

Aspect Workforce Management integrations for Amazon Connect

Configuration Guide

May 2021

Aspect Software, Inc.

Contents

Overview.....	2
Aspect Workforce Management (WFM) on AWS.....	2
Costs and Licenses.....	3
Architecture.....	3
Prerequisites	6
Specialized Knowledge	6
Deployment Options	6
Deployment Steps	7
Step 1. Prepare Your AWS Account.....	7
Step 2. Check Your Amazon Connect Instance.....	7
Step 3. Launch the Quick Start	8
Step 4. Enable Data Streaming	20
Step 5. Test the Deployment	20
Best Practices Using Aspect Workforce Management on AWS	21
Security.....	22
FAQ.....	24
Git Repository	25
Additional Resources	25

Document Revisions	26
--------------------------	----

This Configuration guide was created by Aspect Software, Inc.

[Quick Starts](#) are automated reference deployments that use AWS CloudFormation templates to deploy key technologies on AWS, following AWS best practices.

Overview

This Configuration guide provides step-by-step instructions for deploying Aspect Workforce Management integrations with Amazon Connect on the Amazon Web Services (AWS) Cloud.

This Quick Start is for users of Amazon Connect who want to integrate with Aspect Workforce Management.

Aspect Workforce Management (WFM) on AWS

This AWS Quick Start includes two integration options:

- **Agent productivity** –Enhances the standard Amazon Connect reports with statistics based on agent sign-in/sign-out information. You can designate whether custom agent status labels defined in Amazon Connect should be treated as available, signed in, or signed out. For example, a custom agent status called “Break,” which represents the agent taking time off from work, would typically be treated as signed out, whereas a custom agent status called "Project," which represents some sort of project work, would typically be treated as signed in. Aspect WFM uses this information to produce statistics on average positions staffed (APS) by agent group, available time by agent, unavailable time by agent, and sign-in/sign-out pairs by agent.
- **Real-time adherence** –Provides Amazon Connect agent state change information to the Aspect Real-Time Adherence (RTA) product, which enables you to monitor the activities of your Amazon Connect agents and monitor how well these activities adhere to the agents’ schedules.

Both integration options use an Amazon Connect agent event stream. An Amazon Connect instance supports a single agent event stream. If you have already enabled agent event streams and selected a Kinesis stream for that purpose, choose one of the deployment options that uses an existing Kinesis data stream. If agent event data streaming is not currently enabled for your Amazon Connect instance, choose a deployment option that creates a new Kinesis data stream.

Costs and Licenses

You are responsible for the cost of the AWS services used while running this Quick Start reference deployment. There is no additional cost for using the Quick Start. These integrations require Aspect Workforce Management version 18.1 or higher and Aspect WFM Adapter version 3.3 or higher.

For cost estimates, see the pricing pages for each AWS service you will be using. Prices are subject to change.

Architecture

Deploying an integration option that includes Agent Productivity builds the following environment in the AWS Cloud.

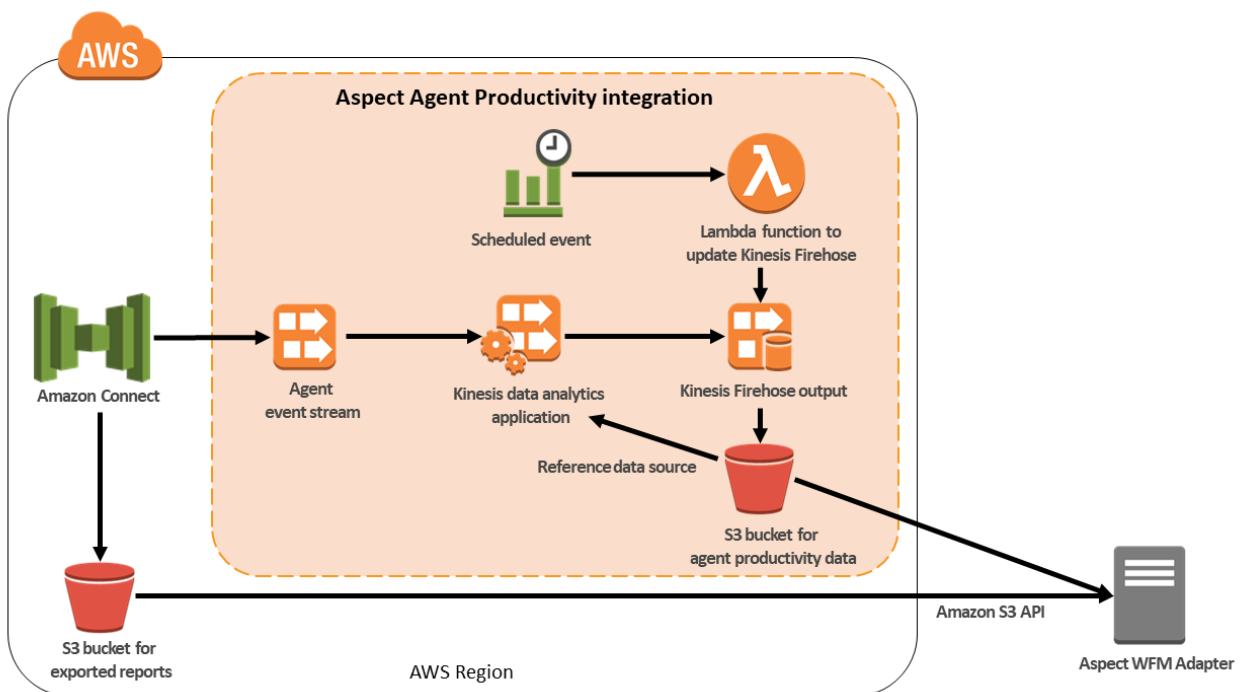


Figure 1: Quick Start architecture for Aspect Agent Productivity on AWS

The Quick Start sets up the following:

- An S3 bucket that holds the Agent Productivity outputs.
- A Kinesis data analytics application that reads agent events from your Amazon Connect instance. The application categorizes those agent events as spans of signed-in or signed-out time.

