

---

# Machine Learning for Telecommunication Implementation Guide



## **Machine Learning for Telecommunication: Implementation Guide**

Copyright © 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

Home .....	1
Overview .....	2
Cost .....	2
Architecture .....	2
Considerations .....	4
The Jupyter Notebook .....	4
AWS Glue .....	5
Demo Data .....	5
SageMaker Instance Types .....	5
Template .....	6
Deployment .....	7
What We'll Cover .....	7
Step 1. Launch the Stack .....	7
Step 2. Run the Notebooks .....	9
Security .....	11
Resources .....	12
Appendix A: Using Your Own Data .....	13
Appendix B: Operational Metrics .....	14
Source Code .....	15
Revisions .....	16
.....	16

# Build a framework for an end-to-end machine learning process including ad-hoc data exploration, data processing and feature engineering, and model training and evaluation

*November 2018 (last update (p. 16): December 2019)*

This implementation guide discusses architectural considerations and configuration steps for deploying the Machine Learning for Telecommunication solution on the Amazon Web Services (AWS) Cloud. It includes links to an [AWS CloudFormation](#) template that launches, configures, and runs the AWS services required to deploy this solution on AWS, using AWS best practices for security and availability.

The guide is intended for data scientists, chief data officers, and data engineers who have practical experience architecting on the AWS Cloud.

# Overview

Machine learning (ML) helps Amazon Web Services (AWS) customers use historical data to predict future outcomes, which can lead to better business decisions. ML techniques are core to the communications service provider (CSP) industry. CSPs can use ML algorithms to construct and refine mathematical models from their business data, and then use those models to help identify fraudulent use of network services, automate network functions (zero touch), and reduce customer churn.

AWS offers several ML services and tools tailored for a variety of use cases and levels of expertise. However, it can be a challenge to understand the mechanics of model training and tuning, identify relevant data features, design a workflow that can perform complex extraction, transformation, load (ETL) activities, and scale to accommodate large datasets.

To help customers get started with a machine learning workflow for CSP use cases, AWS offers the Machine Learning for Telecommunication solution. This solution uses AWS CloudFormation to deploy a scalable, customizable ML architecture that leverages [Amazon SageMaker](#), a fully managed ML service, and [The Jupyter Notebook](#), an open source web application for creating and sharing live code, equations, visualizations, and narrative text.

Additionally, the solution provides a sample demo of ad-hoc data exploration, data processing and feature engineering, and model training and evaluation. It also includes a synthetic telecom IP Data Record (IPDR) dataset to demonstrate how to use ML algorithms to test and train models for predictive analysis in telecommunication. Customers can use the included notebooks as a starting point to develop their own custom ML models, and customize the included notebooks for their own use case.


## Cost

You are responsible for the cost of the AWS services used while running this solution. As of the date of publication, the cost for running this solution with default settings in the US East (N. Virginia) Region is approximately **\$1.21 per hour**. This includes charges for SageMaker, Amazon Simple Storage Service (Amazon S3), and AWS Glue. This estimate does not include variable charges for data processing and data transfer costs.

Prices are subject to change. For full details, see the pricing webpage for each AWS service you will be using in this solution.

## Architecture Overview

Deploying this solution builds the following environment in the AWS Cloud.



**Figure 1: Machine Learning for Telecommunication architecture on AWS**

The AWS CloudFormation template deploys an Amazon Simple Storage Service (Amazon S3) bucket that includes a synthetic IP Data Record (IPDR) dataset in Abstract Syntax Notation One (ASN.1) format or call detail record (CDR) format. The template also deploys an AWS Glue job to convert the dataset from CSV to Parquet compressed format, and an SageMaker instance with Machine Learning (ML) Jupyter notebooks. The solution's Jupyter notebooks include instructions that guide the user through the process of constructing additional data features for future models.

When the workflow is triggered in the Jupyter web interface, the solution ingests data from the Amazon S3 bucket into the SageMaker cluster which enables the user to run Jupyter notebooks on the dataset. Amazon S3 Select reads the Parquet compressed data that was processed by the AWS Glue job. The notebooks preprocess the data, extract features, and divide the data into training and testing. ML algorithms process the training dataset to develop a model to identify anomalies and predict future anomalies. The solution then tests the results of the model's predictions against the testing dataset, identifies false positives and negatives, and retrains to fine-tune the model.

# Implementation Considerations

## The Jupyter Notebook

The Machine Learning for Telecommunication solution uses The Jupyter Notebook, an open source web application that allows you to create and share live code, equations, visualizations, and narrative text.

