

Developer Guide

Amazon Simple Workflow Service



Amazon Simple Workflow Service: Developer Guide

Table of Contents

What is Amazon SWF?	. 1
Workflow components	. 2
Workflow components	. 2
Running your workflow	3
Setting up your development environment	. 4
Develop with Amazon SDKs	. 5
Consider the Amazon Flow Framework	. 5
Getting started	. 6
About the Workflow	. 7
Prerequisites	. 7
Tutorial Steps	. 8
Part 1: Using Amazon SWF with the SDK for Ruby	. 8
Include the Amazon SDK for Ruby	8
Configuring the Amazon Session	9
Registering an Amazon SWF Domain	10
Next Steps	11
Part 2: Implementing the Workflow	12
Designing the Workflow	12
Setting up our Workflow Code	13
Registering the Workflow	14
Polling for Decisions	16
Starting the Workflow Execution	19
Next Steps	21
Part 3: Implementing the Activities	21
Defining a Basic Activity Type	21
Defining GetContactActivity	23
Defining SubscribeTopicActivity	25
Defining WaitForConfirmationActivity	29
Defining SendResultActivity	31
Next Steps	33
Part 4: Implementing the Activities Task Poller	33
Running the Workflow	36
Where Do I Go from Here?	40
Working in the console	41

	Registering a domain	41
	Registering workflow types	42
	Registering activity types	42
	Starting a workflow	43
	To start a workflow execution using the console	43
	Managing workflow executions	44
Ba	asic concepts	47
	Creating a workflow	48
	Modeling Your Workflow and Its Activities	49
	Running workflows	49
	Workflow history	50
	Object identifiers	55
	Domains	56
	Actors	57
	What is an Actor in Amazon SWF?	57
	Workflow Starters	58
	Deciders	58
	Activity Workers	60
	Data Exchange Between Actors	60
	Tasks	61
	Task lists	62
	Decision Task Lists	62
	Activity Task Lists	63
	Task Routing	63
	Workflow execution closure	64
	Workflow execution life cycle	65
	Workflow Execution Life Cycle	65
	Polling for tasks	71
A	dvanced concepts	73
	Versioning	73
	Signals	74
	Child workflows	76
	Markers	77
	Tags	79
	Manage tags	79
	Tag workflow executions	80

Control access to domains with tags	81
Exclusive choice	82
Timers	85
Cancelling activity tasks	85
Security	89
Data Protection	89
Encryption	90
Identity and Access Management	91
Audience	92
Authenticating with identities	92
Managing access using policies	95
Access Control	98
Policy actions	98
Policy resources	99
Policy condition keys	100
ACLs	100
ABAC	101
Temporary credentials	101
Principal permissions	102
Service roles	102
Service-linked roles	102
Identity-based policies	103
Resource-based policies	103
How Amazon Simple Workflow Service works with IAM	104
Identity-based policy examples	105
Basic Principles	107
Amazon SWF IAM Policies	109
API Summary	117
Tag-based Policies	125
Amazon VPC endpoints	126
Troubleshooting	128
Logging and Monitoring	130
Amazon SWF Metrics for CloudWatch	130
Viewing Amazon SWF Metrics	140
Recording to CloudTrail	144
EventBridge for Amazon SWF	151

Using Amazon User Notifications with Amazon SWF	159
Compliance Validation	160
Resilience	160
Infrastructure Security	161
Configuration and Vulnerability Analysis	162
Using the Amazon CLI	163
Working with APIs	165
Making HTTP Requests	165
HTTP Header Contents	166
HTTP Body Content	168
Sample JSON Request and Response	168
Calculating the HMAC-SHA Signature	169
List of Amazon SWF Actions	172
Actions Related to Activities	172
Actions Related to Deciders	172
Actions Related to Workflow Executions	173
Actions Related to Administration	173
Visibility Actions	174
Registering a Domain	175
See Also	176
Setting timeout values	176
Quotas on Timeout Values	176
Workflow Execution and Decision Task Timeouts	176
Activity Task Timeouts	177
See Also	178
Registering a Workflow Type	178
See Also	178
Registering an Activity Type	178
See Also	179
Lambda tasks	179
About Amazon Lambda	179
Benefits and limitations of using Lambda tasks	180
Using Lambda tasks in your workflows	180
Developing an Activity Worker	186
Polling for Activity Tasks	186
Performing the Activity Task	187

Reporting Activity Task Heartbeats	187
Completing or Failing an Activity Task	188
Launching Activity Workers	190
Developing deciders	190
Defining Coordination Logic	191
Polling for Decision Tasks	192
Applying the Coordination Logic	194
Responding with Decisions	194
Closing a Workflow Execution	195
Launching Deciders	197
Starting workflows	198
Setting task priority	199
Setting Task Priority for Workflows	200
Setting Task Priority for Activities	202
Actions that Return Task Priority Information	203
Handling errors	203
Validation Errors	204
Errors in Enacting Actions or Decisions	204
Timeouts	204
Errors raised by user code	205
Errors related to closing a workflow execution	205
Quotas	206
General Account Quotas for Amazon SWF	206
Quotas on Workflow Executions	207
Quotas on Task Executions	207
Amazon SWF throttling quotas	209
Throttling quotas for all Regions	209
Decision quotas for all Regions	211
Workflow-level quotas	212
Requesting a quota increase	212
Additional resources	213
Timeout Types	213
Timeouts in Workflow and Decision Tasks	213
Timeouts in Activity Tasks	214
Endpoints	216
Additional Documentation	216

Amazon Simple Workflow Service API Reference	216
Amazon Flow Framework Documentation	217
Amazon SDK Documentation	217
Amazon CLI Documentation	219
Web Resources	219
Amazon SWF Forum	219
Amazon SWF FAQ	219
Amazon SWF Videos	220
Ruby Flow Options	220
Continue to use the Ruby Flow Framework	
Migrate to the Java Flow Framework	221
Migrate to Step Functions	221
Use the Amazon SWF API directly	222
Document history	223

What is Amazon Simple Workflow Service?

With Amazon Simple Workflow Service (Amazon SWF) you can build, run, and scale background jobs that have parallel or sequential steps. You can coordinate work across distributed components and track the state of tasks.

In Amazon SWF, a *task* represents a logical unit of work that is performed by a component of your application. Coordinating tasks across includes managing inter-task dependencies, scheduling, and concurrency in the flow of your application. With Amazon SWF, you can control and coordinate tasks without worrying about underlying complexities, such as tracking progress and maintaining task state.

When using Amazon SWF, you implement *workers* to perform tasks. Workers can run either on cloud infrastructure, such as Amazon Elastic Compute Cloud (Amazon EC2), or on your own premises. You can create tasks that are long-running, or that may fail, time out, or require restarts —or that may complete with varying throughput and latency. Amazon SWF stores tasks and assigns them to workers when they are ready, tracks progress, and maintains state, including details of task completion.

To coordinate tasks, you write a program that gets the latest task state from Amazon SWF and uses that state to initiate subsequent tasks. Amazon SWF maintains an application's execution state durably, so your application is resilient to individual component failures. With Amazon SWF, you can build, deploy, scale, and modify application components independently.

Other Amazon workflow services

For most use cases, we recommend considering Amazon Step Functions for your workflow and orchestration needs.

With Step Functions, you can create workflows, also called *state machines*, to build distributed applications, automate processes, orchestrate microservices, and create data and machine learning pipelines. In the Step Functions' console or Amazon toolkit in VS Code, you can use the graphical Workflow Studio to visualize, edit, test, and debug your application's workflow.

For more technical information, see the <u>Amazon Step Functions Developer Guide</u>.

Developing workflow components with Amazon SWF

Developing distributed applications requires coordinating many components and dealing with latency and unreliability inherent in remote communication.

With Amazon Simple Workflow Service (Amazon SWF), you can develop asynchronous and distributed applications by providing a programming model and infrastructure for coordinating distributed components and maintaining their execution state in a reliable way. By relying on Amazon SWF, you are freed to focus on building the aspects of your application that differentiate it.

Components of a workflow

<u>Components of a workflow</u> The fundamental concept in Amazon SWF is the *workflow*. A workflow is a set of *activities* that carry out some objective, together with logic that coordinates the activities. For example, a workflow could receive a customer order and take whatever actions are necessary to fulfill the order.

Each workflow runs in a resource called a *domain*, which controls the workflow's scope. An Amazon account can have multiple domains, each of which can contain multiple workflows, but workflows in different domains can't interact.

When designing an Amazon SWF workflow, you define each of the required activities. You then register each activity with Amazon SWF as an activity type. You will provide a name, version, and timeout values. For example, a customer may have an expectation that an order will ship within 24 hours.

In the process of carrying out the workflow, some activities may need to be performed more than once, perhaps with varying inputs. For example, in a customer-order workflow, you might have an activity that handles purchased items. If the customer purchases multiple items, then this activity would have to run multiple times. Amazon SWF has the concept of an *activity task* that represents one invocation of an activity. In our example, the processing of each item would be represented by a single activity task.

An *activity worker* is a program that receives activity tasks, performs them, and provides results. The task might actually be performed by a person. For example, a statistical analyst might receive sets of data, analyze the data, and then send back their analysis.

Workflow components API Version 2012-01-25 2

Activity tasks, and the activity workers that perform them, can run synchronously or asynchronously. Workers can run in one location or be distributed across multiple computers, potentially in different geographic regions. Different activity workers can be written in different programming languages and run on different operating systems. For example, one activity worker might be running on a server in Asia, while another might be running on a mobile device in North America.

The coordination logic in a workflow is contained in a software program called a *decider*. A decider schedules activity tasks, provides input to activity workers, processes events that arrive while the workflow is in progress, and ends (or closes) the workflow after the objective has been met.

The role of the Amazon SWF service is to function as a reliable central hub through which data is exchanged between the decider, the activity workers, and other relevant entities such as the person administering the workflow. Amazon SWF also maintains the state of each workflow execution, which saves your application from having to store the state in a durable way.

The decider directs the workflow by receiving decision tasks from Amazon SWF and responding back to Amazon SWF with decisions. A decision represents an action or set of actions, which are the next steps in the workflow. A typical decision would be to schedule an activity task. Decisions can also be used to delay tasks with timers, request cancellation of in-progress tasks, and to complete workflows.

The mechanism by which both the activity workers and the decider receive their tasks (activity tasks and decision tasks respectively) is by polling the Amazon SWF service.

Amazon SWF informs the decider of the state of the workflow by including, with each decision task, a copy of the current workflow execution history. The workflow execution history is composed of events, where an event represents a significant change in the state of the workflow execution. Examples of events include task completion, task time outs, or the expiration of a timer. The history is a complete, consistent, and authoritative record of the workflow's progress.

Amazon SWF access control uses Amazon Identity and Access Management (IAM), so you can control access to Amazon resources. For example, you can allow a user to access your account, but only to run certain workflows in a particular domain.

Running your workflow

The following provide an overview of the steps necessary to develop and run a workflow in Amazon SWF:

Running your workflow API Version 2012-01-25 3

- 1. Write activity workers to perform the processing steps in your workflow.
- 2. Write a **decider** to handle the coordination logic of your workflow.
- 3. Register your activities and workflow with Amazon SWF.

You can do this step programmatically or by using the Amazon Web Services Management Console.

4. Start your activity workers and decider.

These actors can run on any computing device that can access an Amazon SWF endpoint. For example, you could use compute instances in the cloud, such as Amazon Elastic Compute Cloud (Amazon EC2); servers in your data center; or even a mobile device, to host a decider or activity worker. Once started, the decider and activity workers should start polling Amazon SWF for tasks.

5. Start one or more executions of your workflow.

You can start workflows programmatically or via the Amazon Web Services Management Console.

Each execution runs independently and you can provide each with its own set of input data. When an execution is started, Amazon SWF schedules the initial decision task. In response, your decider begins generating decisions that initiate activity tasks. Execution continues until your decider makes a decision to close the execution.

6. View workflow executions using the Amazon Web Services Management Console.

You can filter and view complete details of running and completed executions. For example, you can select an open execution to see which tasks have been completed and what their results were.

Setting up your development environment

You have the option of developing for Amazon SWF in any of the programming languages supported by Amazon. For Java developers, the Amazon Flow Framework is also available. For more information, see the <u>Amazon Flow Framework</u> website, and see <u>Amazon Flow Framework for Java Developer Guide</u>.

To reduce latency and to store data in a location that meets your requirements, Amazon SWF provides **endpoints** in different Regions.

Each endpoint in Amazon SWF is completely independent. Any domains, workflows, and activities you have registered in one Region will not share data or attributes with those in another Region.

When you register an Amazon SWF domain, workflow, or activity, it exists *only within the Region you registered it in*. For example, you could register a domain named SWF-Flows-1 in two different Regions, but they will share no data or attributes with each other — each acting as a completely independent domain.

For a list of Amazon SWF endpoints, see Regions and Endpoints.

Develop with Amazon SDKs

Amazon SWF is supported by the Amazon SDKs for Java, .NET, Node.js, PHP, Python, and Ruby, providing a convenient way to use the Amazon SWF HTTP API in the programming language of your choice.

You can develop *deciders*, *activity workers*, or *workflow starters* using the API exposed by these libraries. And, you can use visibility operations through these libraries so you can develop your own Amazon SWF monitoring and reporting tools.

To download tools for developing and managing applications on Amazon, including SDKs, go to the Developer Center.

For detailed information about the Amazon SWF operations in each SDK, refer to the languagespecific reference documentation for the SDK.

Consider the Amazon Flow Framework

The Amazon Flow Framework is an enhanced SDK for writing distributed, asynchronous programs that run as workflows on Amazon SWF. The framework is available for the Java programming language and provides classes for writing complex distributed programs.

With the Amazon Flow Framework, you use preconfigured types to map the definition of your workflow directly to methods in your program. The Amazon Flow Framework supports standard object-oriented concepts, such as exception-based error handling. Programs written with the Amazon Flow Framework can be created, run, and debugged entirely within your preferred editor or IDE. For more information, see the Amazon Flow Framework website, and see Amazon Flow Framework for Java Developer Guide.

Develop with Amazon SDKs API Version 2012-01-25 5

Getting started with Amazon SWF

You can get started with the following Amazon Simple Workflow Service workflow application which consists of a set of four activities that operate sequentially. The tutorial also covers the following topics:

- Setting default and execution-time workflow and activity options.
- Polling Amazon SWF for decision and activity tasks.
- Passing data between the activities and the workflow with Amazon SWF.
- Waiting for human tasks and reporting heartbeats to Amazon SWF from an activity task.
- Using Amazon SNS to create a topic, subscribe a user to it, and publish messages to subscribed endpoints.

You can use Amazon SWF and Amazon Simple Notification Service (Amazon SNS) together to emulate a "human task" workflow—one in which a human worker is required to perform some action and then communicate with Amazon SWF to launch the next activity in the workflow.

Because Amazon SWF is a cloud-based web service, communication with Amazon SWF can originate from anywhere a connection to the Internet is available. In this case, we will use Amazon SNS to communicate with the user by either email, an SMS text message, or both.

This tutorial uses the Amazon SDK for Ruby to access Amazon SWF and Amazon SNS, but there are many development options available, including the Amazon Flow Framework for Ruby, which provides easier coordination and communication with Amazon SWF.



Note

This tutorial uses the Amazon SDK for Ruby, but we recommend that you use the Amazon Flow Framework for Java.

Topics

- About the Workflow
- Prerequisites
- Tutorial Steps

- Subscription Workflow Tutorial Part 1: Using Amazon SWF with the Amazon SDK for Ruby
- Subscription Workflow Tutorial Part 2: Implementing the Workflow
- Subscription Workflow Tutorial Part 3: Implementing the Activities
- Subscription Workflow Tutorial Part 4: Implementing the Activities Task Poller
- Subscription Workflow Tutorial: Running the Workflow

About the Workflow

The workflow that we will be developing consists of four major steps:

- 1. Get a subscription address (email or SMS) from the user.
- 2. Create an SNS topic and subscribe the provided endpoints to the topic.
- 3. Wait for the user to confirm the subscription.
- 4. If the user confirms, publish a congratulatory message to the topic.

These steps include activities that are completely automated (steps 2 and 4), and others that require the workflow to wait for a human to provide some data to the activity before the workflow can progress (steps 1 and 3).

Each step relies on data that is generated by the previous step (you must have an endpoint before subscribing it to a topic, and you must have a topic subscription before you can wait for confirmation, etc.) This tutorial will also cover how to provide activity results upon completion, and how to pass input to a task that is being scheduled. Amazon SWF handles coordination and delivery of information between the activities and the workflow, and vice-versa.

We're also using both keyboard input and Amazon SNS to handle communication between Amazon SWF and the human who is providing data to the workflow. In practice, you can use many different techniques to communicate with human users, but Amazon SNS provides a very easy way to use email or text messages to notify the user about events in the workflow.

Prerequisites

To follow along with this tutorial, you will need the following:

Amazon Web Services account

About the Workflow API Version 2012-01-25 7

- Ruby interpreter
- Amazon SDK for Ruby

If you already have these set up, you're ready to continue. If you don't want to run the example, you can still follow the tutorial—much of the content in this tutorial applies to using Amazon SWF and Amazon SNS regardless of the development option you choose.

Tutorial Steps

This tutorial is divided into the following steps:

- 1. Subscription Workflow Tutorial Part 1: Using Amazon SWF with the Amazon SDK for Ruby
- 2. Subscription Workflow Tutorial Part 2: Implementing the Workflow
- 3. Subscription Workflow Tutorial Part 3: Implementing the Activities
- 4. Subscription Workflow Tutorial Part 4: Implementing the Activities Task Poller
- 5. Subscription Workflow Tutorial: Running the Workflow

Subscription Workflow Tutorial Part 1: Using Amazon SWF with the Amazon SDK for Ruby

Topics

- Include the Amazon SDK for Ruby
- Configuring the Amazon Session
- Registering an Amazon SWF Domain
- Next Steps

Include the Amazon SDK for Ruby

Begin by creating a file called utils.rb. The code in this file will obtain, or create if necessary, the Amazon SWF domain used by both the workflow and activities code and will provide a place to put code that is common to all of our classes.

First, we need to include the aws-sdk-v1 library in our code, so that we can use the features provided by the SDK for Ruby.

Tutorial Steps API Version 2012-01-25 8

```
require 'aws-sdk-v1'
```

This gives us access to the Amazon namespace, which provides the ability to set global sessionrelated values, such as your Amazon credentials and region, and also provides access to the Amazon service APIs.

Configuring the Amazon Session

We'll configure the Amazon Session by setting our Amazon credentials (which are needed for accessing Amazon services) and the Amazon region to use.

There are a number of ways to set Amazon credentials in the Amazon SDK for Ruby: by setting them in environment variables (AWS_ACCESS_KEY_ID and AWS_SECRET_ACCESS_KEY) or by setting them with AWS.config. We'll use the latter method, loading them from a YAML configuration file, called aws-config.txt, that looks like this.

```
:access_key_id: REPLACE_WITH_ACCESS_KEY_ID
:secret_access_key: REPLACE_WITH_SECRET_ACCESS_KEY
```

Create this file now, replacing the strings beginning with REPLACE_WITH_ with your Amazon access key ID and secret access key. For information about your Amazon access keys, see How Do I Get Security Credentials? in the Amazon Web Services General Reference.

We also need to set the Amazon region to use. Because we'll be using the Short Message Service (SMS) to send text messages to the user's phone with Amazon SNS, we need to make sure that we're using region supported by Amazon SNS. See Supported Regions and Countries in the Amazon Simple Notification Service Developer Guide.

Note

If you don't have access to us-east-1, or don't care about running the demo with SMS messaging enabled, feel free to use any region you wish to. You can remove the SMS functionality from the sample and use email as the sole endpoint to subscribe to the Amazon SNS topic.

For more information about sending SMS messages, see Sending and Receiving SMS Notifications Using Amazon SNS in the Amazon Simple Notification Service Developer Guide. We'll now add some code to utils.rb to load the config file, get the user's credentials, then provide both the credentials and region to AWS.config.

```
require 'yaml'
# Load the user's credentials from a file, if it exists.
begin
  config_file = File.open('aws-config.txt') { |f| f.read }
  puts "No config file! Hope you set your Amazon credentials in the environment..."
end
if config_file.nil?
  options = { }
  options = YAML.load(config_file)
end
# SMS Messaging (which can be used by Amazon SNS) is available only in the
# `us-east-1` region.
$SMS_REGION = 'us-east-1'
options[:region] = $SMS_REGION
# Now, set the options
AWS.config = options
```

Registering an Amazon SWF Domain

To use Amazon SWF, you need to set up a *domain*: a named entity that will hold your workflows and activities. You can have many Amazon SWF domains registered, but they must all have unique names within your Amazon account, and workflows can't interact across domains: All of the workflows and activities for your application must be in the same domain to interact with one another.

Because we'll be using the same domain throughout our application, we'll create a function in utils.rb called init_domain, that will retrieve the Amazon SWF domain named SWFSampleDomain.

Once you have registered a domain, you can reuse it for many workflow executions. However, it is an error to try to register a domain that already exists, so our code will first check to see if the

domain exists, and will use the existing domain if it can be found. If the domain can't be found, we'll create it.

To work with Amazon SWF domains in the SDK for Ruby, use <u>AWS::SimpleWorkflow.domains</u>, which returns a <u>DomainCollection</u> that can be used to both enumerate and register domains:

- To check to see if a domain is already registered, you can look at the list provided by AWS::Simpleworkflow.domains.registered.
- To register a new domain, use AWS::Simpleworkflow.domains.register.

Here is the code for init_domain in utils.rb.

```
# Registers the domain that the workflow will run in.
def init_domain
  domain_name = 'SWFSampleDomain'
  domain = nil
  swf = AWS::SimpleWorkflow.new
  # First, check to see if the domain already exists and is registered.
  swf.domains.registered.each do | d |
    if(d.name == domain_name)
      domain = d
      break
    end
  end
  if domain.nil?
    # Register the domain for one day.
    domain = swf.domains.create(
      domain_name, 1, { :description => "#{domain_name} domain" })
  end
  return domain
end
```

Next Steps

Next, you will create the workflow and starter code in <u>Subscription Workflow Tutorial Part 2:</u> Implementing the Workflow.

Next Steps API Version 2012-01-25 11

Subscription Workflow Tutorial Part 2: Implementing the Workflow

Up until now, our code has been pretty generic. This is the part where we begin to really define what our workflow does, and what activities we'll need to implement it.

Topics

- Designing the Workflow
- Setting up our Workflow Code
- · Registering the Workflow
- Polling for Decisions
- Starting the Workflow Execution
- Next Steps

Designing the Workflow

If you recall, the initial idea for this workflow consisted of the following steps:

- 1. Get a subscription address (email or SMS) from the user.
- 2. Create an SNS topic and subscribe the provided endpoints to the topic.
- 3. Wait for the user to confirm the subscription.
- 4. If the user confirms, publish a congratulatory message to the topic.

We can think of each step in our workflow as an *activity* that it must perform. Our *workflow* is responsible for scheduling each activity at the appropriate time, and coordinating data transfer between activities.

For this workflow, we'll create a separate activity for each of these steps, naming them descriptively:

- 1. get_contact_activity
- 2. subscribe_topic_activity
- 3. wait_for_confirmation_activity

4. send_result_activity

These activities will be executed in order, and data from each step will be used in the subsequent step.

We could design our application so that all of the code exists in one source file, but this runs contrary to the way that Amazon SWF was designed. It is designed for workflows that can span the entire Internet in scope, so let's at least break the application up into two separate executables:

- swf_sns_workflow.rb Contains the workflow and workflow starter.
- swf_sns_activities.rb Contains the activities and activities starter.

The workflow and activity implementations can be run in separate windows, separate computers, or even different parts of the world. Because Amazon SWF is keeping track of the details of your workflows and activities, your workflow can coordinate scheduling and data transfer of your activities no matter where they are running.

Setting up our Workflow Code

We'll begin by creating a file called swf_sns_workflow.rb. In this file, declare a class called **SampleWorkflow**. Here is the class declaration and its constructor, the initialize method.

```
require_relative 'utils.rb'

# SampleWorkflow - the main workflow for the SWF/SNS Sample

# 
# See the file called `README.md` for a description of what this file does.

class SampleWorkflow

attr_accessor :name

def initialize(workflowId)

# the domain to look for decision tasks in.

@domain = init_domain

# the task list is used to poll for decision tasks.

@workflowId = workflowId

# The list of activities to run, in order. These name/version hashes can be
```

As you can see, we are keeping the following class instance data:

- domain The domain name retrieved from init_domain in utils.rb.
- workflowId The task list passed in to initialize.
- activity_list The activity list, which has the names and versions of the activities we'll run.

The domain name, activity name, and activity version are enough for Amazon SWF to positively identify an activity type, so that is all of the data we need to keep about our activities in order to schedule them.

The task list will be used by the workflow's *decider* code to poll for decision tasks and schedule activities.

At the end of this function, we call a method we haven't yet defined: register_workflow. We'll define this method next.

Registering the Workflow

To use a workflow type, we must first register it. Like an activity type, a workflow type is identified by its domain, name, and version. Also, like both domains and activity types, you can't re-register an existing workflow type. If you need to change anything about a workflow type, you must provide it with a new version, which essentially creates a new type.

Here is the code for register_workflow, which is used to either retrieve the existing workflow type we registered on a previous run or to register the workflow if it has not yet been registered.

```
# Registers the workflow
def register_workflow
workflow_name = 'swf-sns-workflow'
```

Registering the Workflow API Version 2012-01-25 14

```
@workflow_type = nil
   # a default value...
   workflow_version = '1'
   # Check to see if this workflow type already exists. If so, use it.
   @domain.workflow_types.each do | a |
     if (a.name == workflow_name) && (a.version == workflow_version)
       @workflow_type = a
     end
   end
   if @workflow_type.nil?
     options = {
       :default_child_policy => :terminate,
       :default_task_start_to_close_timeout => 3600,
       :default_execution_start_to_close_timeout => 24 * 3600 }
     puts "registering workflow: #{workflow_name}, #{workflow_version},
#{options.inspect}"
     @workflow_type = @domain.workflow_types.register(workflow_name, workflow_version,
options)
   end
   puts "** registered workflow: #{workflow_name}"
 end
```

First, we check to see if the workflow name and version is already registered by iterating through the domain's <u>workflow_types</u> collection. If we find a match, we'll use the workflow type that was already registered.

If we don't find a match, then a new workflow type is registered (by calling <u>register</u> on the same workflow_types collection that we were searching for the workflow in) with the name 'swf-sns-workflow', version '1', and the following options.

```
options = {
  :default_child_policy => :terminate,
  :default_task_start_to_close_timeout => 3600,
  :default_execution_start_to_close_timeout => 24 * 3600 }
```

Options passed in during registration are used to set *default behavior* for our workflow type, so we don't need to set these values every time we start a new workflow execution.

Registering the Workflow API Version 2012-01-25 15

Here, we just set some timeout values: the maximum time it can take from the time a task starts to when it closes (one hour), and the maximum time it can take for the workflow execution to complete (24 hours). If either of these times are exceeded, the task or workflow will timeout.

For more information about timeout values, see Amazon SWF Timeout Types.

Polling for Decisions

At the heart of every workflow execution there is a *decider*. The decider's responsibility is for managing the execution of the workflow itself. The decider receives *decision tasks* and responds to them, either by scheduling new activities, cancelling and restarting activities, or by setting the state of the workflow execution as complete, cancelled, or failed.

The decider uses the workflow execution's *task list* name to receive decision tasks to respond to. To poll for decision tasks, call <u>poll</u> on the domain's <u>decision_tasks</u> collection to loop over available decision tasks. You can then check for new events in the decision task by iterating over its new_events collection.

The returned events are <u>AWS::SimpleWorkflow::HistoryEvent</u> objects, and you can get the type of the event by using the returned event's <u>event_type</u> member. For a list and description of history event types, see <u>HistoryEvent</u> in the *Amazon Simple Workflow Service API Reference*.

Here is the beginning of the decision task poller's logic. A new method in our workflow class called poll_for_decisions.

```
def poll_for_decisions
  # first, poll for decision tasks...
  @domain.decision_tasks.poll(@workflowId) do | task |
    task.new_events.each do | event |
    case event.event_type
```

We'll now branch the execution of our decider based on the event_type that is received. The first one we are likely to receive is **WorkflowExecutionStarted**. When this event is received, it means that Amazon SWF is signaling to your decider that it should begin the workflow execution. We'll begin by scheduling the first activity by calling schedule_activity_task on the task we received while polling.

We'll pass it the first activity we declared in our activity list, which, because we reversed the list so we can use it like a stack, occupies the last position on the list. The "activities" we defined are just

Polling for Decisions API Version 2012-01-25 16

maps consisting of a name and version number, but this is all that Amazon SWF needs to identify the activity for scheduling, assuming that the activity has already been registered.

```
when 'WorkflowExecutionStarted'
    # schedule the last activity on the (reversed, remember?) list to
    # begin the workflow.
    puts "** scheduling activity task: #{@activity_list.last[:name]}"

task.schedule_activity_task( @activity_list.last,
    { :workflowId => "#{@workflowId}-activities" } )
```

When we schedule an activity, Amazon SWF sends an *activity task* to the activity task list that we pass in while scheduling it, signaling the task to begin. We'll deal with activity tasks in <u>Subscription Workflow Tutorial Part 3: Implementing the Activities</u>, but it is worth noting that we don't execute the task here. We only tell Amazon SWF that it should be *scheduled*.

The next activity that we'll need to address is the **ActivityTaskCompleted** event, which occurs when Amazon SWF has received an activity completed response from an activity task.

```
when 'ActivityTaskCompleted'
  # we are running the activities in strict sequential order, and
  # using the results of the previous activity as input for the next
  # activity.
  last_activity = @activity_list.pop
  if(@activity_list.empty?)
    puts "!! All activities complete! Sending complete_workflow_execution..."
    task.complete_workflow_execution
    return true;
  else
    # schedule the next activity, passing any results from the
    # previous activity. Results will be received in the activity
    # task.
    puts "** scheduling activity task: #{@activity_list.last[:name]}"
    if event.attributes.has_key?('result')
      task.schedule_activity_task(
        @activity_list.last,
        { :input => event.attributes[:result],
          :workflowId => "#{@workflowId}-activities" } )
    else
      task.schedule_activity_task(
        @activity_list.last, { :workflowId => "#{@workflowId}-activities" } )
```

Polling for Decisions API Version 2012-01-25 17

end end

Because we are executing our tasks in a linear fashion, and only one activity is executing at once, we'll take this opportunity to pop the completed task from the activity_list stack. If this results in an empty list, then we know that our workflow is complete. In this case, we signal to Amazon SWF that our workflow is complete by calling complete_workflow_execution on the task.

In the event that the list still has entries, we'll schedule the next activity on the list (again, in the last position). This time, however, we'll look to see if the previous activity returned any result data to Amazon SWF upon completion, which is provided to the workflow in the event's attributes, in the optional result key. If the activity generated a result, we'll pass it as the input option to the next scheduled activity, along with the activity task list.

By retrieving the result values of completed activities, and by setting the input values of scheduled activities, we can pass data from one activity to the next, or we can use data from an activity to change behavior in our decider based on the results from an activity.

For the purposes of this tutorial, these two event types are the most important in defining the behavior of our workflow. However, an activity can generate events other than ActivityTaskCompleted. We'll wrap up our decider code by providing demonstration handler code for the ActivityTaskTimedOut and ActivityTaskFailed events, and for the WorkflowExecutionCompleted event, which will be generated when Amazon SWF processes the complete_workflow_execution call that we make when we run out of activities to run.

```
when 'ActivityTaskTimedOut'
   puts "!! Failing workflow execution! (timed out activity)"
   task.fail_workflow_execution
   return false

when 'ActivityTaskFailed'
   puts "!! Failing workflow execution! (failed activity)"
   task.fail_workflow_execution
   return false

when 'WorkflowExecutionCompleted'
   puts "## Yesss, workflow execution completed!"
   task.workflow_execution.terminate
   return false
end
end
```

Polling for Decisions API Version 2012-01-25 18

```
end
end
```

Starting the Workflow Execution

Before any decision tasks will be generated for the workflow to poll for, we need to start the workflow execution.

To start the workflow execution, call <u>start_execution</u> on your registered workflow type (<u>AWS::SimpleWorkflow::WorkflowType</u>). We'll define a small wrapper around this to make use of the workflow_type instance member that we retrieved in the class constructor.

```
def start_execution
  workflow_execution = @workflow_type.start_execution( {
    :workflowId => @workflowId } )
  poll_for_decisions
  end
end
```

Once the workflow is executing, decision events will begin to appear on the workflow's task list, which is passed as a workflow execution option in start_execution.

Unlike options that are provided when the workflow type is registered, options that are passed to start_execution are not considered to be part of the workflow type. You are free to change these per workflow execution without changing the workflow's version.

Because we'd like the workflow to begin executing when we run the file, add some code that instantiates the class and then calls the start_execution method that we just defined.

```
if __FILE__ == $0
  require 'securerandom'

# Use a different task list name every time we start a new workflow execution.
#
# This avoids issues if our pollers re-start before SWF considers them closed,
# causing the pollers to get events from previously-run executions.
workflowId = SecureRandom.uuid

# Let the user start the activity worker first...

puts ""
puts "Amazon SWF Example"
```

```
puts "-----"
 puts ""
 puts "Start the activity worker, preferably in a separate command-line window, with"
 puts "the following command:"
 puts ""
 puts "> ruby swf_sns_activities.rb #{workflowId}-activities"
 puts "You can copy & paste it if you like, just don't copy the '>' character."
 puts ""
 puts "Press return when you're ready..."
 i = gets
 # Now, start the workflow.
 puts "Starting workflow execution."
 sample_workflow = SampleWorkflow.new(workflowId)
 sample_workflow.start_execution
end
```

To avoid any task list naming conflicts, we'll use SecureRandom.uuid to generate a random UUID that we can use as the task list name, guaranteeing that a different task list name is used for each workflow execution.

Note

Task lists are used to record events about a workflow execution, so if you use the same task list for multiple executions of the same workflow type, you might get events that were generated during a previous execution, especially if you are running them in near succession to each other, which is often the case when trying out new code or running tests.

To avoid the issue of having to deal with artifacts from previous executions, we can use a new task list for each execution, specifying it when we begin the workflow execution.

There is also a bit of code here to provide instructions for the person running it (probably you), and to provide the "activity" version of the task list. The decider uses this task list name to schedule activities for the workflow, and the activities implementation will listen for activity events on this task list name to know when to begin the scheduled activities and to provide updates about the activity execution.

The code also waits for the user to start running the activities starter *before* it starts the workflow execution, so the activities starter will be ready to respond when activity tasks begin appearing on the provided task list.

Next Steps

You have implemented the work flow. Next, you will define the activities and an activities starter, in Subscription Workflow Tutorial Part 3: Implementing the Activities.

Subscription Workflow Tutorial Part 3: Implementing the Activities

We'll now implement each of the activities in our workflow, beginning with a base class that provides some common features for the activity code.

