



EMR Serverless API Reference

Amazon EMR Serverless



API Version 2021-07-13

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon EMR Serverless: EMR Serverless API Reference

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

Table of Contents

Welcome	1
Actions	2
CancelJobRun	3
Request Syntax	3
URI Request Parameters	3
Request Body	3
Response Syntax	4
Response Elements	4
Errors	4
See Also	5
CreateApplication	6
Request Syntax	6
URI Request Parameters	8
Request Body	8
Response Syntax	13
Response Elements	13
Errors	14
See Also	15
DeleteApplication	16
Request Syntax	16
URI Request Parameters	16
Request Body	16
Response Syntax	16
Response Elements	16
Errors	16
See Also	17
GetApplication	18
Request Syntax	18
URI Request Parameters	18
Request Body	18
Response Syntax	18
Response Elements	21
Errors	21
See Also	22

GetDashboardForJobRun	23
Request Syntax	23
URI Request Parameters	23
Request Body	24
Response Syntax	24
Response Elements	24
Errors	24
See Also	25
GetJobRun	26
Request Syntax	26
URI Request Parameters	26
Request Body	26
Response Syntax	27
Response Elements	29
Errors	29
See Also	30
GetResourceDashboard	31
Request Syntax	31
URI Request Parameters	31
Request Body	32
Response Syntax	32
Response Elements	32
Errors	32
See Also	33
GetSession	34
Request Syntax	34
URI Request Parameters	34
Request Body	34
Response Syntax	34
Response Elements	36
Errors	36
See Also	36
GetSessionEndpoint	38
Request Syntax	38
URI Request Parameters	38
Request Body	38

Response Syntax	38
Response Elements	39
Errors	40
See Also	40
ListApplications	42
Request Syntax	42
URI Request Parameters	42
Request Body	42
Response Syntax	43
Response Elements	43
Errors	44
See Also	44
ListJobRunAttempts	45
Request Syntax	45
URI Request Parameters	45
Request Body	46
Response Syntax	46
Response Elements	46
Errors	47
See Also	47
ListJobRuns	49
Request Syntax	49
URI Request Parameters	49
Request Body	50
Response Syntax	50
Response Elements	51
Errors	51
See Also	52
ListSessions	53
Request Syntax	53
URI Request Parameters	53
Request Body	54
Response Syntax	54
Response Elements	54
Errors	55
See Also	55

ListTagsForResource	57
Request Syntax	57
URI Request Parameters	57
Request Body	57
Response Syntax	57
Response Elements	58
Errors	58
See Also	59
StartApplication	60
Request Syntax	60
URI Request Parameters	60
Request Body	60
Response Syntax	60
Response Elements	60
Errors	60
See Also	61
StartJobRun	62
Request Syntax	62
URI Request Parameters	63
Request Body	64
Response Syntax	66
Response Elements	66
Errors	67
See Also	68
StartSession	69
Request Syntax	69
URI Request Parameters	69
Request Body	70
Response Syntax	71
Response Elements	72
Errors	73
See Also	73
StopApplication	75
Request Syntax	75
URI Request Parameters	75
Request Body	75

Response Syntax	75
Response Elements	75
Errors	75
See Also	76
TagResource	77
Request Syntax	77
URI Request Parameters	77
Request Body	77
Response Syntax	78
Response Elements	78
Errors	78
See Also	79
TerminateSession	80
Request Syntax	80
URI Request Parameters	80
Request Body	80
Response Syntax	80
Response Elements	81
Errors	81
See Also	82
UntagResource	83
Request Syntax	83
URI Request Parameters	83
Request Body	83
Response Syntax	84
Response Elements	84
Errors	84
See Also	84
UpdateApplication	86
Request Syntax	86
URI Request Parameters	88
Request Body	88
Response Syntax	92
Response Elements	95
Errors	95
See Also	96

Data Types	97
Application	99
Contents	99
See Also	104
ApplicationSummary	105
Contents	105
See Also	107
AutoStartConfig	108
Contents	108
See Also	108
AutoStopConfig	109
Contents	109
See Also	109
CloudWatchLoggingConfiguration	110
Contents	110
See Also	111
Configuration	113
Contents	113
See Also	114
ConfigurationOverrides	115
Contents	115
See Also	115
DiskEncryptionConfiguration	116
Contents	116
See Also	116
Hive	118
Contents	118
See Also	119
IdentityCenterConfiguration	120
Contents	120
See Also	121
IdentityCenterConfigurationInput	122
Contents	122
See Also	122
ImageConfiguration	123
Contents	123

See Also	124
ImageConfigurationInput	125
Contents	125
See Also	125
InitialCapacityConfig	127
Contents	127
See Also	127
InteractiveConfiguration	128
Contents	128
See Also	128
JobDriver	129
Contents	129
See Also	129
JobLevelCostAllocationConfiguration	130
Contents	130
See Also	130
JobRun	131
Contents	131
See Also	137
JobRunAttemptSummary	139
Contents	139
See Also	142
JobRunExecutionIamPolicy	143
Contents	143
See Also	143
JobRunSummary	145
Contents	145
See Also	148
ManagedPersistenceMonitoringConfiguration	149
Contents	149
See Also	149
MaximumAllowedResources	150
Contents	150
See Also	151
MonitoringConfiguration	152
Contents	152

See Also	152
NetworkConfiguration	154
Contents	154
See Also	154
PrometheusMonitoringConfiguration	156
Contents	156
See Also	156
ResourceUtilization	157
Contents	157
See Also	157
RetryPolicy	159
Contents	159
See Also	159
S3MonitoringConfiguration	160
Contents	160
See Also	160
SchedulerConfiguration	161
Contents	161
See Also	161
Session	162
Contents	162
See Also	166
SessionConfigurationOverrides	168
Contents	168
See Also	168
SessionSummary	169
Contents	169
See Also	171
SparkSubmit	173
Contents	173
See Also	174
TotalResourceUtilization	175
Contents	175
See Also	175
WorkerResourceConfig	177
Contents	177

See Also	178
WorkerTypeSpecification	179
Contents	179
See Also	179
WorkerTypeSpecificationInput	180
Contents	180
See Also	180
Common Parameters	181
Common Error Types	184

Welcome

Amazon EMR Serverless is a new deployment option for Amazon EMR. Amazon EMR Serverless provides a serverless runtime environment that simplifies running analytics applications using the latest open source frameworks such as Apache Spark and Apache Hive. With Amazon EMR Serverless, you don't have to configure, optimize, secure, or operate clusters to run applications with these frameworks.

The API reference to Amazon EMR Serverless is `emr-serverless`. The `emr-serverless` prefix is used in the following scenarios:

- It is the prefix in the CLI commands for Amazon EMR Serverless. For example, `aws emr-serverless start-job-run`.
- It is the prefix before IAM policy actions for Amazon EMR Serverless. For example, "Action": `["emr-serverless:StartJobRun"]`. For more information, see [Policy actions for Amazon EMR Serverless](#).
- It is the prefix used in Amazon EMR Serverless service endpoints. For example, `emr-serverless.us-east-2.amazonaws.com`.

This document was last published on September 21, 2026.

Actions

The following actions are supported:

- [CancelJobRun](#)
- [CreateApplication](#)
- [DeleteApplication](#)
- [GetApplication](#)
- [GetDashboardForJobRun](#)
- [GetJobRun](#)
- [GetResourceDashboard](#)
- [GetSession](#)
- [GetSessionEndpoint](#)
- [ListApplications](#)
- [ListJobRunAttempts](#)
- [ListJobRuns](#)
- [ListSessions](#)
- [ListTagsForResource](#)
- [StartApplication](#)
- [StartJobRun](#)
- [StartSession](#)
- [StopApplication](#)
- [TagResource](#)
- [TerminateSession](#)
- [UntagResource](#)
- [UpdateApplication](#)

CancelJobRun

Cancels a job run.

Request Syntax

```
DELETE /applications/applicationId/jobruns/jobRunId?  
shutdownGracePeriodInSeconds=shutdownGracePeriodInSeconds HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application on which the job run will be canceled.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

jobRunId

The ID of the job run to cancel.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

shutdownGracePeriodInSeconds

The duration in seconds to wait before forcefully terminating the job after cancellation is requested.

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "applicationId": "string",
  "jobRunId": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[applicationId](#)

The output contains the application ID on which the job run is cancelled.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

[jobRunId](#)

The output contains the ID of the cancelled job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

CreateApplication

Creates an application.

Request Syntax

POST /applications HTTP/1.1
Content-type: application/json

```
{
  "architecture": "string",
  "autoStartConfiguration": {
    "enabled": boolean
  },
  "autoStopConfiguration": {
    "enabled": boolean,
    "idleTimeoutMinutes": number
  },
  "clientToken": "string",
  "diskEncryptionConfiguration": {
    "encryptionContext": {
      "string" : "string"
    },
    "encryptionKeyArn": "string"
  },
  "identityCenterConfiguration": {
    "identityCenterInstanceArn": "string",
    "userBackgroundSessionsEnabled": boolean
  },
  "imageConfiguration": {
    "applicationLevelDigestResolution": boolean,
    "imageUri": "string"
  },
  "initialCapacity": {
    "string" : {
      "workerConfiguration": {
        "cpu": "string",
        "disk": "string",
        "diskType": "string",
        "memory": "string"
      },
      "workerCount": number
    }
  }
}
```

```

},
"interactiveConfiguration": {
  "livyEndpointEnabled": boolean,
  "sessionEnabled": boolean,
  "studioEnabled": boolean
},
"jobLevelCostAllocationConfiguration": {
  "enabled": boolean
},
"maximumCapacity": {
  "cpu": "string",
  "disk": "string",
  "memory": "string"
},
"monitoringConfiguration": {
  "cloudWatchLoggingConfiguration": {
    "enabled": boolean,
    "encryptionKeyArn": "string",
    "logGroupName": "string",
    "logStreamNamePrefix": "string",
    "logTypes": {
      "string" : [ "string" ]
    }
  },
  "managedPersistenceMonitoringConfiguration": {
    "enabled": boolean,
    "encryptionKeyArn": "string"
  },
  "prometheusMonitoringConfiguration": {
    "remoteWriteUrl": "string"
  },
  "s3MonitoringConfiguration": {
    "encryptionKeyArn": "string",
    "logUri": "string"
  }
},
"name": "string",
"networkConfiguration": {
  "securityGroupIds": [ "string" ],
  "subnetIds": [ "string" ]
},
"releaseLabel": "string",
"runtimeConfiguration": [
  {

```