The custom Jupyter notebooks develop a machine learning (ML) model that can identify and predict insights such as call-characteristic anomalies. The notebooks use a default set of predefined, call detail record (CDR) related features to train the model, test the model, then compute the [area under the curve](#) (AUC) metric based on the model scores on the test data, plot the [receiver operating characteristic](#) (ROC) curve to better understand the model performance, and then report the false positives, false negatives, true positives, and true negatives using the [confusion matrix](#). Users can modify the notebook code and features to create different ML models and related visualizations for their specific use cases.

### Note

Five Jupyter notebooks are included in the solution which can be used with your datasets to solve a business problem or use case.

`ML-Telecom-NaiveBayes.ipynb`: This notebook demonstrates the exploration of CDR data and classification with a Apache Spark ML Naive Bayes Algorithm. The notebook reads the Parquet data with Amazon S3 Select for performance optimization, and showcases visualization of data exploration with bar charts and box plots. Additionally, it classifies and predicts the probability of call disconnect reason code 16, and further displays the ML evaluation steps with the confusion matrix and cross validation.

`ML-Telecom-PCA-KMeans.ipynb`: This notebook demonstrates unsupervised learning with Apache Spark ML on CDR data with principal component analysis (PCA) and k-means. The notebook reads the Parquet data with Amazon S3 Select for performance optimization, showcases feature extraction with PCA and an elbow graph for optimizing the choice of k for k-means clustering. Additionally, it further displays the exploration of the data features with a correlation matrix heat map, visualizes the k-means clustering data in 3D and 2D, and further evaluates the k-means cluster by creating centriods.

`ML-Telecom-RandomForestClassifier.ipynb`: This notebook uses a supervised learning algorithm of the Apache Spark ML RandomForest Classifier for classifying the call disconnect reason in CDR Data. The notebook reads the Parquet data with Amazon S3 Select for performance optimization, showcases exploratory data analysis with the ML concepts of feature selection and correlation heat map matrix. Additionally, it uses ChiSqSelector for figuring feature importances, and further evaluates the model with a confusion matrix, AUC, ROC, and precision recall metrics.

`ML-Telecom-TimeSeries-RandomForestClassifier-DeepAR.ipynb`: This notebook demonstrates the [DeepAR](#) supervised learning algorithm for forecasting Scalar Time-Series with telecom data. The notebook uses a hybrid approach of using Apache Spark ML for classifying the call disconnect reason as an anomaly and Amazon DeepAR for time series prediction of the call disconnect reason anomaly. Additionally, it further demonstrates how to prepare the dataset for the time series training with DeepAR and how to use the trained model for inference, and uses the context length and prediction length in minutes interval to showcase the time series prediction graph.

`ML-Telecom-RandomCutForest.ipynb`: This notebook demonstrates unsupervised anomaly detection on time series data with the SageMaker [Random Cut Forest algorithm](#). The notebook uses the call service duration of CDR data for anomaly detection. Additionally, it further showcases data exploration by plotting the data, hyper parameter tuning and calculating anomaly scores on the data, and displays concepts of data shingling and prediction.

## AWS Glue

The Machine Learning for Telecommunication solution invokes an AWS Glue job during the solution deployment to process the synthetic call detail record (CDR) data or the customer's data to convert from CSV to Parquet format. By default, the AWS Glue job deploys 10 data processing units (DPUs) for preprocessing and can be scheduled with a scheduler. Note that when you use your own dataset, you need to modify the schema definitions to meet your data attributes to enable the AWS Glue job to run successfully.

## Demo Data

This solution includes synthetic demo IP Data Record (IPDR) datasets in Abstract Syntax Notation One (ASN.1) format and call detail record (CDR) format. You can choose to copy these datasets into the source bucket to run the demo of AWS Glue transformations and ML model predictions, or use your own datasets.

The sample datasets descriptions in the [Call Detail Record](#) are as follows:

- **Call Detail Record – Start:** Indicates that the call has been successfully established/connected and a session has successfully started.
- **Call Detail Record – Stop:** Indicates that records that previously established a successful session are now terminated.
- **Call Detail Record - Start Sample:** A subset of the CDR start dataset
- **Call Detail Record – Stop Sample:** A subset of the CDR stop dataset

## SageMaker Instance Types

This solution features three instance types for SageMaker (ml.t2.medium, ml.m4.xlarge, ml.p2.xlarge). By default, the solution deploys an ml.t2.medium instance. Note that the solution can provision an ml.p2.xlarge instance using the **Notebook Instance Type** template parameter. This instance allows Jupyter notebooks with large amounts of data to be processed faster, but will incur additional charges.



