Details the event outcomes as evaluated by the selected detector version and its rules.
- **STATUS** — Indicates whether the event was evaluated successfully. If the event was not evaluated successfully, this column shows a reason code for the failure.
- **RULE_RESULTS** — A list of all the rules that matched, based on the rule execution mode.

Getting batch predictions

The following steps assume that you have already created an event type, trained a model using that event type (optional), and created a detector version for that event type.

To get a batch prediction

1. Sign in to the AWS Management Console and open the Amazon Fraud Detector console at <https://console.aws.amazon.com/frauddetector>.
2. In the left navigation pane of the Amazon Fraud Detector console, choose **Batch Predictions**, and then choose **New batch prediction**.
3. In **Job name**, specify a name for your batch prediction job. If you don't specify a name, Amazon Fraud Detector randomly generates a job name.
4. In **Detector**, choose the detector for this batch prediction.
5. In **Detector version**, choose the detector version for this batch prediction. You can choose a detector version in any status. If your detector has a detector version in **Active** status, that version is automatically selected, but you can also change this selection if needed.

6. In **IAM role**, choose or create a role that has read and write access to your input and output Amazon S3 buckets. See [Guidance on IAM roles](#) for more information.

To get batch predictions, the IAM role that calls the `CreateBatchPredictionJob` operation must have read permissions to your input S3 bucket and write permissions to your output S3 bucket. For more information about bucket permissions, see [User policy examples](#) in the *Amazon S3 User Guide*.

7. In **Input data location**, specify the Amazon S3 location of your input data. If you want the output file in a different S3 bucket, select **Separate data location for output** and provide the Amazon S3 location for your output data.
8. (Optional) Create tags for your batch prediction job.
9. Choose **Start**.

Amazon Fraud Detector creates the batch prediction job, and the job's status is `In progress`. Batch prediction job processing times vary depending on the number of events and your detector version configuration.

To stop a batch prediction job that is in progress, go to the batch prediction job detail page, choose **Actions**, and then choose **Stop batch prediction**. If you stop a batch prediction job, you won't receive any results for the job.

When the batch prediction job's status changes to `Complete`, you can retrieve the job's output from the designated output Amazon S3 bucket. The output file's name is in the format `batch prediction job name_file creation timestamp_output.csv`. For example, the output file from a job named `mybatchjob` is `mybatchjob_1611170650_output.csv`.

To search for specific events evaluated by a batch prediction job, in the left navigation pane of the Amazon Fraud Detector console, choose **Search past predictions**.

To delete a batch prediction job that has completed, go to the batch prediction job detail page, choose **Actions** and then choose **Delete batch prediction**.

Guidance on IAM roles

To get batch predictions, the IAM role that calls the [CreateBatchPredictionJob](#) operation must have read permissions to your input S3 bucket and write permissions to your output S3 bucket. For more information about bucket permissions, see [User policy examples](#) in the *Amazon S3 User Guide*. On

the Amazon Fraud Detector console, you have three options for selecting an IAM role for Batch Predictions:

1. Create a role when creating a new Batch Prediction job.
2. Select an existing IAM role that you have previously created in Amazon Fraud Detector console. Make sure to add the `S3:PutObject` permission to the role before you do this step.
3. Enter a custom ARN for a previously created IAM role.

If you receive an error related to your IAM role, verify the following:

1. Your Amazon S3 input and output buckets are in the same region as your detector.
2. The IAM role you are using has the `s3:GetObject` permission for your input S3 bucket and the `s3:PutObject` permission for your output S3 bucket.
3. The IAM role you are using has a trust policy for service principal `frauddetector.amazonaws.com`.

Get batch fraud predictions using the AWS SDK for Python (Boto3)

The following example shows a sample request for the [CreateBatchPredictionJob](#) API. A batch prediction job must include the following existing resources: detector, detector version, and event type name. The following example assumes you have created an event type `sample_registration`, a detector `sample_detector`, and a detector version 1.

```
import boto3
fraudDetector = boto3.client('frauddetector')

fraudDetector.create_batch_prediction_job (
    jobId = 'sample_batch',
    inputPath = 's3://bucket_name/input_file_name.csv',
    outputPath = 's3://bucket_name/',
    eventName = 'sample_registration',
    detectorName = 'sample_detector',
    detectorVersion = '1',
    iamRoleArn = 'arn:aws:iam::*:role/service-role/AmazonFraudDetector-DataAccessRole-
** '
)
```

Prediction explanations

Prediction explanations provide insight into how each event variable impacted your model's fraud prediction score, and are automatically generated as part of the fraud prediction. Each fraud prediction comes with a risk score between 1 and 1000. Prediction explanations give you details of the influence of each event variable on the risk scores in terms of magnitude (0-5, 5 being highest) and direction (drove score higher or lower). You can also use prediction explanations for the following tasks:

- To identify top risk indicators during manual investigations when an event is flagged for review.
- To narrow down root causes that lead to false positive predictions (for example, high risk scores for legitimate events).
- To analyze fraud patterns across event data and detect bias, if any, in your dataset.

Important

Prediction explanations are automatically generated and available only for models trained on or after *June 30, 2021*. To receive prediction explanations for models trained before *June 30, 2021*, retrain those models.

Prediction explanations provide the following set of values for each event variable that was used to train the model.

Relative impact

Provides a visual reference of the variable's impact in terms of magnitude on the fraud prediction scores. The relative impact values consist of a star rating (0-5, 5 being the highest) and direction (increased/decreased) impact of the fraud risk.

- Variables that increased fraud risk are indicated by red colored stars. The higher the number of red colored stars, the more the variable drove up the fraud score and increased likelihood of fraud.
- Variables that decreased fraud risk are indicated by green colored stars. The higher the number of green colored stars, the more the variable drove down the fraud risk score and decreased likelihood of fraud.

- Zero stars for all variables indicate that none of the variables on their own significantly changed the fraud risk.

Raw explanation value

Provides raw, uninterpreted value represented as log-odds of the fraud. These values are usually between -10 to +10, but range from - infinity to + infinity.

- A positive value indicates that the variable drove the risk score up.
- A negative value indicates that the variable drove the risk score down.

In the Amazon Fraud Detector console, the prediction explanation values are displayed as follows. The colored star ratings and the corresponding raw numerical values make it easy to see the relative influence between variables.

| Prediction explanations - <i>preview</i> | | | |
|---|---------------------------------|-------------------|-------------------------|
| This prediction is based on contribution from each variable to the overall likelihood of a fraudulent event. Prediction explanations give you better understanding of how an event's input variables influence fraud prediction scores. For details on calculations, refer to documentation | | | |
| <input checked="" type="checkbox"/> Show raw prediction explanation value | | | |
| Variables that increased fraud risk | | | |
| Name | Value | Relative impact ⓘ | Raw explanation value ⓘ |
| comp_255 | whatsapp | ★★★★★ | 0.49 |
| req_255 | 0 | ★★★★★ | 0.29 |
| sentiment_description | 0.2 | ★★★★★ | 0.12 |
| desc_255 | this is the company description | ★★★★★ | 0.07 |
| title | king | ★★★★★ | 0.07 |
| required_experience | 5 | ★★★★★ | 0.04 |
| required_education | masters | ★★★★★ | 0.03 |
| has_questions | true | ★★★★★ | 0.01 |
| Variables that decreased fraud risk | | | |
| Name | Value | Relative Impact ⓘ | Raw explanation value ⓘ |
| has_company_logo | true | ★★★★★ | -0.26 |
| req_desc_similarity | 0.3 | ★★★★★ | -0.21 |
| employment_type | temp | ★★★★★ | -0.21 |
| job_location | california | ★★★★★ | -0.11 |
| job_function | engineer | ★★★★★ | -0.06 |
| industry | software | ★★★★★ | -0.05 |
| sentiment_requirements | 0.5 | ★★★★★ | -0.01 |
| telecommuting | yes | ★★★★★ | -0.00 |
| company_desc_similarity | 0.0 | ★★★★★ | -0.00 |

Viewing prediction explanations

After you generate fraud predictions, you can view prediction explanations in the Amazon Fraud Detector console. To view the prediction explanations using APIs from the AWS SDK, you must first call the `ListEventPrediction` API to obtain the prediction timestamp for the event, and then call the `GetEventPredictionMetadata` API to get the prediction explanations.

View prediction explanations using Amazon Fraud Detector console

To view the prediction explanations using the console,

1. Open the AWS Console and sign in to your account. Navigate to Amazon Fraud Detector.
2. In the left navigation pane, choose **Search past predictions**.
3. Use the **Property**, **Operator**, and **Value** filters to select the prediction you want review.
4. In the top **Filter** pane, make sure to select the time period for when the prediction you want to review was generated.
5. The **Results** pane displays a list of all the predictions generated during the specified time period. Click the Event ID of the prediction to view the prediction explanations.
6. Scroll down to the **Prediction explanations** pane.
7. Set the **Show raw prediction explanation value** button **on** to view raw prediction explanation value of all the variables.