- A Kinesis Data Firehose data delivery stream that writes the output of the Kinesis data analytics application to the S3 bucket.
- A Lambda function that also writes to the Kinesis Data Firehose data delivery stream. The Lambda function ensures that output objects are written to the S3 bucket during periods of time when no agents are signed-in.
- A CloudWatch Events rule that calls the Lambda function on a schedule, every five minutes.
- An IAM user to provide read access to objects in the two S3 buckets used by WFM Adapter:
 - The S3 bucket for exported reports, associated with your Amazon Connect instance. You will need to schedule the export of several Amazon Connect reports. For details, see the Aspect WFM Adapter documentation.
 - The S3 bucket created by this integration.

You will need to deploy a separate integration for each Amazon Connect instance for which Agent Productivity support is required.

Deploying an integration option that includes Real-Time Adherence builds the following environment in the AWS Cloud.

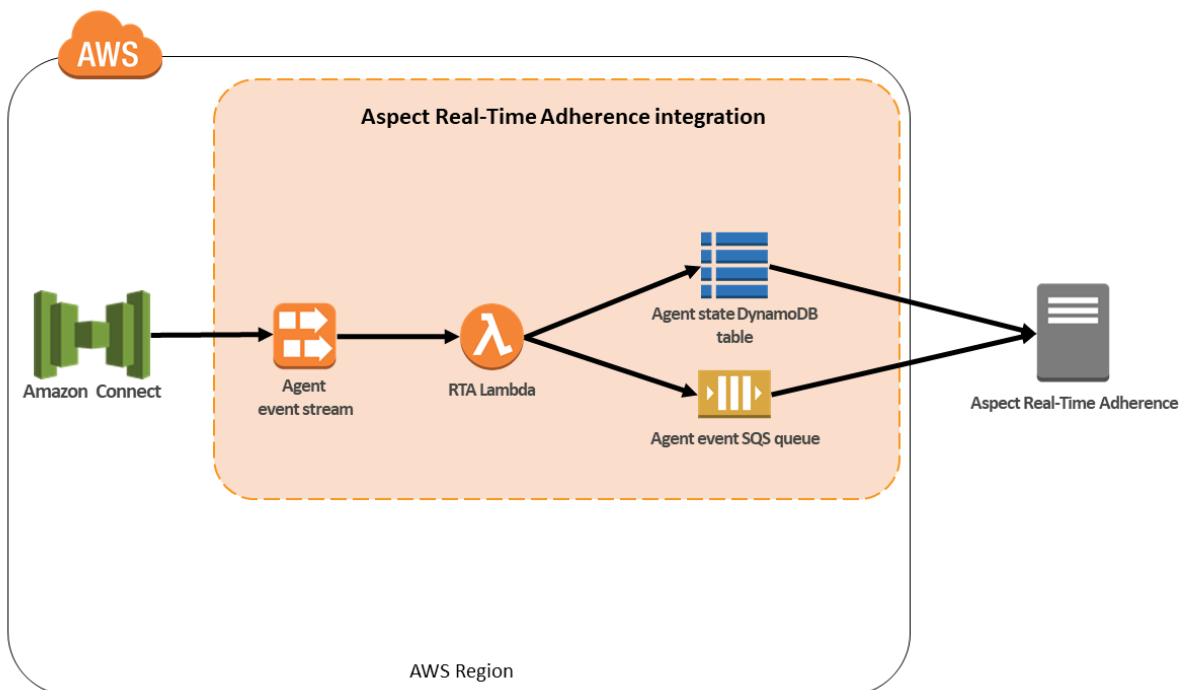


Figure 2: Quick Start architecture for Aspect Real-Time Adherence on AWS

The Quick Start sets up the following:

- A DynamoDB table that holds the latest state change for each agent signed into your Amazon Connect instance.
- An optional DynamoDB table that holds configuration information. If you choose to store RTA configuration in a table, the QuickStart will create this table. If you choose to store RTA configuration in the environment variables of the RTA Lambda, the QuickStart will not create this table.
- One or more SQS queues that are used to relay Amazon Connect agent state changes to the Aspect Real-Time Adherence product. By default, the QuickStart will create a single SQS queue. You can create multiple SQS queues (up to 10) to improve the throughput of state changes.
- A Lambda function that is called whenever your Amazon Connect instance writes events to its agent event Kinesis data stream. The Lambda function is responsible for updating the DynamoDB table and SQS queue(s) created by the integration.
 - In cases where an RTA server is configured to monitor a subset of the agents associated with an Amazon Connect instance, you should configure the QuickStart to filter based on Agent Hierarchy Groups. To do so, identify an Agent Hierarchy Group Level configured in the Amazon Connect instance that can be used to distinguish the agents being monitored on the RTA server. Configure AHGFilterLevel to match the Agent Hierarchy Group Level number (1-5). Configure a list of Agent Hierarchy Group names that are available for the specified level and match the agents to be included. By default, group names should be separated by commas. However, you can specify a different separator in AHGFilterSeparator if some of your group names contain commas. Note: if you want to send all state changes to the RTA server, use the default AHGFilterLevel value of 0.
- An IAM user to provide access to the AWS resources required by Aspect Real-Time Adherence:
 - Read-only access to the DynamoDB table that holds the latest state change for each agent signed into your Amazon Connect instance.
 - Read/write access to the SQS queue(s).
 - If configuration information is stored in a table:
 - Read/write access to the configuration DynamoDB table.
 - Otherwise,
 - Access to read and update the configuration of the Lambda function.

You will need to deploy a separate integration for each Amazon Connect instance connection to an Aspect Real-Time Adherence server. That is, if you wish to monitor agent states from an Amazon Connect instance on two different Aspect Real-Time Adherence servers, then you will need to deploy a separate integration for each RTA server.

If you choose to add Aspect WFM with a new Kinesis data stream, the following resources will be created:

- A Kinesis data stream to be configured as the agent event stream of your Amazon Connect instance. This Kinesis data stream is designated as “Agent event stream” in the figures above.

Prerequisites

Specialized Knowledge

Before you deploy this Quick Start, we recommend that you become familiar with the following AWS services. (If you are new to AWS, see [Getting Started with AWS](#).)