# AWS CloudFormation Template

This solution uses AWS CloudFormation to automate the deployment of the Machine Learning for Telecommunication solution on the AWS Cloud. It includes the following AWS CloudFormation template, which you can download before deployment:

[View  
Template](#)

**machine-learning-for-telecommunication.template:** Use this template to launch the Machine Learning for Telecommunication solution and all associated components. The default configuration deploys an Amazon Simple Storage Service (Amazon S3) bucket, an AWS Glue job, and an SageMaker instance, but you can also customize the template based on your specific network needs.

# Automated Deployment

Before you launch the automated deployment, please review the architecture, security, and other information discussed in this guide. Follow the step-by-step instructions in this section to configure and deploy the Machine Learning for Telecommunication solution into your account.

**Time to deploy:** Approximately five minutes

## What We'll Cover

The procedure for deploying this architecture on AWS consists of the following steps. For detailed instructions, follow the links for each step.

### [Step 1. Launch the Stack \(p. 7\)](#)

- Launch the AWS CloudFormation template into your AWS account.
- Enter values for required parameters: **Source Bucket**, **Destination Bucket**
- Review the other template parameters and adjust if necessary.

### [Step 2. Run the Notebooks \(p. 9\)](#)

- Test the machine learning algorithms and run the notebooks.

## Step 1. Launch the Stack

This automated AWS CloudFormation template deploys the Machine Learning for Telecommunication solution on the AWS Cloud.

### **Note**

You are responsible for the cost of the AWS services used while running this solution. See the [Cost \(p. 2\)](#) section for more details. For full details, see the pricing webpage for each AWS service you will be using in this solution.

1. Sign in to the AWS Management Console and click the button below to launch the `machine-learning-for-telecommunication` AWS CloudFormation template.



You can also [download the template](#) as a starting point for your own implementation.

2. The template is launched in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the region selector in the console navigation bar.

### **Note**

This solution uses SageMaker, which is currently available in specific AWS Regions only. Therefore, you must launch this solution in an AWS Region where SageMaker is available. For the most current service availability by region, see [AWS service offerings by region](#).

Machine Learning for Telecommunication  
Implementation Guide  
Step 1. Launch the Stack

3. On the **Select Template** page, assign a name to your solution stack.
4. Under **Parameters**, review the parameters for the template and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
<b>Industry</b>	telecom	The industry for which the solution is deployed
<b>Source Bucket</b>	<Requires input>	Unique name of the solution-created Amazon S3 bucket where your application artifacts will be stored
<b>Source Prefix</b>	machine-learning-for-all/V1.0.0/data	Folder location of the dataset for the solution to process. By default, this is the location of synthetic dataset
<b>Destination Bucket</b>	<Requires input>	Unique name of the solution-created Amazon S3 bucket where your AWS Glue processed Parquet data will be stored
<b>Detination Prefix</b>	machine-learning-for-all/V1.0/parq-conv	Folder location of the AWS Glue processed dataset where glue job converts from CSV to Parquet
<b>Synthetic Data</b>	Yes	Select <b>Yes</b> to transfer the synthetic demo dataset to the solution created Amazon S3 bucket and use it for modeling in the notebook. Select <b>No</b> to use your own dataset for modeling.  <b>Note</b> You may need to update the schema. For more information, see <a href="#">Appendix A (p. 13)</a> .
<b>Cron Job</b>	cron(0 0 1/1 * ? *)	AWS Glue job scheduler for execution at a specific time. Use this parameter for long running jobs
<b>Start Transformation</b>	Yes	Choose whether to start an ETL job transformation for converting CSV to Parquet format
<b>Deploy Sagemaker</b>	Yes	Choose whether to deploy an SageMaker instance for

Parameter	Default	Description
		execution of machine learning Jupyter notebooks
<b>Notebook Instance Type</b>	ml.t2.medium	Select the instance type of the SageMaker instance  <b>Note</b> We recommend a ml.p2.xlarge instance for notebooks with large amounts of data. For example the ML-telecom-TimeSeries-RandomForestClassifier-DeepAR notebook.
<b>KmsKeyId</b>	<Optional Input >	AWS KMS key ID used to encrypt data at rest on the ML storage volume attached to the notebook instance.

5. Choose **Next**.
6. On the **Options** page, choose **Next**.
7. On the **Review** page, review and confirm the settings. Be sure to check the box acknowledging that the template will create AWS Identity and Access Management (IAM) resources.
8. Choose **Create** to deploy the stack.

You can view the status of the stack in the AWS CloudFormation console in the Status column. You should see a status of **CREATE\_COMPLETE** in approximately five minutes.

**Note**

In addition to the primary AWS Lambda functions, this solution includes the `SolutionHelper` Lambda function, which runs only during initial configuration or when resources are updated or deleted. When running this solution, the `SolutionHelper` function is inactive. However, do not delete this function as it is necessary to manage associated resources.