View prediction explanations using the AWS SDK for Python (Boto3)

The following examples show sample requests for viewing prediction explanations using `ListEventPredictions` and `GetEventPredictionMetadata` APIs from the AWS SDK.

Example 1: Get a list of most recent predictions using `ListEventPredictions` API

```
import boto3
fraudDetector = boto3.client('frauddetector')
fraudDetector.list_event_predictions(
    maxResults = 10,
    predictionTimeRange = {
        end_time: '2022-01-13T23:18:21Z',
        start_time: '2022-01-13T20:18:21Z'
    }
)
```


Example 2; Get a list of past predictions for event type "registration" using ListEventPredictions API

```
import boto3
fraudDetector = boto3.client('frauddetector')
fraudDetector.list_event_predictions(
    eventType = {
        value = 'registration'
    }
    maxResults = 70,
    nextToken = "10",
    predictionTimeRange = {
        end_time: '2021-07-13T23:18:21Z',
        start_time: '2021-07-13T20:18:21Z'
    }
)
```

Example 3: Get details of a past prediction for a specified event ID, event type, detector ID, and detector version ID that was generated in the specified time period using GetEventPredictionMetadata API.

The predictionTimestamp specified for this request is obtained by first calling the ListEventPredictions API.

```
import boto3
fraudDetector = boto3.client('frauddetector')
fraudDetector.get_event_prediction_metadata (
    detectorId = 'sample_detector',
    detectorVersionId = '1',
    eventId = '802454d3-f7d8-482d-97e8-c4b6db9a0428',
    eventTypeName = 'sample_registration',
    predictionTimestamp = '2021-07-13T21:18:21Z'
)
```

Understanding how prediction explanations are calculated

Amazon Fraud Detector uses [SHAP \(SHapeley Additive exPlanations\)](#) to explain individual event predictions by computing the **raw explanation values** of each event variable used for model

training. The raw explanation values are computed by the model as part of the classification algorithm when generating predictions. These raw explanation values represent the contribution of each input to the logarithm of the odds of fraud. The raw explanation values (from $-\infty$ to $+\infty$) are converted to a **relative impact value** (-5 to +5) using a mapping. The relative impact value derived from raw explanation value represents the number of times increase in odds of fraud (positive) or legit (negative), making it easier to understand the prediction explanations.

Security in Amazon Fraud Detector

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to Amazon Fraud Detector, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon Fraud Detector. The following topics show you how to configure Amazon Fraud Detector to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Amazon Fraud Detector resources.

Topics

- [Data Protection in Amazon Fraud Detector](#)
- [Identity and access management for Amazon Fraud Detector](#)
- [Logging and monitoring in Amazon Fraud Detector](#)
- [Compliance validation for Amazon Fraud Detector](#)
- [Resilience in Amazon Fraud Detector](#)
- [Infrastructure Security in Amazon Fraud Detector](#)

Data Protection in Amazon Fraud Detector

The AWS [shared responsibility model](#) applies to data protection in Amazon Fraud Detector. As described in this model, AWS is responsible for protecting the global infrastructure that runs all

of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-2](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Amazon Fraud Detector or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Encrypting data at rest

Amazon Fraud Detector encrypts your data at rest with your choice of an encryption key. You can choose one of the following:

- An AWS owned [KMS key](#). If you don't specify an encryption key your data is encrypted with this key by default.

- A customer managed [KMS key](#). You can control access to your customer managed KMS key using [key policies](#). For information on creating and managing customer managed KMS key, see [Key management](#).

Encrypting data in transit

Amazon Fraud Detector copies data out of your account and processes it in an internal AWS system. By default, Amazon Fraud Detector uses TLS 1.2 with AWS certificates to encrypt data in transit.

Key management

Amazon Fraud Detector encrypts your data using one of two types of keys:

- An AWS owned [KMS key](#). This is the default.
- A customer managed [KMS key](#).

Creating customer managed KMS key

You can create customer managed KMS key using either the AWS KMS console or the [CreateKey](#) API. When creating the key make sure you,

- Select a symmetric encryption customer managed KMS key, Amazon Fraud Detector does not support asymmetric KMS keys. For more information, see [Asymmetric Keys in AWS KMS](#) in the AWS Key Management Service Developer Guide.
- Create single region KMS key. Amazon Fraud Detector does not support multi-region KMS keys. For more information, see [Multi-region keys in AWS KMS](#) in the AWS Key Management Service Developer Guide.
- Provide the following [key policy](#) to grant permissions to Amazon Fraud Detector to use the key.

```
{
  "Effect": "Allow",
  "Principal": {
    "Service": "frauddetector.amazonaws.com"
  },
  "Action": [
    "kms:Encrypt",
```

```
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey",
    "kms:CreateGrant",
    "kms:RetireGrant"
  ],
  "Resource": "*"
}
```

For information on key policies, see [Using Key Policies in AWS KMS](#) in the AWS Key Management Service Developer Guide.

Encrypting data using customer managed KMS key

Use Amazon Fraud Detector's [PutKMSEncryptionKey](#) API to encrypt your Amazon Fraud Detector data at rest using the customer managed KMS key. You can change the encryption configuration at any time using [PutKMSEncryptionKey](#) API.

Important notes about encrypted data

- Data generated after setting up the customer managed KMS key is encrypted. Data generated before setting up the customer managed KMS key will remain unencrypted.
- If customer managed KMS key is changed, the data that was encrypted using the previous encryption configuration will not be re-encrypted.

View data

When you use customer managed KMS key to encrypt your Amazon Fraud Detector data, the data encrypted using this method is not searchable using filters in the **Search Past Predictions** area of the Amazon Fraud Detector console. To ensure complete search results, use one or more of the following properties to filter results:

- Event ID
- Evaluation timestamp
- Detector status
- Detector version
- Model version

- Model type
- Rule evaluation status
- Rule execution mode
- Rule match status
- Rule version
- Variable data source

If customer managed KMS key was either deleted or is scheduled for deletion, your data might not be available. For more information, see [Deleting KMS key](#).

Amazon Fraud Detector and interface VPC endpoints (AWS PrivateLink)

You can establish a private connection between your VPC and Amazon Fraud Detector by creating an *interface VPC endpoint*. Interface endpoints are powered by [AWS PrivateLink](#), a technology that enables you to privately access Amazon Fraud Detector APIs without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC don't need public IP addresses to communicate with Amazon Fraud Detector APIs. Traffic between your VPC and Amazon Fraud Detector does not leave the Amazon network.

Each interface endpoint is represented by one or more [Elastic Network Interfaces](#) in your subnets.

For more information, see [Interface VPC endpoints \(AWS PrivateLink\)](#) in the *Amazon VPC User Guide*.

Considerations for Amazon Fraud Detector VPC endpoints

Before you set up an interface VPC endpoint for Amazon Fraud Detector, ensure that you review [Interface endpoint properties and limitations](#) in the *Amazon VPC User Guide*.

Amazon Fraud Detector supports making calls to all of its API actions from your VPC.

VPC endpoint policies are supported for Amazon Fraud Detector. By default, full access to Amazon Fraud Detector is allowed through the endpoint. For more information, see [Controlling access to services with VPC endpoints](#) in the *Amazon VPC User Guide*.

Creating an interface VPC endpoint for Amazon Fraud Detector

You can create a VPC endpoint for the Amazon Fraud Detector service using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see [Creating an interface endpoint](#) in the *Amazon VPC User Guide*.

Create a VPC endpoint for Amazon Fraud Detector using the following service name:

- `com.amazonaws.region.frauddetector`

If you enable private DNS for the endpoint, you can make API requests to Amazon Fraud Detector using its default DNS name for the Region, for example, `frauddetector.us-east-1.amazonaws.com`.

For more information, see [Accessing a service through an interface endpoint](#) in the *Amazon VPC User Guide*.

Creating a VPC endpoint policy for Amazon Fraud Detector

You can create a policy for interface VPC endpoints for Amazon Fraud Detector to specify the following:

- The principal that can perform actions
- The actions that can be performed
- The resources on which actions can be performed

For more information, see [Controlling Access to Services with VPC Endpoints](#) in the *Amazon VPC User Guide*.

The following example VPC endpoint policy specifies that all users who have access to the VPC interface endpoint are allowed to access the Amazon Fraud Detector detector named `my_detector`.

```
{
  "Statement": [
    {
      "Action": "frauddetector:*Detector",
      "Effect": "Allow",
```



```
        "Resource": "arn:aws:frauddetector:us-east-1:123456789012:detector/  
my_detector",  
        "Principal": "*"   
    }  
]  
}
```

In this example, the following are denied:

- Other Amazon Fraud Detector API actions
- Invoking Amazon Fraud Detector GetEventPrediction API

Note

In this example, users can still take other Amazon Fraud Detector API actions from outside the VPC. For information about how to restrict API calls to those from within the VPC, see [Amazon Fraud Detector identity-based policies](#).

