
Real-Time Analytics with Spark Streaming

Implementation Guide



Real-Time Analytics with Spark Streaming: 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	3
Considerations	4
Application Requirements	4
Processing Engine	4
Application JAR	4
Apache Zeppelin	4
Demo Application	4
Single Application Deployment	4
EMR Web Interfaces	5
Memory-Optimized Instance Types	5
Real-Time Data Visualization	5
Regional Deployment	5
Templates	6
Deployment	7
Prerequisites	7
What We'll Cover	7
Step 1. Launch the Stack	7
Step 2. Stop a Running Application	11
Security	12
Security Groups	12
Security at Rest for Amazon EMR	12
Resources	13
Appendix A	14
Appendix B	15
Appendix C	16
Source Code	17
Revisions	18
.....	18

Deploy a solution designed to work with Spark Streaming applications that processes real-time and batch data to obtain actionable insights

February 2017 (last update (p. 18): April 2020)

This implementation guide discusses architectural considerations and configuration steps for deploying the Real-Time Analytics with Spark Streaming in the Amazon Web Services (AWS) Cloud. It includes links to [AWS CloudFormation](#) templates that launch, configure, and run the AWS services required to deploy this solution using AWS best practices for security and availability.

The guide is intended for IT infrastructure architects who have practical experience architecting in the AWS Cloud, building big data applications, and are familiar with Apache Spark.

Overview

Many Amazon Web Services (AWS) customers use batch data reports to gain strategic insight into long-term business trends, and a growing number of customers also require [streaming data](#) to obtain actionable insights from their data in real time. Batch data is collected over a period of time and processed in batches, and this data can provide snapshots of trends that shape decision-making. Streaming data is generated continuously from thousands of data sources and it can help companies proactively respond to changing conditions.

A *lambda architecture* (not to be confused with the AWS Lambda service) is one way to implement real-time and batch data processing in a single framework. The lambda architecture divides processing into three layers: the batch layer in which new data is appended to the master data set and stored as batch views, the serving layer in which batch views are indexed, and the speed layer in which the real-time data views are produced, continuously updated, and stored for read/write operations.

AWS provides many of the building blocks required to build a secure, flexible, cost-effective lambda architecture in the cloud. (For more information on implementing a lambda architecture on AWS, see [Lambda Architecture for Batch and Real-Time Processing on AWS](#).) These include [Amazon Kinesis Data Streams](#), a platform for processing terabytes of streaming data, [Amazon EMR](#), a service that distributes and processes data across dynamically scalable Amazon Elastic Compute Cloud (Amazon EC2) instances, and Amazon Simple Storage Service (Amazon S3), a secure and durable object store. Customers can combine these AWS services with [Apache Spark Streaming](#), for fault-tolerant stream processing of live-data streams, and [Spark SQL](#), which allows Spark code to execute relational queries, to build a single architecture to process real-time and batch data.

The Real-Time Analytics with Spark Streaming solution is an AWS-provided reference implementation that automatically provisions and configures the AWS services necessary to start processing real-time and batch data in minutes. The solution is designed to work with customers' Spark Streaming applications, and also includes a demo application and data producer to create an example environment. The solution also leverages [Apache Zeppelin](#), a web-based notebook for interactive data analytics, to enable customers to visualize both their real-time and batch data.

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 listed in Table 1. This includes charges for Amazon Kinesis Data Streams, Amazon EC2, and Amazon EMR. Prices are subject to change. For full details, see the pricing webpage for each AWS service you will be using in this solution

AWS Service	Resource Count and Type	Total Cost/Hour
Amazon Kinesis Data Streams	2 shards	\$0.03
Amazon EC2	1 - t2.medium instance	\$0.05
	(2 if using the demo template)	\$0.10
NAT Gateways	2 (one per availability zone)	\$0.09
Amazon EMR	3 - r3.xlarge instances	\$0.76

Note

The solution creates a customer master key (CMK) in [AWS Key Management Service \(AWS KMS\)](#) that is used to encrypt data at rest. Rotation is automatically enabled, so each new key version raises the cost of the CMK by \$1/month (rotation happens [once a year](#)).