```
{
  "classification": "string",
  "configurations": [
    "Configuration"
  ],
  "properties": {
    "string" : "string"
  }
},
"schedulerConfiguration": {
  "maxConcurrentRuns": number,
  "queueTimeoutMinutes": number
},
"tags": {
  "string" : "string"
},
"type": "string",
"workerTypeSpecifications": {
  "string" : {
    "imageConfiguration": {
      "applicationLevelDigestResolution": boolean,
      "imageUri": "string"
    }
  }
}
}
```

URI Request Parameters

The request does not use any URI parameters.

Request Body

The request accepts the following data in JSON format.

architecture

The CPU architecture of an application.

Type: String

Valid Values: ARM64 | X86_64

Required: No

autoStartConfiguration

The configuration for an application to automatically start on job submission.

Type: [AutoStartConfig](#) object

Required: No

autoStopConfiguration

The configuration for an application to automatically stop after a certain amount of time being idle.

Type: [AutoStopConfig](#) object

Required: No

clientToken

The client idempotency token of the application to create. Its value must be unique for each request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [A-Za-z0-9._-]+

Required: Yes

diskEncryptionConfiguration

The configuration object that allows encrypting local disks.

Type: [DiskEncryptionConfiguration](#) object

Required: No

identityCenterConfiguration

The IAM Identity Center Configuration accepts the Identity Center instance parameter required to enable trusted identity propagation. This configuration allows identity propagation between integrated services and the Identity Center instance.

Type: [IdentityCenterConfigurationInput](#) object

Required: No

imageConfiguration

The image configuration for all worker types. You can either set this parameter or `imageConfiguration` for each worker type in `workerTypeSpecifications`.

Type: [ImageConfigurationInput](#) object

Required: No

initialCapacity

The capacity to initialize when the application is created.

Type: String to [InitialCapacityConfig](#) object map

Map Entries: Minimum number of 0 items. Maximum number of 10 items.

Key Length Constraints: Minimum length of 1. Maximum length of 50.

Key Pattern: `[a-zA-Z]+[-_]*[a-zA-Z]+`

Required: No

interactiveConfiguration

The interactive configuration object that enables the interactive use cases to use when running an application.

Type: [InteractiveConfiguration](#) object

Required: No

jobLevelCostAllocationConfiguration

The configuration object that enables job level cost allocation.

Type: [JobLevelCostAllocationConfiguration](#) object

Required: No

maximumCapacity

The maximum capacity to allocate when the application is created. This is cumulative across all workers at any given point in time, not just when an application is created. No new resources will be created once any one of the defined limits is hit.

Type: [MaximumAllowedResources](#) object

Required: No

monitoringConfiguration

The configuration setting for monitoring.

Type: [MonitoringConfiguration](#) object

Required: No

name

The name of the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [A-Za-z0-9._/#-]+

Required: No

networkConfiguration

The network configuration for customer VPC connectivity.

Type: [NetworkConfiguration](#) object

Required: No

releaseLabel

The Amazon EMR release associated with the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [A-Za-z0-9._/-]+

Required: Yes

runtimeConfiguration

The [Configuration](#) specifications to use when creating an application. Each configuration consists of a classification and properties. This configuration is applied to all the job runs submitted under the application.

Type: Array of [Configuration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Required: No

[schedulerConfiguration](#)

The scheduler configuration for batch and streaming jobs running on this application. Supported with release labels emr-7.0.0 and above.

Type: [SchedulerConfiguration](#) object

Required: No

[tags](#)

The tags assigned to the application.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 200 items.

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: [A-Za-z0-9 /_ . : = + @ -] +

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: [A-Za-z0-9 /_ . : = + @ -] *

Required: No

[type](#)

The type of application you want to start, such as Spark or Hive.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: Yes

[workerTypeSpecifications](#)

The key-value pairs that specify worker type to `WorkerTypeSpecificationInput`. This parameter must contain all valid worker types for a Spark or Hive application. Valid worker

types include Driver and Executor for Spark applications and HiveDriver and TezTask for Hive applications. You can either set image details in this parameter for each worker type, or in `imageConfiguration` for all worker types.

Type: String to [WorkerTypeSpecificationInput](#) object map

Key Length Constraints: Minimum length of 1. Maximum length of 50.

Key Pattern: `[a-zA-Z]+[-_]*[a-zA-Z]+`

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "applicationId": "string",
  "arn": "string",
  "name": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[applicationId](#)

The output contains the application ID.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

[arn](#)

The output contains the ARN of the application.

Type: String

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.*:(\d{12}):\/applications\/[0-9a-zA-Z-]+`

name

The output contains the name of the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[A-Za-z0-9._/#-]+`

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

ConflictException

The request could not be processed because of conflict in the current state of the resource.

HTTP Status Code: 409

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

DeleteApplication

Deletes an application. An application has to be in a stopped or created state in order to be deleted.

Request Syntax

```
DELETE /applications/applicationId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application that will be deleted.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetApplication

Displays detailed information about a specified application.

Request Syntax

```
GET /applications/applicationId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application that will be described.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "application": {
    "applicationId": string,
    "architecture": string,
    "arn": string,
    "autoStartConfiguration": {
      "enabled": boolean
    },
    "autoStopConfiguration": {
      "enabled": boolean,
      "idleTimeoutMinutes": number
    }
  }
}
```

```

},
"createdAt": number,
"diskEncryptionConfiguration": {
  "encryptionContext": {
    "string" : "string"
  },
  "encryptionKeyArn": "string"
},
"identityCenterConfiguration": {
  "identityCenterApplicationArn": "string",
  "identityCenterInstanceArn": "string",
  "userBackgroundSessionsEnabled": boolean
},
"imageConfiguration": {
  "applicationLevelDigestResolution": boolean,
  "imageUri": "string",
  "resolvedImageDigest": "string"
},
"initialCapacity": {
  "string" : {
    "workerConfiguration": {
      "cpu": "string",
      "disk": "string",
      "diskType": "string",
      "memory": "string"
    },
    "workerCount": number
  }
},
"interactiveConfiguration": {
  "livyEndpointEnabled": boolean,
  "sessionEnabled": boolean,
  "studioEnabled": boolean
},
"jobLevelCostAllocationConfiguration": {
  "enabled": boolean
},
"maximumCapacity": {
  "cpu": "string",
  "disk": "string",
  "memory": "string"
},
"monitoringConfiguration": {
  "cloudWatchLoggingConfiguration": {

```

```

    "enabled": boolean,
    "encryptionKeyArn": "string",
    "logGroupName": "string",
    "logStreamNamePrefix": "string",
    "logTypes": {
        "string" : [ "string" ]
    }
},
"managedPersistenceMonitoringConfiguration": {
    "enabled": boolean,
    "encryptionKeyArn": "string"
},
"prometheusMonitoringConfiguration": {
    "remoteWriteUrl": "string"
},
"s3MonitoringConfiguration": {
    "encryptionKeyArn": "string",
    "logUri": "string"
}
},
"name": "string",
"networkConfiguration": {
    "securityGroupIds": [ "string" ],
    "subnetIds": [ "string" ]
},
"releaseLabel": "string",
"runtimeConfiguration": [
    {
        "classification": "string",
        "configurations": [
            "Configuration"
        ],
        "properties": {
            "string" : "string"
        }
    }
],
"schedulerConfiguration": {
    "maxConcurrentRuns": number,
    "queueTimeoutMinutes": number
},
"state": "string",
"stateDetails": "string",
"tags": {

```

```
    "string" : "string"
  },
  "type": "string",
  "updatedAt": number,
  "workerTypeSpecifications": {
    "string" : {
      "imageConfiguration": {
        "applicationLevelDigestResolution": boolean,
        "imageUri": "string",
        "resolvedImageDigest": "string"
      }
    }
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

application

The output displays information about the specified application.

Type: [Application](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetDashboardForJobRun

Creates and returns a URL that you can use to access the application UIs for a job run.

For jobs in a running state, the application UI is a live user interface such as the Spark or Tez web UI. For completed jobs, the application UI is a persistent application user interface such as the Spark History Server or persistent Tez UI.

Note

The URL is valid for one hour after you generate it. To access the application UI after that hour elapses, you must invoke the API again to generate a new URL.

Request Syntax

```
GET /applications/applicationId/jobruns/jobRunId/dashboard?  
accessSystemProfileLogs=accessSystemProfileLogs&attempt=attempt HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

accessSystemProfileLogs

Allows access to system profile logs for Lake Formation-enabled jobs. Default is false.

applicationId

The ID of the application.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

attempt

An optional parameter that indicates the amount of attempts for the job. If not specified, this value defaults to the attempt of the latest job.

Valid Range: Minimum value of 1.

jobRunId

The ID of the job run.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "url": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

url

The URL to view job run's dashboard.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetJobRun

Displays detailed information about a job run.

Request Syntax

```
GET /applications/applicationId/jobruns/jobRunId?attempt=attempt HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application on which the job run is submitted.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

attempt

An optional parameter that indicates the amount of attempts for the job. If not specified, this value defaults to the attempt of the latest job.

Valid Range: Minimum value of 1.

jobRunId

The ID of the job run.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "jobRun": {
    "applicationId": "string",
    "arn": "string",
    "attempt": number,
    "attemptCreatedAt": number,
    "attemptUpdatedAt": number,
    "billedResourceUtilization": {
      "memoryGBHour": number,
      "storageGBHour": number,
      "vCPUHour": number
    },
    "configurationOverrides": {
      "applicationConfiguration": [
        {
          "classification": "string",
          "configurations": [
            "Configuration"
          ],
          "properties": {
            "string" : "string"
          }
        }
      ],
      "diskEncryptionConfiguration": {
        "encryptionContext": {
          "string" : "string"
        },
        "encryptionKeyArn": "string"
      },
      "monitoringConfiguration": {
        "cloudWatchLoggingConfiguration": {
          "enabled": boolean,
          "encryptionKeyArn": "string",
          "logGroupName": "string",
          "logStreamNamePrefix": "string",
          "logTypes": {
            "string" : [ "string" ]
          }
        }
      }
    }
  }
}
```

```

    }
  },
  "managedPersistenceMonitoringConfiguration": {
    "enabled": boolean,
    "encryptionKeyArn": "string"
  },
  "prometheusMonitoringConfiguration": {
    "remoteWriteUrl": "string"
  },
  "s3MonitoringConfiguration": {
    "encryptionKeyArn": "string",
    "logUri": "string"
  }
}
},
"createdAt": number,
"createdBy": "string",
"endedAt": number,
"executionIamPolicy": {
  "policy": "string",
  "policyArns": [ "string" ]
},
"executionRole": "string",
"executionTimeoutMinutes": number,
"imageConfiguration": {
  "applicationLevelDigestResolution": boolean,
  "imageUri": "string",
  "resolvedImageDigest": "string"
},
"jobDriver": { ... },
"jobRunId": "string",
"mode": "string",
"name": "string",
"networkConfiguration": {
  "securityGroupIds": [ "string" ],
  "subnetIds": [ "string" ]
},
"queuedDurationMilliseconds": number,
"releaseLabel": "string",
"retryPolicy": {
  "maxAttempts": number,
  "maxFailedAttemptsPerHour": number
},
"startedAt": number,

```

```

    "state": "string",
    "stateDetails": "string",
    "tags": {
      "string" : "string"
    },
    "totalExecutionDurationSeconds": number,
    "totalResourceUtilization": {
      "memoryGBHour": number,
      "storageGBHour": number,
      "vCPUHour": number
    },
    "updatedAt": number,
    "workerTypeSpecifications": {
      "string" : {
        "imageConfiguration": {
          "applicationLevelDigestResolution": boolean,
          "imageUri": "string",
          "resolvedImageDigest": "string"
        }
      }
    }
  }
}

```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[jobRun](#)

The output displays information about the job run.

Type: [JobRun](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetResourceDashboard

Returns a URL that you can use to access the application UIs for a specified resource, such as a session.

For resources in a running state, the application UI is a live user interface such as the Spark web UI. For terminated resources, the application UI is a persistent application user interface such as the Spark History Server.

Note

The URL is valid for one hour after you generate it. To access the application UI after that hour elapses, you must invoke the API again to generate a new URL.

Request Syntax

```
GET /applications/applicationId/dashboard?  
resourceId=resourceId&resourceType=resourceType HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application that the resource belongs to.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

resourceId

The ID of the resource.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

resourceType

The type of resource to access the dashboard for. Currently, only `Session` is supported.

Valid Values: `SESSION`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "url": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

url

A URL to the resource dashboard. For an active resource, this URL opens the live application UI. For a terminated resource, this URL opens the persistent application UI. This value is not included in the response if the URL is not available.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetSession

Displays detailed information about a session.

Request Syntax

```
GET /applications/applicationId/sessions/sessionId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application that the session belongs to.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

Required: Yes

sessionId

The ID of the session.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "session": {
```

```

    "applicationId": "string",
    "arn": "string",
    "billedResourceUtilization": {
        "memoryGBHour": number,
        "storageGBHour": number,
        "vCPUPercent": number
    },
    "configurationOverrides": {
        "runtimeConfiguration": [
            {
                "classification": "string",
                "configurations": [
                    "Configuration"
                ],
                "properties": {
                    "string" : "string"
                }
            }
        ]
    },
    "createdAt": number,
    "createdBy": "string",
    "endedAt": number,
    "executionRoleArn": "string",
    "idleSince": number,
    "idleTimeoutMinutes": number,
    "name": "string",
    "networkConfiguration": {
        "securityGroupIds": [ "string" ],
        "subnetIds": [ "string" ]
    },
    "releaseLabel": "string",
    "sessionId": "string",
    "startedAt": number,
    "state": "string",
    "stateDetails": "string",
    "tags": {
        "string" : "string"
    },
    "totalExecutionDurationSeconds": number,
    "totalResourceUtilization": {
        "memoryGBHour": number,
        "storageGBHour": number,
        "vCPUPercent": number
    }

```

```
    },  
    "updatedAt": number  
  }  
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[session](#)

The output displays information about the session.

Type: [Session](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerError

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

GetSessionEndpoint

Returns the session endpoint URL and a time-limited authentication token for the specified session. Use the endpoint and token to connect a client to the session. Call this operation again when the authentication token expires to obtain a new token.