Opting out of using your data for service improvement

Historical event data you provide to train models and generate predictions is used solely to provide and maintain your service. This data might also be used to improve the quality of Amazon Fraud Detector. Your trust, privacy, and the security of your content are our highest priority and ensure that our use complies with our commitments to you. See [Data Privacy FAQ](#) for more information

You can choose to opt out of having your event data used to develop or improve the quality of Amazon Fraud Detector by visiting the [AI services opt-out policies](#) page in the *AWS Organizations User Guide* and following the process explained there.

Note

Your AWS accounts will need to be centrally managed by AWS Organizations for you to be able to use the opt-out policy. If you have not already created an organization for your AWS accounts, visit [Creating and managing an organization](#) page and follow the process explained there.

Identity and access management for Amazon Fraud Detector

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon Fraud Detector resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How Amazon Fraud Detector works with IAM](#)
- [Amazon Fraud Detector identity-based policy examples](#)
- [Confused deputy prevention](#)
- [Troubleshooting Amazon Fraud Detector identity and access](#)

Audience

How you use AWS Identity and Access Management (IAM) differs, depending on the work that you do in Amazon Fraud Detector.

Service user – If you use the Amazon Fraud Detector service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Amazon Fraud Detector features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Amazon Fraud Detector, see [Troubleshooting Amazon Fraud Detector identity and access](#).

Service administrator – If you're in charge of Amazon Fraud Detector resources at your company, you probably have full access to Amazon Fraud Detector. It's your job to determine which Amazon Fraud Detector features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon Fraud Detector, see [How Amazon Fraud Detector works with IAM](#).

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Amazon Fraud Detector. To view example Amazon Fraud Detector identity-based policies that you can use in IAM, see [Amazon Fraud Detector identity-based policy examples](#).

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be *authenticated* (signed in to AWS) as the AWS account root user, as an IAM user, or by assuming an IAM role.

You can sign in to AWS as a federated identity by using credentials provided through an identity source. AWS IAM Identity Center (IAM Identity Center) users, your company's single sign-on authentication, and your Google or Facebook credentials are examples of federated identities. When you sign in as a federated identity, your administrator previously set up identity federation using IAM roles. When you access AWS by using federation, you are indirectly assuming a role.

Depending on the type of user you are, you can sign in to the AWS Management Console or the AWS access portal. For more information about signing in to AWS, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

If you access AWS programmatically, AWS provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use AWS tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see [Signing AWS API requests](#) in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, AWS recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see [Multi-factor authentication](#) in the *AWS IAM Identity Center User Guide* and [Using multi-factor authentication \(MFA\) in AWS](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity that has complete access to all AWS services and resources in the account. This identity is called the AWS account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your

root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

Users and groups

An [IAM user](#) is an identity within your AWS account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see [Rotate access keys regularly for use cases that require long-term credentials](#) in the *IAM User Guide*.

An [IAM group](#) is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see [When to create an IAM user \(instead of a role\)](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity within your AWS account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. You can temporarily assume an IAM role in the AWS Management Console by [switching roles](#). You can assume a role by calling an AWS CLI or AWS API operation or by using a custom URL. For more information about methods for using roles, see [Using IAM roles](#) in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see [Creating a role for a third-party Identity Provider](#) in the *IAM User Guide*. If you use IAM Identity Center, you configure a permission set. To control what your identities can access after they authenticate, IAM Identity Center correlates the

permission set to a role in IAM. For information about permissions sets, see [Permission sets](#) in the *AWS IAM Identity Center User Guide*.

- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some AWS services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.
- **Cross-service access** – Some AWS services use features in other AWS services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
 - **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in AWS, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other AWS services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).
 - **Service role** – A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Creating a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.
 - **Service-linked role** – A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making AWS CLI or AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see [Using](#)

[an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see [When to create an IAM role \(instead of a user\)](#) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy is an object in AWS that, when associated with an identity or resource, defines their permissions. AWS evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in AWS as JSON documents. For more information about the structure and contents of JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A user with that policy can get role information from the AWS Management Console, the AWS CLI, or the AWS API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Creating IAM policies](#) in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your AWS account. Managed policies include AWS managed policies and customer managed policies. To learn how to choose

between a managed policy or an inline policy, see [Choosing between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Access Control Lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, AWS WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

AWS supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in AWS Organizations. AWS Organizations is a

service for grouping and centrally managing multiple AWS accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each AWS account root user. For more information about Organizations and SCPs, see [How SCPs work](#) in the *AWS Organizations User Guide*.

- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How Amazon Fraud Detector works with IAM

Before you use IAM to manage access to Amazon Fraud Detector, you should understand what IAM features are available to use with Amazon Fraud Detector. To get a high-level view of how Amazon Fraud Detector and other AWS services work with IAM, see [AWS Services That Work with IAM](#) in the *IAM User Guide*.

Topics

- [Amazon Fraud Detector identity-based policies](#)
- [Amazon Fraud Detector resource-based policies](#)
- [Authorization Based on Amazon Fraud Detector Tags](#)
- [Amazon Fraud Detector IAM roles](#)

Amazon Fraud Detector identity-based policies

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. Amazon Fraud Detector supports specific actions, resources, and condition keys. To learn about all of the elements that you use in a JSON policy, see [IAM JSON Policy Elements Reference](#) in the *IAM User Guide*.

To get started with Amazon Fraud Detector, we recommend creating an user with access restricted to Amazon Fraud Detector operations and required permissions. You can add other permissions as needed. The following policies provide the required permission to use Amazon Fraud Detector: `AmazonFraudDetectorFullAccessPolicy` and `AmazonS3FullAccess`. For more information on setting up Amazon Fraud Detector using these policies see [Set up for Amazon Fraud Detector](#).

Actions

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated AWS API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

Policy actions in Amazon Fraud Detector use the following prefix before the action: `frauddetector:`. For example, to create a rule with the Amazon Fraud Detector `CreateRule` API operation, you include the `frauddetector:CreateRule` action in the policy. Policy statements must include either an `Action` or `NotAction` element. Amazon Fraud Detector defines its own set of actions that describe tasks that you can perform with this service.

To specify multiple actions in a single statement, separate them with commas as follows:

```
"Action": [  
    "frauddetector:action1",  
    "frauddetector:action2"
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word `Describe`, include the following action:

```
"Action": "frauddetector:Describe*"
```

To see a list of Amazon Fraud Detector actions, see [Actions Defined by Amazon Fraud Detector](#) in the *IAM User Guide*.

Resources

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*" 
```

[Resource Types Defined by Amazon Fraud Detector](#) lists all Amazon Fraud Detector resource ARNs.

For example, to specify the my_detector detector in your statement, use the following ARN:

```
"Resource": "arn:aws:frauddetector:us-east-1:123456789012:detector/my_detector" 
```

For more information about the format of ARNs, see [Amazon Resource Names \(ARNs\) and AWS Service Namespaces](#).

To specify all detectors that belong to a specific account, use the wildcard (*):

```
"Resource": "arn:aws:frauddetector:us-east-1:123456789012:detector/*" 
```

Some Amazon Fraud Detector actions, such as those for creating resources, cannot be performed on a specific resource. In those cases, you must use the wildcard (*).

```
"Resource": "*" 
```

To see a list of Amazon Fraud Detector resource types and their ARNs, see [Resources Defined by Amazon Fraud Detector](#) in the *IAM User Guide*. To learn which actions you can specify the ARN of each resource, see [Actions Defined by Amazon Fraud Detector](#).

Condition keys

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element (or *Condition block*) lets you specify conditions in which a statement is in effect. The `Condition` element is optional. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple `Condition` elements in a statement, or multiple keys in a single `Condition` element, AWS evaluates them using a logical AND operation. If you specify multiple values for a single condition key, AWS evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see [IAM policy elements: variables and tags](#) in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

Amazon Fraud Detector defines its own set of condition keys and also supports using some global condition keys. To see all AWS global condition keys, see [AWS Global Condition Context Keys](#) in the *IAM User Guide*.

To see a list of Amazon Fraud Detector condition keys, see [Condition Keys for Amazon Fraud Detector](#) in the *IAM User Guide*. To learn which actions and resources you can use a condition key, see [Actions Defined by Amazon Fraud Detector](#).

Examples

To view examples of Amazon Fraud Detector identity-based policies, see [Amazon Fraud Detector identity-based policy examples](#).

Amazon Fraud Detector resource-based policies

Amazon Fraud Detector does not support resource-based policies.

Authorization Based on Amazon Fraud Detector Tags

You can attach tags to Amazon Fraud Detector resources or pass tags in a request to Amazon Fraud Detector. To control access based on tags, you provide tag information in the [condition element](#) of a policy using the `aws:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys.

Amazon Fraud Detector IAM roles

An [IAM role](#) is an entity within your AWS account that has specific permissions.

Using temporary credentials with Amazon Fraud Detector

You can use temporary credentials to sign in with federation, assume an IAM role, or to assume a cross-account role. You obtain temporary security credentials by calling AWS STS API operations such as [AssumeRole](#) or [GetFederationToken](#).

Amazon Fraud Detector supports using temporary credentials.

Service-linked roles

[Service-linked roles](#) allow AWS services to access resources in other services to complete an action on your behalf. Service-linked roles appear in your IAM account and are owned by the service. An IAM administrator can view but not edit the permissions for service-linked roles.