This pricing does not reflect variable charges incurred from Amazon S3, Amazon CloudWatch, AWS KMS, data transfer fees, or the cost of Amazon DynamoDB. For full details, see the pricing webpage for each AWS service you will be using in this solution.

Apache Spark Streaming, Apache Spark SQL, and Apache Zeppelin are open source. There is no additional cost to use these tools.

Architecture Overview

Deploying this solution with the **default parameters** builds the following environment in the AWS Cloud.

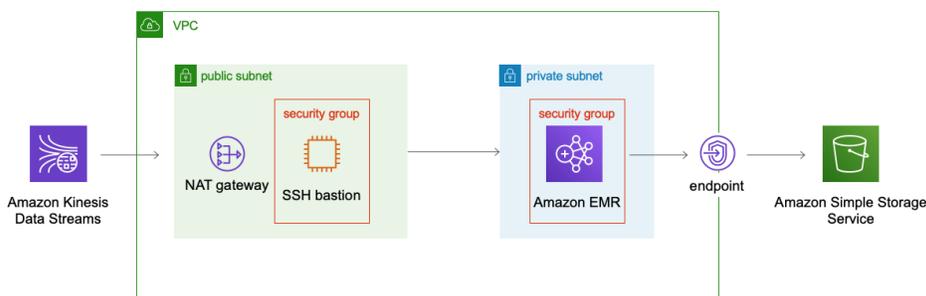


Figure 1: Real-Time Analytics with Spark Streaming default architecture

The AWS CloudFormation template deploys Amazon Kinesis Data Streams which includes Amazon DynamoDB for checkpointing, an Amazon Virtual Private Cloud (Amazon VPC) network with one public and one private subnet, a NAT gateway, a bastion host, an Amazon EMR cluster, and a VPC endpoint to an Amazon S3 bucket.

Amazon Kinesis Data Streams collects data from data sources and sends it through a NAT gateway to the Amazon EMR cluster. Amazon Kinesis Data Streams also includes the Kinesis Client Library (KCL), a pre-built library that helps you easily build Kinesis applications for reading and processing data from an Kinesis stream. The KCL uses a unique Amazon DynamoDB table to keep track of the application's state. Because the KCL uses the name of the Amazon Kinesis Data Streams application to create the name of the table, each application name must be unique.

The private subnet contains an Amazon EMR cluster with Apache Zeppelin. The public subnet contains a NAT gateway to connect Amazon Kinesis Data Streams to the Amazon EMR cluster, and a bastion host that provides SSH access to the Amazon EMR cluster.

The Real-Time Analytics solution is designed to allow you to use your own application, but it also includes a demo application that you can deploy for testing purposes. For more information, see [Appendix A \(p. 14\)](#).

After the Spark Streaming application processes the data, it stores the data in an Amazon S3 bucket.

Implementation Considerations

Application Requirements

The Real-Time Analytics solution requires a working Spark Streaming application written in Java or Scala. We recommend that you use the latest version of Apache Spark for your application. You can choose to deploy your application as a JAR file or a JSON file (see the next section for details).

Processing Engine

When you deploy the solution, you choose the processing engine for your custom Spark Streaming application: a JAR file or an Apache Zeppelin notebook JSON file.

Application JAR

If you choose to package your custom Spark Streaming application as a JAR file, the solution requires a Spark Submit command to launch your application on the Amazon EMR cluster. You can choose to upload a file with the submit script, or you can enter the command directly in the AWS CloudFormation template parameter when you launch the solution.

Apache Zeppelin

If you choose to use Zeppelin, you will upload your custom Spark Streaming application as a `notebook.json` file from your local machine or a URL you specify. The solution automatically creates the dependencies and configurations to visualize your real-time and batch data.

Demo Application

The solution includes an additional AWS CloudFormation template `real-time-analytics-spark-streaming-demo.template` that deploys a demo application for testing purposes.

If you choose to run the demo application, this solution deploys a sample Spark Streaming application on the Amazon EMR cluster, a sample Kinesis stream, and a sample data producer that sends sample data to your Kinesis stream. The demo application is packaged as a JAR file, but it also includes a JSON file you can use to deploy the demo application through the Zeppelin UI. For more information, see [Appendix A \(p. 14\)](#).