Topics

- Defining a Basic Activity Type
- Defining GetContactActivity
- Defining SubscribeTopicActivity
- Defining WaitForConfirmationActivity
- Defining SendResultActivity
- Next Steps

Defining a Basic Activity Type

When designing the workflow, we identified the following activities:

- get_contact_activity
- subscribe_topic_activity
- wait_for_confirmation_activity
- send_result_activity

We'll implement each of these activities now. Because our activities will share some features, let's do a little groundwork and create some common code they can share. We'll call it **BasicActivity**, and define it in a new file called basic_activity.rb.

Next Steps API Version 2012-01-25 21

As with the other source files, we'll include utils.rb to access the init_domain function to set up the sample domain.

```
require_relative 'utils.rb'
```

Next, we'll declare the basic activity class and some common data that we'll be interested in for each activity. We'll save the activity's MYS::SimpleWorkflow::ActivityType instance, name, and results in attributes of the class.

```
class BasicActivity

attr_accessor :activity_type
attr_accessor :name
attr_accessor :results
```

These attributes access instance data that's defined in the class' initialize method, which takes an activity *name*, and an optional *version* and map of *options* to be used when registering the activity with Amazon SWF.

```
def initialize(name, version = 'v1', options = nil)
 @activity_type = nil
 @name = name
 @results = nil
 # get the domain to use for activity tasks.
 @domain = init_domain
 # Check to see if this activity type already exists.
 @domain.activity_types.each do | a |
   if (a.name == @name) && (a.version == version)
      @activity_type = a
    end
 end
 if @activity_type.nil?
   # If no options were specified, use some reasonable defaults.
    if options.nil?
      options = {
        # All timeouts are in seconds.
        :default_task_heartbeat_timeout => 900,
        :default_task_schedule_to_start_timeout => 120,
```

As with workflow type registration, if an activity type is already registered, we can retrieve it by looking at the domain's activity_types collection. If the activity can't be found, it will be registered.

Also, as with workflow types, you can set *default options* that are stored with your activity type when you register it.

The last thing our basic activity gets is a consistent way to run it. We'll define a do_activity method that takes an activity task. As shown, we can use the passed-in activity task to receive data via its input instance attribute.

```
def do_activity(task)
   @results = task.input # may be nil
   return true
  end
end
```

That wraps up the **BasicActivity** class. Now we'll use it to make defining our activities simple and consistent.

Defining GetContactActivity

The first activity that is run during a workflow execution is get_contact_activity, which retrieves the user's Amazon SNS topic subscription information.

Create a new file called get_contact_activity.rb, and require both yaml, which we'll use to prepare a string for passing to Amazon SWF, and basic_activity.rb, which we'll use as the basis for this **GetContactActivity** class.

```
require 'yaml'
require_relative 'basic_activity.rb'

# **GetContactActivity** provides a prompt for the user to enter contact
# information. When the user successfully enters contact information, the
# activity is complete.
```

Defining GetContactActivity API Version 2012-01-25 23

```
class GetContactActivity < BasicActivity</pre>
```

Because we put the activity registration code in **BasicActivity**, the initialize method for **GetContactActivity** is pretty simple. We simply call the base class constructor with the activity name, get_contact_activity. This is all that is required to register our activity.

```
# initialize the activity
def initialize
  super('get_contact_activity')
end
```

We'll now define the do_activity method, which prompts for the user's email and/or phone number.

```
def do_activity(task)
      puts ""
      puts "Please enter either an email address or SMS message (mobile phone) number
to"
      puts "receive SNS notifications. You can also enter both to use both address
types."
      puts ""
      puts "If you enter a phone number, it must be able to receive SMS messages, and
      puts "be 11 digits (such as 12065550101 to represent the number
1-206-555-0101)."
      input_confirmed = false
      while !input_confirmed
        puts ""
        print "Email: "
        email = $stdin.gets.strip
        print "Phone: "
        phone = $stdin.gets.strip
        puts ""
        if (email == '') && (phone == '')
          print "You provided no subscription information. Quit? (y/n)"
           confirmation = $stdin.gets.strip.downcase
           if confirmation == 'y'
             return false
           end
```

Defining GetContactActivity API Version 2012-01-25 24

```
else
         puts "You entered:"
         puts " email: #{email}"
         puts " phone: #{phone}"
         print "\nIs this correct? (y/n): "
         confirmation = $stdin.gets.strip.downcase
         if confirmation == 'y'
           input_confirmed = true
         end
      end
    end
    # make sure that @results is a single string. YAML makes this easy.
    @results = { :email => email, :sms => phone }.to_yaml
    return true
  end
end
```

At the end of do_activity, we take the email and phone number retrieved from the user, place it in a map and then use to_yaml to convert the entire map to a YAML string. There's an important reason for this: any results that you pass to Amazon SWF when you complete an activity must be string data only. Ruby's ability to easily convert objects to YAML strings and then back again into objects is, thankfully, well-suited for this purpose.

That's the end of the get_contact_activity implementation. This data will be used next in the subscribe_topic_activity implementation.

Defining SubscribeTopicActivity

We'll now delve into Amazon SNS and create an activity that uses the information generated by get_contact_activity to subscribe the user to an Amazon SNS topic.

Create a new file called subscribe_topic_activity.rb, add the same requirements that we used for get_contact_activity, declare your class, and provide its initialize method.

```
require 'yaml'
require_relative 'basic_activity.rb'

# **SubscribeTopicActivity** sends an SMS / email message to the user, asking for
# confirmation. When this action has been taken, the activity is complete.
class SubscribeTopicActivity < BasicActivity</pre>
```

```
def initialize
    super('subscribe_topic_activity')
end
```

Now that we have the code in place to get the activity set up and registered, we will add some code to create an Amazon SNS topic. To do so, we'll use the AWS::SNS::Client object's create_topic method.

Add the create_topic method to your class, which takes a passed-in Amazon SNS client object.

```
def create_topic(sns_client)
  topic_arn = sns_client.create_topic(:name => 'SWF_Sample_Topic')[:topic_arn]
 if topic_arn != nil
   # For an SMS notification, setting `DisplayName` is *required*. Note that
   # only the *first 10 characters* of the DisplayName will be shown on the
   # SMS message sent to the user, so choose your DisplayName wisely!
   sns_client.set_topic_attributes( {
      :topic_arn => topic_arn,
      :attribute_name => 'DisplayName',
      :attribute_value => 'SWFSample' } )
 else
   @results = {
      :reason => "Couldn't create SNS topic", :detail => "" }.to_yaml
   return nil
  end
 return topic_arn
end
```

Once we have the topic's Amazon Resource Name (ARN), we can use it with the Amazon SNS client's <u>set_topic_attributes</u> method to set the topic's *DisplayName*, which is required for sending SMS messages with Amazon SNS.

Lastly, we'll define the do_activity method. We'll start by collecting any data that was passed via the input option when the activity was scheduled. As previously mentioned, this must be passed as a string, which we created using to_yaml. When retrieving it, we'll use YAML.load to turn the data into Ruby objects.

Here's the beginning of do_activity, in which we retrieve the input data.

```
def do_activity(task)
```

```
activity_data = {
  :topic_arn => nil,
  :email => { :endpoint => nil, :subscription_arn => nil },
  :sms => { :endpoint => nil, :subscription_arn => nil },
}
if task.input != nil
  input = YAML.load(task.input)
  activity_data[:email][:endpoint] = input[:email]
  activity_data[:sms][:endpoint] = input[:sms]
else
  @results = { :reason => "Didn't receive any input!", :detail => "" }.to_yaml
  puts(" #{@results.inspect}")
  return false
end
# Create an SNS client. This is used to interact with the service. Set the
# region to $SMS_REGION, which is a region that supports SMS notifications
# (defined in the file `utils.rb`).
sns_client = AWS::SNS::Client.new(
  :config => AWS.config.with(:region => $SMS_REGION))
```

If we didn't receive any input, there isn't much to do, so we'll just fail the activity.

Assuming that everything is fine, however, we'll continue filling in our do_activity method, get an Amazon SNS client with the Amazon SDK for Ruby, and pass it to our create_topic method to create the Amazon SNS topic.

```
# Create the topic and get the ARN
activity_data[:topic_arn] = create_topic(sns_client)

if activity_data[:topic_arn].nil?
   return false
end
```

There are a couple of things worth noting here:

- We use <u>AWS.config.with</u> to set the region for our Amazon SNS client. Because we want to send SMS messages, we use the SMS-enabled region that we declared in utils.rb.
- We save the topic's ARN in our activity_data map. This is part of the data that will be passed to the *next* activity in our workflow.

Finally, this activity subscribes the user to the Amazon SNS topic, using the passed-in endpoints (email and SMS). We don't require the user to enter *both* endpoints, but we do need at least one.

```
# Subscribe the user to the topic, using either or both endpoints.
[:email, :sms].each do | x |
    ep = activity_data[x][:endpoint]
    # don't try to subscribe an empty endpoint
    if (ep != nil && ep != "")
     response = sns_client.subscribe( {
        :topic_arn => activity_data[:topic_arn],
        :protocol => x.to_s, :endpoint => ep } )
    activity_data[x][:subscription_arn] = response[:subscription_arn]
    end
end
```

<u>AWS::SNS::Client.subscribe</u> takes the topic ARN, the *protocol* (which, cleverly, we disguised as the activity_data map key for the corresponding endpoint).

Finally, we re-package the information for the next activity in YAML format, so that we can send it back to Amazon SWF.

```
# if at least one subscription arn is set, consider this a success.
   if (activity_data[:email][:subscription_arn] != nil) or (activity_data[:sms]
[:subscription_arn] != nil)
     @results = activity_data.to_yaml
     else
        @results = { :reason => "Couldn't subscribe to SNS topic", :detail =>
     "" }.to_yaml
        puts(" #{@results.inspect}")
        return false
     end
     return true
     end
     end
```

That completes the implementation of the subscribe_topic_activity. Next, we'll define wait_for_confirmation_activity.

Defining WaitForConfirmationActivity

Once a user is subscribed to an Amazon SNS topic, he or she will still need to confirm the subscription request. In this case, we'll be waiting for the user to confirm by either email or an SMS message.

The activity that waits for the user to confirm the subscription is called wait_for_confirmation_activity, and we'll define it here. To begin, create a new file called wait_for_confirmation_activity.rb and set it up as we've set up the previous activities.

```
require 'yaml'
require_relative 'basic_activity.rb'

# **WaitForConfirmationActivity** waits for the user to confirm the SNS
# subscription. When this action has been taken, the activity is complete. It
# might also time out...
class WaitForConfirmationActivity < BasicActivity

# Initialize the class
def initialize
    super('wait_for_confirmation_activity')
end</pre>
```

Next, we'll begin defining the do_activity method and retrieve any input data into a local variable called subscription_data.

```
def do_activity(task)
  if task.input.nil?
    @results = { :reason => "Didn't receive any input!", :detail => "" }.to_yaml
    return false
  end

subscription_data = YAML.load(task.input)
```

Now that we have the topic ARN, we can retrieve the topic by creating a new instance of AWS::SNS::Topic and pass it the ARN.

```
topic = AWS::SNS::Topic.new(subscription_data[:topic_arn])

if topic.nil?
  @results = {
```

```
:reason => "Couldn't get SWF topic ARN",
   :detail => "Topic ARN: #{topic.arn}" }.to_yaml
   return false
end
```

Now, we'll check the topic to see if the user has confirmed the subscription using one of the endpoints. We'll only require that one endpoint has been confirmed to consider the activity a success.

An Amazon SNS topic maintains a list of the <u>subscriptions</u> for that topic, and we can check whether or not the user has confirmed a particular subscription by checking to see if the subscription's ARN is set to anything other than PendingConfirmation.

```
# loop until we get some indication that a subscription was confirmed.
      subscription_confirmed = false
      while(!subscription_confirmed)
        topic.subscriptions.each do | sub |
          if subscription_data[sub.protocol.to_sym][:endpoint] == sub.endpoint
            # this is one of the endpoints we're interested in. Is it subscribed?
            if sub.arn != 'PendingConfirmation'
              subscription_data[sub.protocol.to_sym][:subscription_arn] = sub.arn
              puts "Topic subscription confirmed for (#{sub.protocol}:
#{sub.endpoint})"
              @results = subscription_data.to_yaml
              return true
            else
              puts "Topic subscription still pending for (#{sub.protocol}:
#{sub.endpoint})"
            end
          end
        end
```

If we get an ARN for the subscription, we'll save it in the activity's result data, convert it to YAML, and return true from do_activity, which signals that the activity completed successfully.

Because waiting for a subscription to be confirmed might take a while, we'll occasionally call record_heartbeat on the activity task. This signals to Amazon SWF that the activity is still processing, and can also be used to provide updates about the progress of the activity (if you are doing something, like processing files, that you can report progress for).

```
task.record_heartbeat!(
```

```
{ :details => "#{topic.num_subscriptions_confirmed} confirmed,
#{topic.num_subscriptions_pending} pending" })
    # sleep a bit.
    sleep(4.0)
end
```

This ends our while loop. If we somehow get out of the while loop without success, we'll report failure and finish the do_activity method.

That ends the implementation of wait_for_confirmation_activity. We have only one more activity to define: send_result_activity.

Defining SendResultActivity

If the workflow has progressed this far, we've successfully subscribed the user to an Amazon SNS topic and the user has confirmed the subscription.

Our last activity, send_result_activity, sends the user a confirmation of the successful topic subscription, using the topic that the user subscribed to and the endpoint that the user confirmed the subscription with.

Create a new file called send_result_activity.rb and set it up as we've set up all the activities so far.

```
require 'yaml'
require_relative 'basic_activity.rb'

# **SendResultActivity** sends the result of the activity to the screen, and, if
# the user successfully registered using SNS, to the user using the SNS contact
# information collected.
class SendResultActivity < BasicActivity</pre>
```

Defining SendResultActivity API Version 2012-01-25 31

```
def initialize
   super('send_result_activity')
end
```

Our do_activity method begins similarly, as well, getting the input data from the workflow, converting it from YAML, and then using the topic ARN to create an AWS::SNS::Topic instance.

```
def do_activity(task)
  if task.input.nil?
    @results = { :reason => "Didn't receive any input!", :detail => "" }
    return false
  end

input = YAML.load(task.input)

# get the topic, so we publish a message to it.
topic = AWS::SNS::Topic.new(input[:topic_arn])

if topic.nil?
  @results = {
    :reason => "Couldn't get SWF topic",
    :detail => "Topic ARN: #{topic.arn}" }
  return false
end
```

Once we have the topic, we'll publish a message to it (and echo it to the screen, as well).

```
@results = "Thanks, you've successfully confirmed registration, and your
workflow is complete!"

# send the message via SNS, and also print it on the screen.
topic.publish(@results)
puts(@results)

return true
end
end
```

Publishing to an Amazon SNS topic sends the message that you supply to *all* of the subscribed and confirmed endpoints that exist for that topic. So, if the user confirmed with *both* an email and an SMS number, he or she will receive two confirmation messages, one at each endpoint.

Defining SendResultActivity API Version 2012-01-25 32

Next Steps

This completes the implementation of send_result_activity. Now, you will tie all these activities together in an activity application that handles the activity tasks and can launch activities in response, in Subscription Workflow Tutorial Part 4: Implementing the Activities Task Poller.

Subscription Workflow Tutorial Part 4: Implementing the Activities Task Poller

In Amazon SWF, activity tasks for a running workflow execution appear on the *activity task list*, which is provided when you schedule an activity in the workflow.

We'll implement a basic activity poller to handle these tasks for our workflow, and use it to launch our activities when Amazon SWF places a task on the activity task list to start the activity.

To begin, create a new file called swf_sns_activities.rb. We'll use it to:

- Instantiate the activity classes that we created.
- Register each activity with Amazon SWF.
- Poll for activities and call do_activity for each activity when its name appears on the activity task list.

In swf_sns_activities.rb, add the following statements to require each of the activity classes we defined.

```
require_relative 'get_contact_activity.rb'
require_relative 'subscribe_topic_activity.rb'
require_relative 'wait_for_confirmation_activity.rb'
require_relative 'send_result_activity.rb'
```

Now, we'll create the class and provide some initialization code.

```
class ActivitiesPoller

def initialize(domain, workflowId)
   @domain = domain
   @workflowId = workflowId
```

Next Steps API Version 2012-01-25 33

```
@activities = {}

# These are the activities we'll run
activity_list = [
   GetContactActivity,
   SubscribeTopicActivity,
   WaitForConfirmationActivity,
   SendResultActivity ]

activity_list.each do | activity_class |
   activity_obj = activity_class.new
   puts "** initialized and registered activity: #{activity_obj.name}"
   # add it to the hash
   @activities[activity_obj.name.to_sym] = activity_obj
end
end
```

In addition to saving the passed in *domain* and *task list*, this code instantiates each of the activity classes we created. Because each class registers its associated activity (refer to basic_activity.rb if you need to review that code), this is enough to let Amazon SWF know about all of the activities we'll be running.

For each activity instantiated, we store it on a map using the activity name (such as get_contact_activity) as the key, so we can easily look these up in the activity poller code, which we'll define next.

Create a new method called poll_for_activities and call <u>poll</u> on the <u>activity_tasks</u> held by the domain to get activity tasks.

```
def poll_for_activities
  @domain.activity_tasks.poll(@workflowId) do | task |
    activity_name = task.activity_type.name
```

We can get the activity name from the task's <u>activity_type</u> member. Next, we'll use the activity name associated with this task to look up the class to run do_activity on, passing it the task (which includes any input data that should be transferred to the activity).

```
# find the task on the activities list, and run it.
if @activities.key?(activity_name.to_sym)
  activity = @activities[activity_name.to_sym]
  puts "** Starting activity task: #{activity_name}"
```

```
if activity.do_activity(task)
          puts "++ Activity task completed: #{activity_name}"
          task.complete!({ :result => activity.results })
          # if this is the final activity, stop polling.
          if activity_name == 'send_result_activity'
             return true
          end
        else
          puts "-- Activity task failed: #{activity_name}"
          task.fail!(
            { :reason => activity.results[:reason],
              :details => activity.results[:detail] } )
        end
      else
        puts "couldn't find key in @activities list: #{activity_name}"
        puts "contents: #{@activities.keys}"
      end
    end
  end
end
```

The code just waits for do_activity to complete, and then calls either <u>complete!</u> or <u>fail!</u> on the task based on the return code.

Note

This code exits from the poller once the final activity has been launched, because it has completed its mission and has launched all of the activities. In your own Amazon SWF code, if your activities might be run again, you may want to keep the activity poller running indefinitely.

That's the end of the code for our **ActivitiesPoller** class, but we'll add a little more code at the end of the file to allow the user to run it from the command-line.

```
if __FILE__ == $0
  if ARGV.count < 1
    puts "You must supply a task-list name to use!"
    exit
end
poller = ActivitiesPoller.new(init_domain, ARGV[0])
poller.poll_for_activities</pre>
```

```
puts "All done!"
end
```

If the user runs the file from the command line (passing it an activity task list as the first argument), this code will instantiate the poller class and start it polling for activities. Once the poller completes (after it has launched the final activity), we just print a message and exit.

That's it for the activities poller. All that's left for you to do is to run the code and see how it works, in Subscription Workflow Tutorial: Running the Workflow.

Subscription Workflow Tutorial: Running the Workflow

Now that you've completed the implementation of your workflow, activities, and the workflow and activity pollers, you're ready to run the workflow.

If you haven't done so already, you'll need to provide your Amazon access keys in the aws-config.txt file, like in Configuring the Amazon Session in Part 1 of the tutorial.

Now, go to your command line and change to the directory where the tutorial source files are located. You should have the following files:

```
.
|-- aws-config.txt
|-- basic_activity.rb
|-- get_contact_activity.rb
|-- send_result_activity.rb
|-- subscribe_topic_activity.rb
|-- swf_sns_activities.rb
|-- swf_sns_workflow.rb
|-- utils.rb
`-- wait_for_confirmation_activity.rb
```

Now, start the workflow with the following command.

```
ruby swf_sns_workflow.rb
```

This will begin the workflow, and should print out a message with a line that you can copy and paste into a separate command-line window (or even on another computer, if you've copied the tutorial source files onto it).

Running the Workflow API Version 2012-01-25 36

```
Amazon SWF Example
------
Start the activity worker, preferably in a separate command-line window, with the following command:

> ruby swf_sns_activities.rb 87097e76-7c0c-41c7-817b-92527bb0ea85-activities

You can copy & paste it if you like, just don't copy the '>' character.

Press return when you're ready...
```

The workflow code will wait patiently for you to start the activity poller in a separate window.

Open a new command-line window, change to the directory where the source files are located again, and then use the command provided by the swf_sns_workflow.rb file to start the activity poller. For example, if you received the preceding output, you would type (or paste) the following.

```
ruby swf_sns_activities.rb 87097e76-7c0c-41c7-817b-92527bb0ea85-activities
```

Once you begin running your activity poller, it will start to output information about activities registration.

```
** initialized and registered activity: get_contact_activity

** initialized and registered activity: subscribe_topic_activity

** initialized and registered activity: wait_for_confirmation_activity

** initialized and registered activity: send_result_activity
```

You can now return to your original command-line window, and press return to start your workflow execution. It will register the workflow and schedule the first activity.

```
Starting workflow execution.

** registered workflow: swf-sns-workflow

** scheduling activity task: get_contact_activity
```

Go back to the other window, where your activity poller is running. The result of the first running activity is displayed, providing a prompt for you to enter your email or SMS phone number. Enter either, or both, of these pieces of data, and then confirm your text entry.

Running the Workflow API Version 2012-01-25 37

activity task received: <AWS::SimpleWorkflow::ActivityTask> ** Starting activity task: get_contact_activity

Please enter either an email address or SMS message (mobile phone) number to receive Amazon SNS notifications. You can also enter both to use both address types.

If you enter a phone number, it must be able to receive SMS messages, and must be 11 digits (such as 12065550101 to represent the number 1-206-555-0101).

Email: me@example.com Phone: 12065550101

You entered:

email: me@example.com phone: 12065550101

Is this correct? (y/n): y



(i) Note

The phone number provided is fictitious, and is used only for illustrative purposes. Use your own phone number and email address here!

Soon after entering this information, you should receive an email or text message from Amazon SNS, asking you to confirm your topic subscription. If you entered an SMS number, you will see something like the following appear on your phone.

Would you like to receive messages from SWFSAMPLE? Reply YES SWFSAMPLE to receive messages. Reply HELP or STOP. Msg&data rates may apply.

3:39 PM

Running the Workflow API Version 2012-01-25 38 If you reply to this message with YES, you'll get the response that we provided in send_result_activity.

SWFSAMPLE> Thanks, you've successfully confirmed registration, and your workflow is complete!

3:39 PM

While all of this was happening, did you see what was happening in your command-line window? Both the workflow and activity pollers have been hard at work.

Here's the output from the workflow poller.

```
** scheduling activity task: subscribe_topic_activity

** scheduling activity task: wait_for_confirmation_activity

** scheduling activity task: send_result_activity
!! All activities complete! Sending complete_workflow_execution...
```

Here's the output from the activity poller, which was happening at the same time in another command-line window.

```
++ Activity task completed: get_contact_activity
** Starting activity task: subscribe_topic_activity
++ Activity task completed: subscribe_topic_activity
** Starting activity task: wait_for_confirmation_activity
Topic subscription still pending for (email: me@example.com)
Topic subscription confirmed for (sms: 12065550101)
++ Activity task completed: wait_for_confirmation_activity
** Starting activity task: send_result_activity
Thanks, you've successfully confirmed registration, and your workflow is complete!
++ Activity task completed: send_result_activity
All done!
```

Congratulations, your workflow is complete, and so is this tutorial!

Running the Workflow API Version 2012-01-25 39

You may want to re-run the workflow again to see how timeouts work, or to enter different data. Just remember that once you subscribe to a topic, *you're still subscribed until you unsubscribe*. Rerunning the workflow before unsubscribing to topics will probably result in automatic success, because the wait_for_confirmation_activity will see that your subscription is already confirmed.

To unsubscribe from the Amazon SNS topic

- Respond in the negative (send STOP) to the text message.
- Choose the unsubscribe link that you received in your email.

You're now ready to re-subscribe to the topic again.

Where Do I Go from Here?

This tutorial has covered a lot of ground, but there's still much more you can learn about the Amazon SDK for Ruby, Amazon SWF, or Amazon SNS. For more information and many more examples, see the official documentation for each:

- Amazon SDK for Ruby Documentation
- Amazon Simple Notification Service Documentation
- Amazon Simple Workflow Service Documentation

Where Do I Go from Here?

API Version 2012-01-25 40

Working in the Amazon SWF console

The Amazon SWF console provides options to configure, initiate, and manage workflow executions.

With the Amazon SWF console, you can:

- Register workflow domains.
- Register workflow types, and activity types.
- Start, view, signal, cancel, terminate, and restart workflow executions.

Registering a domain

Workflows run in an Amazon resource called a *domain*, which controls the workflow's scope. An Amazon account can have multiple domains, each of which can contain multiple workflows, but workflows in different domains can't interact.

Domain registration is the only functionality initially available in the console. After at least one domain is registered, you can perform the following actions for the domain:

- Register workflow and activity types.
- Initiate workflow executions.
- Cancel, terminate, and send signals to running workflow executions.
- Restart closed workflow executions.

You can also perform domain management actions, such as deprecating and undeprecating domains.

After you deprecate a domain, you can't use it to create new workflow executions or register new workflows. Deprecating a domain also deprecates all the activity and workflows registered in the domain. Executions that were started before the domain was deprecated continue to run.

After undeprecating a previously deprecated domain, you can resume using the domain to register workflow types and start new workflow executions.

For more information about these domain management actions, see <u>DeprecateDomain</u> and <u>UndeprecateDomain</u>.

Registering a domain API Version 2012-01-25 41

Registering workflow types

You can register workflow types in the Amazon SWF console after you have registered at least one domain.

A workflow type is a set of activity types that carry out an objective and contain the logic that coordinates the activities. Workflow types coordinate and manage the execution of activities that can be run asynchronously across multiple computing devices and feature both sequential and parallel processing methods.

To register an Amazon SWF workflow type using the console

- 1. Open the domain in which you want to register a workflow.
- 2. Choose Register, and then choose Register Workflow.
- On the Register Workflow page, enter the Workflow name and Workflow version. Optionally, you can also specify a <u>Default task list</u> that will be used to schedule decision tasks for executions of this workflow.
- 4. (Optional) Choose Advanced options to specify the following details for your workflow:
 - Default Task priority The default task priority to assign to the workflow.
 - <u>Default Execution start to close timeout</u> The default maximum duration for executions of this workflow.
 - <u>Default Task start to close timeout</u> The default maximum duration of decision tasks for this workflow.
 - **Default Child policy** The default policy to use for the child workflow executions.
 - <u>Default Lambda role</u> The default IAM role attached to this workflow.
- Choose Register workflow.

Registering activity types

Activities are tasks which you want your workflow type to coordinate and execute (for example: verify customer's order, charge credit card etc.). The order in which activities are performed is determined by the workflow type's coordination logic.

You can register activity types after at least one domain is registered.

Registering workflow types API Version 2012-01-25 42

To register an Amazon SWF activity type using the console

- 1. Open the domain in which you want to register an activity.
- 2. Choose Register, and then choose Register Activity.
- 3. On the **Register activity** page, enter the <u>Activity name</u> and <u>Activity version</u>. Optionally, you can also specify a <u>Default task list</u> that will be used to schedule tasks of this activity.
- 4. (Optional) Choose **Advanced options** to specify the following details for your activity:
 - Default Task priority The default task priority to assign to the activity.
 - <u>Default task schedule to start timeout</u> The default maximum duration that a task of this activity can wait before being assigned to a worker.
 - <u>Default Task start to close timeout</u> The default maximum duration that a worker can take to process tasks of this activity.
 - <u>Default task schedule to close timeout</u> The default maximum duration for a task of this activity.
 - <u>Default task heartbeat timeout</u> The default maximum time before which a worker processing a task of this type must report progress by calling RecordActivityTaskHeartbeat.
- 5. Choose **Register activity**.

Starting a workflow

You can start a workflow execution from the Amazon SWF console. You cannot start a workflow execution until you have registered at least one workflow.

To start a workflow execution using the console

- 1. Open the Amazon SWF console, and in the left navigation pane, choose **Domains**.
- 2. Below the domain name, choose Workflows.
- 3. On the **Workflows** page, choose the workflow that you want to execute.
- 4. Choose Start execution.
- On the Start execution page, enter the <u>Workflow name</u> and Execution ID to identify your execution by a name. Optionally, you can also specify a <u>Task list</u> that will be used for the decision tasks generated for this workflow execution.
- 6. (Optional) Choose **Advanced options** to specify the following details for your workflow execution:

Starting a workflow API Version 2012-01-25 43

- Task priority The task priority to use for this workflow execution.
- Execution start to close timeout The total duration for this workflow execution.
- <u>Task start to close timeout</u> The maximum duration of decision tasks for this workflow execution.
- <u>Child policy</u> The policy to use for the child workflow executions of this workflow execution
 if it is terminated, by calling the <u>TerminateWorkflowExecution</u> action explicitly or due to an
 expired timeout.
- Lambda role The IAM role to attach to this workflow execution.
- 7. Choose **Start execution**.

Managing workflow executions

You can filter your workflow executions by name, status, ID, and tag. You can send signals with inputs into active workflow executions. If you need to cancel or terminate a workflow, you can use the **Try-cancel** option. Cancelling is preferable over terminating a workflow execution because canceling gives the workflow an opportunity to perform any clean-up tasks and then close properly.

In the console, you can manage the workflow executions that are currently running and/or closed.

To manage your workflow executions

- 1. Open a domain to manage its workflow executions.
- 2. Choose Find Execution.
- 3. On the **Workflow executions** page, choose **Filter executions by property**, and then under **Properties** choose one of the following filters:

Choose	To apply this filter
Workflow	Choose this filter to list executions of a specific workflow. For example, to view executions of the fiction-books-order-workflow , do the following:
	1. Choose Workflow .
	2. Under Operators , choose Equals .

Choose	To apply this filter				
	3. Under Workflows, choose fiction-books-order-workflow.4. (Optional) Choose Clear filters to remove the filter and start a new search for executions.				
Status	Choose this filter to list executions with a specific status. For example, to view executions with the status Failed, do the following: 1. Choose Status. 2. Under Operators, choose Equals. 3. Under Statuses, choose Failed. 4. (Optional) Choose Clear filters to remove the filter and start a new search for executions.				
Execution ID	Choose this filter to view a workflow execution based on its ID. For example, to view the execution with ID fiction-books-orde r-category1 , do the following: 1. Choose Execution ID. 2. Under Operators, choose Equals. 3. Under Execution IDs, choose fiction-books-order-category1. 4. (Optional) Choose Clear filters to remove the filter and start a new search for executions.				
Tag	Choose this filter to list executions with a specific tag. For example, to view executions with the status purchaseOrder, do the following: 1. Choose Tag. 2. Under Operators, choose Equals. 3. Under Tag, choose purchaseOrder. 4. (Optional) Choose Clear filters to remove the filter and start a new search for executions.				

- 4. (Optional) After applying the required filter to list workflow executions, you can perform the following operations to an **Active** execution:
 - Signal Use this option to send a running workflow execution additional data. To do this:
 - 1. Choose the execution to which you want to send additional data.
 - 2. Choose **Signal**, and then specify the data in the **Signal execution** dialog box.
 - 3. Choose Signal.
 - **Try-cancel** Use this option to try to cancel a workflow execution. It is preferable to cancel a workflow execution rather than terminate it. Canceling provides the workflow execution an opportunity to perform any clean-up tasks and then close properly.
 - 1. Choose the execution which you want to cancel.
 - 2. Choose **Try-cancel**.
 - **Terminate** Use this option to terminate a workflow execution. Note that it is preferable to cancel a workflow execution rather than terminate it.
 - 1. Choose the execution which you want to terminate.
 - 2. For **Child policy**, make sure **Terminate** is selected.
 - 3. (Optional) Specify the **Reason** and **Details** for terminating the execution.
 - 4. Choose Terminate.
- 5. (Optional) **Re-run** Use this option to re-run a closed workflow execution.
 - 1. In the list of workflow executions, select the closed execution to re-run. When you select a closed execution, the **Re-run** button becomes enabled. Choose **Re-run**.
 - 2. On the **Re-run execution** page, specify the details for workflow execution as mentioned in Starting a workflow.

Basic workflow concepts in Amazon SWF



Note


The concepts in this chapter provide an overview of the Amazon Simple Workflow Service and describe its major features. If you are looking for examples, see Working with Amazon **SWF APIs.**

Using the Amazon Simple Workflow Service (Amazon SWF), you can implement distributed, asynchronous applications as workflows. Workflows coordinate and manage the execution of activities that can be run asynchronously across multiple computing devices and that can feature both sequential and parallel processing.

When designing a workflow, you analyze your application to identify its component tasks. In Amazon SWF, these tasks are represented by activities. The order in which activities are performed is determined by the workflow's coordination logic.

Example workflow for an e-commerce application

The following figure shows an e-commerce order-processing workflow involving both people and automated processes:



The e-commerce application workflow starts when a customer places an order, and includes four tasks:

- 1. Verify the order.
- 2. If the order is valid, charge the customer.
- 3. If the payment is made, ship the order.
- 4. If the order is shipped, save the order details.

The tasks in this workflow are *sequential*: an order must be verified before a credit card can be charged; a credit card must be charged successfully before an order can be shipped; and an order must be shipped before it can be recorded. Even so, because Amazon SWF supports distributed processes, these tasks can be carried out in different locations. If the tasks are programmatic in nature, they can also be written in different programming languages or using different tools.

In addition to sequential processing of tasks, Amazon SWF also supports workflows with parallel processing of tasks. Parallel tasks are performed at the same time, and may be carried out independently by different applications or human workers. Your workflow makes decisions about how to proceed once one or more of the parallel tasks have been completed.

Additional concepts

- Creating a workflow in Amazon SWF
- Running workflows in Amazon SWF
- Workflow history in Amazon SWF
- Object identifiers in Amazon SWF
- Domains in Amazon SWF
- Actors in Amazon SWF
- Tasks in Amazon SWF
- Task lists in Amazon SWF
- Workflow execution closure in Amazon SWF
- Life cycle of a Amazon SWF workflow
- Polling for tasks in Amazon SWF

Creating a workflow in Amazon SWF

Creating a basic sequential workflow involves the following stages.

Creating a workflow API Version 2012-01-25 48


- Modeling a workflow, registering its type, and registering its activity types
- Developing and launching activity workers that perform activity tasks
- Developing and launching deciders that use the workflow history to determine what to do next
- Developing and launching workflow starters, that is, applications that start workflow executions

Modeling Your Workflow and Its Activities

To use Amazon SWF, model the logical steps in your application as activities. An activity represents a single logical step or task in your workflow. For example, authorizing a credit card is an activity that involves providing a credit card number and other information, and receiving an approval code or a message that the card was declined.

In addition to defining activities, you also need to define the coordination logic that handles decision points. For example, the coordination logic might schedule a different follow-up activity depending on whether the credit card was approved or declined.

The following figure shows an example of a sequential customer order workflow with four activities (Verify Order, Charge Credit Card, Ship Order, and Record Completion).



Running workflows in Amazon SWF

After the coordination logic and the activities have been designed, you register these components as workflow and activity types with Amazon SWF. During registration, you specify a name, a version, and default configuration values for each type.