Amazon Fraud Detector does not support service-linked roles.

Service roles

This feature allows a service to assume a [service role](#) on your behalf. This role allows the service to access resources in other services to complete an action on your behalf. Service roles appear in your account and are owned by the account. This means that an administrator can change the permissions for this role. However, doing so might break the functionality of the service.

Amazon Fraud Detector supports service roles.

Amazon Fraud Detector identity-based policy examples

By default, users and IAM roles don't have permission to create or modify Amazon Fraud Detector resources. They also can't perform tasks using the AWS Management Console, AWS CLI, or AWS API. An administrator must create IAM policies that grant users and roles permission to perform

specific API operations on the specified resources they need. The administrator must then attach those policies to the users or groups that require those permissions.

To learn how to create an IAM identity-based policy using these example JSON policy documents, see [Creating Policies on the JSON Tab](#) in the *IAM User Guide*.

Topics

- [Policy best practices](#)
- [AWS-managed \(predefined\) policy for Amazon Fraud Detector](#)
- [Allow users to view their own permissions](#)
- [Allow full access to Amazon Fraud Detector resources](#)
- [Allow read-only access to Amazon Fraud Detector resources](#)
- [Allow access to a specific resource](#)
- [Allow access to specific resources when using dual mode API](#)
- [Limiting access based on tags](#)

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon Fraud Detector resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies that are specific to your use cases. For more information, see [AWS managed policies](#) or [AWS managed policies for job functions](#) in the *IAM User Guide*.
- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to

specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as AWS CloudFormation. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.

- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see [IAM Access Analyzer policy validation](#) in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see [Configuring MFA-protected API access](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

AWS-managed (predefined) policy for Amazon Fraud Detector

AWS addresses many common use cases by providing standalone IAM policies that are created and administered by AWS. These AWS managed policies grant necessary permissions for common use cases so that you can avoid having to investigate which permissions are needed. For more information, see [AWS Managed Policies](#) in the *AWS Identity and Access Management Management User Guide*.

The following AWS managed policy, which you can attach to users in your account, is specific to Amazon Fraud Detector:

`AmazonFraudDetectorFullAccess`: Grants full access to Amazon Fraud Detector resources, actions and the supported operations including:

- List and describe all model endpoints in Amazon SageMaker
- List all IAM roles in the account
- List all Amazon S3 buckets
- Allow IAM Pass Role to pass a role to Amazon Fraud Detector

This policy does not provide unrestricted S3 access. If you need to upload model training datasets to S3, the `AmazonS3FullAccess` managed policy (or scoped-down custom Amazon S3 access policy) is also required.

You can review the policy's permissions by signing in to the IAM console and searching by the policy name. You can also create your own custom IAM policies to allow permissions for Amazon Fraud Detector actions and resources as you need them. You can attach these custom policies to the users or groups that require them.

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
      "Effect": "Allow",
      "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsWithUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
      ],
      "Resource": ["arn:aws:iam::*:user/${aws:username}"]
    },
    {
      "Sid": "NavigateInConsole",
      "Effect": "Allow",
      "Action": [
        "iam:GetGroupPolicy",
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedGroupPolicies",
        "iam:ListGroupPolicies",
        "iam:ListPolicyVersions",
        "iam:ListPolicies",
        "iam:ListUsers"
      ]
    }
  ]
}
```

```
    ],
    "Resource": "*"
  }
]
}
```

Allow full access to Amazon Fraud Detector resources

The following example gives a user in your AWS account full access to all Amazon Fraud Detector resources and actions.

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

Allow read-only access to Amazon Fraud Detector resources

In this example, you grant a user in your AWS account read-only access to your Amazon Fraud Detector resources.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "frauddetector:GetEventTypes",
        "frauddetector:BatchGetVariable",
        "frauddetector:DescribeDetector",
        "frauddetector:GetModelVersion",
        "frauddetector:GetEventPrediction",
        "frauddetector:GetExternalModels",
        "frauddetector:GetLabels",

```



```

        "frauddetector:GetVariables",
        "frauddetector:GetDetectors",
        "frauddetector:GetRules",
        "frauddetector:ListTagsForResource",
        "frauddetector:GetKMSEncryptionKey",
        "frauddetector:DescribeModelVersions",
        "frauddetector:GetDetectorVersion",
        "frauddetector:GetPrediction",
        "frauddetector:GetOutcomes",
        "frauddetector:GetEntityTypes",
        "frauddetector:GetModels"
    ],
    "Resource": "*"
}
]
}

```

Allow access to a specific resource

In this example of a resource-level policy, you grant an user in your AWS account access to all actions and resources except for one particular Detector resource.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "frauddetector:*"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": [
        "frauddetector:*Detector"
      ],
      "Resource": "arn:${Partition}:frauddetector:${Region}:${Account}:detector/
${detector-name}"
    }
  ]
}

```

Allow access to specific resources when using dual mode API

Amazon Fraud Detector provides dual mode get APIs that work as both List and Describe operation. A dual mode API when called without any parameters returns a list of the specified resource associated with your AWS account. A dual mode API when called with parameter returns the details of the specified resource. The resource can be models, variables, event types, or entity types.

The dual mode APIs support resource-level permissions in IAM policies. However, the resource level permissions are only applied when one or more parameters are provided as part of the request. For example, if user calls [GetVariables](#) API and provides a variable name and if there is an IAM Deny policy attached to the variable resource or the variable name, user will receive `AccessDeniedException` error. If user calls `GetVariables` API and does not specify a variable name, all variables are returned, which can cause information leak.

To allow users to view details of specific resources only, use an IAM `NotResource` policy element in an IAM Deny policy. After you add this policy element to an IAM Deny policy, users can only view the details of the resources that are specified in the `NotResource` block. For more information, see [IAM JSON policy elements: NotResource](#) in the *IAM User Guide*.

The following example policy allows users to access all of Amazon Fraud Detector's resources. However, the `NotResource` policy element is used to limit [GetVariables](#) API calls to only the variable names with the prefixes `user*`, `job_*`, and `var*`.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "frauddetector:*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": "frauddetector:GetVariables",
      "NotResource": [
        "arn:aws:frauddetector:*:*:variable/user*",
        "arn:aws:frauddetector:*:*:variable/job_*",
        "arn:aws:frauddetector:*:*:variable/var*"
      ]
    }
  ]
}
```

```
]
}
```

Response

For this example policy, the response exhibit the following behavior:

- A `GetVariables` call that doesn't include variable names results in an `AccessDeniedException` error because the request maps to the `Deny` statement.
- A `GetVariables` call that includes a variable name that's not allowed, results in an `AccessDeniedException` error because the variable name doesn't map to the variable name in the `NotResource` block. For example, a `GetVariables` call with a variable name `email_address` results in an `AccessDeniedException` error.
- A `GetVariables` call that includes a variable name that matches a variable name in the `NotResource` block is returned as expected. For example, a `GetVariables` call that includes variable name `job_cpa` returns the details of `job_cpa` variable.

Limiting access based on tags

This example policy demonstrates how to limit access to Amazon Fraud Detector based on resource tags. This example assumes that:

- In your AWS account you have defined two different groups, named `Team1` and `Team2`
- You have created four detectors
- You want to allow members of `Team1` to make API calls on 2 detectors
- You want to allow members of `Team2` to make API calls on the other 2 detectors

To control access to API calls (example)

1. Add a tag with the key `Project` and value `A` to the detectors used by `Team1`.
2. Add a tag with the key `Project` and value `B` to the detectors used by `Team2`.
3. Create an IAM policy with a `ResourceTag` condition that denies access to detectors that have tags with key `Project` and value `B`, and attach that policy to `Team1`.
4. Create an IAM policy with a `ResourceTag` condition that denies access to detectors that have tags with key `Project` and value `A`, and attach that policy to `Team2`.