- [Amazon Connect](#)
- [Amazon S3](#)
- [Amazon DynamoDB](#)
- [Amazon SQS](#)
- [Amazon Kinesis](#)
- [AWS Lambda](#)
- [Amazon CloudWatch](#)
- [AWS CloudFormation](#)

Deployment Options

This Quick Start provides six deployment options divided into two categories:

- **Add Aspect WFM with a new Kinesis data stream.** Choose this category if agent event data streaming is not currently enabled for your Amazon Connect instance. This category includes three deployment options:
 - **Deploy for agent productivity.** Choose this option if you plan to use the Agent Productivity feature of Aspect WFM but not the Real-Time Adherence feature of Aspect WFM.
 - **Deploy for real-time adherence.** Choose this option if you plan to use the Real-Time Adherence feature of Aspect WFM but not the Agent Productivity feature of Aspect WFM.
 - **Deploy for both options.** Choose this option if you plan to use both the Agent Productivity and Real-Time Adherence features of Aspect WFM.
- **Add Aspect WFM with an existing Kinesis data stream.** Choose this category if agent event data streaming is enabled for your Amazon Connect instance, and you have selected a Kinesis stream for that purpose. This category includes three deployment options:

- **Deploy for agent productivity.** Choose this option if you plan to use the Agent Productivity feature of Aspect WFM but not the Real-Time Adherence feature of Aspect WFM. You can choose this option to add support for the Agent Productivity feature of Aspect WFM, if you previously deployed support for Real-Time Adherence.
- **Deploy for real-time adherence.** Choose this option if you plan to use the Real-Time Adherence feature of Aspect WFM but not the Agent Productivity feature of Aspect WFM. You can choose this option to add support for the Real-Time Adherence feature of Aspect WFM, if you previously deployed support for Agent Productivity. You will need to choose this option if you want to monitor your Amazon Connect instance from multiple Aspect Real-Time Adherence servers. You will need to deploy this option once for each additional Real-Time Adherence server.
- **Deploy for both options.** Choose this option if you plan to use both the Agent Productivity and Real-Time Adherence features of Aspect WFM.

The Quick Start provides templates for these options. It also lets you configure Aspect Workforce Management settings, as discussed later in this guide.

Deployment Steps

Step 1. Prepare Your AWS Account

1. If you don't already have an AWS account, create one at <https://aws.amazon.com> by following the on-screen instructions.
2. Use the region selector in the navigation bar to choose the AWS Region where you want to deploy Aspect Workforce Management on AWS. You must select the same AWS Region as your Amazon Connect instance.

Step 2. Check Your Amazon Connect Instance

1. Edit your Amazon Connect instance within the AWS console.
2. Select Data streaming. Verify whether the Enable data streaming checkbox is checked and a Kinesis Stream is selected under Agent Events. If Agent Event data streaming is enabled for your Amazon Connect instance, you should choose one of the deployment options listed under the category "Add Aspect WFM with an existing Kinesis data stream," in the next step. If not, you should choose one of the deployment options listed under the category "Add Aspect WFM with a new Kinesis data stream," in the next step.
3. If Agent Event data streaming is enabled for your Amazon Connect instance, switch to the Amazon Kinesis service, edit the Kinesis data stream that was selected as the Agent

Event stream, and locate the Stream ARN. We recommend that you copy the Stream ARN to the clipboard, so you can paste it as the value of the Agent Event Kinesis Stream ARN in the next step.

Step 3. Launch the Quick Start

Note You are responsible for the cost of the AWS services used while running this Quick Start reference deployment. There is no additional cost for using this Quick Start. For full details, see the pricing pages for each AWS service you will be using in this Quick Start. Prices are subject to change.

1. Choose one of the following options to launch the AWS CloudFormation template into your AWS account. For help choosing an option, see [deployment options](#) earlier in this guide. The deployment options are available [here](#). Each deployment takes about 5 minutes to complete.
2. Check the region that's displayed in the upper-right corner of the navigation bar, and change it if necessary. This is where the services for Aspect Workforce Management will be built. The template is launched in the US East (Ohio) Region by default. You must select the same region as your Amazon Connect instance.
3. On the **Select Template** page, keep the default setting for the template URL, and then choose **Next**.
4. On the **Specify Details** page, change the stack name if needed. Review the parameters for the template. Provide values for the parameters that require input. For all other parameters, review the default settings and customize them as necessary. When you finish reviewing and customizing the parameters, choose **Next**.

In the following tables, parameters are listed by category and described separately for the six deployment options:

- [Parameters for deploying Aspect Agent Productivity with a new Kinesis data stream](#)
- [Parameters for deploying Aspect Real-Time Adherence with a new Kinesis data stream](#)
- [Parameters for deploying Aspect Agent Productivity and Real-Time Adherence with a new Kinesis data stream](#)
- [Parameters for deploying Aspect Agent Productivity with an existing Kinesis data stream](#)
- [Parameters for deploying Aspect Real-Time Adherence with an existing Kinesis data stream](#)

- [Parameters for deploying Aspect Agent Productivity and Real-Time Adherence with an existing Kinesis data stream](#)
- **Option 1: Parameters for deploying Aspect Agent Productivity with a new Kinesis data stream**

[View template](#)

Aspect WFM Configuration:

Parameter label (name)	Default	Description
Aspect Kinesis Application Name (KdaApplicationName)	aspect-wfm-ap	Enter the name of the Kinesis Data Analytics application to create. This string can include numbers, lowercase letters, uppercase letters, and hyphens (-). It cannot start or end with a hyphen (-). If you have multiple Amazon Connect instances in your AWS account, you will need to create a Kinesis Data Analytics application for each Amazon Connect instance and give each application a unique name. The Kinesis Firehose Data Delivery stream that writes agent status reports to S3 is named based on the application name. The name of the Firehose stream determines the filenames of agent productivity report files written to S3. By default, WFM Adapter expects the report files to be named based on an application name of "aspect-wfm-ap". If you specify a different application name, be sure to update the WFM Adapter configuration for the corresponding data source to match your application name, in the Firehose Report Mapping data source parameter. Consult the WFM Adapter help for more information.
Aspect WFM Report Prefix (AspectApDataS3Prefix)	OUTPUT/	Enter the prefix for Aspect WFM Agent Productivity reports written to S3. You should configure the corresponding WFM Adapter data source to match this value, in the Firehose Report Root Path data source parameter.