Single Application Deployment

The solution is designed to work with only one Spark Streaming application at a time. If you want to change applications, you must first [stop the running application \(p. 11\)](#) and then deploy the solution again with a new application. This also applies if you deploy the demo application: you must stop the running demo application before you can deploy the demo application JSON file, a custom JAR file or a custom JSON file.

If you want to upload an application file that uses the same name as an application from a previous deployment (for example, a new version of a JSON template), you must clear the application name from the Amazon DynamoDB table before you deploy solution. See [Tracking Amazon Kinesis Data Streams Application State](#) in the Amazon Kinesis Data Streams Developer Guide for more information.

EMR Web Interfaces

When you launch an Amazon EMR cluster in a public subnet, the master node of the cluster has a public DNS which allows you to create an SSH tunnel and securely access the Amazon EMR web interfaces. Because this solution deploys the Amazon EMR cluster in a *private* subnet, the master node will not have a public DNS for secure SSH access. To allow you to access the Amazon EMR web interfaces, this solution deploys a bastion host with a public IP address. You must configure dynamic port forwarding to connect to the bastion host. For more information, see [View Web Interfaces Hosted on Amazon EMR Clusters](#) in the Amazon EMR Management Guide.

Memory-Optimized Instance Types

We recommend memory-optimized Amazon Elastic Compute Cloud (Amazon EC2) instance types for Apache Spark workloads because Spark attempts to process as much data in memory as possible. By default, this solution deploys an r3.xlarge instance for the Amazon EMR cluster nodes to deliver optimal performance.

Real-Time Data Visualization

For users who choose Zeppelin as their processing engine for this solution, Zeppelin will display visualizations of your real-time data as it flows. These visualizations can be shared or published to external dashboards. For more information, go to the [Apache Zeppelin website](#).

Regional Deployment

The Real-Time Analytics solution uses AWS Lambda during initial configuration or when resources are updated or deleted. Therefore, you must deploy this solution in an AWS Region that supports AWS Lambda. For the most current AWS Lambda availability by region, see [AWS service offerings by region](#).

AWS CloudFormation Templates

This solution uses AWS CloudFormation to automate the deployment of this Real-Time Analytics with Spark Streaming on the AWS Cloud. It includes the following AWS CloudFormation templates, which you can download before deployment:

[View
Template](#)

real-time-analytics-spark-streaming.template: Use this template to launch the Real-Time Analytics solution and all associated components. The default configuration deploys an Amazon VPC network, a bastion host, a VPC endpoint, an Kinesis stream, an Amazon EMR cluster, and an Amazon S3 bucket.

[View
Template](#)

real-time-analytics-spark-streaming-demo.template: Use this template to launch the Real-Time Analytics solution with the demo application. This configuration deploys an Amazon VPC network, a bastion host, a VPC endpoint, an Amazon EMR cluster, an Amazon S3 bucket, a demo Spark Streaming application, a sample Kinesis stream, and a sample data producer.

Automated Deployment

Before you launch the automated deployment, please review the implementation considerations and prerequisites discussed in this guide. Follow the step-by-step instructions in this section to configure and deploy the Real-Time Analytics solution into your account.

Time to deploy: Approximately 15 minutes

Prerequisites

Review the application requirements and processing options in [Implementation Considerations \(p. 4\)](#).

- Before you deploy the solution, you must upload your working Spark Streaming application to an Amazon Simple Storage Service (Amazon S3) bucket. If you are using a Spark Submit script to launch your custom application, you must have a *spark-submit.sh* file with the Spark Submit command in an Amazon S3 bucket.
- Remember that you can only deploy one running application with a unique name at a time with this solution (see [Single Application Deployment \(p. 4\)](#) for detailed information).
- Before you deploy the solution, you must enable the Amazon EMR web interfaces to view Apache Zeppelin, the Spark History UI, and the Resource Manager. For more information, see [EMR Web Interfaces \(p. 5\)](#).
- You must also configure dynamic port forwarding to connect to the bastion host to securely access the Amazon EMR web interfaces. For more information, see [View Web Interfaces Hosted on Amazon EMR Clusters](#) in the Amazon EMR Management Guide.