The following is an example of a policy that denies specific actions on any Amazon Fraud Detector resource that has a tag with a key of `Project` and a value of `B`:

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

      "Action": [

        "frauddetector:CreateModel",
        "frauddetector:CancelBatchPredictionJob",
        "frauddetector:CreateBatchPredictionJob",
        "frauddetector>DeleteBatchPredictionJob",
        "frauddetector>DeleteDetector"
      ],

      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/Project": "B"
        }
      }
    }
  ]
}
```

Confused deputy prevention

The confused deputy problem occurs when an entity that doesn't have permission to perform an action can coerce a more-privileged entity to perform the action. AWS provides tools that help you protect your account if you provide third-parties (called *cross-account*) or other AWS services (called *cross-service*) access to resources in your account.

Cross-service confused deputy problem can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should not otherwise have permission to access. To prevent this, you can create policies that help you protect your data for all services with service principals that have been given access to your service's resources.

Amazon Fraud Detector supports using [service roles](#) in your permission policies to allow a service to access another service's resources on your behalf. A role requires two policies: a role trust policy that specifies the principal that is allowed to assume the role and a permissions policy that specifies what can be done with the role. When a service assumes a role on your behalf, the service principal must be allowed to perform the `sts:AssumeRole` action in the role trust policy. When a service calls `sts:AssumeRole`, AWS STS returns a set of temporary security credentials that the service principal uses to access the resources that are permitted by the role's permissions policy.

To prevent cross-service confused deputy problem, Amazon Fraud Detector recommends using the [aws:SourceArn](#) and [aws:SourceAccount](#) global condition context keys in your role trust policy to limit access to the role to only those requests that are generated by expected resources.

The `aws:SourceAccount` specifies the Account ID and the `aws:SourceArn` specifies the ARN of the resource associated with cross-service access. The `aws:SourceArn` must be specified using the [ARN format](#). Ensure that both `aws:SourceAccount` and `aws:SourceArn` are using the same Account ID when used in the same policy statement.

The most effective way to protect against the confused deputy problem is to use the `aws:SourceArn` global condition context key with the full ARN of the resource. If you don't know the full ARN of the resource or if you are specifying multiple resources, use the `aws:SourceArn` global context condition key with a wildcard (*) for the unknown portions of the ARN. For example, `arn:aws:servicename:*:123456789012:*`. For information about Amazon Fraud Detector resources and actions you can use in your permission policies, see [Actions, resources, and condition keys for Amazon Fraud Detector](#).

The following role trust policy example uses wildcard(*) in the `aws:SourceArn` condition key to allow Amazon Fraud Detector access to multiple resources associated with the Account Id.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
```

```

        "Service": [
            "frauddetector.amazonaws.com"
        ],
        "Action": "sts:AssumeRole",
        "Condition": {
            "StringEquals": {
                "aws:SourceAccount": "123456789012"
            },
            "StringLike": {
                "aws:SourceArn": "arn:aws:frauddetector:us-west-2:123456789012:*"
            }
        }
    }
}

```

The following role trust policy allows Amazon Fraud Detector access to only external-model resource. Notice the `aws:SourceArn` param in Condition block. The resource qualifier is built using the model endpoint that is provided for making the `PutExternalModel` API call.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "frauddetector.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "123456789012"
        },
        "StringLike": {
          "aws:SourceArn": "arn:aws:frauddetector:us-west-2:123456789012:external-
model/MyExternalModeldoNotDelete-ReadOnly"
        }
      }
    }
  ]
}

```

```
]
}
```

Troubleshooting Amazon Fraud Detector identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon Fraud Detector and IAM.

Topics

- [I am not authorized to perform an action in Amazon Fraud Detector](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my AWS account to access my Amazon Fraud Detector resources](#)
- [Amazon Fraud Detector could not assume the given role](#)

I am not authorized to perform an action in Amazon Fraud Detector

If the AWS Management Console tells you that you're not authorized to perform an action, then you must contact your administrator for assistance. Your administrator is the person that provided you with your sign-in credentials.

The following example error occurs when the `mateojackson` user tries to use the console to view details about a *detector* but does not have `frauddetector:GetDetectors` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
frauddetector:GetDetectors on resource: my-example-detector
```

In this case, Mateo asks his administrator to update his policies to allow him to access the *my-example-detector* resource using the `frauddetector:GetDetectors` action.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to Amazon Fraud Detector.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in Amazon Fraud Detector. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my Amazon Fraud Detector resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon Fraud Detector supports these features, see [How Amazon Fraud Detector works with IAM](#).
- To learn how to provide access to your resources across AWS accounts that you own, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see [Providing access to AWS accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Amazon Fraud Detector could not assume the given role

If you receive an error that Amazon Fraud Detector could not assume the given role, then you must update the trust relationship for the specified role. By specifying Amazon Fraud Detector as

a trusted entity, the service can assume the role. When you use Amazon Fraud Detector to create a role, this trust relationship is automatically set. You only need to establish this trust relationship for IAM roles that are not created by Amazon Fraud Detector.

To establish a trust relationship for an existing role to Amazon Fraud Detector

1. Open the IAM console at <https://console.aws.amazon.com/iam/>
2. In the navigation pane choose **Roles**.
3. Choose the name of the role that you want to modify, and choose the **Trust relationships** tab.
4. Choose **Edit trust relationship**.
5. Under **Policy Document**, paste the following, and then choose **Update Trust Policy**.

```
{
  "Version": "2012-10-17",
  "Statement": [ {
    "Effect": "Allow",
    "Principal": {
      "Service": "frauddetector.amazonaws.com"
    },
    "Action": "sts:AssumeRole"
  } ]
}
```

Logging and monitoring in Amazon Fraud Detector

AWS provides the following monitoring tools to watch Amazon Fraud Detector, report when something is wrong, and take automatic actions when appropriate:

- *Amazon CloudWatch* monitors your AWS resources and the applications you run on AWS in real time. For more information on CloudWatch, see the [Amazon CloudWatch User Guide](#).
- *AWS CloudTrail* captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon S3 bucket that you specify. For more information on CloudTrail, see the [AWS CloudTrail User Guide](#).

For more information on monitoring Amazon Fraud Detector, see [Monitor Amazon Fraud Detector](#).

Compliance validation for Amazon Fraud Detector

Third-party auditors assess the security and compliance of AWS services as part of multiple AWS compliance programs, such as SOC, PCI, FedRAMP, and HIPAA.

To learn whether an AWS service is within the scope of specific compliance programs, see [AWS services in Scope by Compliance Program](#) and choose the compliance program that you are interested in. For general information, see [AWS Compliance Programs](#).

You can download third-party audit reports using AWS Artifact. For more information, see [Downloading Reports in AWS Artifact](#).

Your compliance responsibility when using AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- [Security and Compliance Quick Start Guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying baseline environments on AWS that are security and compliance focused.
- [Architecting for HIPAA Security and Compliance on Amazon Web Services](#) – This whitepaper describes how companies can use AWS to create HIPAA-eligible applications.

Note

Not all AWS services are HIPAA eligible. For more information, see the [HIPAA Eligible Services Reference](#).

- [AWS Compliance Resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [AWS Customer Compliance Guides](#) – Understand the shared responsibility model through the lens of compliance. The guides summarize the best practices for securing AWS services and map the guidance to security controls across multiple frameworks (including National Institute of Standards and Technology (NIST), Payment Card Industry Security Standards Council (PCI), and International Organization for Standardization (ISO)).

- [Evaluating Resources with Rules](#) in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [AWS Security Hub](#) – This AWS service provides a comprehensive view of your security state within AWS. Security Hub uses security controls to evaluate your AWS resources and to check your compliance against security industry standards and best practices. For a list of supported services and controls, see [Security Hub controls reference](#).
- [Amazon GuardDuty](#) – This AWS service detects potential threats to your AWS accounts, workloads, containers, and data by monitoring your environment for suspicious and malicious activities. GuardDuty can help you address various compliance requirements, like PCI DSS, by meeting intrusion detection requirements mandated by certain compliance frameworks.
- [AWS Audit Manager](#) – This AWS service helps you continuously audit your AWS usage to simplify how you manage risk and compliance with regulations and industry standards.

Resilience in Amazon Fraud Detector

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).

Infrastructure Security in Amazon Fraud Detector

As a managed service, Amazon Fraud Detector is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see [AWS Cloud Security](#). To design your AWS environment using the best practices for infrastructure security, see [Infrastructure Protection](#) in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access Amazon Fraud Detector through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.

- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the [AWS Security Token Service](#) (AWS STS) to generate temporary security credentials to sign requests.

Monitor Amazon Fraud Detector

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon Fraud Detector and your other AWS solutions. AWS provides the following monitoring tools to watch Amazon Fraud Detector, report when something is wrong, and take automatic actions when appropriate:

- *Amazon CloudWatch* monitors your AWS resources and the applications you run on AWS in real time. You can collect and track metrics, create customized dashboards, and set alarms that notify you or take actions when a specified metric reaches a threshold that you specify. For more information, see the [Amazon CloudWatch User Guide](#).
- *AWS CloudTrail* captures API calls and related events made by or on behalf of your AWS account and delivers the log files to an Amazon S3 bucket that you specify. You can identify which users and accounts called AWS, the source IP address from which the calls were made, and when the calls occurred. For more information, see the [AWS CloudTrail User Guide](#).

Topics

- [Monitoring Amazon Fraud Detector with Amazon CloudWatch](#)
- [Logging Amazon Fraud Detector API Calls with AWS CloudTrail](#)

Monitoring Amazon Fraud Detector with Amazon CloudWatch

You can monitor Amazon Fraud Detector using CloudWatch, which collects raw data and processes it into readable, near real-time metrics. These statistics are kept for 15 months, so that you can access historical information and gain a better perspective on how your web application or service is performing. You can also set alarms that watch for certain thresholds, and send notifications or take actions when those thresholds are met. For more information, see the [Amazon CloudWatch User Guide](#).

Topics

- [Using CloudWatch Metrics for Amazon Fraud Detector](#).
- [Amazon Fraud Detector Metrics](#)

Using CloudWatch Metrics for Amazon Fraud Detector.