## Step 2. Run the Notebooks

1. In the AWS Management Console, navigate to the AWS CloudFormation stack **Outputs** tab.
2. Copy the **SageMakerInstance** name.
3. In the AWS Management Console, open the **SageMaker** instance.
4. Under files, choose **telecom**.
5. Open each of the notebooks: `Ml-Telecom-NaiveBayes.ipynb`, `Ml-Telecom-PCA-KMeans.ipynb`, and `Ml-Telecom-RandomForestClassifier.ipynb`.
6. Choose **Cell**, then select **Run all** in each of the notebooks.
7. Open the `Ml-Telecom-RandomCutForest.ipynb` notebook, and update the **bucket\_name** located in cell 2 with the value you provided in the **Source Bucket** template parameter.
8. Choose **Cell**, then select **Run all**.
9. Open the `Ml-Telecom-TimeSeries-RandomForestClassifier-DeepAR.ipynb` notebook, and update the **bucket\_name** located in cell 18 with the value you provided in the **Source Bucket** template parameter.

10 Choose **Cell**, then select **Run all**.

Now you will be able to use the included Jupyter notebooks and the synthetic telecom dataset to demonstrate how to use machine learning algorithms to test and train models for time series predictive analysis in telecommunications.

**Note**

The AWS Glue job processes the data to convert from CSV to Parquet and make sure the conversion job is complete before execution of the notebook. The AWS Glue job for data conversion of synthetic telecommunication dataset takes approximately one minute.

# Security

When you build systems on AWS infrastructure, security responsibilities are shared between you and AWS. This shared model can reduce your operational burden as AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. For more information about security on AWS, visit the [AWS Security Center](#).

# Additional Resources

## **AWS services documentation**

- [AWS CloudFormation](#)
- [AWS Glue](#)
- [SageMaker](#)
- [Amazon Simple Storage Service](#)

## **Other resources**

- [Jupyter](#)
- [Apache Spark MLlib](#)
- [Call Detail Record](#)

# Appendix A: Using Your Own Data

This solution includes a synthetic dataset for model-training purposes, and is intended for novice machine learning (ML) data scientists. (We would like to thank [Ribbon Communications](#) for providing the synthetic dataset. The data was generated by test data generators, and is not customer or sensitive data). If you choose to use the Jupyter notebooks against your own datasets, we recommend following the AWS best practices for [uploading data into Amazon S3](#) to ensure that your data is uploaded quickly and securely.

Use the following steps to use your own datasets in the Jupyter notebooks:

1. Update the **Source Bucket** and **Source Prefix** location in AWS Cloudformation template to point to the location of your data in Amazon S3.
2. Transferring your synthetic data is optional. Modify the **Synthetic Data** to **No**, if the Demo Data transfer is not required.
3. Select a notebook to run.
4. Modify the value of the **bucket\_name** variable to the Amazon S3 folder location of the Parquet files.
5. Choose **Cell**, then select **Run all**.

## Note

AWS Glue charges by the amount of data that is scanned and processed. Customers who have large datasets can save on costs and achieve better performance if the data is partitioned, compressed, or converted into a columnar format such as, [Apache Parquet](#) format.



# Appendix B: Collection of Operational Metrics

This solution includes an option to send anonymous operational metrics to AWS. We use this data to better understand how customers use this solution and related services and products. When enabled, the following information is collected and sent to AWS:

- **Solution ID:** The AWS solution identifier
- **Unique ID (UUID):** Randomly generated, unique identifier for each deployment
- **Timestamp:** Data-collection timestamp
- **Region:** The AWS Region where the solution is being deployed
- **DPU:** The number of DPUs used in the AWS Glue job
- **Instance Type:** The SageMaker instance type

Note that AWS will own the data gathered via this survey. Data collection will be subject to the [AWS Privacy Policy](#). To opt out of this feature, complete one of the following tasks:

a) Modify the AWS CloudFormation template mapping section as follows:

```
"Send" : {  
  "AnonymousUsage" : { "Data" : "Yes" }  
},
```

to

```
"Send" : {  
  "AnonymousUsage" : { "Data" : "No" }  
},
```

OR

b) After the solution has been launched, find the `SolutionHelper` function in the Lambda console and set the `SEND_ANONYMOUS_DATA` environment variable to `No`.

# Source Code

Visit our [GitHub repository](#) to download the templates and scripts for this solution, and to share your customizations with others.

# Document Revisions

Date	Change	
November 2018	Initial publication	
June 2019	Added information about demo data for modeling	
December 2019	Upgraded the solution to Python 3.7	

## Notices

This implementation guide is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

The Machine Learning for Telecommunication solution is licensed under Apache License Version 2.0 available at <https://www.apache.org/licenses/LICENSE-2.0>