Request Syntax

```
GET /applications/applicationId/sessions/sessionId/endpoint HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application that the session belongs to.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

sessionId

The ID of the session.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Content-type: application/json

```
{
  "applicationId": "string",
  "authToken": "string",
  "authTokenExpiresAt": number,
  "endpoint": "string",
  "sessionId": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[applicationId](#)

The output contains the ID of the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

[authToken](#)

The authentication token for connecting to the session endpoint. Call `GetSessionEndpoint` again to obtain a new token before it expires.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 8000.

[authTokenExpiresAt](#)

The expiration time of the authentication token.

Type: Timestamp

[endpoint](#)

The endpoint URL for connecting to the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

sessionId

The output contains the ID of the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerError

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)

- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListApplications

Lists applications based on a set of parameters.

Request Syntax

```
GET /applications?maxResults=maxResults&nextToken=nextToken&states=states HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

maxResults

The maximum number of applications that can be listed.

Valid Range: Minimum value of 1. Maximum value of 50.

nextToken

The token for the next set of application results.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: [A-Za-z0-9_=-]+

states

An optional filter for application states. Note that if this filter contains multiple states, the resulting list will be grouped by the state.

Array Members: Minimum number of 1 item. Maximum number of 7 items.

Valid Values: CREATING | CREATED | STARTING | STARTED | STOPPING | STOPPED | TERMINATED

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "applications": [
    {
      "architecture": "string",
      "arn": "string",
      "createdAt": number,
      "id": "string",
      "name": "string",
      "releaseLabel": "string",
      "state": "string",
      "stateDetails": "string",
      "type": "string",
      "updatedAt": number
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[applications](#)

The output lists the specified applications.

Type: Array of [ApplicationSummary](#) objects

[nextToken](#)

The output displays the token for the next set of application results. This is required for pagination and is available as a response of the previous request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: [A-Za-z0-9_=-]+

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerError

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListJobRunAttempts

Lists all attempt of a job run.

Request Syntax

```
GET /applications/applicationId/jobruns/jobRunId/attempts?  
maxResults=maxResults&nextToken=nextToken HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application for which to list job runs.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

jobRunId

The ID of the job run to list.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

maxResults

The maximum number of job run attempts to list.

Valid Range: Minimum value of 1. Maximum value of 50.

nextToken

The token for the next set of job run attempt results.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: [A-Za-z0-9_=-]+

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "jobRunAttempts": [
    {
      "applicationId": "string",
      "arn": "string",
      "attempt": number,
      "createdAt": number,
      "createdBy": "string",
      "executionRole": "string",
      "id": "string",
      "jobCreatedAt": number,
      "mode": "string",
      "name": "string",
      "releaseLabel": "string",
      "state": "string",
      "stateDetails": "string",
      "type": "string",
      "updatedAt": number
    }
  ],
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

jobRunAttempts

The array of the listed job run attempt objects.

Type: Array of [JobRunAttemptSummary](#) objects

nextToken

The output displays the token for the next set of application results. This is required for pagination and is available as a response of the previous request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: [A-Za-z0-9_=-]+

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListJobRuns

Lists job runs based on a set of parameters.

Request Syntax

```
GET /applications/applicationId/jobruns?  
createdAtAfter=createdAtAfter&createdAtBefore=createdAtBefore&maxResults=maxResults&mode=mode&n  
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application for which to list the job run.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

createdAtAfter

The lower bound of the option to filter by creation date and time.

createdAtBefore

The upper bound of the option to filter by creation date and time.

maxResults

The maximum number of job runs that can be listed.

Valid Range: Minimum value of 1. Maximum value of 50.

mode

The mode of the job runs to list.

Valid Values: BATCH | STREAMING

nextToken

The token for the next set of job run results.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: [A-Za-z0-9_=-]+

states

An optional filter for job run states. Note that if this filter contains multiple states, the resulting list will be grouped by the state.

Array Members: Minimum number of 0 items. Maximum number of 8 items.

Valid Values: SUBMITTED | PENDING | SCHEDULED | RUNNING | SUCCESS | FAILED | CANCELLING | CANCELLED | QUEUED

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "jobRuns": [
    {
      "applicationId": "string",
      "arn": "string",
      "attempt": number,
      "attemptCreatedAt": number,
      "attemptUpdatedAt": number,
      "createdAt": number,
      "createdBy": "string",
      "executionRole": "string",
      "id": "string",
      "mode": "string",
      "name": "string",
      "releaseLabel": "string",
      "state": "string",
```

```
    "stateDetails": "string",
    "type": "string",
    "updatedAt": number
  },
  "nextToken": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[jobRuns](#)

The output lists information about the specified job runs.

Type: Array of [JobRunSummary](#) objects

[nextToken](#)

The output displays the token for the next set of job run results. This is required for pagination and is available as a response of the previous request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: [A-Za-z0-9_=-]+

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListSessions

Lists sessions for the specified application. You can filter sessions by state and creation time.

Request Syntax

```
GET /applications/applicationId/sessions?  
createdAtAfter=createdAtAfter&createdAtBefore=createdAtBefore&maxResults=maxResults&nextToken=nextToken  
HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application to list sessions for.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

Required: Yes

createdAtAfter

The lower bound of the option to filter by creation date and time.

createdAtBefore

The upper bound of the option to filter by creation date and time.

maxResults

The maximum number of sessions to return in each page of results.

Valid Range: Minimum value of 1. Maximum value of 50.

nextToken

The token for the next set of session results.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: [A-Za-z0-9_=-]+

states

An optional filter for session states. Note that if this filter contains multiple states, the resulting list will be grouped by the state.

Array Members: Minimum number of 1 item. Maximum number of 8 items.

Valid Values: SUBMITTED | STARTING | STARTED | IDLE | BUSY | FAILED |
TERMINATING | TERMINATED

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "nextToken": "string",
  "sessions": [
    {
      "applicationId": "string",
      "arn": "string",
      "createdAt": number,
      "createdBy": "string",
      "executionRoleArn": "string",
      "name": "string",
      "releaseLabel": "string",
      "sessionId": "string",
      "state": "string",
      "stateDetails": "string",
      "updatedAt": number
    }
  ]
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

nextToken

The output displays the token for the next set of session results. This is required for pagination and is available as a response of the previous request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: [A-Za-z0-9_=-]+

sessions

The output lists information about the specified sessions.

Type: Array of [SessionSummary](#) objects

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerError

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

ListTagsForResource

Lists the tags assigned to the resources.

Request Syntax

```
GET /tags/resourceArn HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) that identifies the resource to list the tags for. Currently, the supported resources are Amazon EMR Serverless applications and job runs.

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.*:(\d{12}):\./applications\/[0-9a-zA-Z]+(\./jobruns\/[0-9a-zA-Z]+)?`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "tags": {
    "string" : "string"
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

tags

The tags for the resource.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 200 items.

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: [A-Za-z0-9 /_ . : = + @ -] +

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: [A-Za-z0-9 /_ . : = + @ -] *

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerError

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

StartApplication

Starts a specified application and initializes initial capacity if configured.

Request Syntax

```
POST /applications/applicationId/start HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application to start.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ServiceQuotaExceededException

The maximum number of resources per account has been reached.

HTTP Status Code: 402

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

StartJobRun

Starts a job run.

Request Syntax

POST /applications/*applicationId*/jobruns HTTP/1.1

Content-type: application/json

```
{
  "clientToken": "string",
  "configurationOverrides": {
    "applicationConfiguration": [
      {
        "classification": "string",
        "configurations": [
          "Configuration"
        ],
        "properties": {
          "string" : "string"
        }
      }
    ],
    "diskEncryptionConfiguration": {
      "encryptionContext": {
        "string" : "string"
      },
      "encryptionKeyArn": "string"
    },
    "monitoringConfiguration": {
      "cloudWatchLoggingConfiguration": {
        "enabled": boolean,
        "encryptionKeyArn": "string",
        "logGroupName": "string",
        "logStreamNamePrefix": "string",
        "logTypes": {
          "string" : [ "string" ]
        }
      }
    },
    "managedPersistenceMonitoringConfiguration": {
      "enabled": boolean,
      "encryptionKeyArn": "string"
    }
  }
}
```

```

    "prometheusMonitoringConfiguration": {
      "remoteWriteUrl": "string"
    },
    "s3MonitoringConfiguration": {
      "encryptionKeyArn": "string",
      "logUri": "string"
    }
  },
  "executionIamPolicy": {
    "policy": "string",
    "policyArns": [ "string" ]
  },
  "executionRoleArn": "string",
  "executionTimeoutMinutes": number,
  "jobDriver": { ... },
  "mode": "string",
  "name": "string",
  "retryPolicy": {
    "maxAttempts": number,
    "maxFailedAttemptsPerHour": number
  },
  "tags": {
    "string" : "string"
  }
}

```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application on which to run the job.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

Required: Yes

Request Body

The request accepts the following data in JSON format.

clientToken

The client idempotency token of the job run to start. Its value must be unique for each request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[A-Za-z0-9._-]+`

Required: Yes

configurationOverrides

The configuration overrides for the job run.

Type: [ConfigurationOverrides](#) object

Required: No

executionIamPolicy

You can pass an optional IAM policy. The resulting job IAM role permissions will be an intersection of this policy and the policy associated with your job execution role.

Type: [JobRunExecutionIamPolicy](#) object

Required: No

executionRoleArn

The execution role ARN for the job run.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):iam:([0-9]{12}):role((\u002F|(\u002F[\u0021-\u007F]+\u002F))[\w+=,.\@-]+)`

Required: Yes

executionTimeoutMinutes

The maximum duration, in minutes, for the job run. If the job run exceeds this duration, Amazon EMR Serverless cancels it automatically.

For BATCH mode job runs, the maximum value is 10080 minutes (7 days) starting with Amazon EMR release 7.11. Setting a value of 0 to disable the timeout is no longer supported for BATCH mode job runs.

Type: Long

Valid Range: Minimum value of 0. Maximum value of 1000000.

Required: No

jobDriver

The job driver for the job run.

Type: [JobDriver](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: No

mode

The mode of the job run when it starts.

Type: String

Valid Values: BATCH | STREAMING

Required: No

name

The optional job run name. This doesn't have to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `.*\S.*`

Required: No

[retryPolicy](#)

The retry policy when job run starts.

Type: [RetryPolicy](#) object

Required: No

[tags](#)

The tags assigned to the job run.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 200 items.

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: [A-Za-z0-9 /_ . : = + @ -] +

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: [A-Za-z0-9 /_ . : = + @ -] *

Required: No

Response Syntax

```
HTTP/1.1 200
Content-type: application/json

{
  "applicationId": "string",
  "arn": "string",
  "jobRunId": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

applicationId

This output displays the application ID on which the job run was submitted.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

arn

This output displays the ARN of the job run..

Type: String

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.*:(\d{12}):/applications/[0-9a-zA-Z]+/jobruns/[0-9a-zA-Z]+`

jobRunId

The output contains the ID of the started job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

ConflictException

The request could not be processed because of conflict in the current state of the resource.

HTTP Status Code: 409

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

StartSession

Creates and starts a new session on the specified application. The application must be in the STARTED state or have AutoStart enabled, and have interactive sessions enabled. This operation is supported for EMR release 7.13.0 and later.

Request Syntax

POST /applications/*applicationId*/sessions HTTP/1.1

Content-type: application/json

```
{
  "clientToken": "string",
  "configurationOverrides": {
    "runtimeConfiguration": [
      {
        "classification": "string",
        "configurations": [
          "Configuration"
        ],
        "properties": {
          "string" : "string"
        }
      }
    ]
  },
  "executionRoleArn": "string",
  "idleTimeoutMinutes": number,
  "name": "string",
  "tags": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application on which to start the session.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

Request Body

The request accepts the following data in JSON format.

clientToken

A unique, case-sensitive identifier that you provide to ensure the idempotency of the request. If you retry a request that completed successfully using the same client token, the server returns the successful response without performing the operation again.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[A-Za-z0-9._-]+`

Required: Yes

configurationOverrides

The configuration overrides for the session. Only runtime configuration overrides are supported.

Type: [SessionConfigurationOverrides](#) object

Required: No

executionRoleArn

The execution role ARN for the session. Amazon EMR Serverless uses this role to access AWS resources on your behalf during session execution.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):iam::([0-9]{12}):role((\u002F|(\u002F[\u0021-\u007F]+\u002F))[\w+=,.\@-]+)`

Required: Yes

idleTimeoutMinutes

The idle timeout in minutes for the session. After the session remains idle for this duration, Amazon EMR Serverless automatically terminates it.