Only registered workflow and activity types can be used with Amazon SWF. In the e-commerce example, you would register the CustomerOrder workflow type and the VerifyOrder, ChargeCreditCard, ShipOrder, and RecordCompletion activity types.

After registering your workflow type, you can run it as often you like. A *workflow execution* is a running instance of a workflow.

A workflow execution can be started by any process or application, even another workflow execution. In the e-commerce example, a new workflow execution is started with each customer order. The type of application that initiates the workflow depends on how the customer places the order. The workflow could be initiated by a web site or mobile application or by a customer service representative using an internal company application.

With Amazon SWF, you can associate an identifier—called a workflowId—with your workflow executions, so you can integrate your existing business identifiers into your workflow. In the ecommerce example, each workflow execution might be identified using the customer invoice number.

In addition to the identifier that you provide, Amazon SWF associates a unique system-generated identifier—a runId—with each workflow execution. Amazon SWF allows only one workflow execution with this identifier to run at any given time; although you can have multiple workflows executions of the same workflow type, each workflow execution has a distinct runId.

Workflow history in Amazon SWF

Amazon SWF records the progress of every workflow execution in the *workflow history* - a detailed, complete, and consistent record of every event that occurred since the workflow execution started.

An *event* represents a discrete change in your workflow execution's state, such as a new activity being scheduled or a running activity being completed. The workflow history contains every event that causes the execution state of the workflow execution to change, such as scheduled and completed activities, task timeouts, and signals.

Operations that don't change the state of the workflow execution don't typically appear in the workflow history. For example, the workflow history doesn't show poll attempts or the use of visibility operations.

The workflow history has several key benefits:

- Applications can be stateless, because all information about a workflow execution is stored in its workflow history.
- For each workflow execution, the history provides a record of which activities were scheduled, their current status, and their results. The workflow execution uses this information to determine next steps.
- The history provides a detailed audit trail that you can use to monitor running workflow executions and verify completed workflow executions.

The following is a conceptual view of the e-commerce workflow history:

```
Invoice0001

Start Workflow Execution

Schedule Verify Order
Start Verify Order Activity
Complete Verify Order Activity

Schedule Charge Credit Card
Start Charge Credit Card Activity
Complete Charge Credit Card Activity

Schedule Ship Order
Start Ship Order Activity
```

In the preceding example, the order is waiting to ship. In the following example, the order is complete. Because the workflow history is cumulative, the newer events are appended:

```
Invoice0001

Start Workflow Execution

Schedule Verify Order
Start Verify Order Activity
Complete Verify Order Activity

Schedule Charge Credit Card
Start Charge Credit Card Activity

Complete Charge Credit Card Activity
```

```
Schedule Ship Order
Start Ship Order Activity

Complete Ship Order Activity

Schedule Record Order Completion
Start Record Order Completion Activity

Complete Record Order Completion Activity

Close Workflow
```

Programmatically, the events in the workflow execution history are represented as JavaScript Object Notation (JSON) objects. The history itself is a JSON array of these objects. Each event has the following:

- A type, such as <u>WorkflowExecutionStarted</u> or <u>ActivityTaskCompleted</u>
- A timestamp in Unix time format
- An ID that uniquely identifies the event

In addition, each type of event has a distinct set of descriptive attributes that are appropriate to that type. For example, the ActivityTaskCompleted event has attributes that contain the IDs for the events that correspond to the time that the activity task was scheduled and when it was started, as well as an attribute that holds result data.

You can obtain a copy of the current state of the workflow execution history by using the GetWorkflowExecutionHistory action. In addition, as part of the interaction between Amazon SWF and the decider for your workflow, the decider periodically receives copies of the history.

Below is a section of an example workflow execution history in JSON format.

```
[ {
    "eventId": 11,
    "eventTimestamp": 1326671603.102,
    "eventType": "WorkflowExecutionTimedOut",
    "workflowExecutionTimedOutEventAttributes": {
        "childPolicy": "TERMINATE",
        "timeoutType": "START_TO_CLOSE"
     }
}, {
    "decisionTaskScheduledEventAttributes": {
        "startToCloseTimeout": "600",
```

```
"taskList": {
         "name": "specialTaskList"
      }
   },
   "eventId": 10,
   "eventTimestamp": 1326670566.124,
   "eventType": "DecisionTaskScheduled"
}, {
   "activityTaskTimedOutEventAttributes": {
      "details": "Waiting for confirmation",
      "scheduledEventId": 8,
      "startedEventId": 0,
      "timeoutType": "SCHEDULE_TO_START"
   },
   "eventId": 9,
   "eventTimestamp": 1326670566.124,
   "eventType": "ActivityTaskTimedOut"
}, {
   "activityTaskScheduledEventAttributes": {
      "activityId": "verification-27",
      "activityType": {
         "name": "activityVerify",
         "version": "1.0"
      },
      "control": "digital music",
      "decisionTaskCompletedEventId": 7,
      "heartbeatTimeout": "120",
      "input": "5634-0056-4367-0923,12/12,437",
      "scheduleToCloseTimeout": "900",
      "scheduleToStartTimeout": "300",
      "startToCloseTimeout": "600",
      "taskList": {
         "name": "specialTaskList"
      }
   },
   "eventId": 8,
   "eventTimestamp": 1326670266.115,
   "eventType": "ActivityTaskScheduled"
}, {
   "decisionTaskCompletedEventAttributes": {
      "executionContext": "Black Friday",
      "scheduledEventId": 5,
      "startedEventId": 6
   },
```

```
"eventId": 7,
   "eventTimestamp": 1326670266.103,
   "eventType": "DecisionTaskCompleted"
   "decisionTaskStartedEventAttributes": {
      "identity": "Decider01",
      "scheduledEventId": 5
   },
   "eventId": 6,
   "eventTimestamp": 1326670161.497,
   "eventType": "DecisionTaskStarted"
}, {
   "decisionTaskScheduledEventAttributes": {
      "startToCloseTimeout": "600",
      "taskList": {
         "name": "specialTaskList"
      }
   },
   "eventId": 5,
   "eventTimestamp": 1326668752.66,
   "eventType": "DecisionTaskScheduled"
}, {
   "decisionTaskTimedOutEventAttributes": {
      "scheduledEventId": 2,
      "startedEventId": 3,
      "timeoutType": "START_TO_CLOSE"
   },
   "eventId": 4,
   "eventTimestamp": 1326668752.66,
   "eventType": "DecisionTaskTimedOut"
}, {
   "decisionTaskStartedEventAttributes": {
      "identity": "Decider01",
      "scheduledEventId": 2
   },
   "eventId": 3,
   "eventTimestamp": 1326668152.648,
   "eventType": "DecisionTaskStarted"
}, {
   "decisionTaskScheduledEventAttributes": {
      "startToCloseTimeout": "600",
      "taskList": {
         "name": "specialTaskList"
      }
```

```
},
  "eventId": 2,
  "eventTimestamp": 1326668003.094,
  "eventType": "DecisionTaskScheduled"
}
```

For a detailed list of the different types of events that can appear in the workflow execution history, see the HistoryEvent data type in the Amazon Simple Workflow Service API Reference.

Amazon SWF stores the complete history of all workflow executions for a configurable number of days after the execution closes. This period, which is known as the workflow history retention period, is specified when you register a *Domain* for your workflow. Domains are discussed in greater detail later in this section.

Object identifiers in Amazon SWF

The following list describes how Amazon SWF objects, such as workflow executions, are uniquely identified.

- Workflow Type A registered workflow type is identified by its domain, name, and version.
 Workflow types are specified in the call to RegisterWorkflowType.
- Activity Type A registered activity type is identified by its domain, name, and version. Activity
 types are specified in the call to RegisterActivityType.
- Decision Tasks and Activity Tasks Each decision task and activity task is identified by a unique task token. The task token is generated by Amazon SWF and is returned with other information about the task in the response from PollForDecisionTask or PollForActivityTask.
 Although the token is most commonly used by the process that received the task, that process could pass the token to another process, which could then report the completion or failure of the task.
- **Workflow Execution** A single execution of a workflow is identified by the domain, workflow ID, and run ID. The first two are parameters that are passed to StartWorkflowExecution. The run ID is returned by StartWorkflowExecution.

Object identifiers API Version 2012-01-25 55

Domains in Amazon SWF

Workflows run in an Amazon resource called a *domain* which provide a way of scoping Amazon SWF resources within your Amazon account. All the components of a workflow, such as the workflow type and activity types, must be specified to be in a domain.

An Amazon account can have multiple domains, each of which can contain multiple workflows, but workflows in different domains can't interact.

When setting up a new workflow, before you set up any of the other workflow components you need to register a domain if you have not already done so.

When you register a domain, you specify a *workflow history retention period*. The retention period is the length of time that Amazon SWF will continue to retain information about the workflow execution after the workflow execution is complete.

Domain registration is the only functionality initially available in the console. After at least one domain is registered, you can perform the following actions for the domain:

- Register workflow and activity types.
- Initiate workflow executions.
- Cancel, terminate, and send signals to running workflow executions.
- Restart closed workflow executions.

You can also perform domain management actions, such as deprecating and undeprecating domains.

After you deprecate a domain, you can't use it to create new workflow executions or register new workflows. Deprecating a domain also deprecates all the activity and workflows registered in the domain. Executions that were started before the domain was deprecated continue to run.

After undeprecating a previously deprecated domain, you can resume using the domain to register workflow types and start new workflow executions.

For more information about these domain management actions, see <u>DeprecateDomain</u> and <u>UndeprecateDomain</u>.

Domains API Version 2012-01-25 56

Actors in Amazon SWF

Topics


- What is an Actor in Amazon SWF?
- Workflow Starters
- Deciders
- Activity Workers
- Data Exchange Between Actors

What is an Actor in Amazon SWF?

In the course of its operations, Amazon SWF interacts with a number of different types of programmatic *actors*. Actors can be <u>workflow starters</u>, <u>deciders</u>, or <u>activity workers</u>. These actors communicate with Amazon SWF through its API. You can develop these actors in any programming language.

The following diagram shows the Amazon SWF architecture, including Amazon SWF and its actors.

Actors API Version 2012-01-25 57



Workflow Starters

A workflow starter is any application that can initiate workflow executions. In the e-commerce example, one workflow starter could be the website at which the customer places an order. Another workflow starter could be a mobile application or system used by a customer service representative to place the order on behalf of the customer.

Deciders

A decider is an implementation of a workflow's coordination logic. Deciders control the flow of activity tasks in a workflow execution. Whenever a change occurs during a workflow execution, such as the completion of a task, a decision task including the entire workflow history will be passed to a decider. When the decider receives the decision task from Amazon SWF, it analyzes the

Workflow Starters API Version 2012-01-25 58

workflow execution history to determine the next appropriate steps in the workflow execution. The decider communicates these steps back to Amazon SWF using *decisions*. A decision is an Amazon SWF data type that can represent various next actions. For a list of the possible decisions, go to Decision in the Amazon Simple Workflow Service API Reference.

Here is an example of a decision in JSON format, the format in which it is transmitted to Amazon SWF. This decision schedules a new activity task.

```
{
   "decisionType" : "ScheduleActivityTask",
   "scheduleActivityTaskDecisionAttributes" : {
      "activityType" : {
         "name" : "activityVerify",
         "version" : "1.0"
      },
      "activityId" : "verification-27",
      "control" : "digital music",
      "input": "5634-0056-4367-0923,12/12,437",
      "scheduleToCloseTimeout" : "900",
      "taskList" : {
         "name": "specialTaskList"
      },
      "scheduleToStartTimeout" : "300",
      "startToCloseTimeout" : "600",
      "heartbeatTimeout" : "120"
   }
}
```

A decider receives a decision task when the workflow execution starts and each time a state change occurs in the workflow execution. Deciders continue to move the workflow execution forward by receiving decision tasks and responding to Amazon SWF with more decisions until the decider determines that the workflow execution is complete. It then responds with a decision to close the workflow execution. After the workflow execution closes, Amazon SWF will not schedule additional tasks for that execution.

In the e-commerce example, the decider determines if each step was performed correctly, and then either schedules the next step or manages any error conditions.

A decider represents a single computer process or thread. Multiple deciders can process tasks for the same workflow type.

Deciders API Version 2012-01-25 59

Activity Workers

An activity worker is a process or thread that performs the *activity tasks* that are part of your workflow. The activity task represents one of the tasks that you identified in your application.

To use an activity task in your workflow, you must register it using either the Amazon SWF console or the RegisterActivityType action.

Each activity worker polls Amazon SWF for new tasks that are appropriate for that activity worker to perform; certain tasks can be performed only by certain activity workers. After receiving a task, the activity worker processes the task to completion and then reports to Amazon SWF that the task was completed and provides the result. The activity worker then polls for a new task. The activity workers associated with a workflow execution continue in this way, processing tasks until the workflow execution itself is complete. In the e-commerce example, activity workers are independent processes and applications used by people, such as credit card processors and warehouse employees, that perform individual steps in the process.

An activity worker represents a single computer process (or thread). Multiple activity workers can process tasks of the same activity type.

Data Exchange Between Actors

Input data can be provided to a workflow execution when it is started. Similarly, input data can be provided to activity workers when they schedule activity tasks. When an activity task is complete, the activity worker can return results to Amazon SWF. Similarly, a decider can report the results of a workflow execution when the execution is complete. Each actor can send data to, and receive data from, Amazon SWF through strings, the form of which is user-defined. Depending on the size and sensitivity of the data, you can pass data directly or pass a pointer to data stored on another system or service (such as Amazon S3 or DynamoDB). Both the data passed directly and the pointers to other data stores are recorded in the workflow execution history; however, Amazon SWF doesn't copy or cache any of the data from external stores as part of the history.

Because Amazon SWF maintains the complete execution state of each workflow execution, including the inputs and the results of tasks, all actors can be stateless. As a result, workflow processing is highly scalable. As the load on your system grows, you can simply add more actors to increase capacity.

Activity Workers API Version 2012-01-25 60

Tasks in Amazon SWF

Amazon SWF interacts with activity workers and deciders by providing them with work assignments known as tasks. There are three types of tasks in Amazon SWF:

- Activity task An *Activity* task tells an activity worker to perform its function, such as to check inventory or charge a credit card. The activity task contains all the information that the activity worker needs to perform its function.
- Lambda task A Lambda task is similar to an Activity task, but executes a Lambda function instead of a traditional Amazon SWF activity. For more information about how to define a Lambda task, see Amazon Lambda tasks in Amazon SWF.
- **Decision task** A *Decision* task tells a decider that the state of the workflow execution has changed so that the decider can determine the next activity that needs to be performed. The decision task contains the current workflow history.

Amazon SWF schedules a decision task when the workflow starts and whenever the state of the workflow changes, such as when an activity task completes. Each decision task contains a paginated view of the entire workflow execution history. The decider analyzes the workflow execution history and responds back to Amazon SWF with a set of decisions that specify what should occur next in the workflow execution. Essentially, every decision task gives the decider an opportunity to assess the workflow and provide direction back to Amazon SWF.

To ensure that no conflicting decisions are processed, Amazon SWF assigns each decision task to exactly one decider and allows only one decision task at a time to be active in a workflow execution.

The following table shows the relationship between the different constructs related to workflows and deciders.

Logical Design	Registered As	Performed By	Receives & Performs	Generates
Workflow	Workflow Type	Decider	Decision Tasks	Decisions

When an activity worker has completed the activity task, it reports to Amazon SWF that the task was completed, and it includes any relevant results that were generated. Amazon SWF updates the

Tasks API Version 2012-01-25 61

workflow execution history with an event that indicates the task completed and then schedules a decision task to transmit the updated history to the decider.

Amazon SWF assigns each activity task to exactly one activity worker. Once the task is assigned, no other activity worker can claim or perform that task.

The following table shows the relationship between the different constructs related to activities.

Logical Design	Registered As	Performed By	Receives & Performs	Generates
Activity	Activity Type	Activity Worker	Activity Tasks	Results Data

Task lists in Amazon SWF

Task lists provide a way of organizing the various tasks associated with a workflow. You can think of task lists as similar to dynamic queues. When a task is scheduled in Amazon SWF, you can specify a queue (task list) to put it in. Similarly, when you poll Amazon SWF for a task you say which queue (task list) to get the task from.

Task lists provide a flexible mechanism to route tasks to workers as your use case necessitates. Task lists are dynamic in that you don't need to register a task list or explicitly create it through an action: simply scheduling a task creates the task list if it doesn't already exist.

There are separate lists for *activity* tasks and *decision* tasks. A task is always scheduled on only one task list; tasks are not shared across lists. Furthermore, like activities and workflows, task lists are scoped to a particular Amazon region and Amazon SWF domain.

Topics

- Decision Task Lists
- Activity Task Lists
- Task Routing

Decision Task Lists

Each workflow execution is associated with a specific decision task list. When a workflow type is registered (RegisterWorkflowType action), you can specify a default task list for

Task lists API Version 2012-01-25 62

executions of that workflow type. When the workflow starter initiates the workflow execution (StartWorkflowExecution action), it has the option of specifying a different task list for that workflow execution.

When a decider polls for a new decision task (PollForDecisionTask action), the decider specifies a decision task list to draw from. A single decider could serve multiple workflow executions by calling PollForDecisionTask multiple times, using a different task list in each call, where each task list is specific to a particular workflow execution. Alternatively, the decider could poll a single decision task list that provides decision tasks for multiple workflow executions. You could also have multiple deciders serving a single workflow execution by having all of them poll the task list for that workflow execution.

Activity Task Lists

A single activity task list can contain tasks of different activity types. Tasks are scheduled on the task list in order. Amazon SWF returns the tasks from the list in order on a best effort basis. Under some circumstances, the tasks may not come off the list in order.

When an activity type is registered (RegisterActivityType action), you can specify a default task list for that activity type. By default, activity tasks of this type will be scheduled on the specified task list; however, when the decider schedules an activity task (ScheduleActivityTask decision), the decider can optionally specify a different task list on which to schedule the task. If the decider doesn't specify a task list, the default task list is used. As a result, you can place activity tasks on specific task lists according to attributes of the task. For example, you could place all instances of an activity task for a given credit card type on a particular task list.

Task Routing

When an activity worker polls for a new task (PollForActivityTask action), it can specify an activity task list to draw from. If it does, the activity worker will accept tasks only from that list. In this way, you can ensure that certain tasks get assigned only to particular activity workers. For example, you might create a task list that holds tasks that require the use of a high-performance computer. Only activity workers running on the appropriate hardware would poll that task list. Another example would be to create a task list for a particular geographic region. You could then ensure that only workers deployed in that region would pick up those tasks. Or you could create a task list for high-priority orders and always check that list first.

Activity Task Lists API Version 2012-01-25 63

Assigning particular tasks to particular activity workers in this way is called *task routing*. Task routing is optional; if you don't specify a task list when scheduling an activity task, the task is automatically placed on the default task list.

Workflow execution closure in Amazon SWF

Once you start a workflow execution, it is *open*. An open workflow execution could be closed as completed, canceled, failed, or timed out. It could also be continued as a new execution, or it could be terminated. A workflow execution could be closed by the decider, by the person administering the workflow, or by Amazon SWF.

If the decider determines that the activities of the workflow have finished, it should close the workflow execution as completed by using the RespondDecisionTaskCompleted action and pass the CompleteWorkflowExecution decision.

Alternatively, a decider might close the workflow execution as canceled or failed. In order to cancel the execution, the decider should use the RespondDecisionTaskCompleted action and pass the CancelWorkflowExecution decision.

A decider should fail the workflow execution if it enters a state outside the realm of normal completion. In order to fail the execution, the decider should use the RespondDecisionTaskCompleted action and pass the FailWorkflowExecution decision.

Amazon SWF monitors workflow executions to ensure that they don't exceed any user-specified timeout settings. If a workflow execution times out, Amazon SWF automatically closes it. For more information about timeout values, see the Amazon SWF Timeout Types section.

A decider might also close the execution and logically continue it as a new execution using the RespondDecisionTaskCompleted action and passing the ContinueAsNewWorkflowExecution decision. This is a useful strategy for long-running workflow executions for which the history may grow too large over time.

Finally, you could terminate workflow executions directly from the Amazon SWF console or programmatically by using the Termination forces closure of the workflow execution. Cancellation is preferred over termination, because your deciders can manage closure of the workflow execution.


Amazon SWF terminates a workflow execution if the execution exceeds certain service-defined limits. Amazon SWF terminates a child workflow if the parent workflow has terminated and the applicable child policy indicates that the child workflow should also be terminated.

Workflow execution closure API Version 2012-01-25 64

Life cycle of a Amazon SWF workflow

From the start of a workflow execution to its completion, Amazon SWF interacts with actors by assigning them appropriate tasks, either activity tasks or decision tasks.

The following diagram shows the life cycle of an order-processing workflow execution from the perspective of components that act on it.



Workflow Execution Life Cycle

The following table explains each task in the preceding image.

Description	Action, Decision, or Event
1. The workflow starter calls the appropriate Amazon SWF action to start the workflow execution for an order, providing the order information.	<u>StartWorkflowExecution</u> action.
2. Amazon SWF receives the start workflow execution request and then schedules the first decision task.	WorkflowExecutionStarted event and DecisionTaskSchedu led event.

Workflow execution life cycle API Version 2012-01-25 65

Description	Action, Decision, or Event
3. The decider receives the task from Amazon SWF, reviews the history, applies the coordination logic to determine that no previous activities occurred, makes a decision to schedule the Verify Order activity with the information the activity worker needs to process the task, and returns the decision to Amazon SWF.	PollForDecisionTask action. RespondDecisionTaskCompleted action and ScheduleActivityTask decision.
4. Amazon SWF receives the decision, schedules the Verify Order activity task, and waits for the activity task to complete or time out.	<u>ActivityTaskScheduled</u> event
5. An activity worker that can perform the Verify Order activity receives the task, performs it, and returns the results to Amazon SWF.	<pre>PollForActivityTask action and RespondActivityTas kCompleted action.</pre>

Description	Action, Decision, or Event	
6. Amazon SWF receives the results of the Verify Order activity, adds them to the workflow history, and schedules a decision task.	ActivityTaskCompleted event and DecisionTaskSchedu led event.	
7. The decider receives the task from Amazon SWF, reviews the history, applies the coordinat ion logic, makes a decision to schedule a ChargeCreditCard activity task with the information the activity worker needs to process the task, and returns the decision to Amazon SWF.	PollForDecisionTask action. RespondDecisionTask kCompleted action with ScheduleActivityTask decision.	
8. Amazon SWF receives the decision, schedules the ChargeCreditCard activity task, and waits for it to complete or time out.	<pre>DecisionTaskCompleted event and ActivityTaskSchedu led event.</pre>	

Description	Action, Decision, or Event
9. An activity worker that can perform the ChargeCreditCard activity receives the task, performs it, and returns the results to Amazon SWF.	PollForActivityTask and RespondActivityTaskComplete d_action.
10. Amazon SWF receives the results of the ChargeCre ditCard activity task, adds them to the workflow history, and schedules a decision task.	ActivityTaskCompleted event and DecisionTaskSchedu led event.
11. The decider receives the task from Amazon SWF, reviews the history, applies the coordination logic, makes a decision to schedule a ShipOrder activity task with the information the activity worker needs to perform the task, and returns the decision to Amazon SWF.	PollForDecisionTask action. RespondDecisionTaskCompleted with ScheduleActivityTask decision.

Description	Action, Decision, or Event
12. Amazon SWF receives the decision, schedules a ShipOrder activity task, and waits for it to complete or time out.	<pre>DecisionTaskCompleted event and ActivityTaskSchedu led event.</pre>
13. An activity worker that can perform the ShipOrder activity receives the task, performs it, and returns the results to Amazon SWF.	PollForActivityTask action and RespondActivityTas kCompleted action.
14. Amazon SWF receives the results of the ShipOrder activity task, adds them to the workflow history, and schedules a decision task.	ActivityTaskCompleted event and DecisionTaskSchedu led event.

Description	Action, Decision, or Event
15. The decider receives the task from Amazon SWF, reviews the history, applies the coordinat ion logic, makes a decision to schedule a RecordCompletion activity task with the information the activity worker needs to perform the task, and returns the decision to Amazon SWF.	PollForDecisionTask action. RespondDecisionTask KCompleted action with ScheduleActivityTask decision.
16. Amazon SWF receives the decision, schedules a RecordCompletion activity task, and waits for it to complete or time out.	<pre>DecisionTaskCompleted event and ActivityTaskSchedu led event.</pre>
17. An activity worker that can perform the RecordCompletion activity receives the task, performs it, and returns the results to Amazon SWF.	<pre>PollForActivityTask action and RespondActivityTas kCompleted action.</pre>

Description	Action, Decision, or Event
18. Amazon SWF receives the results of the RecordCom pletion activity task, adds them to the workflow history, and schedules a decision task.	ActivityTaskCompleted event and DecisionTaskSchedu led event.
19. The decider receives the task from Amazon SWF, reviews the history, applies the coordinat ion logic, makes a decision to close the workflow execution and returns the decision along with any results to Amazon SWF.	PollForDecisionTask action. RespondDecisionTas kCompleted action with CompleteWorkflowExecution decision.
20. Amazon SWF closes the workflow execution and archives the history for future reference.	WorkflowExecutionCompleted event.

Polling for tasks in Amazon SWF

Deciders and activity workers communicate with Amazon SWF using *long polling*. The decider or activity worker periodically initiates communication with Amazon SWF, notifying Amazon SWF of its availability to accept a task, and then specifies a task list to get tasks from.

Polling for tasks API Version 2012-01-25 71

If a task is available on the specified task list, Amazon SWF returns it immediately in the response. If no task is available, Amazon SWF holds the TCP connection open for up to 60 seconds so that, if a task becomes available during that time, it can be returned in the same connection. If no task becomes available within 60 seconds, it returns an empty response and closes the connection. (An empty response is a Task structure in which the value of taskToken is an empty string.) If this happens, the decider or activity worker should poll again.

Long polling works well for high-volume task processing. Deciders and activity workers can manage their own capacity, and it is easy to use when the deciders and activity workers are behind a firewall.

For more information, see Polling for Decision Tasks and Polling for Activity Tasks.

Polling for tasks API Version 2012-01-25 72

Advanced workflow concepts in Amazon SWF

The e-commerce example in the <u>???</u> section represents a simplified workflow scenario. In reality, you are likely to want your workflow to do concurrent tasks (send an order confirmation email while authorizing a credit card), record major events (all items are packed), update the order with changes (add or remove an item), and make other more advanced decisions as part of your workflow execution. This section describes advanced workflow concepts that you can use to construct your workflows.

Advanced concepts

- Versioning
- Signals
- Child workflows in Amazon SWF
- Markers in Amazon SWF
- Tags in Amazon SWF
- Implementing exclusive choice with Amazon SWF
- Timers in Amazon SWF
- Cancelling activity tasks in Amazon SWF

Versioning

Business needs often require you to have different implementations or variations of the same workflow or activity running simultaneously. For example, you might want to test a new implementation of a workflow while another one is in production. You might also want to run two different implementations with two different feature sets, such as a basic and premium implementation. Versioning enables you to run multiple implementations of workflows and activities concurrently, for any purpose that meets your requirements.

Workflow and activity types have a version associated with them which is specified at registration time. Version is a free-form string and you can choose your own versioning scheme. In order to create a new version of a registered type, you should register it with the same name and a different version. Task lists in Amazon SWF, described earlier, can further help you to implement versioning. Consider a situation in which you have long-running workflow executions of a given type already in progress, and circumstances require that you revise the workflow, such as to add a new feature. You could implement the new feature by creating new versions of activity types and workers, and a new

Versioning API Version 2012-01-25 73

decider. Then you could launch executions of the new workflow version using a different set of task lists. This way, you could have executions of workflows of different versions running simultaneously without affecting each other.

Signals

Signals enable you to inject information into a running workflow execution. In some scenarios, you might want to add information to a running workflow execution to let it know that something has changed or to inform it of an external event. Any process can send a signal to an open workflow execution. For example, one workflow execution might signal another.



Note

An attempt to send a signal to a workflow execution that isn't open results in SignalWorkflowExecution failing with UnknownResourceFault.

To use signals, define the signal name and data to be passed to the signal—if any. Then, program the decider to recognize the signal event (WorkflowExecutionSignaled) in the history and process it appropriately. When a process wants to signal a workflow execution, it makes a call to Amazon SWF (using the SignalWorkflowExecution action or, in the case of a decider, using the SignalExternalWorkflowExecution decision) that specifies the identifier for the target workflow execution, the signal name, and the signal data. Amazon SWF then receives the signal, records it in the history of the target workflow execution, and schedules a decision task for it. When the decider receives the decision task, it also receives the signal inside the workflow execution history. The decider can then take appropriate actions based on the signal and its data.

Sometimes you might want to wait for a signal. For example, a user could cancel an order by sending a signal, but only within one hour of placing the order. Amazon SWF doesn't have a primitive to enable a decider to wait for a signal from the service. Pause functionality needs to be implemented in the decider itself. In order to pause, the decider should start a timer, using the StartTimer decision, which specifies the duration for which the decider will wait for the signal while continuing to poll for decision tasks. When the decider receives a decision task, it should check the history to see if either the signal has been received or the timer has fired. If the signal has been received, then the decider should cancel the timer. However, if instead, the timer has fired, then it means that the signal did not arrive within the specified time. To summarize, in order to wait for a specific signal, do the following.

Signals API Version 2012-01-25 74

- 1. Create a timer for the amount of time the decider should wait.
- 2. When a decision task is received, check the history to see if the signal has arrived or if the timer has fired.
- 3. If a signal has arrived, cancel the timer using a CancelTimer decision and process the signal. Depending on the timing, the history may contain both TimerFired and WorkflowExecutionSignaled events. In such cases, you can rely on the relative order of the events in the history to determine which occurred first.
- 4. If the timer has fired, before a signal is received, then the decider has timed out waiting for the signal. You can fail the execution or do whatever other logic is appropriate to your use case.

For cases in which a workflow should be canceled—for example, the order itself was canceled by the customer—the RequestCancelWorkflowExecution action should be used rather than sending a signal to the workflow.

Some applications for signals include the following:

- Pausing workflow executions from progressing until a signal is received (e.g., waiting for an inventory shipment).
- Providing information to a workflow execution that might affect the logic of how deciders make decisions. This is useful for workflows affected by external events (e.g., trying to finish the sale of a stock after the market closes).
- Updating a workflow execution when you anticipate that changes might occur (e.g., changing order quantities after an order is placed and before it ships).

In the following example, the workflow execution is sent a signal to cancel an order.

```
https://swf.us-east-1.amazonaws.com.cn
SignalWorkflowExecution
{"domain": "867530901",
    "workflowId": "20110927-T-1",
    "runId": "f5ebbac6-941c-4342-ad69-dfd2f8be6689",
    "signalName": "CancelOrder",
    "input": "order 3553"}
```

If the workflow execution receives the signal, Amazon SWF returns a successful HTTP response similar to the following. Amazon SWF will generate a decision task to inform the decider to process the signal.

Signals API Version 2012-01-25 75

HTTP/1.1 200 OK Content-Length: 0

Content-Type: application/json

x-amzn-RequestId: bf78ae15-3f0c-11e1-9914-a356b6ea8bdf

Child workflows in Amazon SWF

Complicated workflows can be broken into smaller, more manageable, and potentially reusable components by using child workflows. A child workflow is a workflow execution that is initiated by another (parent) workflow execution. To initiate a child workflow, the decider for the parent workflow uses the StartChildWorkflowExecution decision. Input data specified with this decision is made available to the child workflow through its history.

The attributes for the StartChildWorkflowExecution decision also specify the *child policy*, that is, how Amazon SWF should handle the situation in which the parent workflow execution terminates before the child workflow execution. There are three possible values:

- TERMINATE: Amazon SWF will terminate the child executions.
- REQUEST_CANCEL: Amazon SWF will attempt to cancel the child execution by placing a WorkflowExecutionCancelRequested event in the child's workflow execution history.
- ABANDON: Amazon SWF will take no action; the child executions will continue to run.

After the child workflow execution starts, it runs like a regular execution. When it completes, Amazon SWF records the completion, along with its results, in the parent workflow execution's workflow history. Examples of child workflows include the following:


- · Credit card processing child workflow used by workflows in different websites
- Email child workflow that verifies the customer email address, checks the opt-out list, sends the email, and verifies that it didn't bounce or fail.
- Database storage and retrieval child workflow that combines connection, setup, transaction, and verification.
- Source code compilation child workflow that combines building, packaging, and verification.

In the e-commerce example, you might want to make the Charge Credit Card activity a child workflow. To do this, you could register a new Verify Customer workflow, register the Verify

Child workflows API Version 2012-01-25 76

Customer Address and Check Fraud DB activities, and define coordination logic for the tasks. Then, a decider in the Customer Order workflow can initiate a Verify Customer child workflow by scheduling the StartChildWorkflowExecution decision that specifies this workflow type.

The following figure shows a customer order workflow that includes a new Verify Customer child workflow, which checks the customer address, checks the fraud database, and charges the credit card.



Multiple workflows could create child workflow executions using the same workflow type. For example, the Verify Customer child workflow could also be used in other parts of an organization. The events for a child workflow are contained in its own workflow history and are not included in the parent's workflow history.

Because child workflows are simply workflow executions that are initiated by a decider, they could also be started as normal stand-alone workflow executions.

Markers in Amazon SWF

At times, you might want to record information in the workflow history of a workflow execution that is specific to your use case. Markers enable you to record information in the workflow execution history that you can use for any custom or scenario-specific purpose.

To use markers, a decider uses the RecordMarker decision, names the marker, attaches desired data to the decision, and notifies Amazon SWF using the RespondDecisionTaskCompleted action.

Amazon SWF receives the request, records the marker in the workflow history, and enacts any

Markers API Version 2012-01-25 77

other decisions in the request. From that point on, deciders can see the marker in the workflow history and use it in any way that you program.

Recording a marker doesn't, by itself, initiate a decision task. To prevent the workflow execution from becoming stuck, something must occur that continues the execution of the workflow. For example, this might include the decider scheduling another activity task, the workflow execution receiving a signal, or a previously scheduled activity task completing.

Examples of markers include the following:

- A counter that counts the number of loops in a recursive workflow.
- Progress of the workflow execution based on the results of activities.
- Information summarized from earlier workflow history events.

In the e-commerce example, you might add an activity that checks the inventory every day and increments the count in a marker each time. Then, you could add decision logic that emails the customer or notifies a manager when the count exceeds five, without having to review the entire history.

In the following example, the decider completes a decision task and responds with a RespondDecisionTaskCompleted action that contains a RecordMarker decision.

If Amazon SWF successfully records the marker, it returns a successful HTTP response similar to the following.

```
HTTP/1.1 200 OK
```

Markers API Version 2012-01-25 78

Content-Length: 0

Content-Type: application/json

x-amzn-RequestId: 6c0373ce-074c-11e1-9083-8318c48dee96

Tags in Amazon SWF

Amazon SWF supports tagging a workflow execution. This is especially useful when you have many resources.

Amazon SWF supports tagging a workflow execution with up to five tags. Each tag is a free-form string and may be up to 256 characters in length. If you want to use tags, you must assign them when you start a workflow execution. You can't add tags to a workflow execution after it has been started, nor may you edit or remove tags that have been assigned to a workflow execution.

IAM supports controlling access to Amazon SWF domains based on tags. To control access based on tags, provide information about your tags in the condition element of an IAM policy.