To use metrics, you must specify the following information:

- The metric namespace. A *namespace* is a CloudWatch container Amazon Fraud Detector uses to publish its metrics into. If you are using the CloudWatch [ListMetrics](#) API or the [list-metrics](#) command to view the metrics for Amazon Fraud Detector, specify `AWS/FraudDetector` for the namespace.
- The metric dimension. A *dimension* is a name-value pair that helps you to uniquely identify a metric, for example, `DetectorId` can be a dimension name. Specifying a metric dimension is optional.
- The metric name, such as `GetEventPrediction`.

You can get monitoring data for Amazon Fraud Detector by using the AWS Management Console, the AWS CLI, or the CloudWatch API. You can also use the CloudWatch API through one of the Amazon AWS Software Development Kits (SDKs) or the CloudWatch API tools. The console displays a series of graphs based on the raw data from the CloudWatch API. Depending on your needs, you might prefer to use either the graphs displayed in the console or retrieved from the API.

The following list shows some common uses for the metrics. These are suggestions to get you started, not a comprehensive list.

| How Do I? | Relevant Metrics |
|---|--|
| How do I track the number of predictions that have been performed? | Monitor the <code>GetEventPrediction</code> metric. |
| How can I monitor <code>GetEventPrediction</code> errors? | Use the <code>GetEventPrediction5xxError</code> and the <code>GetEventPrediction4xxError</code> metrics. |
| How can I monitor the latency of <code>GetEventPrediction</code> calls? | Use the <code>GetEventPredictionLatency</code> metric. |

You must have the appropriate CloudWatch permissions to monitor Amazon Fraud Detector with CloudWatch. For more information, see [Authentication and Access Control for Amazon CloudWatch](#).

Access Amazon Fraud Detector Metrics

The following steps show how to access Amazon Fraud Detector metrics using the CloudWatch console.

To view metrics (console)

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. Choose **Metrics**, choose the **All Metrics** tab, and then choose **Fraud Detector**.
3. Choose the metric dimension.
4. Choose the desired metric from the list, and choose a time period for the graph.

Create an Alarm

You can create a CloudWatch alarm that sends an Amazon Simple Notification Service (Amazon SNS) message when the alarm changes state. An alarm watches a single metric over a time period that you specify. It performs one or more actions based on the value of the metric relative to a given threshold over a number of time periods. The action is a notification sent to an Amazon SNS topic or an Auto Scaling policy.

Alarms invoke actions for sustained state changes only. CloudWatch alarms don't invoke actions simply because they are in a particular state. The state must have changed and have been maintained for a specified number of time periods.

To set an alarm (console)

1. Sign in to the AWS Management Console and open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, choose **Alarms**, and choose **Create Alarm**. This opens the **Create Alarm Wizard**.
3. Choose **Select metric**.
4. In the **All metrics** tab, choose **Fraud Detector**.

5. Choose **By Detector ID**, and then choose the **GetEventPrediction** metric.
6. Choose the **Graphed metrics** tab.
7. For **Statistic**, choose **Sum**.
8. Choose **Select metric**.
9. For **Conditions**, choose **Static** for **Threshold type** and **Greater** for **Whenever...**, and then enter a maximum value of your choice. Choose **Next**.
10. To send alarms to an existing Amazon SNS topic, for **Send notification to:**, choose an existing SNS topic. To set the name and email addresses for a new email subscription list, choose **New list**. CloudWatch saves the list and displays it in the field so you can use it to set future alarms.

Note

If you use **New list** to create a new Amazon SNS topic, the email addresses must be verified before the intended recipients receive notifications. Amazon SNS sends email only when the alarm enters an alarm state. If this alarm state change happens before the email addresses are verified, intended recipients don't receive a notification.

11. Choose **Next**. Add a name and optional description for your alarm. Choose **Next**.
12. Choose **Create Alarm**.

Amazon Fraud Detector Metrics

Amazon Fraud Detector sends the following metrics to CloudWatch. All metrics support these statistics: Average, Minimum, Maximum, Sum.

| Metric | Description |
|---------------------------|--|
| GetEventPrediction | The number of GetEventPrediction API requests. Valid Dimensions: DetectorID |
| GetEventPredictionLatency | The interval of time taken to respond to a client request from the GetEventPrediction request. Valid Dimensions: DetectorID Unit: Milliseconds |

| Metric | Description |
|----------------------------|---|
| GetEventPrediction4XXError | <p>The number of GetEventPrediction requests where Amazon Fraud Detector returned a 4xx HTTP response code. For each 4xx response, 1 is sent.</p> <p>Valid Dimensions: DetectorID</p> |
| GetEventPrediction5XXError | <p>The number of GetEventPrediction requests where Amazon Fraud Detector returned a 5xx HTTP response code. For each 5xx response, 1 is sent.</p> <p>Valid Dimensions: DetectorID</p> |
| Prediction | <p>The number of predictions. 1 is sent if successful.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID</p> |
| PredictionLatency | <p>The interval of time taken for a prediction operation .</p> <p>Valid Dimensions: DetectorID , DetectorVersionID</p> <p>Unit: Milliseconds</p> |
| PredictionError | <p>The number of predictions where Amazon Fraud Detector encountered an error. 1 is sent if an error is encountered.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID</p> |
| VariableUsed | <p>The number of GetEventPrediction requests where the variable was used as part of the evaluation.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , VariableName</p> |

| Metric | Description |
|-------------------------|---|
| VariableDefaultReturned | <p>The number of GetEventPrediction requests where the variable was not present as part of the Event Attributes and therefore the default value for the variable was used during evaluation.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , VariableName</p> |
| RuleNotEvaluated | <p>The number of GetEventPrediction requests where the rule was not evaluated because a prior rule matched.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , RuleID</p> |
| RuleEvaluateTrue | <p>The number of GetEventPrediction requests where the rule triggered as True and the rule outcome was returned.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , RuleID</p> |
| RuleEvaluateFalse | <p>The number of GetEventPrediction requests where the rule evaluated to False.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , RuleID</p> |
| RuleEvaluateError | <p>The number of GetEventPrediction requests where the rule evaluates in error</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , RuleID</p> |

| Metric | Description |
|--|--|
| OutcomeReturned | <p>The number of GetEventPrediction calls where the specified outcome was returned.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , OutcomeName</p> |
| ModelInvocation (Amazon SageMaker model endpoint) | <p>The number of GetEventPrediction requests where the SageMaker model endpoint was invoked as part of the evaluation.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , ModelEndpoint</p> |
| ModelInvocationError (Amazon SageMaker model endpoint) | <p>The number of GetEventPrediction requests where the invoked SageMaker model endpoint returned an error during evaluation.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , ModelEndpoint</p> |
| ModelInvocationLatency (Amazon SageMaker model endpoint) | <p>The interval of time taken by the Imported Model to respond as viewed from Amazon Fraud Detector. This interval includes only the model invocation.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , ModelEndpoint</p> <p>Unit: Milliseconds</p> |
| ModelInvocation | <p>The number of GetEventPrediction requests where the model was invoked as part of the evaluation.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , ModelType , ModelID</p> |

| Metric | Description |
|------------------------|---|
| ModelInvocationError | <p>The number of GetEventPrediction requests where the Amazon Fraud Detector model returned an error during evaluation.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , ModelType , ModelID</p> |
| ModelInvocationLatency | <p>The interval of time taken by the Amazon Fraud Detector Model to respond as viewed from Amazon Fraud Detector. This interval includes only the model invocation.</p> <p>Valid Dimensions: DetectorID , DetectorVersionID , ModelType , ModelID</p> <p>Unit: Milliseconds</p> |

Logging Amazon Fraud Detector API Calls with AWS CloudTrail

Amazon Fraud Detector is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Amazon Fraud Detector. CloudTrail captures all API calls for Amazon Fraud Detector as events, including calls from the Amazon Fraud Detector console and calls from code to the Amazon Fraud Detector APIs.

If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon Fraud Detector. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Amazon Fraud Detector, the IP address that the request was made from, who made the request, when it was made, and additional details.

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

Amazon Fraud Detector Information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in Amazon Fraud Detector, that activity is recorded in a CloudTrail event along with other AWS service

events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing Events with CloudTrail Event History](#).

For an ongoing record of events in your AWS account, including events for Amazon Fraud Detector, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data that's collected in CloudTrail logs. For more information, see the following:

- [Overview for Creating a Trail](#)
- [CloudTrail Supported Services and Integrations](#)
- [Configuring Amazon SNS Notifications for CloudTrail](#)
- [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#)

Amazon Fraud Detector supports logging every action (API operation) as an event in CloudTrail log files. For more information, see [Actions](#).

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity Element](#).