Type: Long

Valid Range: Minimum value of 0. Maximum value of 1000000.

Required: No

name

The optional name for the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `.*\S.*`

Required: No

tags

The tags to assign to the session.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 200 items.

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: `[A-Za-z0-9 /_ . : = + @ - ] +`

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: `[A-Za-z0-9 /_ . : = + @ - ] *`

Required: No

Response Syntax

```
HTTP/1.1 200
```

```
Content-type: application/json

{
  "applicationId": "string",
  "arn": "string",
  "sessionId": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

applicationId

The output contains the application ID on which the session was started.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

arn

The output contains the ARN of the session.

Type: String

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.*:(\d{12}):\applications\[0-9a-zA-Z]+\sessions\[0-9a-zA-Z]+`

sessionId

The output contains the ID of the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

ConflictException

The request could not be processed because of conflict in the current state of the resource.

HTTP Status Code: 409

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ServiceQuotaExceededException

The maximum number of resources per account has been reached.

HTTP Status Code: 402

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

StopApplication

Stops a specified application and releases initial capacity if configured. All scheduled and running jobs must be completed or cancelled before stopping an application.

Request Syntax

```
POST /applications/applicationId/stop HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application to stop.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

TagResource

Assigns tags to resources. A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value, both of which you define. Tags enable you to categorize your AWS resources by attributes such as purpose, owner, or environment. When you have many resources of the same type, you can quickly identify a specific resource based on the tags you've assigned to it.

Request Syntax

```
POST /tags/resourceArn HTTP/1.1
Content-type: application/json

{
  "tags": {
    "string" : "string"
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) that identifies the resource to list the tags for. Currently, the supported resources are Amazon EMR Serverless applications and job runs.

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.*:(\d{12}):\applications\`
`[0-9a-zA-Z]+(\jobruns\[/code>`

Required: Yes

Request Body

The request accepts the following data in JSON format.

tags

The tags to add to the resource. A tag is an array of key-value pairs.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 200 items.

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: [A-Za-z0-9 /_ . : = + @ -] +

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: [A-Za-z0-9 /_ . : = + @ -] *

Required: Yes

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

TerminateSession

Terminates the specified session. After you terminate a session, it enters the TERMINATING state and then the TERMINATED state. You can still access the Spark History Server for a terminated session through the GetResourceDashboard operation.

Request Syntax

```
DELETE /applications/applicationId/sessions/sessionId HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application that the session belongs to.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

sessionId

The ID of the session to terminate.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

```
Content-type: application/json

{
  "applicationId": "string",
  "sessionId": "string"
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

[applicationId](#)

The output contains the application ID on which the session was terminated.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

[sessionId](#)

The output contains the ID of the terminated session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerError

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UntagResource

Removes tags from resources.

Request Syntax

```
DELETE /tags/resourceArn?tagKeys=tagKeys HTTP/1.1
```

URI Request Parameters

The request uses the following URI parameters.

resourceArn

The Amazon Resource Name (ARN) that identifies the resource to list the tags for. Currently, the supported resources are Amazon EMR Serverless applications and job runs.

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*) :emr-serverless: .+ :(\d{12}) : \backslash applications \backslash [0-9a-zA-Z]+ ( \backslash jobruns \backslash [0-9a-zA-Z]+ ) ?`

Required: Yes

tagKeys

The keys of the tags to be removed.

Array Members: Minimum number of 1 item. Maximum number of 200 items.

Length Constraints: Minimum length of 1. Maximum length of 128.

Pattern: `[A-Za-z0-9 / _ . : = + @ - ] +`

Required: Yes

Request Body

The request does not have a request body.

Response Syntax

```
HTTP/1.1 200
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response with an empty HTTP body.

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerErrorException

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

UpdateApplication

Updates a specified application. An application has to be in a stopped or created state in order to be updated.

Request Syntax

```
PATCH /applications/applicationId HTTP/1.1
```

```
Content-type: application/json
```

```
{
  "architecture": "string",
  "autoStartConfiguration": {
    "enabled": boolean
  },
  "autoStopConfiguration": {
    "enabled": boolean,
    "idleTimeoutMinutes": number
  },
  "clientToken": "string",
  "diskEncryptionConfiguration": {
    "encryptionContext": {
      "string" : "string"
    },
    "encryptionKeyArn": "string"
  },
  "identityCenterConfiguration": {
    "identityCenterInstanceArn": "string",
    "userBackgroundSessionsEnabled": boolean
  },
  "imageConfiguration": {
    "applicationLevelDigestResolution": boolean,
    "imageUri": "string"
  },
  "initialCapacity": {
    "string" : {
      "workerConfiguration": {
        "cpu": "string",
        "disk": "string",
        "diskType": "string",
        "memory": "string"
      },
      "workerCount": number
    }
  }
}
```

```

    }
  },
  "interactiveConfiguration": {
    "livyEndpointEnabled": boolean,
    "sessionEnabled": boolean,
    "studioEnabled": boolean
  },
  "jobLevelCostAllocationConfiguration": {
    "enabled": boolean
  },
  "maximumCapacity": {
    "cpu": "string",
    "disk": "string",
    "memory": "string"
  },
  "monitoringConfiguration": {
    "cloudWatchLoggingConfiguration": {
      "enabled": boolean,
      "encryptionKeyArn": "string",
      "logGroupName": "string",
      "logStreamNamePrefix": "string",
      "logTypes": {
        "string" : [ "string" ]
      }
    },
    "managedPersistenceMonitoringConfiguration": {
      "enabled": boolean,
      "encryptionKeyArn": "string"
    },
    "prometheusMonitoringConfiguration": {
      "remoteWriteUrl": "string"
    },
    "s3MonitoringConfiguration": {
      "encryptionKeyArn": "string",
      "logUri": "string"
    }
  },
  "networkConfiguration": {
    "securityGroupIds": [ "string" ],
    "subnetIds": [ "string" ]
  },
  "releaseLabel": "string",
  "runtimeConfiguration": [
    {

```

```
    "classification": "string",
    "configurations": [
      "Configuration"
    ],
    "properties": {
      "string" : "string"
    }
  },
  "schedulerConfiguration": {
    "maxConcurrentRuns": number,
    "queueTimeoutMinutes": number
  },
  "workerTypeSpecifications": {
    "string" : {
      "imageConfiguration": {
        "applicationLevelDigestResolution": boolean,
        "imageUri": "string"
      }
    }
  }
}
```

URI Request Parameters

The request uses the following URI parameters.

applicationId

The ID of the application to update.

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [0-9a-z]+

Required: Yes

Request Body

The request accepts the following data in JSON format.

architecture

The CPU architecture of an application.

Type: String

Valid Values: ARM64 | X86_64

Required: No

autoStartConfiguration

The configuration for an application to automatically start on job submission.

Type: [AutoStartConfig](#) object

Required: No

autoStopConfiguration

The configuration for an application to automatically stop after a certain amount of time being idle.

Type: [AutoStopConfig](#) object

Required: No

clientToken

The client idempotency token of the application to update. Its value must be unique for each request.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [A-Za-z0-9._-]+

Required: Yes

diskEncryptionConfiguration

The configuration object that allows encrypting local disks.

Type: [DiskEncryptionConfiguration](#) object

Required: No

identityCenterConfiguration

Specifies the IAM Identity Center configuration used to enable or disable trusted identity propagation. When provided, this configuration determines how the application interacts with IAM Identity Center for user authentication and access control.

Type: [IdentityCenterConfigurationInput](#) object

Required: No

imageConfiguration

The image configuration to be used for all worker types. You can either set this parameter or `imageConfiguration` for each worker type in `WorkerTypeSpecificationInput`.

Type: [ImageConfigurationInput](#) object

Required: No

initialCapacity

The capacity to initialize when the application is updated.

Type: String to [InitialCapacityConfig](#) object map

Map Entries: Minimum number of 0 items. Maximum number of 10 items.

Key Length Constraints: Minimum length of 1. Maximum length of 50.

Key Pattern: `[a-zA-Z]+[-_]*[a-zA-Z]+`

Required: No

interactiveConfiguration

The interactive configuration object that contains new interactive use cases when the application is updated.

Type: [InteractiveConfiguration](#) object

Required: No

jobLevelCostAllocationConfiguration

The configuration object that enables job level cost allocation.

Type: [JobLevelCostAllocationConfiguration](#) object

Required: No

maximumCapacity

The maximum capacity to allocate when the application is updated. This is cumulative across all workers at any given point in time during the lifespan of the application. No new resources will be created once any one of the defined limits is hit.

Type: [MaximumAllowedResources](#) object

Required: No

monitoringConfiguration

The configuration setting for monitoring.

Type: [MonitoringConfiguration](#) object

Required: No

networkConfiguration

The network configuration for customer VPC connectivity.

Type: [NetworkConfiguration](#) object

Required: No

releaseLabel

The Amazon EMR release label for the application. You can change the release label to use a different release of Amazon EMR.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [A-Za-z0-9._/-]+

Required: No

runtimeConfiguration

The [Configuration](#) specifications to use when updating an application. Each configuration consists of a classification and properties. This configuration is applied across all the job runs submitted under the application.

Type: Array of [Configuration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Required: No

[schedulerConfiguration](#)

The scheduler configuration for batch and streaming jobs running on this application. Supported with release labels emr-7.0.0 and above.

Type: [SchedulerConfiguration](#) object

Required: No

[workerTypeSpecifications](#)

The key-value pairs that specify worker type to `WorkerTypeSpecificationInput`. This parameter must contain all valid worker types for a Spark or Hive application. Valid worker types include Driver and Executor for Spark applications and HiveDriver and TezTask for Hive applications. You can either set image details in this parameter for each worker type, or in `imageConfiguration` for all worker types.

Type: String to [WorkerTypeSpecificationInput](#) object map

Key Length Constraints: Minimum length of 1. Maximum length of 50.

Key Pattern: `[a-zA-Z]+[-_]*[a-zA-Z]+`

Required: No

Response Syntax

HTTP/1.1 200

Content-type: application/json

```
{
  "application": {
    "applicationId": "string",
    "architecture": "string",
    "arn": "string",
    "autoStartConfiguration": {
      "enabled": boolean
    },
    "autoStopConfiguration": {
      "enabled": boolean,
```

```
    "idleTimeoutMinutes": number
  },
  "createdAt": number,
  "diskEncryptionConfiguration": {
    "encryptionContext": {
      "string" : "string"
    },
    "encryptionKeyArn": "string"
  },
  "identityCenterConfiguration": {
    "identityCenterApplicationArn": "string",
    "identityCenterInstanceArn": "string",
    "userBackgroundSessionsEnabled": boolean
  },
  "imageConfiguration": {
    "applicationLevelDigestResolution": boolean,
    "imageUri": "string",
    "resolvedImageDigest": "string"
  },
  "initialCapacity": {
    "string" : {
      "workerConfiguration": {
        "cpu": "string",
        "disk": "string",
        "diskType": "string",
        "memory": "string"
      },
      "workerCount": number
    }
  },
  "interactiveConfiguration": {
    "livyEndpointEnabled": boolean,
    "sessionEnabled": boolean,
    "studioEnabled": boolean
  },
  "jobLevelCostAllocationConfiguration": {
    "enabled": boolean
  },
  "maximumCapacity": {
    "cpu": "string",
    "disk": "string",
    "memory": "string"
  },
  "monitoringConfiguration": {
```

```

    "cloudWatchLoggingConfiguration": {
      "enabled": boolean,
      "encryptionKeyArn": "string",
      "logGroupName": "string",
      "logStreamNamePrefix": "string",
      "logTypes": {
        "string" : [ "string" ]
      }
    },
    "managedPersistenceMonitoringConfiguration": {
      "enabled": boolean,
      "encryptionKeyArn": "string"
    },
    "prometheusMonitoringConfiguration": {
      "remoteWriteUrl": "string"
    },
    "s3MonitoringConfiguration": {
      "encryptionKeyArn": "string",
      "logUri": "string"
    }
  },
  "name": "string",
  "networkConfiguration": {
    "securityGroupIds": [ "string" ],
    "subnetIds": [ "string" ]
  },
  "releaseLabel": "string",
  "runtimeConfiguration": [
    {
      "classification": "string",
      "configurations": [
        "Configuration"
      ],
      "properties": {
        "string" : "string"
      }
    }
  ],
  "schedulerConfiguration": {
    "maxConcurrentRuns": number,
    "queueTimeoutMinutes": number
  },
  "state": "string",
  "stateDetails": "string",

```

```
"tags": {
  "string" : "string"
},
"type": "string",
"updatedAt": number,
"workerTypeSpecifications": {
  "string" : {
    "imageConfiguration": {
      "applicationLevelDigestResolution": boolean,
      "imageUri": "string",
      "resolvedImageDigest": "string"
    }
  }
}
```

Response Elements

If the action is successful, the service sends back an HTTP 200 response.