Manage tags

Manage Amazon Simple Workflow Service tags using the Amazon SDKs or by interacting directly with the Amazon SWF API. Using the API you can add tags when registering a domain, list tags for an existing domain, and add or delete tags for an existing domain.



There is a limit of 50 tags per resource. See General Account Quotas for Amazon SWF

- RegisterDomain
- ListTagsForResource
- TagResource
- UntagResource

For more information see Working with Amazon SWF APIs, and Amazon Simple Workflow Service API Reference.

Tags API Version 2012-01-25 79

Tag workflow executions

With Amazon SWF, you can associate tags with workflow executions and then query for workflow executions based on these tags. You can filter the listi when you use the visibility operations. By carefully selecting the tags you assign to an execution, you can use them to provide meaningful listings.

For example, suppose you run several fulfillment centers. With tags, you could list the processes occurring in a specific fulfillment center. Or, if a customer is converting different types of media files, tags could indicate different processes when converting video, audio, and image files.

You can associate up to five tags with a workflow execution when you start the execution using the StartWorkflowExecution action, StartChildWorkflowExecution decision, or ContinueAsNewWorkflowExecution decision. When you use visibility actions to list or count workflow executions, you can filter results based on your tags.

To use tagging

- Devise a tagging strategy. Think about your business requirements and create a list of tags
 that are meaningful to you. Determine which executions will get which tags. Even though
 an execution can be assigned a maximum of five tags, your tag library can have any number
 of tags. Because each tag can be any string value up to 256 characters in length, a tag can
 describe almost any business concept.
- 2. Tag an execution with up to five tags when you create it.
- 3. List or count the executions that are tagged with a particular tag by specifying the tagFilter parameter with the ListOpenWorkflowExecutions, ListClosedWorkflowExecutions, CountOpenWorkflowExecutions, and CountClosedWorkflowExecutions actions. The action will filter the executions based on the tags specified.

When you associate a tag with a workflow execution, it is permanently associated with that execution, and can't be removed.

You can specify only one tag in the tagFilter parameter with ListWorkflowExecutions. Also, tag matching is case sensitive, and only exact matches return results.

Assume you have already set up two executions that are tagged as follows.

Tag workflow executions API Version 2012-01-25 80

Execution Name	Assigned Tags
Execution-One	Consumer, 2011-February
Execution-Two	Wholesale, 2011-March

You can filter the list of executions returned by ListOpenWorkflowExecutions on the Consumer tag. The oldestDate and latestDate values are specified as Unix Time values.

```
https://swf.us-east-1.amazonaws.com.cn
RespondDecisionTaskCompleted
{
    "domain":"867530901",
    "startTimeFilter":{
        "oldestDate":1262332800,
        "latestDate":1325348400
},
    "tagFilter":{
        "tag":"Consumer"
     }
}
```

Control access to domains with tags

You can control access to Amazon Simple Workflow Service domains by referencing tags associated with Amazon SWF domains in IAM. For instance, you could restrict domains that include a tag where the key is environment and the value is production:

}

This policy will Deny the access to any domain that has been tagged as environment/production.

For more information, see:

- Controlling Access Using IAM Tags
- Tag-based Policies


Implementing exclusive choice with Amazon SWF

In some scenarios, you might want to schedule a different set of activities based on the outcome of a previous activity. With the exclusive choice pattern, you can create flexible workflows that meet the complex requirements of your application.

Amazon SWF does not have a specific exclusive choice action. To implement exclusive choice, you must write your decider logic to make decisions based on the results of a previous activity. Some applications for exclusive choice include the following:

- Performing cleanup activities if the results of a previous activity were unsuccessful
- Scheduling different activities based on whether the customer purchased a basic or advanced plan
- Performing different customer authentication activities based on the customer's ordering history

In the e-commerce example, you might use exclusive choice to either ship or cancel an order based on the outcome of charging the credit card. In the following figure, the decider schedules the Ship Order and Record Completion activity tasks if the credit card is successfully charged. Otherwise, it schedules the Cancel Order and Email Customer activity tasks.



The decider schedules the ShipOrder activity if the credit card is successfully charged. Otherwise, the decider schedules the CancelOrder activity.

Exclusive choice API Version 2012-01-25 82

In this case, program the decider to interpret the history and determine whether the credit card was successfully charged. To do this, you might have logic similar to the following

```
IF lastEvent = "WorkflowExecutionStarted"
 addToDecisions ScheduleActivityTask(ActivityType = "VerifyOrderActivity")
ELSIF lastEvent = "ActivityTaskCompleted"
     AND ActivityType = "VerifyOrderActivity"
 addToDecisions ScheduleActivityTask(ActivityType = "ChargeCreditCardActivity")
#Successful Credit Card Charge Activities
ELSIF lastEvent = "ActivityTaskCompleted"
      AND ActivityType = "ChargeCreditCardActivity"
  addToDecisions ScheduleActivityTask(ActivityType = "ShipOrderActivity")
ELSIF lastEvent = "ActivityTaskCompleted"
     AND ActivityType = "ShipOrderActivity"
  addToDecisions ScheduleActivityTask(ActivityType = "RecordOrderCompletionActivity")
ELSIF lastEvent = "ActivityTaskCompleted"
     AND ActivityType = "RecordOrderCompletionActivity"
 addToDecisions CompleteWorkflowExecution
#Unsuccessful Credit Card Charge Activities
ELSIF lastEvent = "ActivityTaskFailed"
      AND ActivityType = "ChargeCreditCardActivity"
  addToDecisions ScheduleActivityTask(ActivityType = "CancelOrderActivity")
ELSIF lastEvent = "ActivityTaskCompleted"
     AND ActivityType = "CancelOrderActivity"
 addToDecisions ScheduleActivityTask(ActivityType = "EmailCustomerActivity")
ELSIF lastEvent = "ActivityTaskCompleted"
     AND ActivityType = "EmailCustomerActivity"
 addToDecisions CompleteWorkflowExecution
ENDIF
```

If the credit card was successfully charged, the decider should respond with RespondDecisionTaskCompleted to schedule the ShipOrder activity.

```
https://swf.us-east-1.amazonaws.com.cn
RespondDecisionTaskCompleted
```

Exclusive choice API Version 2012-01-25 83

```
{
  "taskToken": "12342e17-80f6-FAKE-TASK-TOKEN32f0223",
  "decisions":[
      {
          "decisionType": "ScheduleActivityTask",
          "scheduleActivityTaskDecisionAttributes":{
               "control": "OPTIONAL_DATA_FOR_DECIDER",
               "activityType":{
                   "name": "ShipOrder",
                  "version":"2.4"
              },
              "activityId": "3e2e6e55-e7c4-fee-deed-aa815722b7be",
              "scheduleToCloseTimeout": "3600",
               "taskList":{
                   "name": "SHIPPING"
              },
              "scheduleToStartTimeout": "600",
              "startToCloseTimeout": "3600",
              "heartbeatTimeout":"300",
              "input": "123 Main Street, Anytown, United States"
          }
      }
  ]
}
```

If the credit card was not successfully charged, the decider should respond with RespondDecisionTaskCompleted to schedule the CancelOrder activity.

Exclusive choice API Version 2012-01-25 84

If Amazon SWF is able to validate the data in the RespondDecisionTaskCompleted action, Amazon SWF returns a successful HTTP response similar to the following.

```
HTTP/1.1 200 OK
Content-Length: 11
Content-Type: application/json
x-amzn-RequestId: 93cec6f7-0747-11e1-b533-79b402604df1
```

Timers in Amazon SWF

With a timer, you can notify your decider when a certain amount of time has elapsed.

When responding to a decision task, the decider has the option to respond with a StartTimer decision. This decision specifies an amount of time after which the timer should fire. After the specified time has elapsed, Amazon SWF will add a TimerFired event to the workflow execution history and schedule a decision task. The decider can then use this information to inform further decisions. One common application for a timer is to delay the execution of an activity task. For example, a customer might want to take delayed delivery of an item.

Cancelling activity tasks in Amazon SWF

Activity task cancellation informs the decider to end activities that no longer need to be performed. Amazon SWF uses a cooperative cancellation mechanism and doesn't forcibly interrupt running activity tasks. You must program your activity workers to handle cancellation requests.

The decider can decide to cancel an activity task while it is processing a decision task. To cancel an activity task, the decider uses the RespondDecisionTaskCompleted action with the RequestCancelActivityTask decision.

Timers API Version 2012-01-25 85

If the activity task has not yet been acquired by an activity worker, the service will cancel the task. Note that there is a potential race condition in that an activity worker could acquire the task at any time. If the task has already been assigned to an activity worker, then the activity worker will be requested to cancel the task.

In this example, the workflow execution is sent a signal to cancel the order.

```
https://swf.us-east-1.amazonaws.com.cn
SignalWorkflowExecution
{"domain": "867530901",
"workflowId": "20110927-T-1",
"runId": "9ba33198-4b18-4792-9c15-7181fb3a8852",
"signalName": "CancelOrder",
"input": "order 3553"}
```

If the workflow execution receives the signal, Amazon SWF returns a successful HTTP response similar to the following. Amazon SWF will generate a decision task to inform the decider to process the signal.

```
HTTP/1.1 200 0K
Content-Length: 0
Content-Type: application/json
x-amzn-RequestId: 6c0373ce-074c-11e1-9083-8318c48dee96
```

When the decider processes the decision task and sees the signal in the history, the decider attempts to cancel the outstanding activity that has the ShipOrderActivity0001 activity ID. The activity ID is provided in the workflow history from the schedule activity task event.

Cancelling activity tasks API Version 2012-01-25 86

If Amazon SWF successfully receives the cancellation request, it returns a successful HTTP response similar to the following:

HTTP/1.1 200 OK Content-Length: 0

Content-Type: application/json

x-amzn-RequestId: 6c0373ce-074c-11e1-9083-8318c48dee96

The cancellation attempt is recorded in the history as the ActivityTaskCancelRequested event.

If the task is successfully canceled—as indicated by an ActivityTaskCanceled event—program your decider to take the appropriate steps that should follow task cancellation such as closing the workflow execution.

If the activity task could not be canceled—for example, if the task completes, fails, or times out instead of canceling—your decider should accept the results of the activity or perform any cleanup or mitigation necessitated by your use case.

If the activity task has already been acquired by an activity worker, then the request to cancel is transmitted through the task-heartbeat mechanism. Activity workers can periodically use RecordActivityTaskHeartbeat to report to Amazon SWF that the task is still in progress.

Note that activity workers are not required to heartbeat, although it is recommended for long-running tasks. Task cancellation requires periodic heartbeat to be recorded; if the worker doesn't heartbeat, the task can't be canceled.

If the decider requests a cancellation of the task, Amazon SWF sets the value of the cancelRequest object to true. The cancelRequest object is part of the ActivityTaskStatus object which is returned by the service in response to RecordActivityTaskHeartbeat.

Amazon SWF doesn't prevent the successful completion of an activity task whose cancellation has been requested; it is up to the activity to determine how to handle the cancellation request. Depending on your requirements, program the activity worker to either cancel the activity task or ignore the cancellation request.

If you want the activity worker to indicate that the work for the activity task was canceled, program it to respond with a RespondActivityTaskCanceled. If you want the activity worker to complete the task, program it to respond with a standard RespondActivityTaskCompleted.

Cancelling activity tasks API Version 2012-01-25 87

When Amazon SWF receives the RespondActivityTaskCompleted or RespondActivityTaskCanceled request, it updates the workflow execution history and schedules a decision task to inform the decider.

Program the decider to process the decision task and return any additional decisions. If the activity task is successfully canceled, program the decider to perform the tasks needed to continue or close the workflow execution. If the activity task isn't successfully canceled, program the decider to accept the results, ignore the results, or schedule any required cleanup.

Cancelling activity tasks API Version 2012-01-25 88

Security in Amazon Simple Workflow Service

This section provides information about Amazon Simple Workflow Service security and authentication.

Topics

- Data protection in Amazon Simple Workflow Service
- Identity and Access Management in Amazon Simple Workflow Service
- Logging and Monitoring
- Compliance Validation for Amazon Simple Workflow Service
- Resilience in Amazon Simple Workflow Service
- Infrastructure Security in Amazon Simple Workflow Service
- Configuration and Vulnerability Analysis in Amazon Simple Workflow Service

Amazon SWF uses IAM to control access to other Amazon services and resources. For an overview of how IAM works, see Overview of Access Management in the IAM User Guide. For an overview of security credentials, see Amazon Security Credentials in the Amazon Web Services General Reference.

Data protection in Amazon Simple Workflow Service

The Amazon shared responsibility model applies to data protection in Amazon Simple Workflow Service. As described in this model, Amazon is responsible for protecting the global infrastructure that runs all of the Amazon Web Services Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the Amazon Web Services services that you use. For more information about data privacy, see the Data Privacy FAQ.

For data protection purposes, we recommend that you protect Amazon Web Services account credentials and set up individual users with Amazon IAM Identity Center or Amazon Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

Use multi-factor authentication (MFA) with each account.

Data Protection API Version 2012-01-25 89

- Use SSL/TLS to communicate with Amazon resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with Amazon CloudTrail. For information about using CloudTrail trails to capture Amazon activities, see Working with CloudTrail trails in the Amazon CloudTrail User Guide.
- Use Amazon encryption solutions, along with all default security controls within Amazon Web Services services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing Amazon through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-3.

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Amazon SWF or other Amazon Web Services services using the console, API, Amazon CLI, or Amazon SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Encryption in Amazon Simple Workflow Service

Encryption at rest

Amazon SWF always encrypts your data at rest. Data in Amazon Simple Workflow Service is encrypted at rest using transparent server-side encryption. This helps reduce the operational burden and complexity involved in protecting sensitive data. With encryption at rest, you can build security-sensitive applications that meet encryption compliance and regulatory requirements

Encryption in transit

All data that passes between Amazon SWF and other services is encrypted using Transport Layer Security (TLS).

Encryption API Version 2012-01-25 90

Identity and Access Management in Amazon Simple Workflow Service

Access to Amazon SWF requires credentials that Amazon can use to authenticate your requests. These credentials must have permissions to access Amazon resources, such as retrieving event data from other Amazon resources.. The following sections provide details on how you can use Amazon Identity and Access Management (IAM) and Amazon SWF to help secure your resources by controlling access to them.

Amazon Identity and Access Management (IAM) is an Amazon Web Services service that helps an administrator securely control access to Amazon resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon SWF resources. IAM is an Amazon Web Services service that you can use with no additional charge.

Topics

- Audience
- · Authenticating with identities
- Managing access using policies
- Access Control
- Policy actions for Amazon SWF
- Policy resources for Amazon SWF
- Policy condition keys for Amazon SWF
- ACLs in Amazon SWF
- ABAC with Amazon SWF
- Using temporary credentials with Amazon SWF
- Cross-service principal permissions for Amazon SWF
- Service roles for Amazon SWF
- Service-linked roles for Amazon SWF
- Identity-based policies for Amazon SWF
- Resource-based policies within Amazon SWF
- How Amazon Simple Workflow Service works with IAM
- Identity-based policy examples for Amazon Simple Workflow Service
- Basic Principles

- Amazon SWF IAM Policies
- API Summary
- Tag-based Policies
- Amazon VPC endpoints for Amazon SWF
- Troubleshooting Amazon Simple Workflow Service identity and access

Audience

How you use Amazon Identity and Access Management (IAM) differs, depending on the work that you do in Amazon SWF.

Service user – If you use the Amazon SWF service to do your job, then your administrator provides you with the credentials and permissions that you need. As you use more Amazon SWF features to do your work, you might need additional permissions. Understanding how access is managed can help you request the right permissions from your administrator. If you cannot access a feature in Amazon SWF, see Troubleshooting Amazon Simple Workflow Service identity and access.

Service administrator – If you're in charge of Amazon SWF resources at your company, you probably have full access to Amazon SWF. It's your job to determine which Amazon SWF features and resources your service users should access. You must then submit requests to your IAM administrator to change the permissions of your service users. Review the information on this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon SWF, see How Amazon Simple Workflow Service works with IAM.

IAM administrator – If you're an IAM administrator, you might want to learn details about how you can write policies to manage access to Amazon SWF. To view example Amazon SWF identity-based policies that you can use in IAM, see <u>Identity-based policy examples for Amazon Simple Workflow Service</u>.

Authenticating with identities

Authentication is how you sign in to Amazon using your identity credentials. You must be *authenticated* (signed in to Amazon) as the Amazon Web Services account root user, as an IAM user, or by assuming an IAM role.

If you access Amazon programmatically, Amazon provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use Amazon tools, you must sign requests yourself. For more information about using the

Audience API Version 2012-01-25 92

recommended method to sign requests yourself, see <u>Amazon Signature Version 4 for API requests</u> in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, Amazon recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see <u>Amazon Multi-factor authentication in IAM in the IAM User Guide.</u>

Amazon Web Services account root user

When you create an Amazon Web Services account, you begin with one sign-in identity that has complete access to all Amazon Web Services services and resources in the account. This identity is called the Amazon Web Services account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see Tasks that require root user credentials in the *IAM User Guide*.

Federated identity

As a best practice, require human users, including users that require administrator access, to use federation with an identity provider to access Amazon Web Services services by using temporary credentials.

A federated identity is a user from your enterprise user directory, a web identity provider, the Amazon Directory Service, or any user that accesses Amazon Web Services services by using credentials provided through an identity source. When federated identities access Amazon Web Services accounts, they assume roles, and the roles provide temporary credentials.

IAM users and groups

An <u>IAM user</u> is an identity within your Amazon Web Services account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see <u>Rotate access keys regularly for use cases that require long-term credentials</u> in the <u>IAM User Guide</u>.

An <u>IAM group</u> is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier

to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see <u>Use cases for IAM users</u> in the *IAM User Guide*.

IAM roles

An <u>IAM role</u> is an identity within your Amazon Web Services account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. To temporarily assume an IAM role in the Amazon Web Services Management Console, you can <u>switch from a user to an IAM role (console)</u>. You can assume a role by calling an Amazon CLI or Amazon API operation or by using a custom URL. For more information about methods for using roles, see <u>Methods to assume a role</u> in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For information about roles for federation, see <u>Create a role for a third-party identity provider</u> (federation) in the *IAM User Guide*.
- **Temporary IAM user permissions** An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- Cross-account access You can use an IAM role to allow someone (a trusted principal) in a
 different account to access resources in your account. Roles are the primary way to grant crossaccount access. However, with some Amazon Web Services services, you can attach a policy
 directly to a resource (instead of using a role as a proxy). To learn the difference between roles
 and resource-based policies for cross-account access, see Cross account resource access in IAM in
 the IAM User Guide.
- Cross-service access Some Amazon Web Services services use features in other Amazon Web Services services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
 - Forward access sessions (FAS) When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an

action that then initiates another action in a different service. FAS uses the permissions of the principal calling an Amazon Web Services service, combined with the requesting Amazon Web Services service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.

- Service role A service role is an <u>IAM role</u> that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM.
 For more information, see <u>Create a role to delegate permissions to an Amazon Web Services</u> service in the *IAM User Guide*.
- Service-linked role A service-linked role is a type of service role that is linked to an Amazon Web Services service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- Applications running on Amazon EC2 You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making Amazon CLI or Amazon API requests. This is preferable to storing access keys within the EC2 instance. To assign an Amazon role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs that are running on the EC2 instance to get temporary credentials. For more information, see Use an IAM role to grant permissions to applications running on Amazon EC2 instances in the IAM User Guide.

Managing access using policies

You control access in Amazon by creating policies and attaching them to Amazon identities or resources. A policy is an object in Amazon that, when associated with an identity or resource, defines their permissions. Amazon evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in Amazon as JSON documents. For more information about the structure and contents of JSON policy documents, see Overview of JSON policies in the IAM User Guide.

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the iam: GetRole action. A user with that policy can get role information from the Amazon Web Services Management Console, the Amazon CLI, or the Amazon API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see <u>Define custom IAM permissions with customer managed policies</u> in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your Amazon Web Services account. Managed policies include Amazon managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see Choose between managed policies and inline policies in the IAM User Guide.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must <u>specify a principal</u> in a resource-based policy. Principals can include accounts, users, roles, federated users, or Amazon Web Services services.

Resource-based policies are inline policies that are located in that service. You can't use Amazon managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, Amazon WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see <u>Access control list (ACL) overview</u> in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

Amazon supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- Permissions boundaries A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the Principal field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see Permissions boundaries for IAM entities in the IAM User Guide.
- Service control policies (SCPs) SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in Amazon Organizations. Amazon Organizations is a service for grouping and centrally managing multiple Amazon Web Services accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in member accounts, including each Amazon Web Services account root user. For more information about Organizations and SCPs, see Service control policies in the Amazon Organizations User Guide.
- Resource control policies (RCPs) RCPs are JSON policies that you can use to set the maximum available permissions for resources in your accounts without updating the IAM policies attached to each resource that you own. The RCP limits permissions for resources in member accounts and can impact the effective permissions for identities, including the Amazon Web Services account root user, regardless of whether they belong to your organization. For more information about Organizations and RCPs, including a list of Amazon Web Services services that support RCPs, see Resource control policies (RCPs) in the Amazon Organizations User Guide.

• Session policies – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see Session policies in the IAM User Guide.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how Amazon determines whether to allow a request when multiple policy types are involved, see Policy evaluation logic in the *IAM User Guide*.

Access Control

You can have valid credentials to authenticate your requests, but unless you have permissions you cannot create or access Amazon SWF resources. For example, you must have permissions to invoke Amazon Lambda, Amazon Simple Notification Service (Amazon SNS), and Amazon Simple Queue Service (Amazon SQS) targets associated with your Amazon SWF rules.

The following sections describe how to manage permissions for Amazon SWF. We recommend that you read the overview first.

- Basic Principles
- Amazon SWF IAM Policies
- Amazon SWF Policy Examples

Policy actions for Amazon SWF

Supports policy actions: Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Action element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated Amazon API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Access Control API Version 2012-01-25 98

Include actions in a policy to grant permissions to perform the associated operation.

To see a list of Amazon SWF actions, see <u>Resources Defined by Amazon Simple Workflow Service</u> in the *Service Authorization Reference*.

Policy actions in Amazon SWF use the following prefix before the action:

```
swf
```

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [
    "swf:action1",
    "swf:action2"
]
```

To view examples of Amazon SWF identity-based policies, see <u>Identity-based policy examples for</u> Amazon Simple Workflow Service.

Policy resources for Amazon SWF

Supports policy resources: Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its <u>Amazon Resource Name (ARN)</u>. You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

To see a list of Amazon SWF resource types and their ARNs, see <u>Actions Defined by Amazon Simple</u> <u>Workflow Service</u> in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see Resources Defined by Amazon Simple Workflow Service.

Policy resources API Version 2012-01-25 99

To view examples of Amazon SWF identity-based policies, see <u>Identity-based policy examples for</u> Amazon Simple Workflow Service.

Policy condition keys for Amazon SWF

Supports service-specific policy condition keys: Yes

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Condition element (or Condition *block*) lets you specify conditions in which a statement is in effect. The Condition element is optional. You can create conditional expressions that use <u>condition operators</u>, such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple Condition elements in a statement, or multiple keys in a single Condition element, Amazon evaluates them using a logical AND operation. If you specify multiple values for a single condition key, Amazon evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see IAM policy elements: variables and tags in the IAM User Guide.

Amazon supports global condition keys and service-specific condition keys. To see all Amazon global condition keys, see Amazon global condition context keys in the *IAM User Guide*.

To see a list of Amazon SWF condition keys, see <u>Condition Keys for Amazon Simple Workflow</u>
<u>Service</u> in the <u>Service Authorization Reference</u>. To learn with which actions and resources you can use a condition key, see <u>Resources Defined by Amazon Simple Workflow Service</u>.

To view examples of Amazon SWF identity-based policies, see <u>Identity-based policy examples for Amazon Simple Workflow Service</u>.

ACLs in Amazon SWF

Supports ACLs: No

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Policy condition keys API Version 2012-01-25 100

ABAC with Amazon SWF

Supports ABAC (tags in policies): Partial

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. In Amazon, these attributes are called *tags*. You can attach tags to IAM entities (users or roles) and to many Amazon resources. Tagging entities and resources is the first step of ABAC. Then you design ABAC policies to allow operations when the principal's tag matches the tag on the resource that they are trying to access.

ABAC is helpful in environments that are growing rapidly and helps with situations where policy management becomes cumbersome.

To control access based on tags, you provide tag information in the <u>condition element</u> of a policy using the aws:ResourceTag/*key-name*, aws:RequestTag/*key-name*, or aws:TagKeys condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see <u>Define permissions with ABAC authorization</u> in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see <u>Use attribute-based access control</u> (ABAC) in the *IAM User Guide*.

Using temporary credentials with Amazon SWF

Supports temporary credentials: Yes

Some Amazon Web Services services don't work when you sign in using temporary credentials. For additional information, including which Amazon Web Services services work with temporary credentials, see Amazon Web Services services that work with IAM in the IAM User Guide.

You are using temporary credentials if you sign in to the Amazon Web Services Management Console using any method except a user name and password. For example, when you access Amazon using your company's single sign-on (SSO) link, that process automatically creates temporary credentials. You also automatically create temporary credentials when you sign in to the console as a user and then switch roles. For more information about switching roles, see Switch from a user to an IAM role (console) in the IAM User Guide.

ABAC API Version 2012-01-25 101

You can manually create temporary credentials using the Amazon CLI or Amazon API. You can then use those temporary credentials to access Amazon. Amazon recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see Temporary security credentials in IAM.

Cross-service principal permissions for Amazon SWF

Supports forward access sessions (FAS): Yes

When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an Amazon Web Services service, combined with the requesting Amazon Web Services service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see Forward access sessions.

Service roles for Amazon SWF

Supports service roles: Yes

A service role is an IAM role that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see Create a role to delegate permissions to an Amazon Web Services service in the IAM User Guide.



Marning

Changing the permissions for a service role might break Amazon SWF functionality. Edit service roles only when Amazon SWF provides guidance to do so.

Service-linked roles for Amazon SWF

Supports service-linked roles: No

A service-linked role is a type of service role that is linked to an Amazon Web Services service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in

Principal permissions API Version 2012-01-25 102 your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing service-linked roles, see <u>Amazon services that work with IAM</u>. Find a service in the table that includes a Yes in the **Service-linked role** column. Choose the **Yes** link to view the service-linked role documentation for that service.

Identity-based policies for Amazon SWF

Supports identity-based policies: Yes

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see Define custom IAM permissions with customer managed policies in the IAM User Guide.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. You can't specify the principal in an identity-based policy because it applies to the user or role to which it is attached. To learn about all of the elements that you can use in a JSON policy, see IAM JSON policy elements reference in the IAM User Guide.

Identity-based policy examples for Amazon SWF

To view examples of Amazon SWF identity-based policies, see <u>Identity-based policy examples for Amazon Simple Workflow Service</u>.

Resource-based policies within Amazon SWF

Supports resource-based policies: No

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must <u>specify a principal</u> in a resource-based policy. Principals can include accounts, users, roles, federated users, or Amazon Web Services services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. Adding a cross-account principal to a resource-

Identity-based policies API Version 2012-01-25 103

based policy is only half of establishing the trust relationship. When the principal and the resource are in different Amazon Web Services accounts, an IAM administrator in the trusted account must also grant the principal entity (user or role) permission to access the resource. They grant permission by attaching an identity-based policy to the entity. However, if a resource-based policy grants access to a principal in the same account, no additional identity-based policy is required. For more information, see Cross account resource access in IAM in the IAM User Guide.

How Amazon Simple Workflow Service works with IAM

Before you use IAM to manage access to Amazon SWF, learn what IAM features are available to use with Amazon SWF.

IAM features you can use with Amazon Simple Workflow Service

IAM feature	Amazon SWF support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys (service-specific)	Yes
ACLs	No
ABAC (tags in policies)	Partial
Temporary credentials	Yes
Principal permissions	Yes
Service roles	Yes
Service-linked roles	No

To get a high-level view of how Amazon SWF and other Amazon services work with most IAM features, see Amazon services that work with IAM in the IAM User Guide.

Identity-based policy examples for Amazon Simple Workflow Service

By default, users and roles don't have permission to create or modify Amazon SWF resources. They also can't perform tasks by using the Amazon Web Services Management Console, Amazon Command Line Interface (Amazon CLI), or Amazon API. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

To learn how to create an IAM identity-based policy by using these example JSON policy documents, see Create IAM policies (console) in the IAM User Guide.

For details about actions and resource types defined by Amazon SWF, including the format of the ARNs for each of the resource types, see <u>Actions, Resources, and Condition Keys for Amazon Simple Workflow Service</u> in the *Service Authorization Reference*.

Topics

- Policy best practices
- Using the Amazon SWF console
- Allow users to view their own permissions

Policy best practices

Identity-based policies determine whether someone can create, access, or delete Amazon SWF resources in your account. These actions can incur costs for your Amazon Web Services account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- Get started with Amazon managed policies and move toward least-privilege permissions
 - To get started granting permissions to your users and workloads, use the *Amazon managed policies* that grant permissions for many common use cases. They are available in your Amazon Web Services account. We recommend that you reduce permissions further by defining Amazon customer managed policies that are specific to your use cases. For more information, see <u>Amazon managed policies</u> or <u>Amazon managed policies</u> for job functions in the *IAM User Guide*.
- Apply least-privilege permissions When you set permissions with IAM policies, grant only the
 permissions required to perform a task. You do this by defining the actions that can be taken on
 specific resources under specific conditions, also known as least-privilege permissions. For more
 information about using IAM to apply permissions, see Policies and permissions in IAM in the
 IAM User Guide.

- Use conditions in IAM policies to further restrict access You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific Amazon Web Services service, such as Amazon CloudFormation. For more information, see IAM JSON policy elements: Condition in the IAM User Guide.
- Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see <u>Validate policies with IAM Access Analyzer</u> in the *IAM User Guide*.
- Require multi-factor authentication (MFA) If you have a scenario that requires IAM users or a root user in your Amazon Web Services account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see Secure API access with MFA in the IAM User Guide.

For more information about best practices in IAM, see <u>Security best practices in IAM</u> in the *IAM User Guide*.

Using the Amazon SWF console

To access the Amazon Simple Workflow Service console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the Amazon SWF resources in your Amazon Web Services account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (users or roles) with that policy.

You don't need to allow minimum console permissions for users that are making calls only to the Amazon CLI or the Amazon API. Instead, allow access to only the actions that match the API operation that they're trying to perform.

To ensure that users and roles can still use the Amazon SWF console, also attach the Amazon SWF *ConsoleAccess* or *ReadOnly* Amazon managed policy to the entities. For more information, see Adding permissions to a user in the *IAM User Guide*.

Allow users to view their own permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the Amazon CLI or Amazon API.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ViewOwnUserInfo",
            "Effect": "Allow",
            "Action": [
                "iam:GetUserPolicy",
                "iam:ListGroupsForUser",
                "iam:ListAttachedUserPolicies",
                "iam:ListUserPolicies",
                "iam:GetUser"
            ],
            "Resource": ["arn:aws-cn:iam::*:user/${aws:username}"]
        },
        {
            "Sid": "NavigateInConsole",
            "Effect": "Allow",
            "Action": [
                "iam:GetGroupPolicy",
                "iam:GetPolicyVersion",
                "iam:GetPolicy",
                "iam:ListAttachedGroupPolicies",
                "iam:ListGroupPolicies",
                "iam:ListPolicyVersions",
                "iam:ListPolicies",
                "iam:ListUsers"
            ],
            "Resource": "*"
        }
    ]
}
```

Basic Principles

Amazon SWF access control is based primarily on two types of permissions:

Basic Principles API Version 2012-01-25 107

Resource permissions: Which Amazon SWF resources a user can access.

You can express resource permissions only for domains.

• API permissions: Which Amazon SWF actions a user can call.

The simplest approach is to grant full account access—call any Amazon SWF action in any domain—or deny access entirely. However, IAM supports a more granular approach to access control that is often more useful. For example, you could:

- Allow a user to call any Amazon SWF action without restrictions, but only in a specified domain. You could use such a policy to allow workflow applications that are under development to use any action, but only a "sandbox" domain.
- Allow a user to access any domain, but constrain how they use the API. You could use such a policy to allow an "auditor" application to call the API in any domain, but allow only read access.
- Allow a user to call only a limited set of actions in certain domains. You could use such a policy
 to allow a workflow starter to call only the StartWorkflowExecution action in a specified
 domain.

Amazon SWF access control is based on the following principles:

- Access control decisions are based only on IAM policies; all policy auditing and manipulation is done through IAM.
- The access control model uses a deny-by-default policy; any access that isn't explicitly allowed is denied.
- You control access to Amazon SWF resources by attaching appropriate IAM policies to the workflow's actors.
- Resource permissions can be expressed only for domains.
- You can further constrain the usage of some actions by applying conditions to one or more parameters.
- If you grant permission to use <u>RespondDecisionTaskCompleted</u>, you can express permissions for the list of decisions included in that action.

Each of the decisions has one or more parameters, much like a regular API call. To allow for policies to be as readable as possible, you can express permissions on decisions as if they were actual API calls, including applying conditions to some parameters. These types of permissions are called *pseudo API* permissions.

Basic Principles API Version 2012-01-25 108

For a summary of which regular and pseudo API parameters can be constrained by using conditions, see API Summary.

Amazon SWF IAM Policies

An IAM policy contains one or more Statement elements, each of which contains a set of elements that define the policy. For a complete list of elements and a general discussion of how to construct policies, see The Access Policy Language. Amazon SWF access control is based on the following elements:

Effect

(Required) The effect of the statement: deny or allow.



Note

You must explicitly allow access; IAM denies access by default.

Resource

(Required) The resource—an entity in an Amazon service that a user can interact with—that the statement applies to.

You can express resource permissions only for domains. For example, a policy can allow access to only certain domains in your account. To express permissions for a domain, set Resource to the domain's Amazon Resource Name (ARN), which has the format "arn:awscn:swf:Region:AccountID:/domain/DomainName". Region is the Amazon region, AccountID is the account ID with no dashes, and *DomainName* is the domain name.

Action

(Required) The action that the statement applies to, which you refer to by using the following format: serviceId:action. For Amazon SWF, set serviceID to swf. For example, swf:StartWorkflowExecution refers to the StartWorkflowExecution action, and is used to control which users are allowed to start workflows.

If you grant permission to use RespondDecisionTaskCompleted, you can also control access to the included list of decisions by using Action to express permissions for the pseudo API. Because IAM denies access by default, a decider's decision must be explicitly allowed or it will not be accepted. You can use a * value to allow all decisions.

Condition

(Optional) Expresses a constraint on one or more of an action's parameters, which restricts the allowed values.

Amazon SWF actions often have a wide scope, which you can reduce by using IAM conditions. For example, to limit which task lists the PollForActivityTask action is allowed to access, you include a Condition and use the swf:taskList.name key to specify the allowable lists.

You can express constraints for the following entities.

- The workflow type. The name and version have separate keys.
- The activity type. The name and version have separate keys.
- Task lists.
- Tags. You can specify multiple tags for some actions. In that case, each tag has a separate key.