Amazon Connect Configuration:

Parameter label (name)	Default	Description
Amazon Connect Exported Reports Bucket (ConnectS3Bucket)	<i>Requires input</i>	Enter the name of the Exported reports S3 bucket of your Amazon Connect instance.
Amazon Connect Exported Reports Prefix (ConnectS3Prefix)	connect/directory/Reports/Aspect/	Enter the prefix under which the exported reports read by WFM Adapter will be exported, in the Exported reports bucket of your Amazon Connect instance.

Parameter label (name)	Default	Description
Amazon Connect Exported Reports Bucket KMS Key ARN (ConnectS3KmsArn)		Optional Amazon Resource Name (ARN) of the KMS key required to decrypt the exported reports of your Amazon Connect instance. You should supply this ARN if you configured your exported reports bucket to use server-side encryption with a customer-managed customer master key (CMK). Otherwise, you should leave this parameter blank.

AWS Quick Start Configuration:

Parameter label (name)	Default	Description
Quick Start S3 Bucket Name (QSS3BucketName)	quickstart-reference	The S3 bucket you have created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers,
Quick Start S3 Key Prefix (QSS3KeyPrefix)	connect-integration-aspect-wfm/	The S3 key name prefix used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes.

- **Option 2: Parameters for deploying Aspect Real-Time Adherence with a new Kinesis data stream**

[View template](#)

Aspect WFM Configuration:

Parameter label (name)	Default	Description
Number of SQS Queues (NumSQSQueues)	1	The number of SQS queues to create (must be between 1 and 10). Creating additional SQS queues can reduce the latency of state changes.
Write to SQS Flag Location (LambdaConfigLocation)	env	Once RTA reads an initial set of states from the agent event DynamoDB table, it toggles a Write to SQS flag to ensure that agent events are written to the SQS queue(s). This setting controls where that Write to SQS flag is stored either in an environment variable of the RTA Lambda (env) or in a configuration DynamoDB table (table). Storing this flag in an environment variable of the Lambda requires less overhead but requires greater permissions for the RTA user. Storing this flag in a configuration DynamoDB table requires fewer permissions for the RTA user but requires an additional DyanmoDB read each time the RTA Lambda executes. By default, this

Parameter label (name)	Default	Description
		flag will be stored in an environment variable of the RTA Lambda.

Amazon Connect Configuration:

Parameter label (name)	Default	Description
Agent Hierarchy Group Level (AHGFilterLevel)	0	The Agent Hierarchy Group Level on which agent events should be filtered (must be between 0 and 5). Configuring a value of 0 disables filtering. Configuring a value between 1 and 5 means that agent events will be filtered based on that Agent Hierarchy Group Level.
Agent Hierarchy Group filter value separator (AHGFilterSeparator)	,	Character used to separate matching values listed in AHGFilterValues. You should configure a separator character that does not appear in any of the individual matching values.
Agent Hierarchy Group filter matching values (AHGFilterValues)		An optional list of matching Agent Hierarchy Group Level values. If you enable Agent Hierarchy Group filtering, any agent events with an Agent Hierarchy Group Level value matching one of the values listed here will be sent on to Real-Time Adherence and any agent events that do not match the filter will be discarded. Matching values should be separated with the character configured in AHGFilterSeparator.

AWS Quick Start Configuration:

Parameter label (name)	Default	Description
Quick Start S3 Bucket Name (QSS3BucketName)	quickstart-reference	The S3 bucket you have created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers,
Quick Start S3 Key Prefix (QSS3KeyPrefix)	connect-integration-aspect-wfm/	The S3 key name prefix used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes.

- **Option 3: Parameters for deploying Aspect Agent Productivity and Real-Time Adherence with a new Kinesis data stream**

[View template](#)

Aspect WFM Configuration:

Parameter label (name)	Default	Description
Aspect Kinesis Application Name (KdaApplicationName)	aspect-wfm-ap	Enter the name of the Kinesis Data Analytics application to create. This string can include numbers, lowercase letters, uppercase letters, and hyphens (-). It cannot start or end with a hyphen (-). If you have multiple Amazon Connect instances in your AWS account, you will need to create a Kinesis Data Analytics application for each Amazon Connect instance and give each application a unique name. The Kinesis Firehose Data Delivery stream that writes agent status reports to S3 is named based on the application name. The name of the Firehose stream determines the filenames of agent productivity report files written to S3. By default, WFM Adapter expects the report files to be named based on an application name of "aspect-wfm-ap". If you specify a different application name, be sure to update the WFM Adapter configuration for the corresponding data source to match your application name, in the Firehose Report Mapping data source parameter. Consult the WFM Adapter help for more information.
Aspect WFM Report Prefix (AspectApDataS3Prefix)	OUTPUT/	Enter the prefix for Aspect WFM Agent Productivity reports written to S3. You should configure the corresponding WFM Adapter data source to match this value, in the Firehose Report Root Path data source parameter.
Number of SQS Queues (NumSQSQueues)	1	The number of SQS queues to create (must be between 1 and 10). Creating additional SQS queues can reduce the latency of state changes.
Write to SQS Flag Location (LambdaConfigLocation)	env	Once RTA reads an initial set of states from the agent event DynamoDB table, it toggles a Write to SQS flag to ensure that agent events are written to the SQS queue(s). This setting controls where that Write to SQS flag is stored either in an environment variable of the RTA Lambda (env) or in a configuration DynamoDB table (table). Storing this flag in an environment variable of the Lambda requires less overhead but requires greater permissions for the RTA user. Storing this flag in a configuration DynamoDB table requires fewer permissions for the RTA user but requires an additional DyanmoDB read each time the RTA Lambda executes. By default, this flag will be stored in an environment variable of the RTA Lambda.