The following data is returned in JSON format by the service.

application

Information about the updated application.

Type: [Application](#) object

Errors

For information about the errors that are common to all actions, see [Common Error Types](#).

InternalServerError

Request processing failed because of an error or failure with the service.

HTTP Status Code: 500

ResourceNotFoundException

The specified resource was not found.

HTTP Status Code: 404

ValidationException

The input fails to satisfy the constraints specified by an AWS service.

HTTP Status Code: 400

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python \(Boto3\)](#)
- [AWS SDK for Ruby V3](#)

Data Types

The EMR Serverless API contains several data types that various actions use. This section describes each data type in detail.

Note

The order of each element in a data type structure is not guaranteed. Applications should not assume a particular order.

The following data types are supported:

- [Application](#)
- [ApplicationSummary](#)
- [AutoStartConfig](#)
- [AutoStopConfig](#)
- [CloudWatchLoggingConfiguration](#)
- [Configuration](#)
- [ConfigurationOverrides](#)
- [DiskEncryptionConfiguration](#)
- [Hive](#)
- [IdentityCenterConfiguration](#)
- [IdentityCenterConfigurationInput](#)
- [ImageConfiguration](#)
- [ImageConfigurationInput](#)
- [InitialCapacityConfig](#)
- [InteractiveConfiguration](#)
- [JobDriver](#)
- [JobLevelCostAllocationConfiguration](#)
- [JobRun](#)
- [JobRunAttemptSummary](#)
- [JobRunExecutionIamPolicy](#)

- [JobRunSummary](#)
- [ManagedPersistenceMonitoringConfiguration](#)
- [MaximumAllowedResources](#)
- [MonitoringConfiguration](#)
- [NetworkConfiguration](#)
- [PrometheusMonitoringConfiguration](#)
- [ResourceUtilization](#)
- [RetryPolicy](#)
- [S3MonitoringConfiguration](#)
- [SchedulerConfiguration](#)
- [Session](#)
- [SessionConfigurationOverrides](#)
- [SessionSummary](#)
- [SparkSubmit](#)
- [TotalResourceUtilization](#)
- [WorkerResourceConfig](#)
- [WorkerTypeSpecification](#)
- [WorkerTypeSpecificationInput](#)

Application

Information about an application. Amazon EMR Serverless uses applications to run jobs.

Contents

applicationId

The ID of the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

arn

The ARN of the application.

Type: String

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.+:(\d{12}):\applications\[0-9a-zA-Z]+`

Required: Yes

createdAt

The date and time when the application run was created.

Type: Timestamp

Required: Yes

releaseLabel

The Amazon EMR release associated with the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[A-Za-z0-9._/-]+`

Required: Yes

state

The state of the application.

Type: String

Valid Values: CREATING | CREATED | STARTING | STARTED | STOPPING | STOPPED | TERMINATED

Required: Yes

type

The type of application, such as Spark or Hive.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: Yes

updatedAt

The date and time when the application run was last updated.

Type: Timestamp

Required: Yes

architecture

The CPU architecture of an application.

Type: String

Valid Values: ARM64 | X86_64

Required: No

autoStartConfiguration

The configuration for an application to automatically start on job submission.

Type: [AutoStartConfig](#) object

Required: No

autoStopConfiguration

The configuration for an application to automatically stop after a certain amount of time being idle.

Type: [AutoStopConfig](#) object

Required: No

diskEncryptionConfiguration

The configuration object that allows encrypting local disks.

Type: [DiskEncryptionConfiguration](#) object

Required: No

identityCenterConfiguration

The IAM Identity Center configuration applied to enable trusted identity propagation.

Type: [IdentityCenterConfiguration](#) object

Required: No

imageConfiguration

The image configuration applied to all worker types.

Type: [ImageConfiguration](#) object

Required: No

initialCapacity

The initial capacity of the application.

Type: String to [InitialCapacityConfig](#) object map

Map Entries: Minimum number of 0 items. Maximum number of 10 items.

Key Length Constraints: Minimum length of 1. Maximum length of 50.

Key Pattern: `[a-zA-Z]+[-_]*[a-zA-Z]+`

Required: No

interactiveConfiguration

The interactive configuration object that enables the interactive use cases for an application.

Type: [InteractiveConfiguration](#) object

Required: No

jobLevelCostAllocationConfiguration

The configuration object that enables job level cost allocation.

Type: [JobLevelCostAllocationConfiguration](#) object

Required: No

maximumCapacity

The maximum capacity of the application. This is cumulative across all workers at any given point in time during the lifespan of the application is created. No new resources will be created once any one of the defined limits is hit.

Type: [MaximumAllowedResources](#) object

Required: No

monitoringConfiguration

The configuration setting for monitoring.

Type: [MonitoringConfiguration](#) object

Required: No

name

The name of the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[A-Za-z0-9._/#- ]+`

Required: No

networkConfiguration

The network configuration for customer VPC connectivity for the application.

Type: [NetworkConfiguration](#) object

Required: No

runtimeConfiguration

The [Configuration](#) specifications of an application. Each configuration consists of a classification and properties. You use this parameter when creating or updating an application. To see the runtimeConfiguration object of an application, run the [GetApplication](#) API operation.

Type: Array of [Configuration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Required: No

schedulerConfiguration

The scheduler configuration for batch and streaming jobs running on this application. Supported with release labels emr-7.0.0 and above.

Type: [SchedulerConfiguration](#) object

Required: No

stateDetails

The state details of the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `.*\S.*`

Required: No

tags

The tags assigned to the application.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 200 items.

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: `[A-Za-z0-9 /_ . : = + @ - ] +`

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: `[A-Za-z0-9 /_ . : = + @ - ] *`

Required: No

workerTypeSpecifications

The specification applied to each worker type.

Type: String to [WorkerTypeSpecification](#) object map

Key Length Constraints: Minimum length of 1. Maximum length of 50.

Key Pattern: `[a-zA-Z] + [ - _ ] * [a-zA-Z] +`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ApplicationSummary

The summary of attributes associated with an application.

Contents

arn

The ARN of the application.

Type: String

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.\+:(\d{12}):\./applications\./[0-9a-zA-Z-Z]+`

Required: Yes

createdAt

The date and time when the application was created.

Type: Timestamp

Required: Yes

id

The ID of the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z-Z]+`

Required: Yes

releaseLabel

The Amazon EMR release associated with the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [A-Za-z0-9._/-]+

Required: Yes

state

The state of the application.

Type: String

Valid Values: CREATING | CREATED | STARTING | STARTED | STOPPING | STOPPED | TERMINATED

Required: Yes

type

The type of application, such as Spark or Hive.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Required: Yes

updatedAt

The date and time when the application was last updated.

Type: Timestamp

Required: Yes

architecture

The CPU architecture of an application.

Type: String

Valid Values: ARM64 | X86_64

Required: No

name

The name of the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [A-Za-z0-9._/#-]+

Required: No

stateDetails

The state details of the application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: .*S.*

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AutoStartConfig

The configuration for an application to automatically start on job submission.

Contents

enabled

Enables the application to automatically start on job submission. Defaults to true.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

AutoStopConfig

The configuration for an application to automatically stop after a certain amount of time being idle.

Contents

enabled

Enables the application to automatically stop after a certain amount of time being idle.
Defaults to true.

Type: Boolean

Required: No

idleTimeoutMinutes

The amount of idle time in minutes after which your application will automatically stop.
Defaults to 15 minutes.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10080.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

CloudWatchLoggingConfiguration

The Amazon CloudWatch configuration for monitoring logs. You can configure your jobs to send log information to CloudWatch.

Contents

enabled

Enables CloudWatch logging.

Type: Boolean

Required: Yes

encryptionKeyArn

The AWS Key Management Service (KMS) key ARN to encrypt the logs that you store in CloudWatch Logs.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):kms:[a-zA-Z0-9-]*:([0-9]{12}):key\[a-zA-Z0-9-]+\`

Required: No

logGroupName

The name of the log group in Amazon CloudWatch Logs where you want to publish your logs.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[\\.\-_\#A-Za-z0-9]+\`

Required: No

logStreamNamePrefix

Prefix for the CloudWatch log stream name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 512.

Pattern: `[^:]*`

Required: No

logTypes

The types of logs that you want to publish to CloudWatch. If you don't specify any log types, driver STDOUT and STDERR logs will be published to CloudWatch Logs by default. For more information including the supported worker types for Hive and Spark, see [Logging for EMR Serverless with CloudWatch](#).

- **Key Valid Values:** SPARK_DRIVER, SPARK_EXECUTOR, HIVE_DRIVER, TEZ_TASK
- **Array Members Valid Values:** STDOUT, STDERR, HIVE_LOG, TEZ_AM, SYSTEM_LOGS

Type: String to array of strings map

Map Entries: Maximum number of 4 items.

Key Length Constraints: Minimum length of 1. Maximum length of 50.

Key Pattern: `[a-zA-Z]+[-_]*[a-zA-Z]+`

Array Members: Minimum number of 1 item. Maximum number of 5 items.

Length Constraints: Minimum length of 1. Maximum length of 50.

Pattern: `[a-zA-Z]+[-_]*[a-zA-Z]+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Configuration

A configuration specification to be used when provisioning an application. A configuration consists of a classification, properties, and optional nested configurations. A classification refers to an application-specific configuration file. Properties are the settings you want to change in that file.

Contents

classification

The classification within a configuration.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `.*\S.*`

Required: Yes

configurations

A list of additional configurations to apply within a configuration object.

Type: Array of [Configuration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Required: No

properties

A set of properties specified within a configuration classification.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 100 items.

Key Length Constraints: Minimum length of 1. Maximum length of 1024.

Key Pattern: `.*\S.*`

Value Length Constraints: Minimum length of 0. Maximum length of 1024.

Value Pattern: `.*\S.*`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ConfigurationOverrides

A configuration specification to be used to override existing configurations.

Contents

applicationConfiguration

The override configurations for the application.

Type: Array of [Configuration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Required: No

diskEncryptionConfiguration

The override configuration to encrypt local disks.

Type: [DiskEncryptionConfiguration](#) object

Required: No

monitoringConfiguration

The override configurations for monitoring.

Type: [MonitoringConfiguration](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

DiskEncryptionConfiguration

The configuration object that allows encrypting local disks.

Contents

encryptionContext

Specifies the optional encryption context that will be used when encrypting the data. An encryption context is a collection of non-secret key-value pairs that represent additional authenticated data.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 8 items.

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Value Length Constraints: Minimum length of 1. Maximum length of 384.

Required: No

encryptionKeyArn

The KMS key ARN to encrypt local disks.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):kms:[a-zA-Z0-9\-*]*:([0-9]{12}):key\[a-zA-Z0-9-]*`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Hive

The configurations for the Hive job driver.

Contents

query

The query for the Hive job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 10280.

Pattern: `.*\S.*`

Required: Yes

initQueryFile

The query file for the Hive job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `.*\S.*`

Required: No

parameters

The parameters for the Hive job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 102400.

Pattern: `.*\S.*`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IdentityCenterConfiguration

The IAM Identity Center Configuration accepts the Identity Center instance parameter required to enable trusted identity propagation. This configuration allows identity propagation between integrated services and the Identity Center instance.

Contents

identityCenterApplicationArn

The ARN of the EMR Serverless created IAM Identity Center Application that provides trusted-identity propagation.

Type: String

Pattern: `arn:(aws[a-zA-Z0-9-]*):sso::\d{12}:application/(sso)?ins-[a-zA-Z0-9-]{16}/apl-[a-zA-Z0-9-]{16}`

Required: No

identityCenterInstanceArn

The ARN of the IAM Identity Center instance.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):sso:::instance/(sso)?ins-[a-zA-Z0-9-]{16}`

Required: No

userBackgroundSessionsEnabled

Enables user background sessions for this application so Livy sessions can continue running after users log out of their interactive notebook or their Identity Center sessions expire.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

IdentityCenterConfigurationInput

The IAM Identity Center Configuration accepts the Identity Center instance parameter required to enable trusted identity propagation. This configuration allows identity propagation between integrated services and the Identity Center instance.