Note

For Amazon SWF, the values are all strings so you constrain a parameter by using a string operator such as StringEquals, which restricts the parameter to a specified string. However, the regular string comparison operators such as StringEquals require all requests to include the parameter. If you don't include the parameter explicitly, and there is no default value such as the default task list provided during type registration, access will be denied.

It is often useful to treat conditions as optional, so that you can call an action without necessarily including the associated parameter. For example, you might want to allow a decider to specify a set of RespondDecisionTaskCompleted decisions, but also allow it to specify only one of them for any particular call. In that case, you constrain the appropriate parameters by using a StringEqualsIfExists operator, which allows access if the parameter satisfies the condition, but doesn't deny access if the parameter is absent.

For a complete list of constrainable parameters and the associated keys, see API Summary.

The following section provides examples of how to construct Amazon SWF policies. For details, see String Conditions.

Amazon SWF Policy Examples

A workflow consists of multiple actors—activities, deciders, and so on. You can control access for each actor by attaching an appropriate IAM policy. This section provides some examples. The following shows the simplest case:

```
{
   "Version": "2012-10-17",
   "Statement" : [ {
        "Effect" : "Allow",
        "Action" : "swf:*",
        "Resource" : "arn:aws-cn:swf:*:123456789012:/domain/*"
   } ]
}
```

If you attach this policy to an actor, it has full account access across all regions. You can use wildcards to have a single value represent multiple resources, actions, or regions.

- The first wildcard (*) in the Resource value indicates that the resource permissions apply to all regions. To restrict permissions to a single region, replace the wildcard with the appropriate region string, such as us-east-1.
- The second wildcard (*) in the Resource value allows the actor to access any of the account's domains in the specified regions.
- The wildcard (*) in the Action value allows the actor to call any Amazon SWF action.

For details on how to use wildcards, see Element Descriptions

The following sections show examples of policies that grant permissions in a more granular way.

Domain Permissions

If you want to restrict a department's workflows to a particular domain, you can use something like:

```
{
   "Version": "2012-10-17",
   "Statement": [ {
        "Effect" : "Allow",
        "Action" : "swf:*",
```

```
"Resource" : "arn:aws-cn:swf:*:123456789012:/domain/department1"
} ]
}
```

If you attach this policy to an actor, it can call any action, but only for the department1 domain.

If you want an actor to have access to more than one domain, you can express permission for each domain separately, as follows:

If you attach this policy to an actor, it can use any Amazon SWF action in the department1 and department2 domains. You can also sometimes use wildcards to represent multiple domains.

API Permissions and Constraints

You control which actions an actor can use with the Action element. Optionally, you can constrain the action's allowable parameter values by using a Condition element.

If you want to restrict an actor to only certain actions, you can use something like the following:

```
"Version": "2012-10-17",
"Statement": [ {
    "Effect" : "Allow",
    "Action" : "swf:StartWorkflowExecution",
    "Resource" : "arn:aws-cn:swf:*:123456789012:/domain/department2"
} ]
```

}

If you attach this policy to an actor, it can call StartWorkflowExecution to start workflows in the department2 domain. It can't use any other actions or start workflows in any other domains.

You can further restrict which workflows an actor can start by constraining one or more of the StartWorkflowExecution parameter values, as follows:

This policy constrains the StartWorkflowExecution action's name and version parameters. If you attach the policy to an actor, it can run only version2 of workflow1 in the department1 domain and both parameters must be included in the request.

You can constrain a parameter without requiring it to be included in a request by using a StringEqualsIfExists operator, as follows:

```
"Version": "2012-10-17",
"Statement" : [ {
    "Effect" : "Allow",
    "Action" : "swf:StartWorkflowExecution",
    "Resource" : "arn:aws-cn:swf:*:123456789012:/domain/some_domain",
    "Condition" : {
        "StringEqualsIfExists" : { "swf:taskList.name" : "task_list_name" }
    }
}
```

```
} ]
```

This policy allows an actor to optionally specify a task list when starting a workflow execution.

You can constrain a list of tags for some actions. In that case, each tag has a separate key, so you use swf:tagList.member.0 to constrain the first tag in the list, swf:tagList.member.1 to constrain the second tag in the list, and so on, up to a maximum of 5. However, you must be careful how you constrain tag lists. For instance, here is an example of a policy that is *not* recommended:

This policy allows you to optionally specify either some_ok_tag or another_ok_tag. However, this policy constrains only the first element of the tag list. The list could have additional elements with arbitrary values that would all be allowed because this policy doesn't apply any conditions to swf:tagList.member.1, swf:tagList.member.2, and so on .

One way to address this issue is to disallow the use of tag lists. The following policy ensures that only some_ok_tag or another_ok_tag are allowed by requiring the list to have only one element.

```
"Version": "2012-10-17",
"Statement" : [ {
    "Effect" : "Allow",
    "Action" : "swf:StartWorkflowExecution",
    "Resource" : "arn:aws-cn:swf:*:123456789012:/domain/some_domain",
    "Condition" : {
```

```
"StringEqualsIfExists" : {
        "swf:tagList.member.0" : "some_ok_tag", "another_ok_tag"
     },
        "Null" : { "swf:tagList.member.1" : "true" }
     }
}
```

Pseudo API Permissions and Constraints

If you want to restrict the decisions available to RespondDecisionTaskCompleted, you must first allow the actor to call RespondDecisionTaskCompleted. You can then express permissions for the appropriate pseudo API members by using the same syntax as for the regular API, as follows:

```
{
  "Version": "2012-10-17",
  "Statement" : [
     {
       "Resource" : "arn:aws-cn:swf:*:123456789012:/domain/*",
       "Action" : "swf:RespondDecisionTaskCompleted",
       "Effect" : "Allow"
     }, {
       "Resource" : "*",
       "Action" : "swf:ScheduleActivityTask",
       "Effect" : "Allow",
       "Condition" : {
          }
     }
  ]
}
```

If you attach this policy to an actor, the first Statement element allows the actor to call RespondDecisionTaskCompleted. The second element allows the actor to use the ScheduleActivityTask decision to direct Amazon SWF to schedule an activity task. To allow all decisions, replace "swf:ScheduleActivityTask" with "swf:*".

You can use Condition operators to constrain parameters just as with the regular API. The StringEquals operator in this Condition allows RespondDecisionTaskCompleted to schedule an activity task for the SomeActivityType activity, and it must schedule that task. If

you want to allow RespondDecisionTaskCompleted to use a parameter value but not require it to do so, you can instead use the StringEqualsIfExists operator.

Amazon managed policy: SimpleWorkflowFullAccess

You can attach the SimpleWorkflowFullAccess policy to your IAM identities.

This policy provides full access to the Amazon SWF configuration service.

Permission details

This policy includes the following permissions.

Service Model Limitations on IAM Policies

You must consider service model constraints when creating IAM policies. It is possible to create a syntactically valid IAM policy that represents an invalid Amazon SWF request; a request that is allowed in terms of access control can still fail because it is an invalid request.

For instance, the following policy for <u>ListOpenWorkflowExecutions</u> is *not* recommended:

```
"Version": "2012-10-17",
"Statement" : [ {
    "Effect" : "Allow",
    "Action" : "swf:ListOpenWorkflowExecutions",
    "Resource" : "arn:aws-cn:swf:*:123456789012:/domain/domain_name",
    "Condition" : {
```

```
"StringEquals" : {
        "swf:typeFilter.name" : "workflow_name",
        "swf:typeFilter.version" : "workflow_version",
        "swf:tagFilter.tag" : "some_tag"
     }
    }
}
```

The Amazon SWF service model doesn't allow the typeFilter and tagFilter parameters to be used in the same ListOpenWorkflowExecutions request. The policy therefore allows calls that the service will reject—by throwing ValidationException—as an invalid request.

API Summary

This section briefly describes how you can use IAM policies to control how an actor can use each API and pseudo API to access Amazon SWF resources.

- For all actions except RegisterDomain and ListDomains, you can allow or deny access to any or all of an account's domains by expressing permissions for the domain resource.
- You can allow or deny permission for any member of the regular API and, if you grant permission to call RespondDecisionTaskCompleted, any member of the pseudo API.
- You can use a Condition to constrain some parameters' allowable values.

The following sections list the parameters that can be constrained for each member of the regular and pseudo API and provide the associated key, and note any limitations on how you can control domain access.

Regular API

This section lists the regular API members, and briefly describes the parameters that can be constrained and the associated keys. It also notes any limitations on how you can control domain access.

CountClosedWorkflowExecutions

- tagFilter.tag String constraint. The key is swf:tagFilter.tag
- typeFilter.name String constraint. The key is swf:typeFilter.name.

• typeFilter.version – String constraint. The key is swf:typeFilter.version.



Note

CountClosedWorkflowExecutions requires typeFilter and tagFilter to be mutually exclusive.

CountOpenWorkflowExecutions

- tagFilter.tag String constraint. The key is swf:tagFilter.tag
- typeFilter.name String constraint. The key is swf:typeFilter.name.
- typeFilter.version String constraint. The key is swf:typeFilter.version.



Note

CountOpenWorkflowExecutions requires typeFilter and tagFilter to be mutually exclusive.

CountPendingActivityTasks

• taskList.name - String constraint. The key is swf:taskList.name.

CountPendingDecisionTasks

• taskList.name - String constraint. The key is swf:taskList.name.

DeleteActivityType

- activityType.name String constraint. The key is swf:activityType.name.
- activityType.version String constraint. The key is swf:activityType.version.

DeprecateActivityType

• activityType.name - String constraint. The key is swf:activityType.name.

• activityType.version - String constraint. The key is swf:activityType.version.

DeprecateDomain

• You can't constrain this action's parameters.

DeleteWorkflowType

- workflowType.name String constraint. The key is swf:workflowType.name.
- workflowType.version String constraint. The key is swf:workflowType.version.

DeprecateWorkflowType

- workflowType.name String constraint. The key is swf:workflowType.name.
- workflowType.version String constraint. The key is swf:workflowType.version.

DescribeActivityType

- activityType.name String constraint. The key is swf:activityType.name.
- activityType.version String constraint. The key is swf:activityType.version.

<u>DescribeDomain</u>

• You can't constrain this action's parameters.

<u>DescribeWorkflowExecution</u>

• You can't constrain this action's parameters.

DescribeWorkflowType

- workflowType.name String constraint. The key is swf:workflowType.name.
- workflowType.version String constraint. The key is swf:workflowType.version.

<u>GetWorkflowExecutionHistory</u>

• You can't constrain this action's parameters.

ListActivityTypes

• You can't constrain this action's parameters.

ListClosedWorkflowExecutions

- tagFilter.tag String constraint. The key is swf:tagFilter.tag
- typeFilter.name String constraint. The key is swf:typeFilter.name.
- typeFilter.version String constraint. The key is swf:typeFilter.version.

Note

ListClosedWorkflowExecutions requires typeFilter and tagFilter to be mutually exclusive.

ListDomains

• You can't constrain this action's parameters.

<u>ListOpenWorkflowExecutions</u>

- tagFilter.tag String constraint. The key is swf:tagFilter.tag
- typeFilter.name String constraint. The key is swf:typeFilter.name.
- typeFilter.version String constraint. The key is swf:typeFilter.version.

Note

ListOpenWorkflowExecutions requires typeFilter and tagFilter to be mutually exclusive.

ListWorkflowTypes

• You can't constrain this action's parameters.

PollForActivityTask

• taskList.name - String constraint. The key is swf:taskList.name.

PollForDecisionTask

• taskList.name - String constraint. The key is swf:taskList.name.

RecordActivityTaskHeartbeat

• You can't constrain this action's parameters.

RegisterActivityType

- defaultTaskList.name String constraint. The key is swf:defaultTaskList.name.
- name String constraint. The key is swf:name.
- version String constraint. The key is swf:version.

RegisterDomain

• name – The name of the domain being registered is available as the resource of this action.

RegisterWorkflowType

- defaultTaskList.name String constraint. The key is swf:defaultTaskList.name.
- name String constraint. The key is swf:name.
- version String constraint. The key is swf:version.

RequestCancelWorkflowExecution

• You can't constrain this action's parameters.

RespondActivityTaskCanceled

You can't constrain this action's parameters.

RespondActivityTaskCompleted

• You can't constrain this action's parameters.

RespondActivityTaskFailed

• You can't constrain this action's parameters.

RespondDecisionTaskCompleted

 decisions.member.N – Restricted indirectly through pseudo API permissions. For details, see Pseudo API.

SignalWorkflowExecution

• You can't constrain this action's parameters.

StartWorkflowExecution

- tagList.member.0 String constraint. The key is swf:tagList.member.0
- tagList.member.1 String constraint. The key is swf:tagList.member.1
- tagList.member.2 String constraint. The key is swf:tagList.member.2
- tagList.member.3 String constraint. The key is swf:tagList.member.3
- tagList.member.4 String constraint. The key is swf:tagList.member.4
- taskList.name String constraint. The key is swf:taskList.name.
- workflowType.name String constraint. The key is swf:workflowType.name.
- workflowType.version String constraint. The key is swf:workflowType.version.

Note

You can't constrain more than five tags.

TerminateWorkflowExecution

You can't constrain this action's parameters.

Pseudo API

This section lists the members of the pseudo API, which represent the decisions included in RespondDecisionTaskCompleted. If you have granted permission to use RespondDecisionTaskCompleted, your policy can express permissions for the members of this API in the same way as the regular API. You can further restrict some members of the pseudo-API by setting conditions on one or more parameters. This section lists the pseudo API members, and briefly describes the parameters that can be constrained and the associated keys.

Note

The aws:SourceIP, aws:UserAgent, and aws:SecureTransport keys are not available for the pseudo API. If your intended security policy requires these keys to control access to the pseudo API, you can use them with the RespondDecisionTaskCompleted action.

CancelTimer

• You can't constrain this action's parameters.

CancelWorkflowExecution

• You can't constrain this action's parameters.

CompleteWorkflowExecution

You can't constrain this action's parameters.

ContinueAsNewWorkflowExecution

- tagList.member.0 String constraint. The key is swf:tagList.member.0
- tagList.member.1 String constraint. The key is swf:tagList.member.1
- tagList.member.2 String constraint. The key is swf:tagList.member.2

- tagList.member.3 String constraint. The key is swf:tagList.member.3
- tagList.member.4 String constraint. The key is swf:tagList.member.4
- taskList.name String constraint. The key is swf:taskList.name.
- workflowTypeVersion String constraint. The key is swf:workflowTypeVersion.



Note

You can't constrain more than five tags.

FailWorkflowExecution

• You can't constrain this action's parameters.

RecordMarker

• You can't constrain this action's parameters.

RequestCancelActivityTask

You can't constrain this action's parameters.

RequestCancelExternalWorkflowExecution

You can't constrain this action's parameters.

ScheduleActivityTask

- activityType.name String constraint. The key is swf:activityType.name.
- activityType.version String constraint. The key is swf:activityType.version.
- taskList.name String constraint. The key is swf:taskList.name.

SignalExternalWorkflowExecution

You can't constrain this action's parameters.

StartChildWorkflowExecution

- tagList.member.0 String constraint. The key is swf:tagList.member.0
- tagList.member.1 String constraint. The key is swf:tagList.member.1
- tagList.member.2 String constraint. The key is swf:tagList.member.2
- tagList.member.3 String constraint. The key is swf:tagList.member.3
- tagList.member.4 String constraint. The key is swf:tagList.member.4
- taskList.name String constraint. The key is swf:taskList.name.
- workflowType.name String constraint. The key is swf:workflowType.name.
- workflowType.version String constraint. The key is swf:workflowType.version.

Note

You can't constrain more than five tags.

StartTimer

• You can't constrain this action's parameters.

Tag-based Policies

Amazon SWF supports policies based on tags. For instance, you could restrict Amazon SWF domains that include a tag with the key environment and the value production:

Taq-based Policies API Version 2012-01-25 125

}

This policy will Deny the access to any domain that has been tagged as environment/ production.

For more information on tagging, see:

- Tags in Amazon SWF
- Controlling Access Using IAM Tags

Amazon VPC endpoints for Amazon SWF



Note

Amazon PrivateLink support is currently available in the Amazon Top Secret - East, Amazon Secret Region, and China Regions only.

If you use Amazon Virtual Private Cloud (Amazon VPC) to host your Amazon resources, you can establish a connection between your Amazon VPC and Amazon Simple Workflow Service workflows. You can use this connection with your Amazon SWF workflows without crossing the public internet.

Amazon VPC lets you launch Amazon resources in a custom virtual network. You can use a VPC to control your network settings, such as the IP address range, subnets, route tables, and network gateways. For more information about VPCs, see the Amazon VPC User Guide.

To connect your Amazon VPC to Amazon SWF you must first define an interface VPC endpoint, which lets you connect your VPC to other Amazon Web Services services. The endpoint provides reliable, scalable connectivity, without requiring an internet gateway, network address translation (NAT) instance, or VPN connection. For more information, see Interface VPC Endpoints (Amazon PrivateLink) in the Amazon VPC User Guide.

Creating the Endpoint

You can create an Amazon SWF endpoint in your VPC using the Amazon Web Services Management Console, the Amazon Command Line Interface (Amazon CLI), an Amazon SDK, the Amazon SWF API, or Amazon CloudFormation.

Amazon VPC endpoints API Version 2012-01-25 126 For information about creating and configuring an endpoint using the Amazon VPC console or the Amazon CLI, see Creating an Interface Endpoint in the Amazon VPC User Guide.



Note

When you create an endpoint, specify Amazon SWF as the service that you want your VPC to connect to. In the Amazon VPC console, service names vary based on the Amazon Region. For example, in the Amazon Top Secret - East Region, the service name for Amazon SWF is **com.amazonaws.us-iso-east-1.swf**.

For information about creating and configuring an endpoint using Amazon CloudFormation, see the AWS::EC2::VPCEndpoint resource in the Amazon CloudFormation User Guide.

Amazon VPC Endpoint Policies

To control connectivity access to Amazon SWF you can attach an Amazon Identity and Access Management (IAM) endpoint policy while creating an Amazon VPC endpoint. You can create complex IAM rules by attaching multiple endpoint policies. For more information, see:

- Amazon Virtual Private Cloud Endpoint Policies for Amazon SWF
- Controlling Access to Services with VPC Endpoints

Amazon Virtual Private Cloud Endpoint Policies for Amazon SWF

You can create an Amazon VPC endpoint policy for Amazon SWF in which you specify the following:

- The principal that can perform actions.
- The actions that can be performed.
- The resources on which the actions can be performed.

The following example shows an Amazon VPC endpoint policy that allows all Amazon SWF operations on a single domain for a specific IAM role.

```
{
  "Version": "2012-10-17",
  "Statement": [
```

Amazon VPC endpoints API Version 2012-01-25 127

```
{
    "Effect": "Allow",
    "Action": "swf:*",
    "Resource": "arn:aws-cn:swf:*:123456789012:/domain/myDomain",
    "Principal": {
        "AWS": "arn:aws-cn:iam::123456789012:role/MyRole"
     }
}
```

- For more information about creating endpoint policies, see <u>Controlling Access to Services with VPC Endpoints.</u>
- For information about how you can use IAM to control access to your Amazon and Amazon SWF resources, see Identity and Access Management in Amazon Simple Workflow Service.

Troubleshooting Amazon Simple Workflow Service identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon SWF and IAM.

Topics

- I am not authorized to perform an action in Amazon SWF
- I am not authorized to perform iam:PassRole
- I want to allow people outside of my Amazon Web Services account to access my Amazon SWF resources

I am not authorized to perform an action in Amazon SWF

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the mateojackson user tries to use the console to view details about a fictional my-example-widget resource but does not have the fictional swf: GetWidget permissions.

```
User: arn:aws-cn:iam::123456789012:user/mateojackson is not authorized to perform: swf:GetWidget on resource: my-example-widget
```

Troubleshooting API Version 2012-01-25 128

In this case, Mateo's policy must be updated to allow him to access the my-example-widget resource using the swf: GetWidget action.

If you need help, contact your Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the iam: PassRole action, your policies must be updated to allow you to pass a role to Amazon SWF.

Some Amazon Web Services services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named marymajor tries to use the console to perform an action in Amazon SWF. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws-cn:iam::123456789012:user/marymajor is not authorized to perform: iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the iam: PassRole action.

If you need help, contact your Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my Amazon Web Services account to access my Amazon SWF resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

• To learn whether Amazon SWF supports these features, see How Amazon Simple Workflow Service works with IAM.

Troubleshooting API Version 2012-01-25 129

- To learn how to provide access to your resources across Amazon Web Services accounts that you own, see IAM User Guide.
- To learn how to provide access to your resources to third-party Amazon Web Services accounts, see <u>Providing access to Amazon Web Services accounts owned by third parties</u> in the *IAM User Guide*.
- To learn how to provide access through identity federation, see Providing access to externally authenticated users (identity federation) in the IAM User Guide.
- To learn the difference between using roles and resource-based policies for cross-account access, see Cross account resource access in IAM in the IAM User Guide.

Logging and Monitoring

This section provides information about logging and monitoring Amazon SWF.

Topics

- Amazon SWF Metrics for CloudWatch
- Viewing Amazon SWF Metrics for CloudWatch using the Amazon Web Services Management Console
- Recording API calls with Amazon CloudTrail
- EventBridge for Amazon SWF execution status changes
- Using Amazon User Notifications with Amazon Simple Workflow Service

Amazon SWF Metrics for CloudWatch

Amazon SWF now provides metrics for CloudWatch that you can use to track your workflows and activities and set alarms on threshold values that you choose. You can view metrics using the Amazon Web Services Management Console. For more information, see <u>Viewing Amazon SWF</u> Metrics for CloudWatch using the Amazon Web Services Management Console.

Topics

- Reporting Units for Amazon SWF Metrics
- API and Decision Event Metrics
- Amazon SWF Metrics

Logging and Monitoring API Version 2012-01-25 130

Amazon SWF non-ASCII resource names and CloudWatch dimensions

Reporting Units for Amazon SWF Metrics

Metrics that Report a Time Interval

Some of the Amazon SWF metrics for CloudWatch are *time intervals*, always measured in milliseconds. The CloudWatch unit is reported as Time. These metrics generally correspond to stages of your workflow execution for which you can set workflow and activity timeouts, and have similar names.

For example, the DecisionTaskStartToCloseTime metric measures the time it took for the decision task to complete after it began executing, which is the same time period for which you can set a DecisionTaskStartToCloseTimeout value.

For a diagram of each of these workflow stages and to learn when they occur over the workflow and activity lifecycles, see Amazon SWF Timeout Types.

Metrics that Report a Count

Some of the Amazon SWF metrics for CloudWatch report results as a *count*. For example, WorkflowsCanceled, records a result as either *one* or *zero*, indicating whether or not the workflow was canceled. A value of zero doesn't indicate that the metric was not reported, only that the condition described by the metric did not occur.

Some of the Amazon SWF metrics for CloudWatch that report a Count in CloudWatch are a *count* per second. For instance, ProvisionedRefillRate, which is reported as a Count in CloudWatch, represents a rate of the Count of requests per second.

For count metrics, minimum and maximum will always be either zero or one, but average will be a value ranging from zero to one.

API and Decision Event Metrics

You can monitor both API and Decision events in CloudWatch to provide insight into your usage and capacity. See <u>deciders</u> in the <u>Basic workflow concepts in Amazon SWF</u> section, and the <u>Decision</u> topic in the Amazon Simple Workflow Service API Reference.

You can also monitor these limits to alarm when you are approaching your Amazon SWF throttling limits. See Amazon SWF throttling quotas for a description of these limits and their default

settings. These limits are designed to prevent incorrect workflows from consuming excessive system resources. To request an increase to your limits see: ???.

As a best practice, you should configure CloudWatch alarms at around 60% of your API or decision events capacity. This will allow you to either adjust your workflow, or request a service limit increase, before Amazon SWF throttling is enabled. Depending on the <u>burstiness</u> of your calls, you can configure different alarms to notify when you are approaching your service limits:

- If your traffic has significant spikes, set an alarm at 60% of your ProvisionedBucketSize limits.
- If your calls have a relatively steady rate, set an alarm at 60% of your ProvisionedRefillRate limit for your related API and decision events.

Amazon SWF Metrics

The following metrics are available for Amazon SWF:

Metric	Description
DecisionTaskSchedu leToStartTime	The time interval, in milliseconds, between the time that the decision task was scheduled and when it was picked up by a worker and started.
	CloudWatch Units: Time
	Dimensions: Domain, WorkflowTypeName, WorkflowTypeVersion
	Valid statistics: Average, Minimum, Maximum
DecisionTaskStartT oCloseTime	The time interval, in milliseconds, between the time that the decision task was started and when it closed.
	CloudWatch Units: Time
	Dimensions: Domain, WorkflowTypeName, WorkflowTypeVersion
	Valid statistics: Average, Minimum, Maximum

Metric	Description
DecisionTasksCompl eted	The count of decision tasks that have been completed.
	CloudWatch Units: Count
	Dimensions: Domain, WorkflowTypeName, WorkflowTypeVersion
	Valid statistics: Sum
PendingTasks	The count of pending tasks in a 1 minute interval for a specific Task List.
	CloudWatch Units: Count
	Dimensions: Domain, TaskListName
	Valid statistics: Sum
StartedDecisionTas ksTimedOutOnClose	The count of decision tasks that started but timed out on closing.
	CloudWatch Units: Count
	Dimensions: Domain, WorkflowTypeName, WorkflowTypeVersion
	Valid statistics: Sum
WorkflowStartToClo seTime	The time, in milliseconds, between the time the workflow started and when it closed.
	CloudWatch Units: Time
	Dimensions: Domain, WorkflowTypeName, WorkflowTypeVersion
	Valid statistics: Average, Minimum, Maximum

Metric	Description
WorkflowsCanceled	The count of workflows that were canceled.
	CloudWatch Units: Count
	Dimensions: Domain, WorkflowTypeName, WorkflowTypeVersion
	Valid statistics: Sum
WorkflowsCompleted	The count of workflows that completed.
	CloudWatch Units: Count
	Dimensions: Domain, WorkflowTypeName, WorkflowTypeVersion
	Valid statistics: Sum
WorkflowsContinued	The count of workflows that continued as new.
AsNew	CloudWatch Units: Count
	Dimensions: Domain, WorkflowTypeName, WorkflowTypeVersion
	Valid statistics: Sum
WorkflowsFailed	The count of workflows that failed.
	CloudWatch Units: Count
	Dimensions: Domain, WorkflowTypeName, WorkflowTypeVersion
	Valid statistics: Sum

Metric	Description
WorkflowsTerminated	The count of workflows that were terminated.
	CloudWatch Units: Count
	<pre>Dimensions: Cause, Domain, WorkflowTypeName, WorkflowTypeVersion</pre>
	Valid statistics: Sum
WorkflowsTimedOut	The count of workflows that timed out, for any reason.
	CloudWatch Units: Count
	Dimensions: Domain, WorkflowTypeName, WorkflowTypeVersion
	Valid statistics: Sum
ActivityTaskSchedu leToCloseTime	The time interval, in milliseconds, between the time when the activity was scheduled and when it closed.
	CloudWatch Units: Time
	Dimensions: Domain, ActivityTypeName, ActivityTypeVersion
	Valid statistics: Average, Minimum, Maximum
ActivityTaskSchedu leToStartTime	The time interval, in milliseconds, between the time when the activity task was scheduled and when it started.
	CloudWatch Units: Time
	Dimensions: Domain, ActivityTypeName, ActivityTypeVersion
	Valid statistics: Average, Minimum, Maximum

Metric	Description
ActivityTaskStartT oCloseTime	The time interval, in milliseconds, between the time when the activity task started and when it closed.
	CloudWatch Units: Time
	<pre>Dimensions: Domain, ActivityTypeName, ActivityT ypeVersion</pre>
	Valid statistics: Average, Minimum, Maximum
ActivityTasksCance led	The count of activity tasks that were canceled.
	CloudWatch Units: Count
	Dimensions: Domain, ActivityTypeName, ActivityTypeVersion
	Valid statistics: Sum
ActivityTasksCompl eted	The count of activity tasks that completed.
	CloudWatch Units: Count
	<pre>Dimensions: Domain, ActivityTypeName, ActivityT ypeVersion</pre>
	Valid statistics: Sum
ActivityTasksFailed	The count of activity tasks that failed.
	CloudWatch Units: Count
	Dimensions: Domain, ActivityTypeName, ActivityTypeVersion
	Valid statistics: Sum

Metric	Description
ScheduledActivityT asksTimedOutOnClose	The count of activity tasks that were scheduled but timed out on close.
	CloudWatch Units: Count
	Dimensions: Domain, ActivityTypeName, ActivityTypeVersion
	Valid statistics: Sum
ScheduledActivityT asksTimedOutOnStart	The count of activity tasks that were scheduled but timed out on start.
	CloudWatch Units: Count
	Dimensions: Domain, ActivityTypeName, ActivityTypeVersion
	Valid statistics: Sum
<pre>StartedActivityTas ksTimedOutOnClose</pre>	The count of activity tasks that were started but timed out on close.
	CloudWatch Units: Count
	Dimensions: Domain, ActivityTypeName, ActivityTypeVersion
	Valid statistics: Sum
StartedActivityTas ksTimedOutOnHeartb eat	The count of activity tasks that were started but timed out due to a heartbeat timeout.
	CloudWatch Units: Count
	Dimensions: Domain, ActivityTypeName, ActivityTypeVersion
	Valid statistics: Sum

Metric	Description
ThrottledEvents	The count of requests that have been throttled.
	CloudWatch Units: Count
	Dimensions: APIName, DecisionName, Throttlin gScope
	Valid statistics: Sum
ProvisionedBucketS ize	The count of available requests per second.
	Dimensions: APIName, DecisionName
	Valid statistics: Minimum
ConsumedCapacity	The count of requests per second.
	CloudWatch Units: Count
	Dimensions: APIName, DecisionName
	Valid statistics: Sum
ConsumedLimit	The amount of general limit that has been consumed.
	Dimensions: GeneralLimitType
ProvisionedRefillR ate	The count of requests per second that are allowed into the bucket.
	Dimensions: APIName, DecisionName
	Valid statistics: Minimum
ProvisionedLimit	The amount of general limit that is provisioned to the account.
	Dimensions: GeneralLimitType

Dimension	Description
Domain	Filters data to the Amazon SWF domain that the workflow or activity is running in.
ActivityTypeName	Filters data to the name of the activity type.
ActivityTypeVersion	Filters data to the version of the activity type.
WorkflowTypeName	Filters data to the name of the workflow type for this workflow execution.
WorkflowTypeVersion	Filters data to the version of the workflow type for this workflow execution.
APIName	Filters data to an API of the specified API name.
DecisionName	Filters data to the specified Decision name.
TaskListName	Filters data to the specified Task List name.
TaskListClassifica tion	Filters data to the classification of the task list. Value is "D" for Decision Task Lists and "A" for Activity Task Lists.
ThrottlingScope	Filters data to the specified throttling scope. Value is "Account" when exceeding account-level quota, or "Workflow" when exceeding workflow-level quota.

Amazon SWF non-ASCII resource names and CloudWatch dimensions

Amazon SWF allows non-ASCII characters in resource names such as TaskList and DomainName. However, the dimension values of CloudWatch metrics can only contain printable ASCII characters. To ensure that Amazon SWF uses dimension values that are compatible with CloudWatch requirements, Amazon SWF resource names that do not meet these requirements are converted and will have a checksum appended as follows:

- Any non-ASCII character is replaced with ?.
- The input string or converted string will, if necessary, be truncated. This ensures that when the checksum is appended, the new string length will not exceed the CloudWatch maximum.

• Because any non-ASCII characters are converted to ?, some CloudWatch metric dimension values that were different before conversion may appear to be the same after conversion. To help differentiate between them, an underscore (_) followed by the first 16 characters of the SHA256 checksum of the original resource name is appended to the resource name.

Conversion examples:

- test àpple would be converted to test ?pple_82cc5b8e3a771d12
- àòà would be converted to ???_2fec5edbb2c05c22.
- The TaskList names applé and apple would both be converted to ?ppl?, and would be identical. Appending the checksum returns distinct values, ?ppl?_f39a36df9d85a69d and ? ppl? da3efb4f11dd0f7f.



You can generate your own SHA256 checksum. For example, to use the shasum command line tool:

echo -n "<the original resource name>" | shasum -a 256 | cut -c1-16

Viewing Amazon SWF Metrics for CloudWatch using the Amazon Web **Services Management Console**

Amazon CloudWatch provides a number of viewable metrics for Amazon SWF workflows and activities. You can view the metrics and set alarms for your Amazon SWF workflow executions using the Amazon Web Services Management Console. You must be logged in to the console to proceed.

For a description of each of the available metrics, see Amazon SWF Metrics for CloudWatch.


Topics

- Viewing Metrics
- Setting Alarms


Viewing Metrics

To view your metrics for Amazon SWF

- 1. Sign in to the Amazon Web Services Management Console and open the CloudWatch console at https://console.amazonaws.cn/cloudwatch/.
- 2. In the navigation pane, under **Metrics**, choose **SWF**.



If you have run any workflow executions recently, you will see two lists of metrics presented: **Workflow Type Metrics** and **Activity Type Metrics**.





Note


Initially you might only see the Workflow Type Metrics; Activity Type Metrics are presented in the same view, but you may need to scroll down to see them.

Up to 50 of the most recent metrics will be shown at a time, divided among workflow and activity metrics.

You can use the interactive headings above each column in the list to sort your metrics using any of the provided dimensions. For workflows, the dimensions are **Domain**, **WorkflowTypeName**, WorkflowTypeVersion, and Metric Name. For activities, the dimensions are Domain, ActivityTypeName, ActivityTypeVersion, and Metric Name.

The various types of metrics are described in Amazon SWF Metrics for CloudWatch.

You can view graphs for metrics by choosing the boxes next to the metric row in the list, and change the graph parameters using the **Time Range** controls to the right of the graph view.



For details about any point on the graph, place your cursor over the graph point. A detail of the point's dimensions will be shown.


For more information about working with CloudWatch metrics, see Viewing, Graphing, and Publishing Metrics in the Amazon CloudWatch User Guide.

Setting Alarms

You can use CloudWatch alarms to perform actions such as notifying you when an alarm threshold is reached. For example, you can set an alarm to send a notification to an SNS topic or to send an email when the WorkflowsFailed metric rises above a certain threshold.

To set an alarm on any of your metrics

- 1. Choose a single metric by choosing its box.
- 2. To the right of the graph, in the **Tools** controls, choose **Create Alarm**.
- 3. On the **Define Alarm** screen, enter the alarm threshold value, period parameters, and actions to take.



For more information about setting and using CloudWatch alarms, see <u>Creating Amazon</u> <u>CloudWatch Alarms</u> in the *Amazon CloudWatch User Guide*.

Recording API calls with Amazon CloudTrail

Amazon Simple Workflow Service is integrated with <u>Amazon CloudTrail</u>, a service that provides a record of actions taken by a user, role, or an Amazon Web Services service. CloudTrail captures all API calls for Amazon SWF as events. The calls captured include calls from the Amazon SWF console and code calls to the Amazon SWF API operations. Using the information collected by CloudTrail, you can determine the request that was made to Amazon SWF, the IP address from which the request was made, when it was made, and additional details.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root user or user credentials.
- Whether the request was made on behalf of an IAM Identity Center user.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another Amazon Web Services service.