Amazon Connect Configuration:

Parameter label (name)	Default	Description
Amazon Connect Exported Reports Bucket (ConnectS3Bucket)	<i>Requires input</i>	Enter the name of the Exported reports S3 bucket of your Amazon Connect instance.
Amazon Connect Exported Reports Prefix (ConnectS3Prefix)	connect/directory/Reports/Aspect/	Enter the prefix under which the exported reports read by WFM Adapter will be exported, in the Exported reports bucket of your Amazon Connect instance.
Amazon Connect Exported Reports Bucket KMS Key ARN (ConnectS3KmsArn)		Optional Amazon Resource Name (ARN) of the KMS key required to decrypt the exported reports of your Amazon Connect instance. You should supply this ARN if you configured your exported reports bucket to use server-side encryption with a customer-managed customer master key (CMK). Otherwise, you should leave this parameter blank.
Agent Hierarchy Group Level (AHGFilterLevel)	0	The Agent Hierarchy Group Level on which agent events should be filtered (must be between 0 and 5). Configuring a value of 0 disables filtering. Configuring a value between 1 and 5 means that agent events will be filtered based on that Agent Hierarchy Group Level.
Agent Hierarchy Group filter value separator (AHGFilterSeparator)	,	Character used to separate matching values listed in AHGFilterValues. You should configure a separator character that does not appear in any of the individual matching values.
Agent Hierarchy Group filter matching values (AHGFilterValues)		An optional list of matching Agent Hierarchy Group Level values. If you enable Agent Hierarchy Group filtering, any agent events with an Agent Hierarchy Group Level value matching one of the values listed here will be sent on to Real-Time Adherence and any agent events that do not match the filter will be discarded. Matching values should be separated with the character configured in AHGFilterSeparator.

AWS Quick Start Configuration:

Parameter label (name)	Default	Description
Quick Start S3 Bucket Name (QSS3BucketName)	quickstart-reference	The S3 bucket you have created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers,
Quick Start S3 Key Prefix (QSS3KeyPrefix)	connect-integration-aspect-wfm/	The S3 key name prefix used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes.

- **Option 4: Parameters for deploying Aspect Agent Productivity with an existing Kinesis data stream**

[View template](#)

Aspect WFM Configuration:

Parameter label (name)	Default	Description
Aspect Kinesis Application Name (KdaApplicationName)	aspect-wfm-ap	Enter the name of the Kinesis Data Analytics application to create. This string can include numbers, lowercase letters, uppercase letters, and hyphens (-). It cannot start or end with a hyphen (-). If you have multiple Amazon Connect instances in your AWS account, you will need to create a Kinesis Data Analytics application for each Amazon Connect instance and give each application a unique name. The Kinesis Firehose Data Delivery stream that writes agent status reports to S3 is named based on the application name. The name of the Firehose stream determines the filenames of agent productivity report files written to S3. By default, WFM Adapter expects the report files to be named based on an application name of "aspect-wfm-ap". If you specify a different application name, be sure to update the WFM Adapter configuration for the corresponding data source to match your application name, in the Firehose Report Mapping data source parameter. Consult the WFM Adapter help for more information.
Aspect WFM Report Prefix (AspectApDataS3Prefix)	OUTPUT/	Enter the prefix for Aspect WFM Agent Productivity reports written to S3. You should configure the corresponding WFM Adapter data source to match this value, in the Firehose Report Root Path data source parameter.

Amazon Connect Configuration:

Parameter label (name)	Default	Description
Amazon Connect Exported Reports Bucket (ConnectS3Bucket)	<i>Requires input</i>	Enter the name of the Exported reports S3 bucket of your Amazon Connect instance.
Amazon Connect Exported Reports Prefix (ConnectS3Prefix)	connect/directory/Reports/Aspect/	Enter the prefix under which the exported reports read by WFM Adapter will be exported, in the Exported reports bucket of your Amazon Connect instance.
Amazon Connect Exported Reports Bucket KMS Key ARN (ConnectS3KmsArn)		Optional Amazon Resource Name (ARN) of the KMS key required to decrypt the exported reports of your Amazon Connect instance. You should supply this ARN if you configured your exported reports bucket to use server-side encryption with a customer-managed customer master key (CMK). Otherwise, you should leave this parameter blank.
Agent Event Kinesis Stream ARN (AgentEventStreamArn)	<i>Requires input</i>	Enter the Amazon Resource Name (ARN) of the Amazon Kinesis Data Stream to which Amazon Connect delivers Agent Events.
Agent Event Kinesis Stream KMS Key ARN (AgentEventStreamKmsArn)		Optional Amazon Resource Name (ARN) of the KMS key required to decrypt agent events. You should supply this ARN if you configured your Agent Event Kinesis stream to use server-side encryption with a customer-managed customer master key (CMK). Otherwise, you should leave this parameter blank.

AWS Quick Start Configuration:

Parameter label (name)	Default	Description
Quick Start S3 Bucket Name (QSS3BucketName)	quickstart-reference	The S3 bucket you have created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers,
Quick Start S3 Key Prefix (QSS3KeyPrefix)	connect-integration-aspect-wfm/	The S3 key name prefix used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes.

- **Option 5: Parameters for deploying Aspect Real-Time Adherence with an existing Kinesis data stream**

[View template](#)

Aspect WFM Configuration:

Parameter label (name)	Default	Description
Number of SQS Queues (NumSQSQueues)	1	The number of SQS queues to create (must be between 1 and 10). Creating additional SQS queues can reduce the latency of state changes.
Write to SQS Flag Location (LambdaConfigLocation)	env	Once RTA reads an initial set of states from the agent event DynamoDB table, it toggles a Write to SQS flag to ensure that agent events are written to the SQS queue(s). This setting controls where that Write to SQS flag is stored either in an environment variable of the RTA Lambda (env) or in a configuration DynamoDB table (table). Storing this flag in an environment variable of the Lambda requires less overhead but requires greater permissions for the RTA user. Storing this flag in a configuration DynamoDB table requires fewer permissions for the RTA user but requires an additional DyanmoDB read each time the RTA Lambda executes. By default, this flag will be stored in an environment variable of the RTA Lambda.