Contents

identityCenterInstanceArn

The ARN of the IAM Identity Center instance.

Type: String

Length Constraints: Minimum length of 10. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):sso:::instance/(sso)?ins-[a-zA-Z0-9-]{16}`

Required: No

userBackgroundSessionsEnabled

Enables user background sessions for this application so Livy sessions can continue running after users log out of their interactive notebook or their Identity Center sessions expire.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ImageConfiguration

The applied image configuration.

Contents

imageUri

The image URI.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: (`[0-9]{12}`)\.`dkr`\.`ecr`\.(`[a-z0-9-]+`).(`[a-z0-9._-]+`)\(/((?:`[a-z0-9]+(?:[-._][a-z0-9]+)*/`)*`[a-z0-9]+(?:[-._][a-z0-9]+)*`)(?::(`[a-zA-Z0-9_]+[a-zA-Z0-9-._]*`)|`@(sha256:[0-9a-f]{64})`))

Required: Yes

applicationLevelDigestResolution

Boolean value indicating if the digest resolution is application level or workload level. If true, a custom image URI is resolved at application start time and all workloads submitted will use that image digest. If false, the custom image URI is resolved at the workload submission time.

Type: Boolean

Required: No

resolvedImageDigest

The SHA256 digest of the image URI. This indicates which specific image the application is configured for. The image digest doesn't exist until an application has started.

Type: String

Pattern: `sha256:[0-9a-f]{64}`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ImageConfigurationInput

The image configuration.

Contents

applicationLevelDigestResolution

Boolean value indicating if the digest resolution is application level or workload level. If true, a custom image URI is resolved at application start time and all workloads submitted will use that image digest. If false, the custom image URI is resolved at the workload submission time.

Type: Boolean

Required: No

imageUri

The URI of an image in the Amazon ECR registry. This field is required when you create a new application. If you leave this field blank in an update, Amazon EMR will remove the image configuration.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: (`[0-9]{12}`)\.`dkr`\.`ecr`\.(`[a-z0-9-]+`).(`[a-z0-9._-]+`)\(/((?:`[a-z0-9-]+`(?:`[-._]``[a-z0-9-]+`)\*|/)\*`[a-z0-9-]+`(?:`[-._]``[a-z0-9-]+`)\*)(?::(`[a-zA-Z0-9_]+``[a-zA-Z0-9-._]*`)|@(`sha256:[0-9a-f]{64}`))

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InitialCapacityConfig

The initial capacity configuration per worker.

Contents

workerCount

The number of workers in the initial capacity configuration.

Type: Long

Valid Range: Minimum value of 1. Maximum value of 1000000.

Required: Yes

workerConfiguration

The resource configuration of the initial capacity configuration.

Type: [WorkerResourceConfig](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

InteractiveConfiguration

The configuration to use to enable the different types of interactive use cases in an application.

Contents

livyEndpointEnabled

Enables an Apache Livy endpoint that you can connect to and run interactive jobs.

Type: Boolean

Required: No

sessionEnabled

Enables interactive sessions on the application. When set to `true`, you can start interactive sessions using the `StartSession` operation.

Type: Boolean

Required: No

studioEnabled

Enables you to connect an application to Amazon EMR Studio to run interactive workloads in a notebook.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobDriver

The driver that the job runs on.

Contents

Important

This data type is a UNION, so only one of the following members can be specified when used or returned.

hive

The job driver parameters specified for Hive.

Type: [Hive](#) object

Required: No

sparkSubmit

The job driver parameters specified for Spark.

Type: [SparkSubmit](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobLevelCostAllocationConfiguration

The configuration object that enables job level cost allocation.

Contents

enabled

Enables job level cost allocation for the application.

Type: Boolean

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobRun

Information about a job run. A job run is a unit of work, such as a Spark JAR, Hive query, or SparkSQL query, that you submit to an Amazon EMR Serverless application.

Contents

applicationId

The ID of the application the job is running on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

arn

The execution role ARN of the job run.

Type: String

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.*:(\d{12}):\/applications\/[0-9a-zA-Z]+\/jobruns\/[0-9a-zA-Z]+`

Required: Yes

createdAt

The date and time when the job run was created.

Type: Timestamp

Required: Yes

createdBy

The user who created the job run.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):(iam|sts)::(\d{12})?:[\w/+=, .@-]+`

Required: Yes

executionRole

The execution role ARN of the job run.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):iam::([0-9]{12}):(role((\u002F|(\u002F[\u0021-\u007F]+\u002F))[\w+=, .@-]+)`

Required: Yes

jobDriver

The job driver for the job run.

Type: [JobDriver](#) object

Note: This object is a Union. Only one member of this object can be specified or returned.

Required: Yes

jobRunId

The ID of the job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

releaseLabel

The Amazon EMR release associated with the application your job is running on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [A-Za-z0-9._/-]+

Required: Yes

state

The state of the job run.

Type: String

Valid Values: SUBMITTED | PENDING | SCHEDULED | RUNNING | SUCCESS | FAILED | CANCELLING | CANCELLED | QUEUED

Required: Yes

stateDetails

The state details of the job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: .*\\S.*

Required: Yes

updatedAt

The date and time when the job run was updated.

Type: Timestamp

Required: Yes

attempt

The attempt of the job run.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

attemptCreatedAt

The date and time of when the job run attempt was created.

Type: Timestamp

Required: No

attemptUpdatedAt

The date and time of when the job run attempt was last updated.

Type: Timestamp

Required: No

billedResourceUtilization

The aggregate vCPU, memory, and storage that AWS has billed for the job run. The billed resources include a 1-minute minimum usage for workers, plus additional storage over 20 GB per worker. Note that billed resources do not include usage for idle pre-initialized workers.

Type: [ResourceUtilization](#) object

Required: No

configurationOverrides

The configuration settings that are used to override default configuration.

Type: [ConfigurationOverrides](#) object

Required: No

endedAt

The date and time when the job was terminated.

Type: Timestamp

Required: No

executionIamPolicy

Optional IAM policy. The resulting job IAM role permissions will be an intersection of the policies passed and the policy associated with your job execution role.

Type: [JobRunExecutionIamPolicy](#) object

Required: No

executionTimeoutMinutes

Returns the job run timeout value from the `StartJobRun` call. If you didn't specify a timeout, this value defaults to 720 minutes.

For BATCH mode job runs, the maximum value is 10080 minutes (7 days) starting with Amazon EMR release 7.11.

Type: Long

Valid Range: Minimum value of 0. Maximum value of 1000000.

Required: No

imageConfiguration

The applied image configuration.

Type: [ImageConfiguration](#) object

Required: No

mode

The mode of the job run.

Type: String

Valid Values: BATCH | STREAMING

Required: No

name

The optional job run name. This doesn't have to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `.*\S.*`

Required: No

networkConfiguration

The network configuration for customer VPC connectivity.

Type: [NetworkConfiguration](#) object

Required: No

queuedDurationMilliseconds

The total time for a job in the QUEUED state in milliseconds.

Type: Long

Required: No

retryPolicy

The retry policy of the job run.

Type: [RetryPolicy](#) object

Required: No

startedAt

The date and time when the job moved to the RUNNING state.

Type: Timestamp

Required: No

tags

The tags assigned to the job run.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 200 items.

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: `[A-Za-z0-9 /_ . : = + @ - ] +`

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: `[A-Za-z0-9 /_ . : = + @ - ] *`

Required: No

totalExecutionDurationSeconds

The job run total execution duration in seconds. This field is only available for job runs in a COMPLETED, FAILED, or CANCELLED state.

Type: Integer

Required: No

totalResourceUtilization

The aggregate vCPU, memory, and storage resources used from the time the job starts to execute, until the time the job terminates, rounded up to the nearest second.

Type: [TotalResourceUtilization](#) object

Required: No

workerTypeSpecifications

The specification applied to each worker type. Includes the JobRun-level ImageConfiguration when the applicationLevelDigestResolution is false for the application.

Type: String to [WorkerTypeSpecification](#) object map

Key Length Constraints: Minimum length of 1. Maximum length of 50.

Key Pattern: `[a-zA-Z]+[-_]*[a-zA-Z]+`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobRunAttemptSummary

The summary of attributes associated with a job run attempt.

Contents

applicationId

The ID of the application the job is running on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

arn

The Amazon Resource Name (ARN) of the job run.

Type: String

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.*:(\d{12}):/applications/[0-9a-zA-Z]+/jobruns/[0-9a-zA-Z]+`

Required: Yes

createdAt

The date and time when the job run attempt was created.

Type: Timestamp

Required: Yes

createdBy

The user who created the job run.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):(iam|sts)::(\d{12})?:[\w/+=, .@-]+`

Required: Yes

executionRole

The Amazon Resource Name (ARN) of the execution role of the job run..

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):iam::([0-9]{12}):(role(\u002F|(\u002F[\u0021-\u007F]+\u002F))[\w+=, .@-]+)`

Required: Yes

id

The ID of the job run attempt.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

jobCreatedAt

The date and time of when the job run was created.

Type: Timestamp

Required: Yes

releaseLabel

The Amazon EMR release label of the job run attempt.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: [A-Za-z0-9._/-]+

Required: Yes

state

The state of the job run attempt.

Type: String

Valid Values: SUBMITTED | PENDING | SCHEDULED | RUNNING | SUCCESS | FAILED | CANCELLING | CANCELLED | QUEUED

Required: Yes

stateDetails

The state details of the job run attempt.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: .*\\S.*

Required: Yes

updatedAt

The date and time of when the job run attempt was last updated.

Type: Timestamp

Required: Yes

attempt

The attempt number of the job run execution.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

mode

The mode of the job run attempt.

Type: String

Valid Values: BATCH | STREAMING

Required: No

name

The name of the job run attempt.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `.*\S.*`

Required: No

type

The type of the job run, such as Spark or Hive.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobRunExecutionIamPolicy

Optional IAM policy. The resulting job IAM role permissions will be an intersection of the policies passed and the policy associated with your job execution role.

Contents

policy

An IAM inline policy to use as an execution IAM policy.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 2048.

Pattern: ([-~]+)

Required: No

policyArns

A list of Amazon Resource Names (ARNs) to use as an execution IAM policy.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 10 items.

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: ([-~... -##-#0-#FF]+)

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

JobRunSummary

The summary of attributes associated with a job run.

Contents

applicationId

The ID of the application the job is running on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

arn

The ARN of the job run.

Type: String

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.+:(\d{12}):\applications\[0-9a-zA-Z]+\jobruns\[0-9a-zA-Z]+`

Required: Yes

createdAt

The date and time when the job run was created.

Type: Timestamp

Required: Yes

createdBy

The user who created the job run.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):(iam|sts)::(\d{12})?:[\w/+=, .@-]+`

Required: Yes

executionRole

The execution role ARN of the job run.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):iam::([0-9]{12}):(role((\u002F|(\u002F[\u0021-\u007F]+\u002F))[\w+=, .@-]+)`

Required: Yes

id

The ID of the job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

releaseLabel

The Amazon EMR release associated with the application your job is running on.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[A-Za-z0-9._/-]+`

Required: Yes

state

The state of the job run.

Type: String

Valid Values: SUBMITTED | PENDING | SCHEDULED | RUNNING | SUCCESS | FAILED | CANCELLING | CANCELLED | QUEUED

Required: Yes

stateDetails

The state details of the job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: .*\\S.*

Required: Yes

updatedAt

The date and time when the job run was last updated.

Type: Timestamp

Required: Yes

attempt

The attempt number of the job run execution.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

attemptCreatedAt

The date and time of when the job run attempt was created.

Type: Timestamp

Required: No

attemptUpdatedAt

The date and time of when the job run attempt was last updated.

Type: Timestamp

Required: No

mode

The mode of the job run.

Type: String

Valid Values: BATCH | STREAMING

Required: No

name

The optional job run name. This doesn't have to be unique.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: .*\\S.*

Required: No

type

The type of job run, such as Spark or Hive.

Type: String

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ManagedPersistenceMonitoringConfiguration

The managed log persistence configuration for a job run.

Contents

enabled

Enables managed logging and defaults to true. If set to false, managed logging will be turned off.