CloudTrail is active in your Amazon Web Services account when you create the account and you automatically have access to the CloudTrail **Event history**. The CloudTrail **Event history** provides a viewable, searchable, downloadable, and immutable record of the past 90 days of recorded management events in an Amazon Web Services Region. For more information, see <u>Working with CloudTrail Event history</u> in the *Amazon CloudTrail User Guide*. There are no CloudTrail charges for viewing the **Event history**.

For an ongoing record of events in your Amazon Web Services account past 90 days, create a trail or a <u>CloudTrail Lake</u> event data store.

CloudTrail trails

A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. All trails created using the Amazon Web Services Management Console are multi-Region. You can create a single-Region or a multi-Region trail by using the Amazon CLI. Creating a multi-Region trail is recommended because you capture activity in all Amazon Web Services Regions in your account. If you create a single-Region trail, you can view only the events logged in the trail's Amazon Web Services Region. For more information about trails, see <u>Creating a trail for your Amazon Web Services</u> account and <u>Creating a trail for an organization</u> in the <u>Amazon CloudTrail User Guide</u>.

You can deliver one copy of your ongoing management events to your Amazon S3 bucket at no charge from CloudTrail by creating a trail, however, there are Amazon S3 storage charges. For more information about CloudTrail pricing, see Amazon CloudTrail Pricing. For information about Amazon S3 pricing, see Amazon S3 Pricing.

CloudTrail Lake event data stores

CloudTrail Lake lets you run SQL-based queries on your events. CloudTrail Lake converts existing events in row-based JSON format to Apache ORC format. ORC is a columnar storage format that is optimized for fast retrieval of data. Events are aggregated into event data stores, which are immutable collections of events based on criteria that you select by applying advanced event selectors. The selectors that you apply to an event data store control which events persist and are available for you to query. For more information about CloudTrail Lake, see Working with Amazon CloudTrail Lake in the Amazon CloudTrail User Guide.

CloudTrail Lake event data stores and queries incur costs. When you create an event data store, you choose the <u>pricing option</u> you want to use for the event data store. The pricing option determines the cost for ingesting and storing events, and the default and maximum retention period for the event data store. For more information about CloudTrail pricing, see <u>Amazon</u> <u>CloudTrail Pricing</u>.

Data events in CloudTrail

<u>Data events</u> provide information about the resource operations performed on or in a resource (for example, reading or writing to an Amazon S3 object). These are also known as data plane operations. Data events are often high-volume activities. By default, CloudTrail doesn't log data events. The CloudTrail **Event history** doesn't record data events.

Additional charges apply for data events. For more information about CloudTrail pricing, see Amazon CloudTrail Pricing.

You can log data events for the Amazon SWF resource types by using the CloudTrail console, Amazon CLI, or CloudTrail API operations. For more information about how to log data events, see Logging data events with the Amazon Command Line Interface in the Amazon CloudTrail User Guide.

The following table lists the Amazon SWF resource types for which you can log data events. The **Data event type** column shows the value to choose from the **Data event type** list on the CloudTrail console. The **resources.type value** column shows the resources.type value, which you would specify when configuring advanced event selectors using the Amazon CLI or CloudTrail APIs. The

Data APIs logged to CloudTrail column shows the API calls logged to CloudTrail for the resource type.

You can configure advanced event selectors to filter on the eventName, readOnly, and resources. ARN fields to log only those events that are important to you. For more information about these fields, see AdvancedFieldSelector in the Amazon CloudTrail API Reference.

Data event type	resources.type value	Data APIs logged to CloudTrail
SWF Domain	AWS::SWF::Domain	CountClosedWorkflowExecutions CountOpenWorkflowExecutions DescribeWorkflowExecution ListClosedWorkflowExecution ListOpenWorkflowExecutions ListOpenWorkflowExecutions GetWorkflowExecutionHistory RequestCancelWorkflowExecution SignalWorkflowExecution StartWorkflowExecution TerminateWorkflowExecution TerminateWorkflowExecution Task Events
		 CountPendingActivityTasks PollForDecisionTask PollForActivityTask

Data event type	resources.type value	Data APIs logged to CloudTrail
Data event type	resources.type value	RecordActivityTask Heartbeat RespondActivityTas kCanceled RespondActivityTas kCompleted RespondActivityTaskFailed RespondDecisionTas kCompleted RespondDecisionTas kCompleted CancelTimer CancelWorkflowExecution CompleteWorkflowEx ecution ContinueAsNewWorkflowExecution FailWorkflowExecution RecordMarker
		 RequestCancelActivityTask RequestCancelExter nalWorkflowExecution ScheduleActivityTask ScheduleLambdaFunction SignalExternalWork flowExecution StartChildWorkflow Execution StartTimer

(1) CloudTrail events and RespondDecisionTaskCompleted

The RespondDecisionTaskCompleted action takes a list of decisions in the request payload. A completed call will emit N+1 CloudTrail data events, one for each decision plus one for the API call itself. The data events and API event will all have the same request id.

Management events in CloudTrail

Management events provide information about management operations that are performed on resources in your Amazon Web Services account. These are also known as control plane operations. By default, CloudTrail logs management events.

Amazon Simple Workflow Service logs the following control plane operations to CloudTrail as management events.

Domain Events

- RegisterDomain
- DescribeDomain
- ListDomains
- DeprecateDomain
- UndeprecateDomain

Activity Events

- RegisterActivityType
- DescribeActivityType
- ListActivityTypes
- DeprecateActivityType
- UndeprecateActivityType
- DeleteActivityType

WorkflowType Events

RegisterWorkflowType

- DescribeWorkflowType
- ListWorkflowTypes
- DeprecateWorkflowType
- UndeprecateWorkflowType
- DeleteWorkflowType

Tag Events

- TagResource
- UntagResource
- ListTagsforResource

Example event

An event represents a single request from any source and includes information about the requested API operation, the date and time of the operation, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so events don't appear in any specific order.

The following example shows a CloudTrail event that demonstrates the CountClosedWorkflowExecutions operation.

```
{
    "eventVersion": "1.09",
    "userIdentity": {
      "type": "AssumedRole",
      "principalId": "1234567890abcdef02345:admin",
      "arn": "arn:aws:sts::111122223333:assumed-role/Admin/admin",
      "accountId": "111122223333",
      "accessKeyId": "abcdef01234567890abc",
      "sessionContext": {
        "sessionIssuer": {
          "type": "Role",
          "principalId": "1234567890abcdef02345",
          "arn": "arn:aws:iam::111122223333:role/Admin",
          "accountId": "111122223333",
          "userName": "Admin"
        },
```

```
"attributes": {
        "creationDate": "2023-11-23T16:37:38Z",
        "mfaAuthenticated": "false"
      }
    }
  },
  "eventTime": "2023-11-23T17:52:46Z",
  "eventSource": "swf.amazonaws.com",
  "eventName": "CountClosedWorkflowExecutions",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "198.51.100.42",
  "userAgent": "aws-internal/3 aws-sdk-java/1.11.42",
  "requestParameters": {
    "domain": "nsg-domain",
    "closeTimeFilter": {
      "oldestDate": "Nov 23, 2023 5:52:46 PM",
      "latestDate": "Nov 23, 2023 5:52:46 PM"
    }
  },
  "responseElements": null,
  "requestID": "a1b2c3d4-5678-90ab-cdef-EXAMPLEaaaaa",
  "eventID": "a1b2c3d4-5678-90ab-cdef-EXAMPLEbbbbb",
  "readOnly": true,
  "resources": [
      "accountId": "111122223333",
      "type": "AWS::SWF::Domain",
      "ARN": "arn:aws:swf:us-east-1:111122223333:/domain/nsg-domain"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": false,
  "recipientAccountId": "111122223333",
  "eventCategory": "Data",
  "tlsDetails": {
    "clientProvidedHostHeader": "swf.example.amazondomains.com"
  }
}
```

For information about CloudTrail record contents, see <u>CloudTrail record contents</u> in the *Amazon CloudTrail User Guide*.

EventBridge for Amazon SWF execution status changes

You use Amazon EventBridge to respond to state changes or events in an Amazon resource. When Amazon SWF emits an event, it always goes to the default EventBridge event bus for your account. You can create a rule for events, associate it with the default event bus, and specify a target action to take when EventBridge receives an event that matches the rule. In this way, you can monitor your workflows without having to constantly poll using the GetWorkflowExecutionHistory API. Based on changes in workflow executions, you can use an EventBridge target to call Amazon Lambda functions, publish messages to Amazon Simple Notification Service (Amazon SNS) topics, and more.

You can see the full contents of an execution status change event using DescribeWorkflowExecution.

For more information, see the Amazon EventBridge User Guide.

EventBridge events

The history event types contain the execution state changes. The detail section of each event contains at least the following parameters:

- eventId: the event ID shown by GetWorkflowExecutionHistory.
- workflowExecutionDetail: the state of the workflow when the event was emitted.
- eventType: the history event type, one of the following:
 - ActivityTaskCanceled
 - ActivityTaskFailed
 - ActivityTaskTimedOut
 - WorkflowExecutionCanceled
 - WorkflowExecutionCompleted
 - WorkflowExecutionFailed
 - WorkflowExecutionStarted
 - WorkflowExecutionTerminated
 - WorkflowExecutionTimedOut
 - WorkflowExecutionContinuedAsNew
 - CancelTimerFailed

- CancelWorkflowExecutionFailed
- ChildWorkflowExecutionFailed
- ChildWorkflowExecutionTimedOut
- CompleteWorkflowExecutionFailed
- ContinueAsNewWorkflowExecutionFailed
- DecisionTaskTimedOut
- FailWorkflowExecutionFailed
- RecordMarkerFailed
- RequestCancelActivityTaskFailed
- RequestCancelExternalWorkflowExecutionFailed
- ScheduleActivityTaskFailed
- SignalExternalWorkflowExecutionFailed
- StartActivityTaskFailed
- StartChildWorkflowExecutionFailed
- StartTimerFailed
- TimerCanceled
- LambdaFunctionFailed
- LambdaFunctionTimedOut
- StartLambdaFunctionFailed
- ScheduleLambdaFunctionFailed

Amazon SWF event examples

The following are examples of Amazon SWF sending events to EventBridge:

Topics

- Execution started
- Execution completed
- Execution failed
- Execution timed out

In each case, the detail section in the event data provides the same information as the DescribeWorkflowExecution API. The executionStatus field indicates the status of the execution at the time the event was sent, either OPEN or CLOSED.

Execution started

```
{
  "version": "0",
  "id": "44444444444",
  "detail-type": "Simple Workflow Execution State Change",
  "source": "aws.swf",
  "account": "444444444444",
  "time": "2020-05-08T15:57:38Z",
  "region": "us-east-1",
  "resources": [
    "arn:aws:swf:us-east-1:4444444444444(/domain/SimpleWorkflowUserSimulator"
  ],
  "detail": {
    "eventId": 1,
    "eventType": "WorkflowExecutionStarted",
    "workflowExecutionDetail": {
      "executionInfo": {
        "execution": {
          "workflowId": "123456789012",
          "runId": "AKIAIOSFODNN7EXAMPLE"
        },
        "workflowType": {
          "name": "SimpleWorkflowUserSimulator",
          "version": "myWorkflow"
        },
        "startTimestamp": 1588953458484,
        "closeTimestamp": null,
        "executionStatus": "OPEN",
        "closeStatus": null,
        "parent": null,
        "parentExecutionArn": null,
        "tagList": null,
        "cancelRequested": false
      },
      "executionConfiguration": {
        "taskStartToCloseTimeout": "60",
        "executionStartToCloseTimeout": "1000",
        "taskList": {
```

```
"name": "44444444444"
        },
        "taskPriority": null,
        "childPolicy": "ABANDON",
        "lambdaRole": "arn:aws:iam::44444444444:role/BasicSWFLambdaExecution"
      },
      "openCounts": {
        "openActivityTasks": 0,
        "openDecisionTasks": 1,
        "openTimers": 0,
        "openChildWorkflowExecutions": 0,
        "openLambdaFunctions": 0
      },
      "latestActivityTaskTimestamp": null,
    }
  }
}
```

Execution completed

```
"version": "0",
"id": "1111-2222-3333",
"detail-type": "Simple Workflow Execution State Change",
"source": "aws.swf",
"account": "444455556666",
"time": "2020-05-08T15:57:39Z",
"region": "us-east-1",
"resources": [
  "arn:aws:swf:us-east-1:444455556666:/domain/SimpleWorkflowUserSimulator"
],
"detail": {
  "eventId": 35,
  "eventType": "WorkflowExecutionCompleted",
  "workflowExecutionDetail": {
    "executionInfo": {
      "execution": {
        "workflowId": "1234-5678-9012",
        "runId": "777788889999"
      },
      "workflowType": {
        "name": "SimpleWorkflowUserSimulator",
        "version": "myWorkflow"
```

```
},
        "startTimestamp": 1588953458820,
        "closeTimestamp": 1588953459448,
        "executionStatus": "CLOSED",
        "closeStatus": "COMPLETED",
        "parent": null,
        "parentExecutionArn": null,
        "tagList": null,
        "cancelRequested": false
      },
      "executionConfiguration": {
        "taskStartToCloseTimeout": "60",
        "executionStartToCloseTimeout": "1000",
        "taskList": {
          "name": "1111-1111-1111"
        },
        "taskPriority": null,
        "childPolicy": "ABANDON",
        "lambdaRole": "arn:aws:iam::444455556666:role/BasicSWFLambdaExecution"
      },
      "openCounts": {
        "openActivityTasks": 0,
        "openDecisionTasks": 0,
        "openTimers": 0,
        "openChildWorkflowExecutions": 0,
        "openLambdaFunctions": 0
      },
      "latestActivityTaskTimestamp": 1588953459402,
    }
  }
}
```

Execution failed

```
"version": "0",
"id": "1111-2222-3333",
"detail-type": "Simple Workflow Execution State Change",
"source": "aws.swf",
"account": "444455556666",
"time": "2020-05-08T15:57:38Z",
"region": "us-east-1",
"resources": [
```

```
"arn:aws:swf:us-east-1:444455556666:/domain/SimpleWorkflowUserSimulator"
],
"detail": {
  "eventId": 11,
  "eventType": "WorkflowExecutionFailed",
  "workflowExecutionDetail": {
    "executionInfo": {
      "execution": {
        "workflowId": "1234-5678-9012",
        "runId": "777788889999"
      },
      "workflowType": {
        "name": "SimpleWorkflowUserSimulator",
        "version": "myWorkflow"
      },
      "startTimestamp": 1588953158481,
      "closeTimestamp": 1588953458560,
      "executionStatus": "CLOSED",
      "closeStatus": "FAILED",
      "parent": null,
      "parentExecutionArn": null,
      "tagList": null,
      "cancelRequested": false
    },
    "executionConfiguration": {
      "taskStartToCloseTimeout": "60",
      "executionStartToCloseTimeout": "1000",
      "taskList": {
        "name": "1111-1111-1111"
      },
      "taskPriority": null,
      "childPolicy": "ABANDON",
      "lambdaRole": "arn:aws:iam::444455556666:role/BasicSWFLambdaExecution"
    },
    "openCounts": {
      "openActivityTasks": 0,
      "openDecisionTasks": 0,
      "openTimers": 0,
      "openChildWorkflowExecutions": 0,
      "openLambdaFunctions": 0
    },
    "latestActivityTaskTimestamp": null,
  }
}
```

}

Execution timed out

```
{
 "version": "0",
 "id": "1111-2222-3333",
 "detail-type": "Simple Workflow Execution State Change",
 "source": "aws.swf",
  "account": "444455556666",
 "time": "2020-05-05T17:26:30Z",
  "region": "us-east-1",
 "resources": [
    "arn:aws:swf:us-east-1:444455556666:/domain/SimpleWorkflowUserSimulator"
 ],
 "detail": {
   "eventId": 6,
    "eventType": "WorkflowExecutionTimedOut",
    "workflowExecutionDetail": {
      "executionInfo": {
        "execution": {
          "workflowId": "1234-5678-9012",
          "runId": "777788889999"
        },
        "workflowType": {
          "name": "SimpleWorkflowUserSimulator",
          "version": "myWorkflow"
        },
        "startTimestamp": 1588698073748,
        "closeTimestamp": 1588699590745,
        "executionStatus": "CLOSED",
        "closeStatus": "TIMED_OUT",
        "parent": null,
        "parentExecutionArn": null,
        "tagList": null,
        "cancelRequested": false
      },
      "executionConfiguration": {
        "taskStartToCloseTimeout": "60",
        "executionStartToCloseTimeout": "1000",
        "taskList": {
          "name": "1111-1111"
        },
```

```
"taskPriority": null,
    "childPolicy": "ABANDON",
    "lambdaRole": "arn:aws:iam::444455556666:role/BasicSWFLambdaExecution"
},
    "openCounts": {
        "openActivityTasks": 1,
        "openDecisionTasks": 0,
        "openTimers": 0,
        "openChildWorkflowExecutions": 0,
        "openLambdaFunctions": 0
},
    "latestActivityTaskTimestamp": 1588699585802,
}
}
```

Execution terminated

```
{
  "version": "0",
  "id": "1111-2222-3333",
  "detail-type": "Simple Workflow Execution State Change",
  "source": "aws.swf",
  "account": "444455556666",
  "time": "2020-05-08T22:37:26Z",
  "region": "us-east-1",
  "resources": [
    "arn:aws:swf:us-east-1:444455556666:/domain/canary"
  ],
  "detail": {
    "eventId": 48,
    "eventType": "WorkflowExecutionTerminated",
    "workflowExecutionDetail": {
      "executionInfo": {
        "execution": {
          "workflowId": "1234-5678-9012",
          "runId": "777788889999"
        },
        "workflowType": {
          "name": "1111-1111-1111",
          "version": "1.3"
        },
        "startTimestamp": 1588977445279,
```

```
"closeTimestamp": 1588977446062,
        "executionStatus": "CLOSED",
        "closeStatus": "TERMINATED",
        "parent": null,
        "parentExecutionArn": null,
        "tagList": null,
        "cancelRequested": false
      },
      "executionConfiguration": {
        "taskStartToCloseTimeout": "60",
        "executionStartToCloseTimeout": "120",
        "taskList": {
          "name": "1111-1111-1111-2222-2222-2222"
        },
        "taskPriority": null,
        "childPolicy": "TERMINATE",
        "lambdaRole": null
      },
      "openCounts": {
        "openActivityTasks": 0,
        "openDecisionTasks": 1,
        "openTimers": 0,
        "openChildWorkflowExecutions": 0,
        "openLambdaFunctions": 0
      "latestActivityTaskTimestamp": 1588977445882,
    }
  }
}
```

Using Amazon User Notifications with Amazon Simple Workflow Service

You can use <u>Amazon User Notifications</u> to set up delivery channels to get notified about Amazon Simple Workflow Service events. You receive a notification when an event matches a rule that you specify. You can receive notifications for events through multiple channels, including email, <u>Amazon Q Developer in chat applications</u> chat notifications, or <u>Amazon Console Mobile Application</u> push notifications. You can also see notifications in the <u>Console Notifications Center</u>. User Notifications supports aggregation, which can reduce the number of notifications you receive during specific events.

Compliance Validation for Amazon Simple Workflow Service

Third-party auditors assess the security and compliance of Amazon Simple Workflow Service as part of multiple Amazon compliance programs. These include SOC, PCI, FedRAMP, HIPAA, and others.

For a list of Amazon services in scope of specific compliance programs, see <u>Amazon Services in Scope by Compliance Program</u>. For general information, see <u>Amazon Compliance Programs</u>.

You can download third-party audit reports using Amazon Artifact. For more information, see Downloading Reports in Amazon Artifact.

Your compliance responsibility when using Amazon SWF is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. Amazon provides the following resources to help with compliance:

- <u>Security and Compliance Quick Start GuidesSecurity and Compliance Quick Start Guides</u> These deployment guides discuss architectural considerations and provide steps for deploying security-and compliance-focused baseline environments on Amazon.
- <u>Architecting for HIPAA Security and Compliance Whitepaper</u> This whitepaper describes how companies can use Amazon to create HIPAA-compliant applications.
- <u>Amazon Compliance Resources</u> This collection of workbooks and guides might apply to your industry and location.
- <u>Evaluating Resources with Rules</u> in the *Amazon Config Developer Guide* The Amazon Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- <u>Amazon Security Hub</u> This Amazon service provides a comprehensive view of your security state within Amazon that helps you check your compliance with security industry standards and best practices.

Resilience in Amazon Simple Workflow Service

The Amazon global infrastructure is built around Amazon Regions and Availability Zones. Amazon Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones

Compliance Validation API Version 2012-01-25 160

without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about Amazon Regions and Availability Zones, see <u>Amazon Global</u> Infrastructure.

In addition to the Amazon global infrastructure, Amazon SWF offers several features to help support your data resiliency and backup needs.

Infrastructure Security in Amazon Simple Workflow Service

As a managed service, is protected by Amazon global network security. For information about Amazon security services and how Amazon protects infrastructure, see Amazon Cloud Security. To design your Amazon environment using the best practices for infrastructure security, see Infrastructure Protection in Security Pillar Amazon Well-Architected Framework.

You use Amazon published API calls to access through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the <u>Amazon Security Token Service</u> (Amazon STS) to generate temporary security credentials to sign requests.

You can call these API operations from any network location, but Amazon SWF does support resource-based access policies, which can include restrictions based on the source IP address. You can also use Amazon SWF policies to control access from specific Amazon Virtual Private Cloud (Amazon VPC) endpoints or specific VPCs. Effectively, this isolates network access to a given Amazon SWF resource from only the specific VPC within the Amazon network.

Infrastructure Security API Version 2012-01-25 161

Configuration and Vulnerability Analysis in Amazon Simple Workflow Service

Configuration and IT controls are a shared responsibility between Amazon and you, our customer. For more information, see the Amazon shared responsibility model.

Using the Amazon CLI with Amazon Simple Workflow Service

Many of the features of Amazon Simple Workflow Service can be accessed from the Amazon CLI. The Amazon CLI provides an alternative to using Amazon SWF with the Amazon Web Services Management Console or in some cases, to programming with the Amazon SWF API and the Amazon Flow Framework.

For example, you can use the Amazon CLI to register a new workflow type:

```
aws swf register-workflow-type --domain MyDomain --name "MySimpleWorkflow" --workflow-version "v1"
```

You can also list your registered workflow types:

```
aws swf list-workflow-types --domain MyDomain --registration-status REGISTERED
```

The following shows an example of the default output in JSON:

```
{
    "typeInfos": [
        {
            "status": "REGISTERED",
            "creationDate": 1377471607.752,
            "workflowType": {
                "version": "v1",
                "name": "MySimpleWorkflow"
            }
        },
            "status": "REGISTERED",
            "creationDate": 1371454149.598,
            "description": "MyDomain subscribe workflow",
            "workflowType": {
                "version": "v3",
                "name": "subscribe"
            }
```

}

The Amazon SWF commands in Amazon CLI provide the ability to start and manage workflow executions, poll for activity tasks, record task heartbeats, and more! For a complete list of Amazon SWF commands, with descriptions of the available arguments and examples showing their use, see Amazon SWF commands in the Amazon CLI Command Reference.

The Amazon CLI commands follow the Amazon SWF API closely, so you can use the Amazon CLI to learn about the underlying Amazon SWF API. You can also use your existing API knowledge to prototype code or perform Amazon SWF actions on the command line.

To learn more about the Amazon CLI, see the Amazon Command Line Interface User Guide.

Working with Amazon SWF APIs

In addition to using the Amazon SDKs that are described in <u>Develop with Amazon SDKs</u>, you can use the HTTP API directly.

To use the API, you send HTTP requests to the <u>SWF endpoint</u> that matches the region that you want to use for your domains, workflows and activities. For more information about making HTTP requests for Amazon SWF, see <u>Making HTTP Requests to Amazon SWF</u>.

This section provides basic information about using the HTTP API to develop your workflows with Amazon SWF. More advanced features, such as using timers, logging with CloudTrail and tagging your workflows are provided in the section, Basic workflow concepts in Amazon SWF.

Topics

- Making HTTP Requests to Amazon SWF
- List of Amazon SWF Actions by Category
- Registering a Domain with Amazon SWF
- Setting timeout values in Amazon SWF
- Registering a Workflow Type with Amazon SWF
- Registering an Activity Type with Amazon SWF
- Amazon Lambda tasks in Amazon SWF
- Developing an Activity Worker in Amazon SWF
- Developing deciders in Amazon SWF
- Starting workflows in Amazon SWF
- Setting task priority in Amazon SWF
- Handling errors in Amazon SWF

Making HTTP Requests to Amazon SWF

If you don't use one of the Amazon SDKs, you can perform Amazon Simple Workflow Service (Amazon SWF) operations over HTTP using the POST request method. The POST method requires that you specify the operation in the header of the request and provide the data for the operation in JSON format in the body of the request.

Making HTTP Requests API Version 2012-01-25 165

HTTP Header Contents

Amazon SWF requires the following information in the header of an HTTP request:

- host The Amazon SWF endpoint.
- x-amz-date You must provide the time stamp in either the HTTP Date header or the Amazon x-amz-date header (some HTTP client libraries don't let you set the Date header). When an x-amz-date header is present, the system ignores any Date header when authenticating the request.

The date must be specified in one of the following three formats, as specified in the HTTP/1.1 RFC:

- Sun, 06 Nov 1994 08:49:37 GMT (RFC 822, updated by RFC 1123)
- Sunday, 06-Nov-94 08:49:37 GMT (RFC 850, obsoleted by RFC 1036)
- Sun Nov 6 08:49:37 1994 (ANSI C's asctime() format)
- x-amzn-authorization The signed request parameters in the format:

AWS3 AWSAccessKeyId=####,Algorithm=HmacSHA256, [,SignedHeaders=Header1;Header2;...] Signature=S(StringToSign)

AWS3 – This is an Amazon implementation-specific tag that denotes the authentication version used to sign the request (currently, for Amazon SWF this value is always AWS3).

AWSAccessKeyId – Your Amazon Access Key ID.

Algorithm – The algorithm used to create the HMAC-SHA value of the string-to-sign, such as HmacSHA256 or HmacSHA1.

Signature – Base64(Algorithm(StringToSign, SigningKey)). For details see <u>Calculating the HMAC-SHA Signature for Amazon SWF</u>

SignedHeaders – (Optional) If present, must contain a list of all the HTTP Headers used in the Canonicalized HttpHeaders calculation. A single semicolon character (;) (ASCII character 59) must be used as the delimiter for list values.

 x-amz-target – The destination service of the request and the operation for the data, in the format

com.amazonaws.swf.service.model.SimpleWorkflowService. + <action>

HTTP Header Contents API Version 2012-01-25 166

For example,

com.amazonaws.swf.service.model.SimpleWorkflowService.RegisterDomain

 content-type – The type needs to specify JSON and the character set, as application/ json; charset=UTF-8

The following is an example header for an HTTP request to create a domain.

```
POST http://swf.us-east-1.amazonaws.com.cn/ HTTP/1.1
Host: swf.us-east-1.amazonaws.com.cn
User-Agent: Mozilla/5.0 (Windows; U; Windows NT 6.1; en-US; rv:1.9.2.25) Gecko/20111212
 Firefox/3.6.25 ( .NET CLR 3.5.30729; .NET4.0E)
Accept: application/json, text/javascript, */*
Accept-Language: en-us, en; q=0.5
Accept-Encoding: gzip, deflate
Accept-Charset: ISO-8859-1, utf-8; q=0.7, *; q=0.7
Keep-Alive: 115
Connection: keep-alive
Content-Type: application/json; charset=UTF-8
X-Requested-With: XMLHttpRequest
X-Amz-Date: Fri, 13 Jan 2012 18:42:12 GMT
X-Amz-Target: com.amazonaws.swf.service.model.SimpleWorkflowService.RegisterDomain
Content-Encoding: amz-1.0
X-Amzn-Authorization: AWS3
 AWSAccessKeyId=AKIAIOSFODNN7EXAMPLE, Algorithm=HmacSHA256, SignedHeaders=Host; X-Amz-
Date;X-Amz-Target;Content-Encoding,Signature=tzjkF551xAxPhzp/BRGFYQRQRq6CqrM254dTDE/
EncI=
Referer: http://swf.us-east-1.amazonaws.com.cn/explorer/index.html
Content-Length: 91
Pragma: no-cache
Cache-Control: no-cache
{"name": "867530902",
 "description": "music",
 "workflowExecutionRetentionPeriodInDays": "60"}
```

Here is an example of the corresponding HTTP response.

```
HTTP/1.1 200 OK
Content-Length: 0
Content-Type: application/json
```

HTTP Header Contents API Version 2012-01-25 167

```
x-amzn-RequestId: 4ec4ac3f-3e16-11e1-9b11-7182192d0b57
```

HTTP Body Content

The body of an HTTP request contains the data for the operation specified in the header of the HTTP request. Use the JSON data format to convey data values and data structure, simultaneously. Elements can be nested within other elements using bracket notation. For example, the following shows a request to list all workflow executions that started between two specified points in time—using Unix Time notation.

```
{
  "domain": "867530901",
  "startTimeFilter":
  {
    "oldestDate": 1325376070,
    "latestDate": 1356998399
  },
    "tagFilter":
  {
      "tag": "music purchase"
  }
}
```

Sample Amazon SWF JSON Request and Response

The following example shows a request to Amazon SWF for a description of the domain that we created previously. Then it shows the Amazon SWF response.

HTTP POST Request

```
POST http://swf.us-east-1.amazonaws.com.cn/ HTTP/1.1
Host: swf.us-east-1.amazonaws.com.cn
User-Agent: Mozilla/5.0 (Windows; U; Windows NT 6.1; en-US; rv:1.9.2.25) Gecko/20111212
Firefox/3.6.25 ( .NET CLR 3.5.30729; .NET4.0E)
Accept: application/json, text/javascript, */*
Accept-Language: en-us,en;q=0.5
Accept-Encoding: gzip,deflate
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7
Keep-Alive: 115
Connection: keep-alive
Content-Type: application/json; charset=UTF-8
```

HTTP Body Content API Version 2012-01-25 168

```
X-Requested-With: XMLHttpRequest
X-Amz-Date: Sun, 15 Jan 2012 03:13:33 GMT
X-Amz-Target: com.amazonaws.swf.service.model.SimpleWorkflowService.DescribeDomain
Content-Encoding: amz-1.0
X-Amzn-Authorization: AWS3
AWSAccessKeyId=AKIAIOSFODNN7EXAMPLE,Algorithm=HmacSHA256,SignedHeaders=Host;X-Amz-Date;X-Amz-Target;Content-
Encoding,Signature=IFJtq3M366CHqMlTpyqYqd9z0ChCoKDC5SCJBsLifu4=
Referer: http://swf.us-east-1.amazonaws.com.cn/explorer/index.html
Content-Length: 21
Pragma: no-cache
Cache-Control: no-cache

{"name": "867530901"}
```

Amazon SWF Response

```
HTTP/1.1 200 OK
Content-Length: 137
Content-Type: application/json
x-amzn-RequestId: e86a6779-3f26-11e1-9a27-0760db01a4a8

{"configuration":
    {"workflowExecutionRetentionPeriodInDays": "60"},
    "domainInfo":
    {"description": "music",
        "name": "867530901",
        "status": "REGISTERED"}
}
```

Notice the protocol (HTTP/1.1) is followed by a status code (200). A code value of 200 indicates a successful operation.

Amazon SWF doesn't serialize null values. If your JSON parser is set to serialize null values for requests, Amazon SWF ignores them.

Calculating the HMAC-SHA Signature for Amazon SWF

Every request to Amazon SWF must be authenticated. The Amazon SDKs automatically sign your requests and manage your token-based authentication. However, if you want to write your own HTTP POST requests, you need to create an x-amzn-authorization value for the HTTP POST Header content as part of your request authentication.

For more information about formatting headers, see HTTP Header Contents. For the Amazon SDK for Java implementation of Amazon Version 3 signing, see the AWSSigner.java class.

Creating a Request Signature

Before you create an HMAC-SHA request signature, you must get your Amazon credentials (the Access Key ID and the Secret Key).



Important

You can use either SHA1 or SHA256 to sign your requests. However, make sure that you use the same method throughout the signing process. The method you choose must match the value of the Algorithm name in the HTTP header.

To create the request signature

- Create a canonical form of the HTTP request headers. The canonical form of the HTTP header includes the following:
 - host
 - Any header element starting with x-amz-

For more information about the included headers, see HTTP Header Contents.

- For each header name-value pair, convert the header name (but not the header value) to lowercase.
- Build a map of the header name to comma-separated header values.

```
x-amz-example: value1
x-amz-example: value2 => x-amz-example:value1,value2
```

For more information, see Section 4.2 of RFC 2616.

For each header name-value pair, convert the name-value pair into a string in the format c. headerName: headerValue. Trim any whitespace from the beginning and end of both headerName and headerValue, with no spaces on each side of the colon.

```
x-amz-example1:value1,value2
```

x-amz-example2:value3

- Insert a new line (U+000A) after each converted string, including the last string.
- Sort the collection of converted strings alphabetically, by header name.
- Create a string-to-sign value that includes the following items: 2.
 - Line 1: The HTTP method (POST), followed by a newline.
 - Line 2: The request URI (/), followed by a newline.
 - Line 3: An empty string followed by a newline.



Note

Typically, the query string appears here, but Amazon SWF doesn't use a query string.

- Lines 4-n: The string representing the canonicalized request headers that you computed in Step 1, followed by a newline. This newline creates a blank line between the headers and the body of the HTTP request. For more information, see RFC 2616.
- The request body, *not* followed by a newline.
- 3. Compute the SHA256 or SHA1 digest of the string-to-sign value. Use the same SHA method throughout the process.
- Compute and Base64-encode the HMAC-SHA using either a SHA256 or a SHA1 digest (depending on the method you used) of the value resulting from the previous step and the temporary secret access key from the Amazon Security Token Service using the GetSessionToken API action.



(i) Note

Amazon SWF expects an equals sign (=) at the end of the Base64-encoded HMAC-SHA value. If your Base64 encoding routine doesn't include the appended equals sign, append one to the end of the value.

For more information about using temporary security credentials with Amazon SWF and other Amazon services, see Amazon Services That Work with IAM in the IAM User Guide.

5. Place the resulting value as the value for the Signature name in the x-amznauthorization header of the HTTP request to Amazon SWF.

6. Amazon SWF verifies the request and performs the specified operation.

List of Amazon SWF Actions by Category

This section lists the reference topics for Amazon SWF actions in the Amazon SWF application programming interface (API). These are listed by *functional category*.

For an alphabetic list of actions, see the Amazon Simple Workflow Service API Reference.

Topics

- Actions Related to Activities
- Actions Related to Deciders
- Actions Related to Workflow Executions
- Actions Related to Administration
- Visibility Actions

Actions Related to Activities

Activity workers use PollForActivityTask to get new activity tasks. After a worker receives an activity task from Amazon SWF, it performs the task and responds using RespondActivityTaskCompleted if successful or RespondActivityTaskFailed if unsuccessful.

The following are actions that are performed by activity workers.