Amazon Connect Configuration:

Parameter label (name)	Default	Description
Agent Event Kinesis Stream ARN (AgentEventStreamArn)	<i>Requires input</i>	Enter the Amazon Resource Name (ARN) of the Amazon Kinesis Data Stream to which Amazon Connect delivers Agent Events.
Agent Event Kinesis Stream KMS Key ARN (AgentEventStreamKmsArn)		Optional Amazon Resource Name (ARN) of the KMS key required to decrypt agent events. You should supply this ARN if you configured your Agent Event Kinesis stream to use server-side encryption with a customer-managed customer master key (CMK). Otherwise, you should leave this parameter blank.
Agent Hierarchy Group Level (AHGFilterLevel)	0	The Agent Hierarchy Group Level on which agent events should be filtered (must be between 0 and 5). Configuring a value of 0 disables filtering. Configuring a value between 1 and 5 means that agent events will be filtered based on that Agent Hierarchy Group Level.
Agent Hierarchy Group filter value separator (AHGFilterSeparator)	,	Character used to separate matching values listed in AHGFilterValues. You should configure a separator

Parameter label (name)	Default	Description
		character that does not appear in any of the individual matching values.
Agent Hierarchy Group filter matching values (AHGFilterValues)		An optional list of matching Agent Hierarchy Group Level values. If you enable Agent Hierarchy Group filtering, any agent events with an Agent Hierarchy Group Level value matching one of the values listed here will be sent on to Real-Time Adherence and any agent events that do not match the filter will be discarded. Matching values should be separated with the character configured in AHGFilterSeparator.

AWS Quick Start Configuration:

Parameter label (name)	Default	Description
Quick Start S3 Bucket Name (QSS3BucketName)	quickstart-reference	The S3 bucket you have created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers,
Quick Start S3 Key Prefix (QSS3KeyPrefix)	connect-integration-aspect-wfm/	The S3 key name prefix used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes.

- **Option 6: Parameters for deploying Aspect Agent Productivity and Real-Time Adherence with an existing Kinesis data stream**

[View template](#)

Aspect WFM Configuration:

Parameter label (name)	Default	Description
Aspect Kinesis Application Name (KdaApplicationName)	aspect-wfm-ap	Enter the name of the Kinesis Data Analytics application to create. This string can include numbers, lowercase letters, uppercase letters, and hyphens (-). It cannot start or end with a hyphen (-). If you have multiple Amazon Connect instances in your AWS account, you will need to create a Kinesis Data Analytics application for each Amazon Connect instance and give each application a unique name. The Kinesis Firehose Data Delivery stream that writes agent status reports to S3 is named based on the application name. The name of the Firehose stream determines the filenames of agent productivity report files written to S3. By default, WFM Adapter expects the report files to be named based on an application name of "aspect-wfm-ap". If you specify a different

Parameter label (name)	Default	Description
		application name, be sure to update the WFM Adapter configuration for the corresponding data source to match your application name, in the Firehose Report Mapping data source parameter. Consult the WFM Adapter help for more information.
Aspect WFM Report Prefix (AspectApDataS3Prefix)	OUTPUT/	Enter the prefix for Aspect WFM Agent Productivity reports written to S3. You should configure the corresponding WFM Adapter data source to match this value, in the Firehose Report Root Path data source parameter.
Number of SQS Queues (NumSQSQueues)	1	The number of SQS queues to create (must be between 1 and 10). Creating additional SQS queues can reduce the latency of state changes.
Write to SQS Flag Location (LambdaConfigLocation)	env	Once RTA reads an initial set of states from the agent event DynamoDB table, it toggles a Write to SQS flag to ensure that agent events are written to the SQS queue(s). This setting controls where that Write to SQS flag is stored either in an environment variable of the RTA Lambda (env) or in a configuration DynamoDB table (table). Storing this flag in an environment variable of the Lambda requires less overhead but requires greater permissions for the RTA user. Storing this flag in a configuration DynamoDB table requires fewer permissions for the RTA user but requires an additional DyanmoDB read each time the RTA Lambda executes. By default, this flag will be stored in an environment variable of the RTA Lambda.

Amazon Connect Configuration:

Parameter label (name)	Default	Description
Amazon Connect Exported Reports Bucket (ConnectS3Bucket)	<i>Requires input</i>	Enter the name of the Exported reports S3 bucket of your Amazon Connect instance.
Amazon Connect Exported Reports Prefix (ConnectS3Prefix)	connect/directory/Reports/Aspect/	Enter the prefix under which the exported reports read by WFM Adapter will be exported, in the Exported reports bucket of your Amazon Connect instance.
Amazon Connect Exported Reports Bucket KMS Key ARN (ConnectS3KmsArn)		Optional Amazon Resource Name (ARN) of the KMS key required to decrypt the exported reports of your Amazon Connect instance. You should supply this ARN if you configured your exported reports bucket to use server-side encryption with a customer-managed customer master key

Parameter label (name)	Default	Description
		(CMK). Otherwise, you should leave this parameter blank.
Agent Event Kinesis Stream ARN (AgentEventStreamArn)	<i>Requires input</i>	Enter the Amazon Resource Name (ARN) of the Amazon Kinesis Data Stream to which Amazon Connect delivers Agent Events.
Agent Event Kinesis Stream KMS Key ARN (AgentEventStreamKmsArn)		Optional Amazon Resource Name (ARN) of the KMS key required to decrypt agent events. You should supply this ARN if you configured your Agent Event Kinesis stream to use server-side encryption with a customer-managed customer master key (CMK). Otherwise, you should leave this parameter blank.
Agent Hierarchy Group Level (AHGFilterLevel)	0	The Agent Hierarchy Group Level on which agent events should be filtered (must be between 0 and 5). Configuring a value of 0 disables filtering. Configuring a value between 1 and 5 means that agent events will be filtered based on that Agent Hierarchy Group Level.
Agent Hierarchy Group filter value separator (AHGFilterSeparator)	,	Character used to separate matching values listed in AHGFilterValues. You should configure a separator character that does not appear in any of the individual matching values.
Agent Hierarchy Group filter matching values (AHGFilterValues)		An optional list of matching Agent Hierarchy Group Level values. If you enable Agent Hierarchy Group filtering, any agent events with an Agent Hierarchy Group Level value matching one of the values listed here will be sent on to Real-Time Adherence and any agent events that do not match the filter will be discarded. Matching values should be separated with the character configured in AHGFilterSeparator.