Understanding Amazon Fraud Detector Log File Entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested operation, the date and time of the operation, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `GetDetectors` operation.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "principal-id",
    "arn": "arn:aws:iam::user-arn",
    "accountId": "account-id",
    "accessKeyId": "access-key",
    "userName": "user-name"
  },
  "eventTime": "2019-11-22T02:18:03Z",
  "eventSource": "frauddetector.amazonaws.com",
  "eventName": "GetDetectors",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "source-ip-address",
  "userAgent": "aws-cli/1.11.16 Python/2.7.11 Darwin/15.6.0 boto3/1.4.73",
  "requestParameters": null,
  "responseElements": null,
  "requestID": "request-id",
  "eventID": "event-id",
  "eventType": "AwsApiCall",
  "recipientAccountId": "recipient-account-id"
}
```




Troubleshoot

The following sections help you troubleshoot issues that you might encounter when working with Amazon Fraud Detector

Troubleshoot training data issues

Use information in this section to help diagnose and resolve issues you might see in the **Model training diagnostic** pane in the Amazon Fraud Detector console when you train your model.

The issues displayed in the **Model training diagnostic** pane are categorized as follows. The requirement to address the issue is dependent on the category of the issue.

-  **Error**- causes the model training to fail. These issues must be addressed for the model to train successfully.
-  **Warning**- causes the model training to continue, however, some of the variables might be getting excluded in the training process. Check for the relevant guidance in this section to improve the quality of your dataset.
-  **Information (Info)**- has no impact on model training and all the variables are used for training. We recommend that you check the relevant guidance in this section to further improve the quality of your dataset and model performance.

Topics

- [Unstable fraud rate in the given dataset](#)
- [Insufficient data](#)
- [Missing or different EVENT_LABEL values](#)
- [Missing or incorrect EVENT_TIMESTAMP values](#)
- [Data not ingested](#)
- [Insufficient variables](#)
- [Missing or incorrect variable type](#)
- [Missing variable values](#)

- [Insufficient unique variable values](#)
- [Incorrect variable expression](#)
- [Insufficient unique entities](#)

Unstable fraud rate in the given dataset

Issue type : Error

Description

Fraud rate in the given data is too unstable through time. Please make sure your fraud and legitimate events are sampled uniformly over time.

Cause

This error occurs if the fraud and legitimate events in your dataset are distributed unevenly and are taken from different time slots. Amazon Fraud Detector model training process samples and partitions your dataset based on `EVENT_TIMESTAMP`. For example, if your dataset consists of fraud events pulled from last 6 months, but only the last month of legitimate events are included, the dataset is considered unstable. An unstable dataset might lead to biases in model performance evaluation.

Solution

Make sure to provide the fraudulent and legitimate events data from same time slot and the fraud rate does not change dramatically over time.

Insufficient data

1. Issue type : Error

Description

Fewer than 50 rows are labeled as fraudulent events. Ensure that both fraudulent and legitimate events exceed the minimum count of 50 and re-train the model.

Cause

This error occurs if your dataset has fewer events labeled as fraudulent than required for model training. Amazon Fraud Detector requires at least 50 fraudulent events to train your model.

Solution

Make sure that your dataset includes a minimum of 50 fraudulent events. You can ensure this by covering a longer time period, if needed.

2. Issue type : Error

Description

Fewer than 50 rows are labeled as legitimate events. Ensure that both fraudulent and legitimate events exceed the minimum count of \$threshold and re-train the model.

Cause

This error occurs if your dataset has fewer events labeled as legitimate than required for model training. Amazon Fraud Detector requires at least 50 legitimate events to train your model.

Solution

Make sure that your dataset includes a minimum of 50 legitimate events. You can ensure this by covering a longer time period, if needed.

3. Issue type : Error

Description

The number of unique entities associated with fraud is less than 100. Consider including more examples of fraudulent entities to improve performance.

Cause

This error occurs if your dataset has fewer entities with fraudulent events than required for model training. The Transaction Fraud Insights (TFI) model requires at least 100 entities with fraud events to ensure maximum coverage of the fraud space. The model may not generalize well if all fraud events are performed by a small group of entities.

Solution

Make sure that your dataset includes at least 100 entities with fraudulent events. You can ensure this by covering a longer time period, if needed.

4. Issue type : Error

Description

The number of unique entities associated with legitimate is less than 100. Consider including more examples of legitimate entities to improve performance.

Cause

This error occurs if your dataset has fewer entities with legitimate events than required for model training. The Transaction Fraud Insights (TFI) model requires at least 100 entities with legitimate events to ensure maximum coverage of the fraud space. The model may not generalize well if all legitimate events are performed by a small group of entities.

Solution

Make sure that your dataset includes at least 100 entities with legitimate events. You can ensure this by covering a longer time period, if needed.

5. Issue type : Error

Description

Less than 100 rows are in the dataset. Ensure there are more than 100 rows in the total dataset and at least 50 rows are labeled as fraudulent.

Cause

This error occurs if your dataset contains fewer than 100 records. Amazon Fraud Detector requires data from at least 100 events (records) in your dataset for model training.

Solution

Make sure that you have data from more than 100 events in your dataset.

Missing or different EVENT_LABEL values

1. Issue type : Error

Description

Greater than 1% of your EVENT_LABEL column are null or are values other than those defined in the model configuration **\$label_values**. Ensure you have less than 1% of missing values

in your `EVENT_LABEL` column and the values are those defined in the model configuration `$label_values`.

Cause

This error occurs because of one of the following reasons:

- More than 1% of the records in the CSV file containing your training data have missing values in the `EVENT_LABEL` column.
- More than 1% of the records in the CSV file containing your training data have values in the `EVENT_LABEL` column that are different than those associated with your event type.

Online Fraud Insights (OFI) model requires that the `EVENT_LABEL` column in each record be populated with one of the labels that's associated with your event type (or, mapped in `CreateModelVersion`).

Solution

If this error is due to the missing `EVENT_LABEL` values, consider assigning proper labels to those records or dropping those records from your dataset. If this error is because labels of some records are not among `label_values`, make sure to add all the values in `EVENT_LABEL` column to labels of the event type and mapped to either fraudulent or legitimate (fraud, legit) in model creation.

2. Issue type : Information

Description

Your `EVENT_LABEL` column contains null values or label values other than those defined in the model configuration `$label_values`. These inconsistent values were converted to 'not fraud' prior to training.

Cause

You get this information because of one of the following reasons:

- Less than 1% of the records in the CSV file containing your training data have missing values in the `EVENT_LABEL` column
- Less than 1% of the records in the CSV file containing your training data have values in the `EVENT_LABEL` column that are different than those associated with your event type.

The model training in both the cases will succeed. However, the label values of those events that have missing or unmapped label values are converted to legitimate. If you consider this to be an issue, follow solution provided below.

Solution

If there are missing EVENT_LABEL values in your dataset, consider dropping those records from your dataset. If the values provided for those EVENT_LABELS are not mapped, make sure that all those values are mapped to either fraudulent or legitimate (fraud, legit) for each event.

Missing or incorrect EVENT_TIMESTAMP values

1. Issue type : Error

Description

Your training data set contains EVENT_TIMESTAMP with timestamps that do not conform to accepted formats. Ensure the format is one of the accepted date/timestamp formats.

Cause

This error occurs if the EVENT_TIMESTAMP column contains value that doesn't comply with the [timestamp formats](#) that are supported by Amazon Fraud Detector.

Solution

Ensure that the values provided for the EVENT_TIMESTAMP column is compliant with the supported [timestamp formats](#). If you have missing values in the EVENT_TIMESTAMP column, you can either backfill those with values using the supported timestamp format or consider dropping the event completely instead of entering strings such as none, null, or missing.

2. Issue type : Error

Your training data set contains EVENT_TIMESTAMP with missing values. Ensure you have no missing values.

Cause

This error occurs if the EVENT_TIMESTAMP column in your dataset has missing values. Amazon Fraud Detector requires that the EVENT_TIMESTAMP column in your dataset have values.

Solution

Ensure that the `EVENT_TIMESTAMP` column in your dataset has values and those values are compliant with the supported [timestamp formats](#). If you have missing values in the `EVENT_TIMESTAMP` column, you can either backfill those with values using the supported timestamp format or consider dropping the event completely instead of entering strings such as `none`, `null`, or `missing`.

Data not ingested

Issue type : Error

Description

No ingested events found for training, please check your training configuration.

Cause

This error occurs if you are creating a model with event data stored with Amazon Fraud Detector but did not import your dataset to Amazon Fraud Detector before you started to train your model.

Solution

Use the `SendEvent` API operation, the `CreateBatchImportJob` API operation, or batch import feature in the Amazon Fraud Detector console, to first import your event data and then train your model. See [Stored event datasets](#) for more information.

Note

We recommend waiting 10 minutes after you have finished importing your data before using it to train your model.

You can use Amazon Fraud Detector console to check number of events already stored for each event type. See [Viewing metrics of your stored events](#) for more information.

Insufficient variables

Issue type : Error

Description

Dataset must contain at least 2 variables suitable for training.

Cause

This error occurs if your dataset contains less than 2 variables that are suitable for model training. Amazon Fraud Detector considers a variable suitable for model training only if it passes all validations. If a variable fails validation, it is excluded in model training and you will see a message in **Model training diagnostic**.