- PollForActivityTask
- RespondActivityTaskCompleted
- RespondActivityTaskFailed
- RespondActivityTaskCanceled
- RecordActivityTaskHeartbeat

Actions Related to Deciders

Deciders use PollForDecisionTask to get decision tasks. After a decider receives a decision task from Amazon SWF, it examines its workflow execution history and decides what to do next. It calls

List of Amazon SWF Actions API Version 2012-01-25 172

RespondDecisionTaskCompleted to complete the decision task and provides zero or more next decisions.

The following are actions that are performed by deciders.

- PollForDecisionTask
- RespondDecisionTaskCompleted

Actions Related to Workflow Executions

The following actions operate on a workflow execution.

- RequestCancelWorkflowExecution
- StartWorkflowExecution
- SignalWorkflowExecution
- TerminateWorkflowExecution

Actions Related to Administration

Although you can perform administrative tasks from the Amazon SWF console, you can use the actions in this section to automate functions or build your own administrative tools.

Activity Management

- RegisterActivityType
- <u>DeprecateActivityType</u>
- UndeprecateActivityType
- DeleteActivityType

Workflow Management

- RegisterWorkflowType
- DeprecateWorkflowType
- <u>UndeprecateWorkflowType</u>

DeleteWorkflowType

Domain Management

These actions allow you to register and deprecate Amazon SWF domains.

- RegisterDomain
- DeprecateDomain
- UndeprecateDomain

For more information and examples of these domain management actions, see <u>Registering a</u> Domain with Amazon SWF.

Workflow Execution Management

- RequestCancelWorkflowExecution
- TerminateWorkflowExecution

Visibility Actions

Although you can perform visibility actions from the Amazon SWF console, you can use the actions in this section to build your own console or administrative tools.

Activity Visibility

- ListActivityTypes
- DescribeActivityType

Workflow Visibility

- ListWorkflowTypes
- DescribeWorkflowType

Workflow Execution Visibility

• DescribeWorkflowExecution

Visibility Actions API Version 2012-01-25 174

- ListOpenWorkflowExecutions
- ListClosedWorkflowExecutions
- CountOpenWorkflowExecutions
- CountClosedWorkflowExecutions
- GetWorkflowExecutionHistory

Domain Visibility

- ListDomains
- DescribeDomain

Task List Visibility

- CountPendingActivityTasks
- CountPendingDecisionTasks

Registering a Domain with Amazon SWF

Your workflow and activity types and the workflow execution itself are all scoped to a *domain*. Domains isolate a set of types, executions, and task lists from others within the same account.

You can register a domain by using the Amazon Web Services Management Console or by using the RegisterDomain action in the Amazon SWF API. The following example uses the API.

```
https://swf.us-east-1.amazonaws.com.cn
RegisterDomain
{
    "name" : "867530901",
    "description" : "music",
    "workflowExecutionRetentionPeriodInDays" : "60"
}
```

The parameters are specified in JavaScript Object Notation (JSON) format. Here, the retention period is set to 60 days. During the retention period, all information about the workflow execution is available through visibility operations using either the Amazon Web Services Management Console or the Amazon SWF API.

Registering a Domain API Version 2012-01-25 175

After registering the domain, you should register the workflow type and the activity types used by the workflow. You need to register the domain first because a registered domain name is part of the required information for registering workflow and activity types.

See Also

RegisterDomain in the Amazon Simple Workflow Service API Reference

Setting timeout values in Amazon SWF

Topics

- Quotas on Timeout Values
- Workflow Execution and Decision Task Timeouts
- Activity Task Timeouts
- See Also

Quotas on Timeout Values

Timeout values are always declared in seconds, and can be set to any number of seconds up to a year (31536000 seconds)—the maximum execution limit for any workflow or activity. The special value NONE is used to set a timeout parameter to "no timeout", or infinite, but the maximum limit of a year still applies.

Workflow Execution and Decision Task Timeouts

You can set timeout values for your Workflow and Decision tasks when registering the workflow type. For example:

```
https://swf.us-east-1.amazonaws.com.cn
RegisterWorkflowType
{
    "domain": "867530901",
    "name": "customerOrderWorkflow",
    "version": "1.0",
    "description": "Handle customer orders",
    "defaultTaskStartToCloseTimeout": "600",
    "defaultExecutionStartToCloseTimeout": "3600",
```

See Also API Version 2012-01-25 176

```
"defaultTaskList": { "name": "mainTaskList" },
   "defaultChildPolicy": "TERMINATE"
}
```

This workflow type registration sets the <u>defaultTaskStartToCloseTimeout</u> to 600 seconds (10 minutes), and <u>defaultExecutionStartToCloseTimeout</u> to 3600 seconds (1 hour).

For more information about workflow type registration, see <u>Registering a Workflow Type with</u> Amazon SWF, and RegisterWorkflowType in the *Amazon Simple Workflow Service API Reference*.

You can override the value set for defaultExecutionStartToCloseTimeout by specifying executionStartToCloseTimeout i.

Activity Task Timeouts

You can set timeout values for your activity tasks when registering the activity type. For example:

```
https://swf.us-east-1.amazonaws.com.cn
RegisterActivityType
{
    "domain": "867530901",
    "name": "activityVerify",
    "version": "1.0",
    "description": "Verify the customer credit",
    "defaultTaskStartToCloseTimeout": "600",
    "defaultTaskHeartbeatTimeout": "120",
    "defaultTaskList": { "name": "mainTaskList" },
    "defaultTaskScheduleToStartTimeout": "1800",
    "defaultTaskScheduleToCloseTimeout": "5400"
}
```

This activity type registration sets the <u>defaultTaskStartToCloseTimeout</u> to 600 seconds (10 minutes), the <u>defaultTaskHeartbeatTimeout</u> to 120 seconds (2 minutes), the <u>defaultTaskScheduleToStartTimeout</u> to 1800 seconds (30 minutes) and defaultTaskScheduleToCloseTimeout to 5400 seconds (1.5 hours).

For more information about activity type registration, see <u>Registering an Activity Type with</u>
<u>Amazon SWF</u>, and <u>RegisterActivityType</u> in the *Amazon Simple Workflow Service API Reference*.

You can override the value set for defaultTaskStartToCloseTimeout by specifying taskStartToCloseTimeout when scheduling the activity task.

Activity Task Timeouts API Version 2012-01-25 177

See Also

Amazon SWF Timeout Types

Registering a Workflow Type with Amazon SWF

The example discussed in this section registers a workflow type using the Amazon SWF API. The name and version that you specify during registration form a unique identifier for the workflow type. The specified domain must have already been registered using the RegisterDomain API action.

The timeout parameters in the following example are duration values specified in seconds. For the defaultTaskStartToCloseTimeout parameter, you can use the duration specifier NONE to indicate no timeout. However, you can't specify a value of NONE for defaultExecutionStartToCloseTimeout; there is a one-year maximum limit on the time that a workflow execution can run. Exceeding this limit always causes the workflow execution to time out. If you specify a value for defaultExecutionStartToCloseTimeout that is greater than one year, the registration will fail.

```
https://swf.us-east-1.amazonaws.com.cn
RegisterWorkflowType
{
    "domain" : "867530901",
    "name" : "customerOrderWorkflow",
    "version" : "1.0",
    "description" : "Handle customer orders",
    "defaultTaskStartToCloseTimeout" : "600",
    "defaultExecutionStartToCloseTimeout" : "3600",
    "defaultTaskList" : { "name": "mainTaskList" },
    "defaultChildPolicy" : "TERMINATE"
}
```

See Also

RegisterWorkflowType in the Amazon Simple Workflow Service API Reference

Registering an Activity Type with Amazon SWF

The following example registers an activity type by using the Amazon SWF API. The name and version that you specify during registration form a unique identifier for the activity type within

See Also API Version 2012-01-25 178

the domain. The specified domain must have already been registered using the RegisterDomain action.

The timeout parameters in this example are duration values specified in seconds. You can use the duration specifier NONE to indicate no timeout.

```
https://swf.us-east-1.amazonaws.com.cn
RegisterActivityType
{
    "domain" : "867530901",
    "name" : "activityVerify",
    "version" : "1.0",
    "description" : "Verify the customer credit",
    "defaultTaskStartToCloseTimeout" : "600",
    "defaultTaskHeartbeatTimeout" : "120",
    "defaultTaskList" : { "name" : "mainTaskList" },
    "defaultTaskScheduleToStartTimeout" : "1800",
    "defaultTaskScheduleToCloseTimeout" : "5400"
}
```

See Also

RegisterActivityType in the Amazon Simple Workflow Service API Reference

Amazon Lambda tasks in Amazon SWF

Topics

- About Amazon Lambda
- Benefits and limitations of using Lambda tasks
- Using Lambda tasks in your workflows

About Amazon Lambda

Amazon Lambda is a fully managed compute service that runs your code in response to events generated by custom code or from various Amazon services such as Amazon S3, DynamoDB, Amazon Kinesis, Amazon SNS, and Amazon Cognito. For more information about Lambda, see the Amazon Lambda Developer Guide.

See Also API Version 2012-01-25 179

Amazon Simple Workflow Service provides a Lambda task so that you can run Lambda functions in place of, or alongside traditional Amazon SWF activities.

Important

Your Amazon account will be charged for Lambda executions (requests) executed by Amazon SWF on your behalf. For details about Lambda pricing, see http:// www.amazonaws.cn/lambda/pricing/.

Benefits and limitations of using Lambda tasks

There are a number of benefits of using Lambda tasks in place of a traditional Amazon SWF activity:

- Lambda tasks don't need to be registered or versioned like Amazon SWF activity types.
- You can use any existing Lambda functions that you've already defined in your workflows.
- · Lambda functions are called directly by Amazon SWF; there is no need for you to implement a worker program to execute them as you must do with traditional activities.
- Lambda provides you with metrics and logs for tracking and analyzing your function executions.

There are also a number of limitations regarding Lambda tasks that you should be aware of:

- Lambda tasks can only be run in Amazon regions that provide support for Lambda. See Lambda Regions and Endpoints in the Amazon Web Services General Reference for details about the currently-supported regions for Lambda.
- Lambda tasks are currently supported only by the base SWF HTTP API and in the Amazon Flow Framework for Java. There is currently no support for Lambda tasks in the Amazon Flow Framework for Ruby.

Using Lambda tasks in your workflows

To use Lambda tasks in your Amazon SWF workflows, you will need to:

- 1. Set up IAM roles to provide Amazon SWF with permission to invoke Lambda functions.
- 2. Attach the IAM roles to your workflows.

3. Call your Lambda function during a workflow execution.

Set up an IAM role

Before you can invoke Lambda functions from Amazon SWF you must provide an IAM role that provides access to Lambda from Amazon SWF. You can either:

- choose a pre-defined role, *AWSLambdaRole*, to give your workflows permission to invoke any Lambda function associated with your account.
- define your own policy and associated role to give workflows permission to invoke particular Lambda functions, specified by their Amazon Resource Names (ARNs).

Limit permissions on an IAM role

You can limit permissions on an IAM role you provide to Amazon SWF by using the SourceArn and SourceAccount context keys in your resource trust policy. These keys limit the usage of an IAM policy so that it is used only from Amazon Simple Workflow Service executions that belong in the specified domain ARN. If you use both global condition context keys, the aws:SourceAccount value and the account referenced in the aws:SourceArn value must use the same account ID when used in the same policy statement.

In the following trust policy example, we use the SourceArn context key to restrict the IAM service role to only be used in Amazon Simple Workflow Service executions that belong to someDomain in the account, 123456789012.

```
}
]
}
```

In the following trust policy example, we use the SourceAccount context key to restrict the IAM service role to only be used in Amazon Simple Workflow Service executions in the account, 123456789012.

```
{
  "Version": "2012-10-17",
  "Statement": [
     {
       "Sid": "",
       "Effect": "Allow",
       "Principal": {
         "Service": "swf.amazonaws.com"
       },
       "Action": "sts:AssumeRole",
       "Condition": {
         "StringLike": {
           "aws:SourceAccount": "123456789012"
       }
      }
  ]
}
```

Providing Amazon SWF with access to invoke any Lambda role

You can use the pre-defined role, AWSLambdaRole, to give your Amazon SWF workflows the ability to invoke any Lambda function associated with your account.

To use AWSLambdaRole to give Amazon SWF access to invoke Lambda functions

- 1. Open the <u>Amazon IAM console</u>.
- 2. Choose **Roles**, then **Create New Role**.
- 3. Give your role a name, such as swf-lambda and choose **Next Step**.
- 4. Under Amazon Service Roles, choose Amazon SWF, and choose Next Step.
- 5. On the Attach Policy screen, choose AWSLambdaRole from the list.
- 6. Choose **Next Step** and then **Create Role** once you've reviewed the role.

Defining an IAM role to provide access to invoke a specific Lambda function

If you want to provide access to invoke a specific Lambda function from your workflow, you will need to define your own IAM policy.

To create an IAM policy to provide access to a particular Lambda function

- 1. Open the Amazon IAM console.
- 2. Choose **Policies**, then **Create Policy**.
- 3. Choose **Copy an Amazon Managed Policy** and select **AWSLambdaRole** from the list. A policy will be generated for you. Optionally edit its name and description to suit your needs.
- 4. In the Resource field of the Policy Document, add the ARN of your Lambda function(s). For example:

Note

For a complete description of how to specify resources in an IAM role, see <u>Overview of IAM Policies</u> in *Using IAM*.

5. Choose **Create Policy** to finish creating your policy.

You can then select this policy when creating a new IAM role, and use that role to give invoke access to your Amazon SWF workflows. This procedure is very similar to creating a role with the *AWSLambdaRole* policy. instead, choose your own policy when creating the role.

To create a Amazon SWF role using your Lambda policy

- 1. Open the Amazon IAM console.
- 2. Choose Roles, then Create New Role.
- 3. Give your role a name, such as swf-lambda-function and choose **Next Step**.
- 4. Under Amazon Service Roles, choose Amazon SWF, and choose Next Step.
- 5. On the Attach Policy screen, choose your Lambda function-specific policy from the list.
- 6. Choose **Next Step** and then **Create Role** once you've reviewed the role.

Attach the IAM role to your workflow

Once you've defined your IAM role, you will need to attach it to the workflow that will be using it to call the Lambda functions you provided Amazon SWF with access to.

There are two places where you can attach the role to your workflow:

- During workflow type registration. This role then may be used as the default Lambda role for every execution of that workflow type.
- When starting a workflow execution. This role will be used only during this workflow's execution (and throughout the entire execution).

To provide a default Lambda role for a workflow type

• When calling RegisterWorkflowType, set the defaultLambdaRole field to the ARN of the role that you defined.

To provide a Lambda role to be used during a workflow execution

• When calling StartWorkflowExecution, set the lambdaRole field to the ARN of the role that you defined.

Note

if the account calling RegisterWorkflowType or StartWorkflowExecution doesn't have permission to use the given role, then the call will fail with an OperationNotPermittedFault.

Call your Lambda function from a Amazon SWF workflow

You can use the ScheduleLambdaFunctionDecisionAttributes data type to identify the Lambda function to call during a workflow execution.

During a call to RespondDecisionTaskCompleted, provide a ScheduleLambdaFunctionDecisionAttributes to your decisions list. For example:

```
{
  "decisions": [{
     "ScheduleLambdaFunctionDecisionAttributes": {
        "id": "lambdaTaskId",
        "name": "myLambdaFunctionName",
        "input": "inputToLambdaFunction",
        "startToCloseTimeout": "30"
     },
     }],
}
```

Set the following parameters:

- *id* with an identifier for the Lambda task. This must be a string from 1-256 characters and must not contain the characters: (colon), / (slash), | (vertical bar), nor any control characters (\u00000 \u0001f and \u0007f \u0009f), nor the literal string arn.
- *name* with the name of your Lambda function. Your Amazon SWF workflow must be provided with an IAM role that gives it access to call the Lambda function. The name provided must follow the constraints for the FunctionName parameter like in the Lambda Invoke action.
- *input* with optional input data for the function. If set, this must follow the constraints for the ClientContext parameter like in the Lambda Invoke action.
- *startToCloseTimeout* with an optional maximum period, in seconds, that the function can take to execute before the task fails with a timeout exception. The value NONE can be used to specify unlimited duration.

For more information, see Implementing Amazon Lambda Tasks

Developing an Activity Worker in Amazon SWF

An activity worker provides the implementation of one or more activity types. An activity worker communicates with Amazon SWF to receive activity tasks and perform them. You can have a fleet of multiple activity workers performing activity tasks of the same activity type.

Amazon SWF makes an activity task available to activity workers when the decider schedules the activity task. When a decider schedules an activity task, it provides the data (which you determine) that the activity worker needs to perform the activity task. Amazon SWF inserts this data into the activity task before sending it to the activity worker.

Activity workers are managed by you. They can be written in any language. A worker can be run anywhere, as long as it can communicate with Amazon SWF through the API. Because Amazon SWF provides all the information needed to perform an activity task, all activity workers can be stateless. Statelessness enables your workflows to be highly scalable; to handle increased capacity requirements, simply add more activity workers.

This section explains how to implement an activity worker. The activity workers should repeatedly do the following.

- 1. Poll Amazon SWF for an activity task.
- 2. Begin performing the task.
- 3. Periodically report a heartbeat to Amazon SWF if the task is long-lived.
- 4. Report that the task completed or failed and return the results to Amazon SWF.

Topics

- Polling for Activity Tasks
- Performing the Activity Task
- Reporting Activity Task Heartbeats
- Completing or Failing an Activity Task
- Launching Activity Workers

Polling for Activity Tasks

To perform activity tasks, each activity worker must poll Amazon SWF by periodically calling the PollForActivityTask action.

In the following example, the activity worker ChargeCreditCardWorker01 polls for a task on the task list, ChargeCreditCard-v0.1. If no activity tasks are available, after 60 seconds, Amazon SWF sends back an empty response. An empty response is a Task structure in which the value of the taskToken is an empty string.

```
https://swf.us-east-1.amazonaws.com.cn
PollForActivityTask
{
    "domain" : "867530901",
    "taskList" : { "name": "ChargeCreditCard-v0.1" },
    "identity" : "ChargeCreditCardWorker01"
}
```

If an activity task becomes available, Amazon SWF returns it to the activity worker. The task contains the data that the decider specifies when it schedules the activity.

After an activity worker receives an activity task, it is ready to perform the work. The next section provides information about performing an activity task.

Performing the Activity Task

After receiving an activity task, the activity worker is ready to perform it.

To perform an activity task

- 1. Program your activity worker to interpret the content in the input field of the task. This field contains the data specified by the decider when the task was scheduled.
- 2. Program the activity worker to begin processing the data and executing your logic.

The next section describes how to program your activity workers to provide status updates to Amazon SWF for long running activities.

Reporting Activity Task Heartbeats

If a heartbeat timeout was registered with the activity type, then the activity worker must record a heartbeat before the heartbeat timeout is exceeded. If an activity task doesn't provide a heartbeat within the timeout, the task times out, Amazon SWF closes it and schedules a new decision task to inform a decider of the timeout. The decider can then reschedule the activity task or take another action.

If, after timing out, the activity worker attempts to contact Amazon SWF, such as by calling RespondActivityTaskCompleted, Amazon SWF will return an UnknownResource fault.

This section describes how to provide an activity heartbeat.

To record an activity task heartbeat, program your activity worker to call the RecordActivityTaskHeartbeat action. This action also provides a string field that you can use to store free-form data to quantify progress in whatever way works for your application.

In this example, the activity worker reports heartbeat to Amazon SWF and uses the details field to report that the activity task is 40 percent complete. To report heartbeat, the activity worker must specify the task token of the activity task.

```
https://swf.us-east-1.amazonaws.com.cn
RecordActivityTaskHeartbeat
{
    "taskToken" : "12342e17-80f6-FAKE-TASK-TOKEN32f0223",
    "details" : "40"
}
```

This action doesn't in itself create an event in the workflow execution history; however, if the task times out, the workflow execution history will contain a ActivityTaskTimedOut event that contains the information from the last heartbeat generated by the activity worker.

Completing or Failing an Activity Task

After executing a task, the activity worker should report whether the activity task completed or failed.

Completing an Activity Task

To complete an activity task, program the activity worker to call the RespondActivityTaskCompleted action after it successfully completes an activity task, specifying the task token.

In this example, the activity worker indicates that the task completed successfully.

```
https://swf.us-east-1.amazonaws.com.cn
RespondActivityTaskCompleted
{
    "taskToken": "12342e17-80f6-FAKE-TASK-TOKEN32f0223",
    "results": "40"
```

}

When the activity completes, Amazon SWF schedules a new decision task for the workflow execution with which the activity is associated.

Program the activity worker to poll for another activity task after it has completed the task at hand. This creates a loop where the activity worker continuously polls for and completes tasks.

If the activity doesn't respond within the *StartToCloseTimeout* period, or if *ScheduleToCloseTimeout* has been exceeded, Amazon SWF times out the activity task and schedules a decision task. This enables a decider to take an appropriate action, such as rescheduling the task.

For example, if an Amazon EC2 instance is executing an activity task and the instance fails before the task is complete, the decider receives a timeout event in the workflow execution history. If the activity task is using a heartbeat, the decider receives the event when the task fails to deliver the next heartbeat after the Amazon EC2 instance fails. If not, the decider eventually receives the event when the activity task fails to complete before it hits one of its overall timeout values. It is then up to the decider to re-assign the task or take some other action.

Failing an Activity Task

If an activity worker can't perform an activity task for some reason, but it can still communicate with Amazon SWF, you can program it to fail the task.

To program an activity worker to fail an activity task, program the activity worker to call the RespondActivityTaskFailed action that specifies the task token of the task.

```
https://swf.us-east-1.amazonaws.com.cn
RespondActivityTaskFailed
{
    "taskToken" : "12342e17-80f6-FAKE-TASK-TOKEN32f0223",
    "reason" : "CC-Invalid",
    "details" : "Credit Card Number Checksum Failed"
}
```

As the developer, you define the values that are stored in the reason and details fields. These are free-form strings; you can use any error code conventions that serve your application. Amazon SWF doesn't process these values. However, Amazon SWF may display these values in the console.

When an activity task is failed, Amazon SWF schedules a decision task for the workflow execution with which the activity task is associated to inform the decider of the failure. Program your decider

to handle failed activities, such as by rescheduling the activity or failing the workflow execution, depending on the nature of the failure.

Launching Activity Workers

To launch activity workers, package your logic into an executable that you can use on your activity worker platform. For example, you might package your activity code as a Java executable that you can run on both Linux and Windows servers.

Once launched, your workers start polling for tasks. Until the decider schedules activity tasks, though, these polls time out with no tasks and your workers just continue polling.

Because polls are outbound requests, activity worker can run on any network that has access to the Amazon SWF endpoint.

You can launch as many activity workers as you like. As the decider schedules activity tasks, Amazon SWF automatically distributes the activity tasks to the polling activity workers.

Developing deciders in Amazon SWF

A decider is an implementation of the coordination logic of your workflow type that runs during the execution of your workflow. You can run multiple deciders for a single workflow type.

Because the execution state for a workflow execution is stored in its workflow history, deciders can be stateless. Amazon SWF maintains the workflow execution history and provides it to a decider with each decision task. This enables you to dynamically add and remove deciders as necessary, which makes the processing of your workflows highly scalable. As the load on your system grows, you simply add more deciders to handle the increased capacity. Note, however, that there can be only one decision task open at any time for a given workflow execution.

Every time a state change occurs for a workflow execution, Amazon SWF schedules a decision task. Each time a decider receives a decision task, it does the following:

- Interprets the workflow execution history provided with the decision task
- Applies the coordination logic based on the workflow execution history and makes decisions on what to do next. Each decision is represented by a Decision structure
- Completes the decision task and provides a list of decisions to Amazon SWF.

This section describes how to develop a decider, which involves:

Launching Activity Workers API Version 2012-01-25 190

- Programming your decider to poll for decision tasks
- Programming your decider to interpret the workflow execution history and make decisions
- Programming your decider to respond to a decision task.

The examples in this section show how you might program a decider for the e-commerce example workflow.

You can implement the decider in any language that you like and run it anywhere, as long as it can communicate with Amazon SWF through its service API.

Topics

- Defining Coordination Logic
- Polling for Decision Tasks
- Applying the Coordination Logic
- Responding with Decisions
- Closing a Workflow Execution
- Launching Deciders

Defining Coordination Logic

The first thing to do when developing a decider is to define the coordination logic. In the e-commerce example, coordination logic that schedules each activity after the previous activity completes might look similar to the following:

```
IF lastEvent = "StartWorkflowInstance"
  addToDecisions ScheduleVerifyOrderActivity

ELSIF lastEvent = "CompleteVerifyOrderActivity"
  addToDecisions ScheduleChargeCreditCardActivity

ELSIF lastEvent = "CompleteChargeCreditCardActivity"
  addToDecisions ScheduleCompleteShipOrderActivity

ELSIF lastEvent = "CompleteShipOrderActivity"
  addToDecisions ScheduleRecordOrderCompletion
ELSIF lastEvent = "CompleteRecordOrderCompletion"
  addToDecisions CloseWorkflow
```

Defining Coordination Logic API Version 2012-01-25 191

ENDIF

The decider applies the coordination logic to the workflow execution history, and creates a list of decisions when completing the decision task using the RespondDecisionTaskCompleted action.

Polling for Decision Tasks

Each decider polls for decision tasks. The decision tasks contain the information that the decider uses to generate decisions such as scheduling activity tasks. To poll for decision tasks, the decider uses the PollForDecisionTask action.

In this example, the decider polls for a decision task, specifying the customerOrderWorkflow-0.1 tasklist.

```
https://swf.us-east-1.amazonaws.com.cn
PollForDecisionTask
{
    "domain": "867530901",
    "taskList": {"name": "customerOrderWorkflow-v0.1"},
    "identity": "Decider01",
    "maximumPageSize": 50,
    "reverseOrder": true
}
```

If a decision task is available from the task list specified, Amazon SWF returns it immediately. If no decision task is available, Amazon SWF holds the connection open for up to 60 seconds, and returns a task as soon as it becomes available. If no task becomes available, Amazon SWF returns an empty response. An empty response is a Task structure in which the value of taskToken is an empty string. Make sure to program your decider to poll for another task if it receives an empty response.

If a decision task is available, Amazon SWF returns a response that contains the decision task as well as a paginated view of the workflow execution history.

In this example, the type of the most recent event indicates the workflow execution started and the input element contains the information needed to perform the first task.

```
{
    "events": [
        {
```

Polling for Decision Tasks API Version 2012-01-25 192

```
"decisionTaskStartedEventAttributes": {
        "identity": "Decider01",
        "scheduledEventId": 2
      },
      "eventId": 3,
      "eventTimestamp": 1326593394.566,
      "eventType": "DecisionTaskStarted"
    }, {
      "decisionTaskScheduledEventAttributes": {
        "startToCloseTimeout": "600",
        "taskList": { "name": "specialTaskList" }
      },
      "eventId": 2,
      "eventTimestamp": 1326592619.474,
      "eventType": "DecisionTaskScheduled"
    }, {
      "eventId": 1,
      "eventTimestamp": 1326592619.474,
      "eventType": "WorkflowExecutionStarted",
      "workflowExecutionStartedEventAttributes": {
        "childPolicy" : "TERMINATE",
        "executionStartToCloseTimeout" : "3600",
        "input" : "data-used-decider-for-first-task",
        "parentInitiatedEventId": 0,
        "tagList" : ["music purchase", "digital", "ricoh-the-dog"],
        "taskList": { "name": "specialTaskList" },
        "taskStartToCloseTimeout": "600",
        "workflowType": {
          "name": "customerOrderWorkflow",
          "version": "1.0"
        }
      }
    }
  ],
}
```

After receiving the workflow execution history, the decider interprets history and makes decisions based on its coordination logic.

Because the number of workflow history events for a single workflow execution might be large, the result returned might be split up across a number of pages. To retrieve subsequent pages, make additional calls to PollForDecisionTask using the *nextPageToken* returned by the initial call.

Polling for Decision Tasks API Version 2012-01-25 193

Note that you do *not* call GetWorkflowExecutionHistory with this *nextPageToken*. Instead, call PollForDecisionTask again.

Applying the Coordination Logic

After the decider receives a decision task, program it to interpret the workflow execution history to determine what has happened so far. Based on this, it should generate a list of decisions.

In the e-commerce example, we are concerned only with the last event in the workflow history, so we define the following logic.

```
IF lastEvent = "StartWorkflowInstance"
   addToDecisions ScheduleVerifyOrderActivity

ELSIF lastEvent = "CompleteVerifyOrderActivity"
   addToDecisions ScheduleChargeCreditCardActivity

ELSIF lastEvent = "CompleteChargeCreditCardActivity"
   addToDecisions ScheduleCompleteShipOrderActivity

ELSIF lastEvent = "CompleteShipOrderActivity"
   addToDecisions ScheduleRecordOrderCompletion

ELSIF lastEvent = "CompleteRecordOrderCompletion"
   addToDecisions CloseWorkflow
ENDIF
```

If the lastEvent is CompleteVerifyOrderActivity, you would add the ScheduleChargeCreditCardActivity activity to the list of decisions.

After the decider determines the decision(s) to make, it can respond to Amazon SWF with appropriate decisions.

Responding with Decisions

After interpreting the workflow history and generating a list of decisions, the decider is ready to respond back to Amazon SWF with those decisions.

Program your decider to extract the data that it needs from the workflow execution history, then create decisions that specify the next appropriate actions for the workflow. The decider transmits

these decision back to Amazon SWF using the RespondDecisionTaskCompleted action. See the *Amazon Simple Workflow Service API Reference* for a list of the available decision types.

In the e-commerce example, when the decider responds with the set of decisions that it generated, it also includes the credit card input from the workflow execution history. The activity worker then has the information it needs to perform the activity task.

When all activities in the workflow execution are complete, the decider closes the workflow execution.

```
https://swf.us-east-1.amazonaws.com.cn
RespondDecisionTaskCompleted
{
  "taskToken": "12342e17-80f6-FAKE-TASK-TOKEN32f0223",
  "decisions" : [
    {
      "decisionType" : "ScheduleActivityTask",
      "scheduleActivityTaskDecisionAttributes" : {
        "control" :"OPTIONAL_DATA_FOR_DECIDER",
        "activityType" : {
          "name" : "ScheduleChargeCreditCardActivity",
          "version" :"1.1"
        },
        "activityId" : "3e2e6e55-e7c4-beef-feed-aa815722b7be",
        "scheduleToCloseTimeout" : "360",
        "taskList" : { "name" : "CC_TASKS" },
        "scheduleToStartTimeout" : "60",
        "startToCloseTimeout" :"300",
        "heartbeatTimeout" :"60",
        "input" : "4321-0001-0002-1234: 0212 : 234"
      }
    }
  ]
}
```

Closing a Workflow Execution

When the decider determines that the business process is complete, that is, that there are no more activities to perform, the decider generates a decision to close the workflow execution.

To close a workflow execution, program your decider to interpret the events in the workflow history to determine what has happened in the execution so far and see if the workflow execution should be closed.

If the workflow has completed successfully, then close the workflow execution by calling RespondDecisionTaskCompleted with the CompleteWorkflowExecution decision. Alternatively, you can fail an erroneous execution using the FailWorkflowExecution decision.

In the e-commerce example, the decider reviews the history and based on the coordination logic adds a decision to close the workflow execution to its list of decisions, and initiates a RespondDecisionTaskCompleted action with a close workflow decision.

Note

There are some cases where closing a workflow execution fails. For example, if a signal is received while the decider is closing the workflow execution, the close decision will fail. To handle this possibility, ensure that the decider continues polling for decision tasks. Also, ensure that the decider that receives the next decision task responds to the event—in this case, a signal—that prevented the execution from closing.

You might also support cancellation of workflow executions. This could be especially useful for long running workflows. To support cancellation, your decider should handle the WorkflowExecutionCancelRequested event in the history. This event indicates that cancellation of the execution has been requested. Your decider should perform appropriate clean-up actions, such as canceling ongoing activity tasks, and closing the workflow calling the RespondDecisionTaskCompleted action with the CancelWorkflowExecution decision.

The following example calls RespondDecisionTaskCompleted to specify that the current workflow execution is canceled.

```
https://swf.us-east-1.amazonaws.com.cn
RespondDecisionTaskCompleted
{
  "taskToken": "12342e17-80f6-FAKE-TASK-TOKEN32f0223",
  "decisions" : [
    {
      "decisionType": "CancelWorkflowExecution",
      "CancelWorkflowExecutionAttributes":{
        "Details": "Customer canceled order"
```

```
}
}
]
}
```

Amazon SWF checks to ensure that the decision to close or cancel the workflow execution is the last decision sent by the decider. That is, it isn't valid to have a set of decisions in which there are decisions after the one that closes the workflow.

Launching Deciders

After completing decider development, you are ready to launch one or more deciders.

To launch deciders, package your coordination logic into an executable that you can use on your decider platform. For example, you might package your decider code as a Java executable that you can run on both Linux and Windows computers.

Once launched, your deciders should start polling Amazon SWF for tasks. Until you start workflow executions and Amazon SWF schedules decision tasks, these polls will time out and get empty responses. An empty response is a Task structure in which the value of taskToken is an empty string. Your deciders should simply continue to poll.

Amazon SWF ensures that only one decision task can be active for a workflow execution at any time. This prevents issues such as conflicting decisions. Additionally, Amazon SWF ensures that a single decision task is assigned to a single decider, regardless of the number of deciders that are running.

If something occurs that generates a decision task while a decider is processing another decision task, Amazon SWF queues the new task until the current task completes. After the current task completes, Amazon SWF makes the new decision task available. Also, decision tasks are batched in the sense that, if multiple activities complete while a decider is processing a decision task, Amazon SWF will create only a single new decision task to account for the multiple task completions. However, each task completion will receive an individual event in the workflow execution history.

Because polls are outbound requests, deciders can run on any network that has access to the Amazon SWF endpoint.

In order for workflow executions to progress, one or more deciders must be running. You can launch as many deciders as you like. Amazon SWF supports multiple deciders polling on the same task list.

Launching Deciders API Version 2012-01-25 197

Starting workflows in Amazon SWF

You can start a workflow execution of a registered workflow type from any application using the StartWorkflowExecution action. When you start the execution you associate an identifier, called the workflowId, with it. The workflowId can be any string that is appropriate for your application, such as the order number in an order processing application. You can't use the same workflowId for multiple open workflow executions within the same domain. For example, if you start two workflow executions with the workflowId Customer Order 01, the second workflow execution will not start and the request will fail. You can, however, reuse the workflowId of a closed execution. Amazon SWF also associates a unique system generated identifier, called the runId, with each workflow execution.

After the workflow and activity types are registered, start the workflow by calling the StartWorkflowExecution action. The value of the input parameter can be any string specified by the application that is starting the workflow. The executionStartToCloseTimeout is the length of time in seconds that the workflow execution can consume from start to close. Exceeding this limit causes the workflow execution to time out. Unlike some of the other timeout parameters in Amazon SWF, you can't specify a value of NONE for this timeout; there is a one-year maximum limit on the time that a workflow execution can run. Similarly, the taskStartToCloseTimeout is the length of time in seconds that a decision task associated with this workflow execution can take before timing out.

```
https://swf.us-east-1.amazonaws.com.cn
StartWorkflowExecution
{
    "domain" : "867530901",
    "workflowId" : "20110927-T-1",
    "workflowType" : {
        "name" : "customerOrderWorkflow", "version" : "1.1"
    },
    "taskList" : { "name" : "specialTaskList" },
    "input" : "arbitrary-string-that-is-meaningful-to-the-workflow",
    "executionStartToCloseTimeout" : "1800",
    "tagList" : [ "music purchase", "digital", "ricoh-the-dog" ],
    "taskStartToCloseTimeout" : "1800",
    "childPolicy" : "TERMINATE"
}
```

Starting workflows API Version 2012-01-25 198

If the StartWorkflowExecution action is successful, Amazon SWF returns the runId for the workflow execution. The runId for a workflow execution is unique within a specific region. Save the runId in case you later need to specify this workflow execution in a call to Amazon SWF. For example, you would use the runId if you later needed to send a signal to the workflow execution.