AWS Quick Start Configuration:

Parameter label (name)	Default	Description
Quick Start S3 Bucket Name (QSS3BucketName)	quickstart-reference	The S3 bucket you have created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers,

Parameter label (name)	Default	Description
Quick Start S3 Key Prefix (QSS3KeyPrefix)	connect- integration-aspect- wfm/	The S3 key name prefix used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes.

- On the **Options** page, you can [specify tags](#) (key-value pairs) for resources in your stack and [set advanced options](#). When you're done, choose **Next**.
- On the **Review** page, review and confirm the template settings. Under **Capabilities**, select the check box to acknowledge that the template will create IAM resources.
- Choose **Create** to deploy the stack.
- Monitor the status of the stack. When the status is **CREATE_COMPLETE**, the Aspect Workforce Management integration is ready.
- Use the URLs displayed in the **Outputs** tab for the stack to view the resources that were created.

Step 4. Enable Data Streaming

If you selected a deployment option that created a Kinesis data stream, follow the steps in the Amazon Connect documentation to set up [data streaming](#). Enable agent event streaming to the Kinesis stream created by the integration. The Amazon Resource Name (ARN) of the Kinesis stream created will be listed in the Outputs of the CloudFormation stack for the integration. The ARN will be listed as the value of the output with a key of "AgentEventStreamARN."

Step 5. Test the Deployment

If you chose an option that includes Agent Productivity, you should follow these steps to test your deployment.

First, you should populate the reference table that describes the Agent Statuses configured in your Amazon Connect instance. The reference table is a file named ASPECT_AGENT_STATUS_INFO.csv. This CSV file has the following form:

```
StateName,StateType
Available,A
Offline,0
```

This file should contain the Agent Status with a Type of Routable (by default, this Agent Status is named “Available”) marked with a StateType of A. This file should also contain any Agent Statuses that should be treated as signed-out, including the Agent Status with a Type of Offline (by default, this Agent Status is named “Offline”) marked with a StateType of O. Any Agent Statuses that are not listed in the file will be treated as signed-in, which is equivalent to a StateType of 1.

Once you have created the appropriate reference table for your Amazon Connect instance, you should add it to the S3 bucket that was created when you deployed this Quick Start. Add the reference table at the root level of the S3 bucket (no prefix) and make sure that you use an object key of ASPECT_AGENT_STATUS_INFO.csv.

Now that you have created the reference table, you can start the Kinesis Data Analytics Application. To do so, select the application in the Kinesis Analytics applications list and choose and Action of Start Application. It will take a few minutes to start the application. Once the application has started, you can test it using the Real time analytics feature of Kinesis Data Analytics.

Note You will need to login to Amazon Connect as an agent and change state to see any results, if there are currently no agents signed into Amazon Connect.

If you chose an option that includes Real-Time Adherence, you will need to configure the Aspect Real-Time Adherence product to test your installation. Consult the Aspect Real-Time Adherence product documentation for details.

Best Practices Using Aspect Workforce Management on AWS

Deployment options that include Agent Productivity create an S3 bucket to hold agent sign-in/sign-out information. Once this bucket is created, you should configure lifecycle rules that are consistent with your policies for other S3 buckets in your AWS account. Consult the Amazon S3 documentation for more information.

Keep in mind that lifecycle rules may affect that your ability to reship reports in Aspect WFM Adapter. You will be limited to the date/time range for which objects are stored in both the S3 bucket created by Aspect Agent Productivity and the Exported reports bucket associated with your Amazon Connect instance. Make sure that your lifecycle rules persist objects during the window of time for which you would like to be able to reship reports.

Both the Agent Productivity and Real-Time Adherence features read from the Agent Event Kinesis data stream of your Amazon Connect instance. If you selected a deployment option that created this Kinesis data stream, the Kinesis data stream will be created with a default shard count of 1. You may need to increase the shard count for the stream to support the Agent Productivity and/or Real-Time Adherence features of Aspect Workforce Management, whether the stream was created by this integration or not. When you size the stream, you need to consider both:

- The amount of data you expect to be written. Amazon Connect writes to the stream when agents change state.
- The amount of data you expect to be read. If you included Agent Productivity support, the corresponding Kinesis Data Analytics application will read from the stream. If you included Real-Time Adherence support, the corresponding Lambda function will read from the stream. You may need to deploy multiple instances of the RTA Lambda function, if you wish to monitor agents an Amazon Connect instance on multiple Real-Time Adherence servers. In addition, you may have other custom applications that read from the stream.

Consult the Amazon Kinesis documentation for more information.

Security

The Aspect WFM Adapter product will need access to two S3 buckets:

- The Exported reports bucket of your Amazon Connect instance
- The S3 bucket created by this deployment

If the integration includes Agent Productivity, it will create an IAM user with the minimum permissions required by Aspect WFM Adapter. To grant access, you will need to create an access key for this user. To do so:

1. Select the IAM Service in the AWS Management Console.
2. Select Users.
3. Select the User created by this integration. To determine the user name, check the Outputs of the CloudFormation stack for the integration. This user name will be listed as the value of the output with a key of “WFMAadapterUserName.”
4. Select the Security credentials tab.
5. Press the Create access key button.

6. Supply the Access key ID and Secret access key from the Create access key dialog when you configure a Data Source for WFM Adapter. Consult the Aspect WFM Adapter documentation for more information.

If you rotate the access keys for this user, you will need to update the corresponding WFM Adapter Data Source with the new key information.