Type: Boolean

Required: No

encryptionKeyArn

The KMS key ARN to encrypt the logs stored in managed log persistence.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*) : kms : [a-zA-Z0-9-]* : ([0-9]{12}) : key\[a-zA-Z0-9-]+\`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MaximumAllowedResources

The maximum allowed cumulative resources for an application. No new resources will be created once the limit is hit.

Contents

cpu

The maximum allowed CPU for an application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 15.

Pattern: `[1-9][0-9]*(\s)?(vCPU|vcpu|VCPU)?`

Required: Yes

memory

The maximum allowed resources for an application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 15.

Pattern: `[1-9][0-9]*(\s)?(GB|gb|gB|Gb)?`

Required: Yes

disk

The maximum allowed disk for an application.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 15.

Pattern: `[1-9][0-9]*(\s)?(GB|gb|gB|Gb)`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

MonitoringConfiguration

The configuration setting for monitoring.

Contents

cloudWatchLoggingConfiguration

The Amazon CloudWatch configuration for monitoring logs. You can configure your jobs to send log information to CloudWatch.

Type: [CloudWatchLoggingConfiguration](#) object

Required: No

managedPersistenceMonitoringConfiguration

The managed log persistence configuration for a job run.

Type: [ManagedPersistenceMonitoringConfiguration](#) object

Required: No

prometheusMonitoringConfiguration

The monitoring configuration object you can configure to send metrics to Amazon Managed Service for Prometheus for a job run.

Type: [PrometheusMonitoringConfiguration](#) object

Required: No

s3MonitoringConfiguration

The Amazon S3 configuration for monitoring log publishing.

Type: [S3MonitoringConfiguration](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

NetworkConfiguration

The network configuration for customer VPC connectivity.

Contents

securityGroupIds

The array of security group Ids for customer VPC connectivity.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 5 items.

Length Constraints: Minimum length of 1. Maximum length of 32.

Pattern: `[-0-9a-zA-Z]+.*`

Required: No

subnetIds

The array of subnet Ids for customer VPC connectivity.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 16 items.

Length Constraints: Minimum length of 1. Maximum length of 32.

Pattern: `[-0-9a-zA-Z]+.*`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

PrometheusMonitoringConfiguration

The monitoring configuration object you can configure to send metrics to Amazon Managed Service for Prometheus for a job run.

Contents

remoteWriteUrl

The remote write URL in the Amazon Managed Service for Prometheus workspace to send metrics to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 10280.

Pattern: `https://aps-workspaces.([a-z]{2}-[a-z-]{1,20}-[1-9]).amazonaws([0-9A-Za-z]{2,4})+/workspaces/[-_.0-9A-Za-z]{1,100}/api/v1/remote_write`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

ResourceUtilization

The resource utilization for memory, storage, and vCPU for jobs.

Contents

memoryGBHour

The aggregated memory used per hour from the time the job starts executing until the job is terminated.

Type: Double

Required: No

storageGBHour

The aggregated storage used per hour from the time the job starts executing until the job is terminated.

Type: Double

Required: No

vCPUPerHour

The aggregated vCPU used per hour from the time the job starts executing until the job is terminated.

Type: Double

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

RetryPolicy

The retry policy to use for a job run.

Contents

maxAttempts

Maximum number of attempts for the job run. This parameter is only applicable for BATCH mode.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

maxFailedAttemptsPerHour

Maximum number of failed attempts per hour. This parameter is only applicable for STREAMING mode.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

S3MonitoringConfiguration

The Amazon S3 configuration for monitoring log publishing. You can configure your jobs to send log information to Amazon S3.

Contents

encryptionKeyArn

The KMS key ARN to encrypt the logs published to the given Amazon S3 destination.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):kms:[a-zA-Z0-9-]*:([0-9]{12}):key\[a-zA-Z0-9-]+\`

Required: No

logUri

The Amazon S3 destination URI for log publishing.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 10280.

Pattern: `.*[\u0020-\uD7FF\uE000-\uFFFD\uD800\uDBFF-\uDC00\uDFFF\u0000-\u001F\u0020-\u007F\u0080-\u00FF\u0100-\u017F\u0180-\u01FF\u0200-\u02FF\u0300-\u037F\u0380-\u03FF\u0400-\u047F\u0480-\u04FF\u0500-\u057F\u0580-\u05FF\u0600-\u06FF\u0700-\u077F\u0780-\u07FF\u0800-\u087F\u0880-\u08FF\u0900-\u097F\u0980-\u09FF\u0A00-\u0A7F\u0A80-\u0AFF\u0B00-\u0B7F\u0B80-\u0BFF\u0C00-\u0C7F\u0C80-\u0CFF\u0D00-\u0D7F\u0D80-\u0DBF\u0E00-\u0E7F\u0E80-\u0EFF\u0F00-\u0F7F\u0F80-\u0FFF\u1000-\u107F\u1080-\u10FF\u1100-\u117F\u1180-\u11FF\u1200-\u127F\u1280-\u12FF\u1300-\u137F\u1380-\u13FF\u1400-\u147F\u1480-\u14FF\u1500-\u157F\u1580-\u15FF\u1600-\u167F\u1680-\u16FF\u1700-\u177F\u1780-\u17FF\u1800-\u187F\u1880-\u18FF\u1900-\u197F\u1980-\u19FF\u1A00-\u1A7F\u1A80-\u1AFF\u1B00-\u1B7F\u1B80-\u1BFF\u1C00-\u1C7F\u1C80-\u1CFF\u1D00-\u1D7F\u1D80-\u1DBF\u1E00-\u1E7F\u1E80-\u1EFF\u1F00-\u1F7F\u1F80-\u1FFF\u2000-\u207F\u2080-\u20FF\u2100-\u217F\u2180-\u21FF\u2200-\u227F\u2280-\u22FF\u2300-\u237F\u2380-\u23FF\u2400-\u247F\u2480-\u24FF\u2500-\u257F\u2580-\u25FF\u2600-\u267F\u2680-\u26FF\u2700-\u277F\u2780-\u27FF\u2800-\u287F\u2880-\u28FF\u2900-\u297F\u2980-\u29FF\u2A00-\u2A7F\u2A80-\u2AFF\u2B00-\u2B7F\u2B80-\u2BFF\u2C00-\u2C7F\u2C80-\u2CFF\u2D00-\u2D7F\u2D80-\u2DBF\u2E00-\u2E7F\u2E80-\u2EFF\u2F00-\u2F7F\u2F80-\u2FFF\u3000-\u307F\u3080-\u30FF\u3100-\u317F\u3180-\u31FF\u3200-\u327F\u3280-\u32FF\u3300-\u337F\u3380-\u33FF\u3400-\u347F\u3480-\u34FF\u3500-\u357F\u3580-\u35FF\u3600-\u367F\u3680-\u36FF\u3700-\u377F\u3780-\u37FF\u3800-\u387F\u3880-\u38FF\u3900-\u397F\u3980-\u39FF\u3A00-\u3A7F\u3A80-\u3AFF\u3B00-\u3B7F\u3B80-\u3BFF\u3C00-\u3C7F\u3C80-\u3CFF\u3D00-\u3D7F\u3D80-\u3DBF\u3E00-\u3E7F\u3E80-\u3EFF\u3F00-\u3F7F\u3F80-\u3FFF\u4000-\u407F\u4080-\u40FF\u4100-\u417F\u4180-\u41FF\u4200-\u427F\u4280-\u42FF\u4300-\u437F\u4380-\u43FF\u4400-\u447F\u4480-\u44FF\u4500-\u457F\u4580-\u45FF\u4600-\u467F\u4680-\u46FF\u4700-\u477F\u4780-\u47FF\u4800-\u487F\u4880-\u48FF\u4900-\u497F\u4980-\u49FF\u4A00-\u4A7F\u4A80-\u4AFF\u4B00-\u4B7F\u4B80-\u4BFF\u4C00-\u4C7F\u4C80-\u4CFF\u4D00-\u4D7F\u4D80-\u4DBF\u4E00-\u4E7F\u4E80-\u4EFF\u4F00-\u4F7F\u4F80-\u4FFF\u5000-\u507F\u5080-\u50FF\u5100-\u517F\u5180-\u51FF\u5200-\u527F\u5280-\u52FF\u5300-\u537F\u5380-\u53FF\u5400-\u547F\u5480-\u54FF\u5500-\u557F\u5580-\u55FF\u5600-\u567F\u5680-\u56FF\u5700-\u577F\u5780-\u57FF\u5800-\u587F\u5880-\u58FF\u5900-\u597F\u5980-\u59FF\u5A00-\u5A7F\u5A80-\u5AFF\u5B00-\u5B7F\u5B80-\u5BFF\u5C00-\u5C7F\u5C80-\u5CFF\u5D00-\u5D7F\u5D80-\u5DBF\u5E00-\u5E7F\u5E80-\u5EFF\u5F00-\u5F7F\u5F80-\u5FFF\u6000-\u607F\u6080-\u60FF\u6100-\u617F\u6180-\u61FF\u6200-\u627F\u6280-\u62FF\u6300-\u637F\u6380-\u63FF\u6400-\u647F\u6480-\u64FF\u6500-\u657F\u6580-\u65FF\u6600-\u667F\u6680-\u66FF\u6700-\u677F\u6780-\u67FF\u6800-\u687F\u6880-\u68FF\u6900-\u697F\u6980-\u69FF\u6A00-\u6A7F\u6A80-\u6AFF\u6B00-\u6B7F\u6B80-\u6BFF\u6C00-\u6C7F\u6C80-\u6CFF\u6D00-\u6D7F\u6D80-\u6DBF\u6E00-\u6E7F\u6E80-\u6EFF\u6F00-\u6F7F\u6F80-\u6FFF\u7000-\u707F\u7080-\u70FF\u7100-\u717F\u7180-\u71FF\u7200-\u727F\u7280-\u72FF\u7300-\u737F\u7380-\u73FF\u7400-\u747F\u7480-\u74FF\u7500-\u757F\u7580-\u75FF\u7600-\u767F\u7680-\u76FF\u7700-\u777F\u7780-\u77FF\u7800-\u787F\u7880-\u78FF\u7900-\u797F\u7980-\u79FF\u7A00-\u7A7F\u7A80-\u7AFF\u7B00-\u7B7F\u7B80-\u7BFF\u7C00-\u7C7F\u7C80-\u7CFF\u7D00-\u7D7F\u7D80-\u7DBF\u7E00-\u7E7F\u7E80-\u7EFF\u7F00-\u7F7F\u7F80-\u7FFF\u8000-\u807F\u8080-\u80FF\u8100-\u817F\u8180-\u81FF\u8200-\u827F\u8280-\u82FF\u8300-\u837F\u8380-\u83FF\u8400-\u847F\u8480-\u84FF\u8500-\u857F\u8580-\u85FF\u8600-\u867F\u8680-\u86FF\u8700-\u877F\u8780-\u87FF\u8800-\u887F\u8880-\u88FF\u8900-\u897F\u8980-\u89FF\u8A00-\u8A7F\u8A80-\u8AFF\u8B00-\u8B7F\u8B80-\u8BFF\u8C00-\u8C7F\u8C80-\u8CFF\u8D00-\u8D7F\u8D80-\u8DBF\u8E00-\u8E7F\u8E80-\u8EFF\u8F00-\u8F7F\u8F80-\u8FFF\u9000-\u907F\u9080-\u90FF\u9100-\u917F\u9180-\u91FF\u9200-\u927F\u9280-\u92FF\u9300-\u937F\u9380-\u93FF\u9400-\u947F\u9480-\u94FF\u9500-\u957F\u9580-\u95FF\u9600-\u967F\u9680-\u96FF\u9700-\u977F\u9780-\u97FF\u9800-\u987F\u9880-\u98FF\u9900-\u997F\u9980-\u99FF\u9A00-\u9A7F\u9A80-\u9AFF\u9B00-\u9B7F\u9B80-\u9BFF\u9C00-\u9C7F\u9C80-\u9CFF\u9D00-\u9D7F\u9D80-\u9DBF\u9E00-\u9E7F\u9E80-\u9EFF\u9F00-\u9F7F\u9F80-\u9FFF\uA000-\uA07F\uA080-\uA0FF\uA100-\uA17F\uA180-\uA1FF\uA200-\uA27F\uA280-\uA2FF\uA300-\uA37F\uA380-\uA3FF\uA400-\uA47F\uA480-\uA4FF\uA500-\uA57F\uA580-\uA5FF\uA600-\uA67F\uA680-\uA6FF\uA700-\uA77F\uA780-\uA7FF\uA800-\uA87F\uA880-\uA8FF\uA900-\uA97F\uA980-\uA9FF\uAA00-\uAA7F\uAA80-\uAAFF\uAB00-\uAB7F\uAB80-\uABFF\uAC00-\uAC7F\uAC80-\uACFF\uAD00-\uAD7F\uAD80-\uADFF\uAE00-\uAE7F\uAE80-\uAEFF\uAF00-\uAF7F\uAF80-\uAFFF\uB000-\uB07F\uB080-\uB0FF\uB100-\uB17F\uB180-\uB1FF\uB200-\uB27F\uB280-\uB2FF\uB300-\uB37F\uB380-\uB3FF\uB400-\uB47F\uB480-\uB4FF\uB500-\uB57F\uB580-\uB5FF\uB600-\uB67F\uB680-\uB6FF\uB700-\uB77F\uB780-\uB7FF\uB800-\uB87F\uB880-\uB8FF\uB900-\uB97F\uB980-\uB9FF\uBA00-\uBA7F\uBA80-\uBAFF\uBB00-\uBB7F\uBB80-\uBBFF\uBC00-\uBC7F\uBC80-\uBCFF\uBD00-\uBD7F\uBD80-\uBDBF\uBE00-\uBE7F\uBE80-\uBEFF\uBF00-\uBF7F\uBF80-\uBFFF\uC000-\uC07F\uC080-\uC0FF\uC100-\uC17F\uC180-\uC1FF\uC200-\uC27F\uC280-\uC2FF\uC300-\uC37F\uC380-\uC3FF\uC400-\uC47F\uC480-\uC4FF\uC500-\uC57F\uC580-\uC5FF\uC600-\uC67F\uC680-\uC6FF\uC700-\uC77F\uC780-\uC7FF\uC800-\uC87F\uC880-\uC8FF\uC900-\uC97F\uC980-\uC9FF\uCA00-\uCA7F\uCA80-\uCAFF\uCB00-\uCB7F\uCB80-\uCBFF\uCC00-\uCC7F\uCC80-\uCCFF\uCD00-\uCD7F\uCD80-\uCDFF\uCE00-\uCE7F\uCE80-\uCEFF\uCF00-\uCF7F\uCF80-\uCFFF\uD000-\uD07F\uD080-\uD0FF\uD100-\uD17F\uD180-\uD1FF\uD200-\uD27F\uD280-\uD2FF\uD300-\uD37F\uD380-\uD3FF\uD400-\uD47F\uD480-\uD4FF\uD500-\uD57F\uD580-\uD5FF\uD600-\uD67F\uD680-\uD6FF\uD700-\uD77F\uD780-\uD7FF\uD800-\uD87F\uD880-\uD8FF\uD900-\uD97F\uD980-\uD9FF\uDA00-\uDA7F\uDA80-\uDAFF\uDB00-\uDB7F\uDB80-\uDBFF\uDC00-\uDC7F\uDC80-\uDCFF\uDD00-\uDD7F\uDD80-\uDDFF\uDE00-\uDE7F\uDE80-\uDEFF\uDF00-\uDF7F\uDF80-\uDFFF\uE000-\uE07F\uE080-\uE0FF\uE100-\uE17F\uE180-\uE1FF\uE200-\uE27F\uE280-\uE2FF\uE300-\uE37F\uE380-\uE3FF\uE400-\uE47F\uE480-\uE4FF\uE500-\uE57F\uE580-\uE5FF\uE600-\uE67F\uE680-\uE6FF\uE700-\uE77F\uE780-\uE7FF\uE800-\uE87F\uE880-\uE8FF\uE900-\uE97F\uE980-\uE9FF\uEA00-\uEA7F\uEA80-\uEAFF\uEB00-\uEB7F\uEB80-\uEBFF\uEC00-\uEC7F\uEC80-\uECFF\uED00-\uED7F\uED80-\uEDFF\uEE00-\uEE7F\uEE80-\uEEFF\uEF00-\uEF7F\uEF80-\uEFFF\uF000-\uF07F\uF080-\uF0FF\uF100-\uF17F\uF180-\uF1FF\uF200-\uF27F\uF280-\uF2FF\uF300-\uF37F\uF380-\uF3FF\uF400-\uF47F\uF480-\uF4FF\uF500-\uF57F\uF580-\uF5FF\uF600-\uF67F\uF680-\uF6FF\uF700-\uF77F\uF780-\uF7FF\uF800-\uF87F\uF880-\uF8FF\uF900-\uF97F\uF980-\uF9FF\uFA00-\uFA7F\uFA80-\uFAFF\uFB00-\uFB7F\uFB80-\uFBFF\uFC00-\uFC7F\uFC80-\uFCFF\uFD00-\uFD7F\uFD80-\uFDFF\uFE00-\uFE7F\uFE80-\uFEFF\uFF00-\uFF7F\uFF80-\uFFFF]`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SchedulerConfiguration