{"runId": "9ba33198-4b18-4792-9c15-7181fb3a8852"}

Setting task priority in Amazon SWF

By default, tasks on a task list are delivered based upon their arrival time: tasks that are scheduled first are generally run first, as far as possible. By setting an optional task priority, you can give priority to certain tasks: Amazon SWF will attempt to deliver higher-priority tasks on a task list before those with lower priority.



Note

Tasks that are scheduled first generally run first, but this is not guaranteed.

You can set task priorities for both workflows and activities. A workflow's task priority doesn't affect the priority of any activity tasks it schedules, nor does it affect any child workflows it starts. The default priority for an activity or workflow is set (either by you or by Amazon SWF) during registration, and the registered task priority is always used unless it is overridden while scheduling the activity or starting a workflow execution.

Task priority values can range from "-2147483648" to "2147483647", with higher numbers indicating higher priority. If you don't set the task priority for an activity or workflow, it will be assigned a priority of zero ("0").

Topics

- Setting Task Priority for Workflows
- Setting Task Priority for Activities
- Actions that Return Task Priority Information

Setting task priority API Version 2012-01-25 199

Setting Task Priority for Workflows

You can set the task priority for a workflow when you register it or start it. The task priority that is set when the workflow type is registered is used as the default for any workflow executions of that type, unless it is overridden when starting the workflow execution.

To register a workflow type with a default task priority, set the *defaultTaskPriority* option when using the RegisterWorkflowType action:

```
{
  "domain": "867530901",
  "name": "expeditedOrderWorkflow",
  "version": "1.0",
  "description": "Expedited customer orders workflow",
  "defaultTaskStartToCloseTimeout": "600",
  "defaultExecutionStartToCloseTimeout": "3600",
  "defaultTaskList": {"name": "mainTaskList"},
  "defaultTaskPriority": "10",
  "defaultChildPolicy": "TERMINATE"
}
```

You can override a workflow type's registered task priority when you start a workflow execution with StartWorkflowExecution:

```
{
  "childPolicy": "TERMINATE",
  "domain": "867530901",
  "executionStartToCloseTimeout": "1800",
  "input": "arbitrary-string-that-is-meaningful-to-the-workflow",
  "tagList": ["music purchase", "digital", "ricoh-the-dog"],
  "taskList": {"name": "specialTaskList"},
  "taskPriority": "-20",
  "taskStartToCloseTimeout": "600",
  "workflowId": "20110927-T-1",
  "workflowType": {"name": "customerOrderWorkflow", "version": "1.0"},
}
```

You can also override the registered task priority when starting a child workflow or when continuing a workflow as new, such as when responding to a decision with RespondDecisionTaskCompleted.

To set a child workflow's task priority, provide the value in startChildWorkflowExecutionDecisionAttributes:

```
"taskToken": "AAAAKgAAAAAAAAAAAAA...",
  "decisions": [
    {
      "decisionType": "StartChildWorkflowExecution",
      "startChildWorkflowExecutionDecisionAttributes": {
        "childPolicy": "TERMINATE",
        "control": "digital music",
        "executionStartToCloseTimeout": "900",
        "input": "201412-Smith-011x",
        "taskList": {"name": "specialTaskList"},
        "taskPriority": "5",
        "taskStartToCloseTimeout": "600",
        "workflowId": "verification-workflow",
        "workflowType": {
          "name": "MyChildWorkflow",
          "version": "1.0"
        }
      }
    }
  ]
}
```

When continuing a workflow as new, set the task priority in continueAsNewWorkflowExecutionDecisionAttributes:

```
}
]
}
```

Setting Task Priority for Activities

You can set the task priority for an activity either when registering it or when scheduling it. The task priority that is set when registering an activity type is used as the default priority when the activity is run, unless it is overridden when scheduling the activity.

To set task priority when registering an activity type, set the *defaultTaskPriority* option when using the RegisterActivityType action:

```
{
  "defaultTaskHeartbeatTimeout": "120",
  "defaultTaskList": {"name": "mainTaskList"},
  "defaultTaskPriority": "10",
  "defaultTaskScheduleToCloseTimeout": "900",
  "defaultTaskScheduleToStartTimeout": "300",
  "defaultTaskStartToCloseTimeout": "600",
  "description": "Verify the customer credit card",
  "domain": "867530901",
  "name": "activityVerify",
  "version": "1.0"
}
```

To schedule a task with a task priority, use the *taskPriority* option when scheduling the activity with the <u>RespondDecisionTaskCompleted</u> action:

```
"input": "abab-101",
    "taskList": {"name": "mainTaskList"},
    "taskPriority": "15"
    }
    }
}
```

Actions that Return Task Priority Information

You can get information about the set task priority (or set default task priority) from the following Amazon SWF actions:

- <u>DescribeActivityType</u> returns the *defaultTaskPriority* of the activity type in the configuration section of the response.
- <u>DescribeWorkflowExecution</u> returns the *taskPriority* of the workflow execution in the executionConfiguration section of the response.
- <u>DescribeWorkflowType</u> returns the *defaultTaskPriority* of the workflow type in the configuration section of the response.
- <u>GetWorkflowExecutionHistory</u> and <u>PollForDecisionTask</u> provide task priority information in the activityTaskScheduledEventAttributes, decisionTaskScheduledEventAttributes, workflowExecutionContinuedAsNewEventAttributes, and workflowExecutionStartedEventAttributes sections of the response.

Handling errors in Amazon SWF

There are a number of different types of errors that can occur during the course of a workflow execution.

Topics

- Validation Errors
- Errors in Enacting Actions or Decisions
- Timeouts
- · Errors raised by user code
- Errors related to closing a workflow execution

Validation Errors

Validation errors occur when a request to Amazon SWF fails because it isn't properly formed or it contains invalid data. In this context, a request could be an action such as DescribeDomain or it could be a decision such as StartTimer. If the request is an action, Amazon SWF returns an error code in the response. Check this error code as it may provide information about what aspect of the request caused the failure. For example, one or more of the arguments passed with the request might be invalid. For a list of common error codes, go to the topic for the action in the Amazon Simple Workflow Service API Reference.

If the request that failed is a decision, an appropriate event will be listed in the workflow execution history. For example, if the StartTimer decision failed, you would see a StartTimerFailed event in the history. The decider should check for these events when it receives the history in response to PollForDecisionTask or GetWorkflowExecutionHistory. Below is a list of possible decision failure events that can occur when the decision isn't correctly formed or contains invalid data.

Errors in Enacting Actions or Decisions

Even if the request is properly formed, errors may occur when Amazon SWF attempts to carry out the request. In these cases, one of the following events in the history will indicate that an error occurred. Look at the reason field of the event to determine the cause of failure.

- CancelTimerFailed
- RequestCancelActivityTaskFailed
- RequestCancelExternalWorkflowExecutionFailed
- ScheduleActivityTaskFailed
- <u>SignalExternalWorkflowExecutionFailed</u>
- StartChildWorkflowExecutionFailed
- <u>StartTimerFailed</u>

Timeouts

<u>Deciders</u>, <u>activity workers</u>, and <u>workflow executions</u> all operate within the constraints of timeout periods. In this type of error, a task or a child workflow times out. An event will appear in the history that describes the timeout. The decider should handle this event by, for example,

Validation Errors API Version 2012-01-25 204

rescheduling the task or restarting the child workflow. For more information about timeouts, see Amazon SWF Timeout Types

- ActivityTaskTimedOut
- ChildWorkflowExecutionTimedOut
- DecisionTaskTimedOut
- WorkflowExecutionTimedOut

Errors raised by user code

Examples of this type of error condition are activity task failures and child-workflow failures. As with timeout errors, Amazon SWF adds an appropriate event to the workflow execution history. The decider should handle this event, possibly by rescheduling the task or restarting the child workflow.

- ActivityTaskFailed
- ChildWorkflowExecutionFailed

Errors related to closing a workflow execution

Deciders may also see the following events if they attempt to close a workflow that has a pending decision task.

- FailWorkflowExecutionFailed
- <u>CompleteWorkFlowExecutionFailed</u>
- <u>ContinueAsNewWorkflowExecutionFailed</u>
- <u>CancelWorkflowExecutionFailed</u>

For more information about any of the events listed above, see <u>History Event</u> in the Amazon SWF API Reference.

Errors raised by user code API Version 2012-01-25 205

Amazon SWF Quotas

Amazon SWF places quotas on the sizes of certain workflow parameters, such as on the number of domains per account and on the size of the workflow execution history. These quotas are designed to prevent erroneous workflows from consuming all of the resources of the system, but are not hard limits. If you find that your application is frequently exceeding these quotas, you can <u>request a service quota increase</u>.

Contents

- General Account Quotas for Amazon SWF
- Quotas on Workflow Executions
- Quotas on Task Executions
- Amazon SWF throttling quotas
 - Throttling quotas for all Regions
 - Decision quotas for all Regions
 - Workflow-level quotas
- Requesting a quota increase

General Account Quotas for Amazon SWF

• Maximum registered domains - 100

This quota includes both registered and deprecated domains.

• Maximum workflow and activity types – 10,000 each per domain

This quota includes both registered and deprecated types.

- API call quota Beyond infrequent spikes, applications may be throttled if they make a large number of API calls in a very short period of time.
- Maximum request size 1 MB per request

This is the *total* data size per Amazon SWF API request, including the request header and all other associated request data.

• Truncated responses for *Count* APIs – Indicates that an internal quota was reached and that the response is not the full count.

Some queries will internally reach the 1 MB quota mentioned above before returning a full response. The following can return a truncated response instead of the full count.

- CountClosedWorkflowExecutions
- CountOpenWorkflowExecutions
- CountPendingActivityTasks
- CountPendingDecisionTasks

For each of these, if the truncated response is set to true, the count is less than the full amount. This internal quota can not be increased.

• Maximum number of tags – 50 tags per resource.

Attempting to add tags beyond 50 will result in a 400 error, TooManyTagsFault.

Quotas on Workflow Executions

• Maximum open workflow executions – 100,000 per domain

This count includes child workflow executions.

- Maximum workflow execution time 1 year. This is a hard quota that can't be changed.
- Maximum workflow execution history size 25,000 events. This is a hard quota that can't be changed.

Best practice is to structure each workflow such that its history does not grow beyond 10,000 events. Because the decider has to fetch the workflow history, a smaller history allows the decider to complete more quickly. If using the Flow Framework, you can use ContinueAsNew to continue a workflow with a fresh history.

• Maximum open child workflow executions – 1,000 per workflow execution

If your use case requires you to go beyond these quotas, you can use features Amazon SWF provides to continue executions and structure your applications using child workflow executions. If you find that you still need a quota increase, see Requesting a quota increase.

Quotas on Task Executions

• Maximum pollers per task list – 1,000 per task list

You can have a maximum of 1,000 pollers which simultaneously poll a particular task list. If you go over 1,000, you receive a LimitExceededException.



Note

While the maximum is 1,000, you might encounter LimitExceededException errors well before this quota. This error does not mean your tasks are being delayed. Instead, it means that you have the maximum amount of idle pollers on a task list. Amazon SWF sets this limit to save resources on both the client and server side. Setting the limit prevents an excessive number of pollers from waiting unnecessarily. You can reduce the LimitExceededException errors by using multiple task lists to distribute polling.

Maximum tasks scheduled per second – 2,000 per task list

You can schedule a maximum of 2,000 tasks per second on a particular task list. If you exceed 2,000, your ScheduleActivityTask decisions will fail with ACTIVITY CREATION RATE EXCEEDED error.



Note

While the maximum is 2,000, you might encounter ACTIVITY_CREATION_RATE_EXCEEDED errors well before this quota. To reduce these errors, use multiple task lists to distribute the load.

• Maximum task execution time – 1 year (constrained by workflow execution time maximum)

You can configure activity timeouts to cause a timeout event to occur if a particular stage of your activity task execution takes too long.

 Maximum time SWF will keep a task in the queue – 1 year (constrained by workflow execution) time quota)

You can configure default activity timeouts during activity registration that will cause a timeout event to occur if a particular stage of your activity task execution takes too long. You can also override the default activity timeouts when you schedule an activity task in your decider code.

• Maximum open activity tasks – 1,000 per workflow execution.

This quota includes both activity tasks that have been scheduled and those being processed by workers.

Quotas on Task Executions API Version 2012-01-25 208

- Maximum open timers 1,000 per workflow execution
- Maximum input/result data size 32,768 characters

This quota affects activity or workflow execution result data, input data when scheduling activity tasks or workflow executions, and input sent with a workflow execution signal.

Maximum decisions in a decision task response – varies

Due to the 1 MB quota on the <u>maximum API request size</u>, the number of decisions returned in a single call to <u>RespondDecisionTaskCompleted</u> will be limited according to the size of the data used by each decision, including the size of any input data provided to scheduled activity tasks or to workflow executions.

Amazon SWF throttling quotas

In addition to the service quotas described previously, certain Amazon SWF API calls and decision events are throttled to maintain service bandwidth, using a <u>token bucket</u> scheme. If your rate of requests consistently exceeds the rates that are listed here, you can <u>request a throttle quota</u> increase.

The throttling and decision quotas are same across all regions.

Throttling quotas for all Regions

The following quotas are applicable at individual account-levels. You can also request an increase to the following quotas. For information about doing this, see Requesting a quota increase.

API name	Bucket size	Refill rate per second
CountClosedWorkflowExecutions	2000	6
CountOpenWorkflowExecutions	2000	6
CountPendingActivityTasks	200	6
CountPendingDecisionTasks	200	6
DeleteActivityType	200	6

API name	Bucket size	Refill rate per second
DeleteWorkflowType	200	6
DeprecateActivityType	200	6
DeprecateDomain	100	6
DeprecateWorkflowType	200	6
DescribeActivityType	2000	6
DescribeDomain	200	6
DescribeWorkflowExecution	2000	6
DescribeWorkflowType	2000	6
GetWorkflowExecutionHistory	2000	60
ListActivityTypes	200	6
ListClosedWorkflowExecutions	200	6
ListDomains	100	6
ListOpenWorkflowExecutions	200	48
ListTagsForResource	50	30
ListWorkflowTypes	200	6
PollForActivityTask	2000	200
PollForDecisionTask	2000	200
RecordActivityTaskHeartbeat	2000	160
RegisterActivityType	200	60
RegisterDomain	100	6

API name	Bucket size	Refill rate per second
RegisterWorkflowType	200	60
RequestCancelWorkflowExecution	2000	30
RespondActivityTaskCanceled	2000	200
RespondActivityTaskCompleted	2000	200
RespondActivityTaskFailed	2000	200
RespondDecisionTaskCompleted	2000	200
SignalWorkflowExecution	2000	30
StartWorkflowExecution	2000	200
TagResource	50	30
TerminateWorkflowExecution	2000	60
UndeprecateActivityType	200	6
UndeprecateDomain	100	6
UndeprecateWorkflowType	200	6
UntagResource	50	30

Decision quotas for all Regions

The following quotas are applicable at individual account-levels. You can also request an increase to the following quotas. For information about doing this, see Requesting a quota increase.

API name	Bucket size	Refill rate per second
RequestCancelExternalWorkfl owExecution	1200	120
ScheduleActivityTask	1000	200
SignalExternalWorkflowExecution	1200	120
StartChildWorkflowExecution	500	12
StartTimer	2000	200

Workflow-level quotas

The following quotas are applicable at workflow-levels and can't be increased.

API name	Bucket size	Refill rate per second
GetWorkflowExecuti onHistory	400	200
SignalWorkflowExec ution	1000	1000
RecordActivityTask Heartbeat	1000	1000
RequestCancelWorkf lowExecution	200	200

Requesting a quota increase

For more information, see Amazon service quotas in the Amazon Web Services General Reference.

Workflow-level quotas API Version 2012-01-25 212

Additional resources and reference info for Amazon SWF

This chapter provides additional resources and reference information that is useful when developing workflows with Amazon SWF.

Topics

- Amazon SWF Timeout Types
- Amazon Simple Workflow Service Endpoints
- Additional Documentation for the Amazon Simple Workflow Service
- Web Resources for the Amazon Simple Workflow Service
- Migration options for Ruby Flow

Amazon SWF Timeout Types

To ensure that workflow executions run correctly, you can set different types of timeouts with Amazon SWF. Some timeouts specify how long the workflow can run in its entirety. Other timeouts specify how long activity tasks can take before being assigned to a worker and how long they can take to complete from the time they are scheduled. All timeouts in the Amazon SWF API are specified in seconds. Amazon SWF also supports the string NONE as a timeout value, which indicates no timeout.


For timeouts related to decision tasks and activity tasks, Amazon SWF adds an event to the workflow execution history. The attributes of the event provide information about what type of timeout occurred and which decision task or activity task was affected. Amazon SWF also schedules a decision task. When the decider receives the new decision task, it will see the timeout event in the history and take an appropriate action by calling the RespondDecisionTaskCompleted action.

A task is considered open from the time that it is scheduled until it is closed. Therefore a task is reported as open while a worker is processing it. A task is closed when a worker reports it as <u>completed</u>, <u>canceled</u>, or <u>failed</u>. A task may also be closed by Amazon SWF as the result of a timeout.

Timeouts in Workflow and Decision Tasks

The following diagram shows how workflow and decision timeouts are related to the lifetime of a workflow:

Timeout Types API Version 2012-01-25 213




There are two timeout types that are relevant to workflow and decision tasks:

- Workflow Start to Close (timeoutType: START_TO_CLOSE) This timeout specifies the maximum time that a workflow execution can take to complete. It is set as a default during workflow registration, but it can be overridden with a different value when the workflow is started. If this timeout is exceeded, Amazon SWF closes the workflow execution and adds an event of type WorkflowExecutionTimedOut to the workflow execution history. In addition to the timeoutType, the event attributes specify the childPolicy that is in effect for this workflow execution. The child policy specifies how child workflow executions are handled if the parent workflow execution times out or otherwise terminates. For example, if the childPolicy is set to TERMINATE, then child workflow executions will be terminated. Once a workflow execution has timed out, you can't take any action on it other than visibility calls.
- Decision Task Start to Close (timeoutType: START_TO_CLOSE) This timeout specifies the maximum time that the corresponding decider can take to complete a decision task. It is set during workflow type registration. If this timeout is exceeded, the task is marked as timed out in the workflow execution history, and Amazon SWF adds an event of type DecisionTaskTimedOut to the workflow history. The event attributes will include the IDs for the events that correspond to when this decision task was scheduled (scheduledEventId) and when it was started (startedEventId). In addition to adding the event, Amazon SWF also schedules a new decision task to alert the decider that this decision task timed out. After this timeout occurs, an attempt to complete the timed-out decision task using RespondDecisionTaskCompleted will fail.

Timeouts in Activity Tasks

The following diagram shows how timeouts are related to the lifetime of an activity task:

Timeouts in Activity Tasks API Version 2012-01-25 214



There are four timeout types that are relevant to activity tasks:

- Activity Task Start to Close (timeoutType: START_TO_CLOSE) This timeout specifies the maximum time that an activity worker can take to process a task after the worker has received the task. Attempts to close a timed out activity task using RespondActivityTaskCanceled, RespondActivityTaskCompleted, and RespondActivityTaskFailed will fail.
- Activity Task Heartbeat (timeoutType: HEARTBEAT) This timeout specifies the maximum time that a task can run before providing its progress through the RecordActivityTaskHeartbeat action.
- Activity Task Schedule to Start (timeoutType: SCHEDULE_TO_START) This timeout specifies how long Amazon SWF waits before timing out the activity task if no workers are available to perform the task. Once timed out, the expired task will not be assigned to another worker.
- Activity Task Schedule to Close (timeoutType: SCHEDULE_TO_CLOSE) This timeout specifies how long the task can take from the time it is scheduled to the time it is complete. As a best practice, this value should not be greater than the sum of the task schedule-to-start timeout and the task start-to-close timeout.



(i) Note

Each of the timeout types has a default value, which is generally set to NONE (infinite). The maximum time for any activity execution is limited to one year, however.

Timeouts in Activity Tasks API Version 2012-01-25 215 You set default values for these during activity type registration, but you can override them with new values when you <u>schedule</u> the activity task. When one of these timeouts occurs, Amazon SWF will add an <u>event</u> of type <u>ActivityTaskTimedOut</u> to the workflow history. The timeoutType value attribute of this event will specify which of these timeouts occurred. For each of the timeouts, the value of timeoutType is shown in parentheses. The event attributes will also include the IDs for the events that correspond to when the activity task was scheduled (scheduledEventId) and when it was started (startedEventId). In addition to adding the event, Amazon SWF also schedules a new decision task to alert the decider that the timeout occurred.

Amazon Simple Workflow Service Endpoints

A list of the current <u>Amazon SWF Regions and Endpoints</u> are provided in the *Amazon Web Services General Reference*, along with the endpoints for other services.

Amazon SWF domains and all related workflows and activities must exist within the same region to communicate with each other. Furthermore, any registered domains, workflows and activities within a region don't exist in other regions. For example, if you create a domain named "MySampleDomain" in both *us-east-1* and in *us-west-2*, they exist as *separate domains*: none of the workflows, task lists, activities, or data associated with your domains are shared across regions.

If you use other Amazon resources in your workflows, such as Amazon EC2 instances, these must also exist in the same region as your Amazon SWF resources. The only exceptions to this are services that span regions, such as Amazon S3 and IAM. You can access these services from workflows that exist in any region that supports them.

Additional Documentation for the Amazon Simple Workflow Service

In addition to this Developer Guide, you may find the following documentation useful.

Amazon Simple Workflow Service API Reference

The <u>Amazon Simple Workflow Service API Reference</u> provides detailed information about the Amazon SWF HTTP API, including actions, request and response structures and error codes.

Endpoints API Version 2012-01-25 216

Amazon Flow Framework Documentation

The Amazon Flow Framework is a programming framework that simplifies the process of implementing distributed asynchronous applications that use Amazon SWF to manage their workflows and activities, so you can focus on implementing your workflow logic.

Each Amazon Flow Framework is designed to work idiomatically in the language for which it is designed, so you can work naturally with your language of choice to implement workflows with all of the benefits of Amazon SWF.

There is an Amazon Flow Framework for Java. The Amazon Flow Framework for Java Developer Guide provides information about how to obtain, set up and use the Amazon Flow Framework for Java.

Amazon SDK Documentation

The Amazon Software Development Kits (SDKs) provide access to Amazon SWF in many different programming languages. The SDKs follow the HTTP API closely, but also provide language-specific programming interfaces for some Amazon SWF features. You can find out more information about each SDK by visiting the following links.



Note

Only SDKs that have support for Amazon SWF at the time of writing are listed here. For a full list of the available Amazon SDKs, visit the Tools for Amazon Web Services page.

Java

The Amazon SDK for Java provides a Java API for Amazon infrastructure services.

To view the available documentation, see the Amazon SDK for Java Documentation page. You can also go directly to the Amazon SWF sections in the SDK reference by following these links:

- Class: AmazonSimpleWorkflowClient
- Class: AmazonSimpleWorkflowAsyncClient
- Interface: AmazonSimpleWorkflow
- Interface: AmazonSimpleWorkflowAsync

JavaScript

The Amazon SDK for JavaScript allows developers to build libraries or applications that make use of Amazon services using a simple and easy-to-use API available both in the browser or inside of Node.js applications on the server.

To view the available documentation, see the <u>Amazon SDK for JavaScript Documentation</u> page. You can also go directly to the Amazon SWF section in the SDK reference by following this link:

Class: AWS.SimpleWorkflow

.NET

The Amazon SDK for .NET is a single, downloadable package that includes Visual Studio project templates, the Amazon .NET library, C# code samples, and documentation. The Amazon SDK for .NET makes it easier for Windows developers to build .NET applications for Amazon SWF and other services.

To view the available documentation, see the <u>Amazon SDK for .NET Documentation</u> page. You can also go directly to the Amazon SWF sections in the SDK reference by following these links:

- Namespace: Amazon.SimpleWorkflow
- Namespace: Amazon.SimpleWorkflow.Model

PHP

The Amazon SDK for PHP provides a PHP programming interface to Amazon SWF.

To view the available documentation, see the <u>Amazon SDK for PHP Documentation</u> page. You can also go directly to the Amazon SWF section in the SDK reference by following this link:

• Class: SwfClient

Python

The Amazon SDK for Python (Boto) provides a Python programming interface to Amazon SWF.

To view the available documentation, see the <u>boto: A Python interface to Amazon Web Services</u> page. You can also go directly to the Amazon SWF sections in the documentation by following these links:

- Amazon SWF Tutorial
- Amazon SWF Reference

Amazon SDK Documentation API Version 2012-01-25 218

Ruby

The Amazon SDK for Ruby provides a Ruby programming interface to Amazon SWF.

To view the available documentation, see the <u>Amazon SDK for Ruby Documentation</u> page. You can also go directly to the Amazon SWF section in the SDK reference by following this link:

Class: AWS::SimpleWorkflow

Amazon CLI Documentation

The Amazon Command Line Interface (Amazon CLI) is a unified tool to manage your Amazon services. With just one tool to download and configure, you can control multiple Amazon services from the command line and automate them through scripts.

For more information about the Amazon CLI, see the Amazon Command Line Interface page.

For an overview of the available commands for Amazon SWF, see <u>swf</u> in the *Amazon CLI Command Reference*.

Web Resources for the Amazon Simple Workflow Service

There are a number of Web resources that you can use to learn more about Amazon SWF or to get help with using the service and developing workflows.

Amazon SWF Forum

The Amazon SWF forum provides a place for you to communicate with other Amazon SWF developers and members of the Amazon SWF development team at Amazon to ask questions and to get answers.

You can visit the forum at: Forum: Amazon Simple Workflow Service.

Amazon SWF FAQ

The Amazon SWF FAQ provide answers to frequently-asked questions about Amazon SWF, including an overview of common use cases, differences between Amazon SWF and other services, and more.

Amazon CLI Documentation API Version 2012-01-25 219

You can access the FAQ here: Amazon SWF FAQ.

Amazon SWF Videos

The <u>Amazon Web Services</u> channel on YouTube provides video training for all of Amazon's Web Services, including Amazon SWF. For a full list of Amazon SWF-related videos, use the following query: <u>Simple Workflow</u> in Amazon Web Services

Migration options for Ruby Flow

The Amazon Flow Framework for Ruby is no longer under active development. While existing code will continue to work indefinitely, there will be no new features or versions. This topic will cover usage and migration options to continue working with Amazon SWF, and information on how to migrate to Step Functions.

Option	Description
Continue to use the Ruby Flow Framework	For now, the Ruby Flow Framework will continue to work. If you do nothing, your code will continue to function as it is. Plan to migrate off of the Amazon Flow Framework for Ruby in the near future.
Migrate to the Java Flow Framework	The Java Flow Framework remains in active development and will continue to receive new features and updates.
Migrate to Step Functions	Step Functions provides a way to coordinate the components of distributed applications using visual workflows controlled by a state machine.
Use the SWF API directly, without the Flow Framework	You can continue to work in Ruby and use the SWF API directly instead of the Ruby Flow Framework.

The advantage the Flow Framework provides, for either Ruby or Java, is that it allows you to focus on your workflow logic. The framework handles much of the details of communication and coordination, and some of the complexity is abstracted. You can continue to have the same level of abstraction by migrating to the Java Flow Framework, or you can directly interact with Amazon SWF SDK directly.

Amazon SWF Videos API Version 2012-01-25 220

Continue to use the Ruby Flow Framework

The Amazon Flow Framework for Ruby will continue to function as it does now in the short term. If you have workflows written in the Amazon Flow Framework for Ruby, these will continue to work. Without updates, support, or security fixes, it is best to have a firm plan to migrate off of the Amazon Flow Framework for Ruby in the near future.

Migrate to the Java Flow Framework

The Amazon Flow Framework for Java will remain in active development. Conceptually, the Amazon Flow Framework for Java is similar to Amazon Flow Framework for Ruby: you can still focus on your workflow logic, and the framework will help manage your decider logic, and will make other aspects of Amazon SWF easier to manage.

- Amazon Flow Framework for Java
- Amazon Flow Framework for Java API Reference

Migrate to Step Functions

Amazon Step Functions provides a service that is similar to Amazon SWF, but where your workflow logic is controlled by a state machine. Step Functions enables you to coordinate the components of distributed applications and microservices using visual workflows. You build applications from individual components that each perform a discrete function, or *task*, allowing you to scale and change applications quickly. Step Functions provides a reliable way to coordinate components and step through the functions of your application. A graphical console provides a way to visualize the components of your application as a series of steps. It automatically triggers and tracks each step, and retries when there are errors, so your application executes in order and as expected, every time. Step Functions logs the state of each step, so when things do go wrong, you can diagnose and debug problems quickly.

In Step Functions, you manage the coordination of your tasks using a state machine, written in declarative JSON, that is defined using the <u>Amazon States Language</u>. By using a state machine, you don't have to write and maintain a decider program to control your application logic. Step Functions provides an intuitive, productive, and agile approach to coordinating application components using visual workflows. You should consider using Amazon Step Functions for all your new applications, and Step Functions provides an excellent platform to migrate to for the workflows you currently have implemented in the Amazon Flow Framework for Ruby.

To help migrate your tasks to Step Functions, while continuing to leverage your Ruby language skills, Step Functions provides an example Ruby activity worker. This example uses best practices for implementing an activity worker, and can be used as a template to migrate your task logic to Step Functions. For more information, see the *Example Activity Worker in Ruby* topic in the Amazon Step Functions Developer Guide.

Note

For many customers, migrating to Step Functions from the Amazon Flow Framework for Ruby is the best option. But, if you require that signals intervene in your processes, or if you need to launch child processes that return a result to a parent, consider using the Amazon SWF API directly, or migrating to the Amazon Flow Framework for Java.

For more information on Amazon Step Functions, see:

- Amazon Step Functions Developer Guide
- Amazon Step Functions API Reference
- Amazon Step Functions Command Line Reference

Use the Amazon SWF API directly

While the Amazon Flow Framework for Ruby manages some of the complexity of Amazon SWF, you can also use the Amazon SWF API directly. Using the API directly allows you to build workflows where you have full control over implementing tasks and coordinating them, without worrying about underlying complexities such as tracking their progress and maintaining their state.

- Amazon Simple Workflow Service Developer Guide
- Amazon Simple Workflow Service API Reference

Document history

The following table describes the important changes to the documentation since the last release of the *Amazon Simple Workflow Service Developer Guide*.

Change	Description	Date Changed
Documentation- only update	Amazon SWF now includes a section about Amazon User Notifications, an Amazon Web Services service that acts as a central location for your Amazon notifications in the Amazon Web Services Management Console. For more information, see Using Amazon User Notifications with Amazon Simple Workflow Service .	May 4, 2023
New feature	Amazon SWF now includes Amazon PrivateLink support in the China Region. For more information, see Amazon VPC endpoints for Amazon SWF.	March 31, 2023
Update	Updated the <u>Quotas on Task Executions</u> section to include Maximum tasks scheduled per second, and the <u>Amazon SWF Metrics for CloudWatch</u> page to include information about <u>using non-ASCII resource names</u> with CloudWatch.	May 12, 2021
New feature	Amazon Simple Workflow Service now supports Amazon EventBridge. For more information, see: • EventBridge for Amazon SWF • EventBridge User Guide	December 18, 2020
New feature	Amazon Simple Workflow Service supports IAM permissio ns using tags. For more information, see the following. • <u>Tags in Amazon SWF</u> • <u>Manage tags</u> • <u>Tag workflow executions</u>	June 20, 2019

Change	Description	Date Changed
	 Control access to domains with tags TagResource UntagResource ListTagsForResource RegisterDomain 	
New feature	Amazon Simple Workflow Service is now available the Europe (Stockholm) region.	December 12, 2018
Update	Improved the Amazon Simple Workflow Service topic on CloudTrail integration. See Recording API calls with Amazon CloudTrail.	August 7, 2018
Update	Added information on the new PendingTasks metric for CloudWatch. For more information, see Amazon SWF Metrics .	June 18, 2018
Update	Improved syntax highlighting in code samples.	March 29, 2018
Update	Added a topic describing options for Ruby Flow users to migrate off of that platform. For more information, see <u>Migration options for Ruby Flow</u> .	March 9, 2018
Update	Improved navigation on advanced concepts topic. See Advanced workflow concepts in Amazon SWF.	February 19, 2018
Update	Improved CloudWatch metrics documentation by adding valid statistics information. See Amazon SWF Metrics for CloudWatch.	December 4, 2017
Update	Changed the TOC to improve the document structure . Added new information on <u>API and Decision Event Metrics</u> .	November 9, 2017

Change	Description	Date Changed
Update	Updated the <u>Amazon SWF Quotas</u> section to include throttling limits for all regions.	October 18, 2017
Update	Changed task_list to workflowId in the <u>Getting</u> <u>started with Amazon SWF</u> to avoid confusion with activity_list .	July 25, 2017
Update	Cleaned up the code examples throughout this guide.	June 5, 2017
Update	Simplified and improved the organization and contents of this guide.	May 19, 2017
Update	Updates and link fixes.	May 16, 2017
Update	Updates and link fixes.	October 1, 2016
Lambda task support	You can specify Lambda tasks in addition to traditional Activity tasks in your workflows. For more information, see Amazon Lambda tasks in Amazon SWF .	July 21, 2015
Support for setting task priority	Amazon SWF now includes support for setting the priority of tasks on a task list, and will attempt to deliver those with higher priority before tasks with lower priority. Information about how to set the task priority for workflows and for activities is provided in Setting task priority in Amazon SWF.	December 17, 2014
Update	Added a new topic that describes how to log Amazon SWF API calls using CloudTrail: Recording API calls with Amazon CloudTrail.	May 8, 2014

Change	Description	Date Changed
Update	Two new topics related to CloudWatch metrics for Amazon SWF have been added: Amazon SWF Metrics for CloudWatch, which provides a list and descriptions of the supported metrics, and Viewing Amazon SWF Metrics for CloudWatch using the Amazon Web Services Managemen tonsole, which provides information about how to view metrics and set alarms with the Amazon Web Services Management Console.	April 28, 2014
Update	Added a new section: Additional resources and reference info for Amazon SWF. This section provides some service reference information and provides information about additional documentation, samples, code and other web resources for Amazon SWF developers.	March 19, 2014
Update	Added a workflow tutorial. See <u>Getting started with</u> <u>Amazon SWF</u> .	October 25, 2013
Update	Added Amazon CLI information and example.	August 26, 2013
Update	Updates and fixes.	August 1, 2013
Update	Updated the document to describe how to use IAM for access control.	February 22, 2013
Initial Release	This is the first release of the <i>Amazon Simple Workflow Service Developer Guide</i> .	October 16, 2012