The Aspect Real-Time Adherence product will need access to the following:

- The DynamoDB table that holds the last state change of each signed-in agent.
- The SQS queue(s) that are used to relay agent events to Aspect Real-Time Adherence.
- If configuration information is stored in a table:
 - The configuration DynamoDB table.
- Otherwise,
 - The configuration of the Lambda function that updates the DynamoDB table and SQS queue(s).

Aspect Real-Time Adherence updates the configuration DynamoDB table or an environment variable to instruct the Lambda function to start writing agent events to the SQS queue(s), after it reads the contents of the DynamoDB table.

If the integration includes Real-Time Adherence, it will create an IAM user with the minimum permissions required by Aspect Real-Time Adherence. To grant access, you will need to create an access key for this user. To do so:

1. Select the IAM Service in the AWS Management Console.
2. Select Users.
3. Select the User created by this integration. To determine the user name, check the Outputs of the CloudFormation stack for the integration. This user name will be listed as the value of the output with a key of “RTAUserName.”
4. Select the Security credentials tab.
5. Press the Create access key button.
6. Supply the Access key ID and Secret access key from the Create access key dialog when you add an Amazon Connect ACD instance in the RTA Server Configuration utility. Consult the Aspect Real-Time Adherence documentation for more information.

If you rotate the access keys for this user, you will need to update Aspect Real-Time Adherence with the new key information.

If you enable server-side encryption on either the Exported reports S3 bucket of your Amazon Connect or the Agent event Kinesis stream of your Amazon Connect, you may need to supply the ARN of the corresponding KMS key during deployment. Server-side encryption can be configured in several ways. If you use a customer managed KMS master key, additional permissions will be required to grant the AWS QuickStart object access to decrypt data from that resource. If AWS is responsible for managing the KMS master key, no additional permissions are required.

By default, the S3 bucket created when you create an Amazon Connect instance is configured to use server-side encryption with an AWS-managed KMS master key (called `aws/connect`). If you use the default AWS-managed KMS master key, you do not need to supply a value for the Amazon Connect Exported Reports Bucket KMS Key ARN parameter when you deploy the AWS QuickStart. If you configure your Amazon Connect to encrypt exported reports using a KMS master key that you created, you should supply the ARN of that KMS master key as the value of the Amazon Connect Exported Reports Bucket KMS Key ARN parameter when you deploy the AWS QuickStart.

By default, Amazon Connect does not support encrypting the contents of the Agent event Kinesis stream with a customer managed KMS master key, since, by default, Amazon Connect would not have the right to encrypt agent events using that key. However, you can work around this issue by using the same customer managed KMS master key for the Kinesis stream that you chose to use for encrypting exported reports, call recordings, or chat transcripts. In that case, you should supply the ARN of the KMS master key as the value of the Agent Event Kinesis Stream KMS Key ARN parameter when you deploy the AWS QuickStart. If your Agent event Kinesis stream is not configured to use server-side encryption, or you use an AWS-managed KMS master key, you do not need to supply a value for the Agent Event Kinesis Stream KMS Key ARN parameter when you deploy the AWS QuickStart.

Note: if you choose one of the AWS QuickStart deployment options where the AWS QuickStart is responsible for creating a Kinesis stream, the Kinesis stream will not be configured to use server-side encryption. If you later choose to enable server-side encryption with a customer managed KMS master key, you will need to grant the required permissions by hand. If you plan configure your Agent event Kinesis stream to use server-side encryption with a customer managed KMS master key, Aspect recommends that you create the Kinesis stream before deploying the AWS QuickStart and then use one of the deployment options for an existing Kinesis stream.

FAQ

Q. I encountered a `CREATE_FAILED` error when I launched the Quick Start.

A. If AWS CloudFormation fails to create the stack, we recommend that you relaunch the template with **Rollback on failure** set to **No**. (This setting is under **Advanced** in the AWS CloudFormation console, **Options** page.) With this setting, the stack's state will be retained and the instance will be left running, so you can troubleshoot the issue. (Look at the log files in %ProgramFiles%\Amazon\EC2ConfigService and C:\cfn\log.)

Important When you set **Rollback on failure** to **No**, you will continue to incur AWS charges for this stack. Please make sure to delete the stack when you finish troubleshooting.

For additional information, see [Troubleshooting AWS CloudFormation](#) on the AWS website.

Q. I encountered a size limitation error when I deployed the AWS Cloudformation templates.

A. We recommend that you launch the Quick Start templates from the location we've provided or from another S3 bucket. If you deploy the templates from a local copy on your computer or from a non-S3 location, you might encounter template size limitations when you create the stack. For more information about AWS CloudFormation limits, see the [AWS documentation](#).

Git Repository

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

Additional Resources

AWS services

- Amazon Connect
<https://aws.amazon.com/documentation/connect/>
- Amazon S3
<https://aws.amazon.com/documentation/s3/>
- Amazon DynamoDB
<https://aws.amazon.com/documentation/dynamodb/>
- Amazon SQS
<https://aws.amazon.com/documentation/sqs/>

- Amazon Kinesis
<https://aws.amazon.com/documentation/kinesis/>
- AWS Lambda
<https://aws.amazon.com/documentation/lambda/>
- Amazon CloudWatch
<https://aws.amazon.com/documentation/cloudwatch/>
- AWS CloudFormation
<https://aws.amazon.com/documentation/cloudformation/>

Quick Start reference deployments

- AWS Quick Start home page
<https://aws.amazon.com/quickstart/>

Document Revisions

Date	Change	In sections
August 2018	Initial publication	—
November 2020	Updates for multiple SQS queues and Agent Hierarchy Group filtering	Architecture, Deployment Steps, Security
December 2020	Updates to support KMS encryption	Deployment Steps, Security
May 2021	Updates to support the Contact Data by Routing Profile report in WFM Adapter security changes for RTA	Overview, Deployment Steps, Security

© 2021, Aspect Software, Inc. All rights reserved.

Notices

The software included with this paper is licensed under the Apache License, Version 2.0 (the "License"). You may not use this file except in compliance with the License. A copy of the License is located at <http://aws.amazon.com/apache2.0/> or in the "license" file accompanying this file. This code is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.