Solution

Ensure that your dataset has at least two variables populated with values and passed all data validations. Note that the event metadata row where you have provided your column headers (EVENT_TIMESTAMP, EVENT_ID, ENTITY_ID, EVENT_LABEL, etc.) aren't considered as variable.

Missing or incorrect variable type

Issue type : Warning

Description

The expected data type for **\$variable_name** is NUMERIC. Review and update **\$variable_name** in your dataset and re-train the model.

Cause

You get this warning if a variable is defined as a NUMERIC variable, but in the dataset, it has values that can't be converted to NUMERIC. As a result, that variable is excluded in model training.

Solution

If you want to keep it as a NUMERIC variable, make sure that values you provide can be converted to float number. Note that if the variable contains missing values, don't fill them with strings such as nonene, null, or missing. If the variable does contain non-numeric values, re-create it as a CATEGORICAL or FREE_FORM_TEXT variable type.

Missing variable values

Issue type : Warning

Description

Greater than **\$threshold** values for **\$variable_name** are missing from your training dataset. Consider modifying **\$variable_name** in your dataset and re-training to improve performance.

Cause

You get this warning if the specified variable is being dropped due to too many missing values. Amazon Fraud Detector allows missing values for a variable. However, if one variable has too many missing values, it doesn't contribute much to the model and that variable is dropped in model training.

Solution

First, verify that those missing values aren't due to mistakes in data collection and preparation. If they are mistakes, then you can consider dropping them from your model training. However, if you do believe those missing values are valuable and still want to keep that variable, you can manually fill missing values with a constant in both model training and real-time inference.

Insufficient unique variable values

Issue type : Warning

Description

The count of unique values of **\$variable_name** is lower than 100. Review and update **\$variable_name** in your dataset and re-train the model.

Cause

You get this warning if the number of unique values of the specified variable is less than the 100. The thresholds differ depending on the variable type. With very few unique values, there's a risk that the dataset isn't general enough to cover the feature space of that variable. As a result, the model might not generalize well on real-time predictions.

Solution

First, make sure the variable distribution is representative of the real business traffic. Then, you can either adopt more fine-trained variables with higher cardinality, such as using `full_customer_name` instead of `first_name` and `last_name` separately or change the variable type to CATEGORICAL, which allows lower cardinality.

Incorrect variable expression

1. Issue type : Information

Description

Greater than 50% of **\$email_variable_name** values do not match the expected regular expression `http://emailregex.com`. Consider modifying **\$email_variable_name** in your dataset and re-training to improve performance.

Cause

This information is displayed if more than 50% records in your dataset has email values that do not comply with a regular email expression and are therefore failing validation.

Solution

Format the email variable values to comply with the regular expression. If there are missing email values, we recommend to leave them empty instead of filling them with strings such as `none`, `null`, or `missing`.

2. Issue type : Information

Description

Greater than 50% of **\$IP_variable_name** values do not match regular expression for IPv4 or IPv6 addresses `https://digitalfortress.tech/tricks/top-15-commonly-used-regex/`. Consider modifying **\$IP_variable_name** in your dataset and re-training to improve performance.

Cause

This information is displayed if more than 50% records in your dataset has IP values that do not comply with a regular IP expression and are therefore failing validation.

Solution

Format the IP values to comply with the regular expression. If there are missing IP values, we recommend to leave them empty instead of filling them with strings such as `none`, `null`, or `missing`.

3. Issue type : Information

Description

Greater than 50% of `$phone_variable_name` values do not match basic phone regular expression `/$pattern/`. Consider modifying `$phone_variable_name` in your dataset and re-training to improve performance .

Cause

This information is displayed if more than 50% records in your dataset has phone numbers that do not comply with a regular phone number expression and are therefore failing validation.

Solution

Format the phone numbers to comply with the regular expression. If there are missing phone numbers, we recommend to leave them empty instead of filling them with strings such as `none`, `null`, or `missing`.

Insufficient unique entities

Issue type : Information

Description

The number of unique entities is less than 1500. Consider including more data to improve performance.

Cause

This information is displayed if your dataset has a smaller number of unique entities than the recommended number. The Transaction Fraud Insights (TFI) model uses both time-series aggregates and generic transaction features to provide the best performance. If your dataset has too few unique entities, then most of your generic data such as `IP_ADDRESS`, `EMAIL_ADDRESS`, might not have unique values. Then, there's also a risk that this dataset isn't general enough to cover the feature space of that variable. As a result, the model might not generalize well on transactions from fresh new entities.

Solution

Include more entities. Extend your training data time range, if needed.

Quotas

Your AWS account has default quotas, formerly referred to as *limits*, for each Amazon Web Service. Unless otherwise noted, each quota is Region-specific. You can request a quota increase for all adjustable quotas mentioned in the tables below. For more information, see [Requesting a quota increase](#)

The following tables outline Amazon Fraud Detector quotas by component.

Amazon Fraud Detector models

| Quota name | Default quota | Adjustable |
|--------------------------------------|---------------|------------|
| Training data size | 5 GB | No |
| Models per account | 50 | No |
| Versions per model | 200 | No |
| Deployed model versions per account | 5 | No |
| Concurrent training jobs per account | 3 | No |
| Concurrent training jobs per model | 1 | No |

Amazon Fraud Detector detectors / variables / outcomes / rules

| Quota name | Default quota | Adjustable |
|-----------------------|---------------|------------|
| Variables per account | 5000 | No |
| Rules per account | 5000 | No |

| Quota name | Default quota | Adjustable |
|-----------------------------|---------------|------------|
| Lists per rule | 3 | No |
| Outcomes per account | 5000 | No |
| Detectors per account | 100 | No |
| Lists per detector | 30 | No |
| Draft versions per detector | 100 | No |
| Models per detector version | 10 | No |
| Labels per account | 100 | No |
| Event types per account | 100 | No |
| Entity types per account | 100 | No |

Amazon Fraud Detector API

| Quota name | Default quota | Adjustable |
|--|---------------|------------|
| GetEventPrediction API calls per second | 200 TPS | Yes |
| Size of payload per GetEventPrediction API call | 256 KB | No |
| Number of inputs per GetEventPrediction API call | 5000 | No |

Document history

The following table describes important changes in Amazon Fraud Detector User Guide. We also update the Amazon Fraud Detector User Guide frequently to address the feedback that you send us.

| Change | Description | Date |
|---|--|-------------------|
| New variable and data types | Amazon Fraud Detector introduces new variable types and a datatype you can use to extract useful information. | June 5, 2023 |
| Event orchestration | The Event orchestration makes it easy for you to send events to AWS services for downstream processing, using Amazon EventBridge. | May 30, 2023 |
| Lists | The Lists resource enables you to reference a set of values such as IP addresses or email addresses, as part of a rule. Use lists in a rule to allow or deny access or a transaction. | February 14, 2023 |
| Data Models Explorer | The Data Models Explorer provides insights into the data elements required by Amazon Fraud Detector to create your fraud detection model. Use data models explorer before you prepare your event dataset. | December 15, 2022 |

| | | |
|---|---|------------------|
| Account Takeover Insights model | Use Account takeover insights (ATI) model to detect accounts that are compromised through malicious takeovers, phishing, or from credentials being stolen. | July 21, 2022 |
| Chapter update | Updated the introductory chapter with additional information about Amazon Fraud Detector | April 11, 2022 |
| Variable enrichment | Enable enrichment of some of the raw data you provide to boost performance for the models that use these data elements and that were trained before February 8, 2022. | February 8, 2022 |
| Opt-out policies | Use opt-out policies to opt out of having your event data used to develop or improve the quality of Amazon Fraud Detector. | January 6, 2022 |
| Confused deputy prevention | Create policies to prevent a third-party or a cross-service entity from manipulating an entity with permissions to act on its behalf to gain access to resources in your account. | December 6, 2021 |

| | | |
|--|--|-------------------|
| Create event dataset | Use the guidance provided in Create event dataset to prepare and gather data for training your model. | November 22, 2021 |
| Prediction explanations | Use Prediction explanations to get insight into how each event variable impacted your model's fraud prediction scores. | November 10, 2021 |
| Troubleshoot | Use information in Troubleshoot training data issues to help diagnose and resolve issues you might see in Amazon Fraud Detector console when you train your model. | October 11, 2021 |
| Transaction fraud insights model | Use Transaction fraud insights (TFI) model to detect online or card-not-present transaction fraud. | October 11, 2021 |
| Stored events | Store your event data in Amazon Fraud Detector and use the stored data to later train your models. By storing event data in Amazon Fraud Detector, you can train models that use auto-computed variables to improve performance, simplify model retraining, and update fraud labels to close the machine learning feedback loop. | October 11, 2021 |

| | | |
|---|---|------------------|
| Model variable importance | Use Model variable importance to gain insight into what is driving performance of your model up or down and which of your model variables contribute the most. And then tweak your model to improve overall performance. | July 9, 2021 |
| Integration with AWS CloudFormation | Use AWS CloudFormation to manage your Amazon Fraud Detector resources. | May 10, 2021 |
| Batch predictions | Use Batch predictions to get predictions for a set of events that do not require real-time scoring. | March 31, 2021 |
| Chapter rework | Rework of Get started and other sections | July 17, 2020 |
| Initial release | Initial release | December 2, 2019 |