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.
- Review the other template parameters, and adjust if necessary.

[Step 2. Stop a Running Application \(p. 11\)](#)

- Navigate to the EMR cluster's Resource Manager.
- Stop the running application.

Step 1. Launch the Stack

The automated AWS CloudFormation templates deploy Real-Time Analytics with Spark Streaming on the AWS Cloud using either your own Spark Streaming application or the AWS-provided [demo application \(p. 14\)](#).

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. Log in to the AWS Management Console and click the button below to launch the solution or demo application 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 Real-Time Analytics solution in a different AWS Region, use the region selector in the console navigation bar.

Note

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

3. On the **Select Template** page, verify that you selected the correct template and choose **Next**.
4. On the **Specify Details** page, assign a name to your Real-Time Analytics solution stack.
5. Under **Parameters**, review the parameters for the template, and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
Key Name	<Requires input>	Public and private key pair, which allows you to connect securely to the bastion host. When you created an AWS account, this is the key pair you created in your preferred AWS Region.
Remote Access CIDR	<Requires input>	The IP address range that can be used to SSH to the bastion host
Availability Zones	<Requires input>	The list of Availability Zones to use for the Amazon VPC subnets
Number of AZs	2	The number of Availability Zones to use in your VPC Note This number must match the number of selections you chose for the Availability Zone parameter.
VPC CIDR	10.0.0.0/16	The VPC CIDR block

Real-Time Analytics with Spark
Streaming Implementation Guide
Step 1. Launch the Stack

Parameter	Default	Description
Private Subnet 1A CIDR	10.0.0.0/19	The CIDR block for the private subnet located in AZ1
Private Subnet 2A CIDR	10.0.32.0/19	The CIDR block for the private subnet located in AZ2
Public Subnet 1 CIDR	10.0.128.0/20	The CIDR block for the public DMZ subnet located in AZ1
Public Subnet 2 CIDR	10.0.144.0/20	The CIDR block for the public DMZ subnet located in AZ2
Kinesis Stream	default-data-stream	The name of the source Amazon Kinesis stream the template will create
Shard count	2	The number of shards for your Amazon Kinesis stream
Master	r5.xlarge	The EMR master node Amazon EC2 instance type
Core	r5.xlarge	The EMR core node Amazon EC2 instance type
Artifact Bucket	<Requires input>	The Amazon S3 bucket location where your application artifacts are stored.
Submit Mode	AppJar	<p>The processing engine of the Spark Streaming application. Choose the appropriate processing engine (Zeppelin for JSON templates; AppJar for JAR files).</p> <p>Note This parameter will be set to DemoApp if you are using the demo template.</p>
Type	None	<p>The submit type of the Spark Streaming application. You can specify a submit script or a submit command.</p> <p>Note Use this parameter only if you choose AppJar as your Submit Mode. This parameter will not show if you are using the demo template.</p>

Real-Time Analytics with Spark
 Streaming Implementation Guide
 Step 1. Launch the Stack

Parameter	Default	Description
Script	<Optional input>	<p>The Amazon S3 bucket where your script with the Spark submit command is stored. For example, <code>s3://{bucket_location/spark_submit.sh}</code></p> <p>Note Use this parameter only if you choose <code>AppJar</code> as your Submit Mode. If you choose <code>Command</code> as your Submit Type, leave this parameter blank. This parameter will not show if you are using the demo template.</p>
Command	<Optional input>	<p>A comma-delimited Spark submit command. For example, <code>--deploy-mode, {cluster/client}, --class {className}, --master, {yarn/local[?]},{s3://AppLocation/AppJar}, {Appname}, {StreamName}, {OutputLoc}</code></p> <p>Note Use this parameter only if you choose <code>AppJar</code> as your Submit Mode. If you choose <code>Script</code> as your Submit Type, leave this parameter blank. This parameter will not show if you are using the demo template.</p>

6. Select **Next**.
7. On the **Options** page, choose **Next**.
8. On the **Review** page, review and confirm the settings. Check the box acknowledging that the template will create AWS Identity and Access Management (IAM) resources and might require the `CAPABILITY_AUTO_EXPAND` capability.
9. 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 roughly 15-20 minutes.

Note

This solution includes two AWS Lambda functions that run only during initial configuration or when resources are updated or deleted.

When running this solution, you will see both Lambda functions in the AWS Lambda console. Do not delete the functions as they are necessary to manage associated resources.

Step 2. Stop a Running Application

The solution is designed to work with only one Spark Streaming application at a time. If you want to change applications, you must first stop the running application and then deploy the solution again with a new application.

Note

To view the EMR Resource Manager web interface, you must enable the EMR web interfaces. For more information, see [EMR Web Interfaces \(p. 5\)](#).

1. On the Amazon EMR console, select the cluster name.
2. On the EMR cluster details page, for **Connections**, choose **Resource Manager**.
3. On the resource manager, select the application **ID**.
4. On the application details page, select **Kill Application**.
5. Select **OK**.

Note

Another option to stop a running application is to use the YARN command line (this approach does not require port forwarding). You must SSH into the Master node (via bastion) and run the following command: `sudo yarn application -kill <application-ID>`.

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](#).

Security Groups

The security groups created in this solution are designed to control and isolate network traffic between the Kinesis sample data producer, the Amazon EMR cluster and the bastion host. We recommend that you review the security groups and further restrict access as needed once the deployment is up and running.

Security at Rest for Amazon EMR

[Security configurations](#) are Amazon EMR templates that can be used to configure data encryption, Kerberos authentication, and Amazon Simple Storage Service (Amazon S3) authorization for EMRFS. The solution creates a security configuration that enables:

- At-rest encryption for EMRFS data in Amazon S3 with Amazon S3-managed encryption keys (SSE – S3)
- At-rest encryption for local disks (Amazon Elastic Block Store root device and storage volumes) with an AWS Key Management Service customer master key (CMK)

However, the solution-created configuration does not enable in-transit encryption because this setting requires either a PEM file or a custom certificate-provider JAR file. You can use a trusted certification authority (CA) to issue certificates. For information about security certificates, see [Providing Certificates for Encrypting Data in Transit with Amazon EMR Encryption](#) in the *Amazon EMR Management Guide*.

Additional Resources

AWS service documentation

- [AWS CloudFormation](#)
- [Amazon DynamoDB](#)
- [Amazon EMR](#)
- [Kinesis](#)
- [Amazon S3](#)
- [Amazon VPC](#)

Associated AWS Whitepaper

- [Lambda Architecture for Batch and Real-Time Processing on AWS with Spark Streaming and Spark SQL](#)

Associated AWS Big Data blog posts

- [Securely Access Web Interfaces on Amazon EMR Launched in a Private Subnet](#)
- [Analyze Realtime Data from Amazon Kinesis Data Streams Using Zeppelin and Spark Streaming](#)
- [The Impact of Using Latest-Generation Instances for Your Amazon EMR Job](#)

Appendix A: Demo Application

The Real-Time Analytics with Spark Streaming solution includes a demo application for testing purposes. Deploying this solution with the demo application builds the following environment in the AWS Cloud.

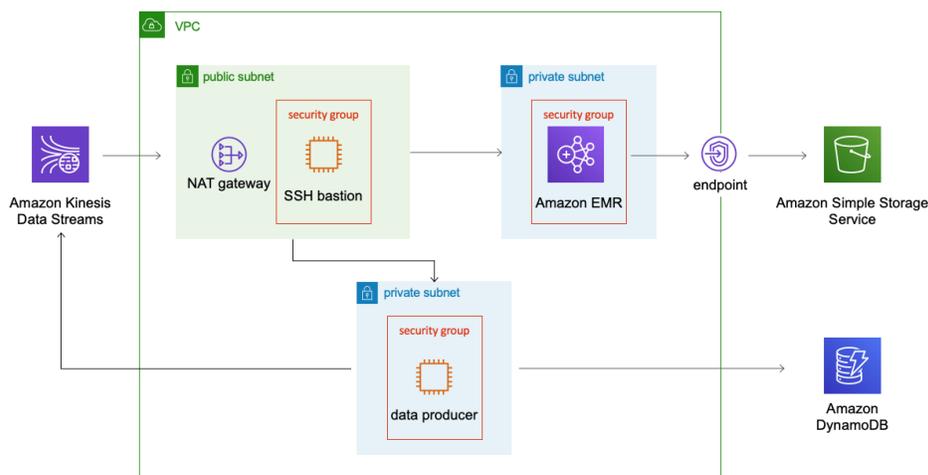


Figure 2: Real-Time Analytics with Spark Streaming architecture with demo application

When you use the demo application, the solution deploys an additional private subnet with a sample data producer and uploads the demo application code to your existing Amazon Simple Storage Service (Amazon S3) bucket. The demo application sends sample data through the NAT gateway to Amazon Kinesis Data Streams. In this architecture, the bastion host provides SSH access to the Amazon EMR cluster and the sample data producer, and the filtered data is stored in Amazon S3.

By default, the data producer will only run for a few seconds (creating one file on Amazon S3). To run the producer manually, take the following steps to replace the necessary values including *<mykeypair>*, *<bastion-host-dns>*, *<kinesis-producer-IP>* address, *<default-data-stream>*, and *<your-AWS-Region>*.

```
# Copy pem file to bastion host using SCP
scp -i "<mykeypair>.pem" <mykeypair>.pem ec2-user@<bastion-host-dns>:<mykeypair>.pem

# Login to bastion host
ssh -i "<mykeypair>.pem" ec2-user@<bastion-host-dns>

# Login to Kinesis Producer
ssh -i "<mykeypair>.pem" ec2-user@<kinesis-producer-IP>

# Run producer for 600 seconds (10 minutes)
sudo java -jar /home/ec2-user/kinesis-producer.jar <default-data-stream> <your-AWS-Region>
600
```

Appendix B: Known Issues

When deleting the solution stack, the Amazon EMR cluster does not delete the following security groups (since they reference one another): `ElasticMapReduce-Slave-Private`, `ElasticMapReduce-ServiceAccess`, and `ElasticMapReduce-Master-Private`. This will cause the VPC stack deletion to fail, since a VPC cannot be deleted while there are still available security groups.

To remove all inbound and outbound rules for a security group, use the AWS Command Line Interface (AWS CLI) terminal and enter the following commands, which must be executed for each group listed above. Using this method requires `jq`, a lightweight command-line JSON processor, to be installed.

```
SECURITY_GROUP_ID=<ID-placeholder>
aws ec2 describe-security-groups --group-ids $SECURITY_GROUP_ID --output json | jq
'.SecurityGroups[0].IpPermissions' > IpPermissions.json

aws ec2 describe-security-groups --group-ids $SECURITY_GROUP_ID --output json | jq
'.SecurityGroups[0].IpPermissionsEgress' > IpPermissionsEgress.json

aws ec2 revoke-security-group-ingress --group-id $SECURITY_GROUP_ID --ip-permissions
file://IpPermissions.json

aws ec2 revoke-security-group-egress --group-id $SECURITY_GROUP_ID --ip-permissions file://
IpPermissionsEgress.json
```

Appendix C: 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 to improve the services and products that we offer. When enabled, the following information is collected and sent to AWS during the initial stack creation:

- **Solution ID:** The AWS solution identifier
- **Unique ID (UUID):** Randomly generated, unique identifier for each Real-Time Analytics with Spark Streaming deployment
- **Timestamp:** Data-collection timestamp
- **EMR Data:** Type of instances and count of the number of instances

Example data:

```
{"Master": "1", "InstanceType": "r3.xlarge", "CoreInstance": "2",  
"InstanceType": "r3.xlarge", "Region": "us-east-1"}
```

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 the following task.

Modify the AWS CloudFormation template mapping section as follows:

```
AnonymousData:  
  SendAnonymousData:  
    Data: Yes
```

to

```
AnonymousData:  
  SendAnonymousData:  
    Data: No
```

Source Code

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

Document Revisions

Date	Change
February 2017	Initial release
March 2017	Corrected demo application architecture diagram
April 2020	Added demo producer application, upgraded the solution to Python 3.8, bug fixes

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. 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 Real-Time Analytics with Spark Streaming is licensed under the terms of the Apache License Version 2.0 available at <https://www.apache.org/licenses/LICENSE-2.0>