The scheduler configuration for batch and streaming jobs running on this application. Supported with release labels emr-7.0.0 and above.

Contents

maxConcurrentRuns

The maximum concurrent job runs on this application. If scheduler configuration is enabled on your application, the default value is 15. The valid range is 1 to 1000.

Type: Integer

Required: No

queueTimeoutMinutes

The maximum duration in minutes for the job in QUEUED state. If scheduler configuration is enabled on your application, the default value is 360 minutes (6 hours). The valid range is from 15 to 720.

Type: Integer

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Session

Information about a session, including the session state, configuration, and timestamps.

Contents

applicationId

The ID of the application that the session belongs to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

arn

The Amazon Resource Name (ARN) of the session.

Type: String

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.*:(\d{12}):\/applications\/[0-9a-zA-Z]+\/sessions\/[0-9a-zA-Z]+`

Required: Yes

createdAt

The date and time that the session was created.

Type: Timestamp

Required: Yes

createdBy

The IAM principal that created the session.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):(iam|sts)::(\d{12})?:[\w/+=, .@-]+`

Required: Yes

executionRoleArn

The Amazon Resource Name (ARN) of the execution role for the session.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):iam::([0-9]{12}):(role((\u002F|(\u002F[\u0021-\u007F]+\u002F))[\w+=, .@-]+)`

Required: Yes

releaseLabel

The Amazon EMR release label associated with the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[A-Za-z0-9._/-]+`

Required: Yes

sessionId

The ID of the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

state

The state of the session.

Type: String

Valid Values: SUBMITTED | STARTING | STARTED | IDLE | BUSY | FAILED |
TERMINATING | TERMINATED

Required: Yes

stateDetails

Additional details about the current state of the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: .*\\S.*

Required: Yes

updatedAt

The date and time that the session was last updated.

Type: Timestamp

Required: Yes

billedResourceUtilization

The aggregate vCPU, memory, and storage that AWS has billed for the session. The billed resources include a 1-minute minimum usage for workers, plus additional storage over 20 GB per worker. Note that billed resources do not include usage for idle pre-initialized workers.

Type: [ResourceUtilization](#) object

Required: No

configurationOverrides

The configuration overrides for the session, including runtime configuration properties.

Type: [SessionConfigurationOverrides](#) object

Required: No

endedAt

The date and time that the session was terminated or failed.

Type: Timestamp

Required: No

idleSince

The date and time that the session became idle.

Type: Timestamp

Required: No

idleTimeoutMinutes

The idle timeout in minutes for the session. After the session remains idle for this duration, it is automatically terminated.

Type: Long

Valid Range: Minimum value of 0. Maximum value of 1000000.

Required: No

name

The optional name of the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `.*\S.*`

Required: No

networkConfiguration

The network configuration for customer VPC connectivity for the session.

Type: [NetworkConfiguration](#) object

Required: No

startedAt

The date and time that the session moved to a running state.

Type: Timestamp

Required: No

tags

The tags assigned to the session.

Type: String to string map

Map Entries: Minimum number of 0 items. Maximum number of 200 items.

Key Length Constraints: Minimum length of 1. Maximum length of 128.

Key Pattern: [A-Za-z0-9 /_ . : = + @ -] +

Value Length Constraints: Minimum length of 0. Maximum length of 256.

Value Pattern: [A-Za-z0-9 /_ . : = + @ -] *

Required: No

totalExecutionDurationSeconds

The total execution duration of the session in seconds.

Type: Long

Required: No

totalResourceUtilization

The aggregate vCPU, memory, and storage resources used from the time the session starts to execute, until the time the session terminates, rounded up to the nearest second.

Type: [TotalResourceUtilization](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SessionConfigurationOverrides

The configuration overrides for a session.

Contents

runtimeConfiguration

The runtime configuration for the session. Contains Spark configuration properties specified at session creation time.

Type: Array of [Configuration](#) objects

Array Members: Minimum number of 0 items. Maximum number of 100 items.

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SessionSummary

The summary of attributes associated with a session.

Contents

applicationId

The ID of the application that the session belongs to.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

arn

The Amazon Resource Name (ARN) of the session.

Type: String

Length Constraints: Minimum length of 60. Maximum length of 1024.

Pattern: `arn:(aws[a-zA-Z0-9-]*):emr-serverless:.*:(\d{12}):\./applications\/[0-9a-zA-Z]+\./sessions\/[0-9a-zA-Z]+`

Required: Yes

createdAt

The date and time that the session was created.

Type: Timestamp

Required: Yes

createdBy

The IAM principal that created the session.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):(iam|sts)::(\d{12})?:[\w/+=, .@-]+`

Required: Yes

executionRoleArn

The Amazon Resource Name (ARN) of the execution role for the session.

Type: String

Length Constraints: Minimum length of 20. Maximum length of 2048.

Pattern: `arn:(aws[a-zA-Z0-9-]*):iam::([0-9]{12}):(role((\u002F|(\u002F[\u0021-\u007F]+\u002F))[\w+=, .@-]+)`

Required: Yes

releaseLabel

The Amazon EMR release label associated with the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[A-Za-z0-9._/-]+`

Required: Yes

sessionId

The ID of the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 64.

Pattern: `[0-9a-z]+`

Required: Yes

state

The state of the session.

Type: String

Valid Values: SUBMITTED | STARTING | STARTED | IDLE | BUSY | FAILED |
TERMINATING | TERMINATED

Required: Yes

stateDetails

Additional details about the current state of the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: .*\\S.*

Required: Yes

updatedAt

The date and time that the session was last updated.

Type: Timestamp

Required: Yes

name

The optional name of the session.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: .*\\S.*

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)

- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

SparkSubmit

The configurations for the Spark submit job driver.

Contents

entryPoint

The entry point for the Spark submit job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 4096.

Pattern: `.*\S.*`

Required: Yes

entryPointArguments

The arguments for the Spark submit job run.

Type: Array of strings

Array Members: Minimum number of 0 items. Maximum number of 1024 items.

Length Constraints: Minimum length of 0.

Required: No

sparkSubmitParameters

The parameters for the Spark submit job run.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 102400.

Pattern: `.*\S.*`

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

TotalResourceUtilization

The aggregate vCPU, memory, and storage resources used from the time job start executing till the time job is terminated, rounded up to the nearest second.

Contents

memoryGBHour

The aggregated memory used per hour from the time job start executing till the time job is terminated.

Type: Double

Required: No

storageGBHour

The aggregated storage used per hour from the time job start executing till the time job is terminated.

Type: Double

Required: No

vCPUPerHour

The aggregated vCPU used per hour from the time job start executing till the time job is terminated.

Type: Double

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)

- [AWS SDK for Ruby V3](#)

WorkerResourceConfig

The cumulative configuration requirements for every worker instance of the worker type.

Contents

cpu

The CPU requirements for every worker instance of the worker type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 15.

Pattern: `[1-9][0-9]*(\s)?(vCPU|vcpu|VCPU)?`

Required: Yes

memory

The memory requirements for every worker instance of the worker type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 15.

Pattern: `[1-9][0-9]*(\s)?(GB|gb|gB|Gb)?`

Required: Yes

disk

The disk requirements for every worker instance of the worker type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 15.

Pattern: `[1-9][0-9]*(\s)?(GB|gb|gB|Gb)`

Required: No

diskType

The disk type for every worker instance of the work type. Shuffle optimized disks have higher performance characteristics and are better for shuffle heavy workloads. Default is STANDARD.

Type: String

Pattern: (SHUFFLE_OPTIMIZED|[Ss]huffle_[Oo]ptimized|STANDARD|[Ss]tandard)

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

WorkerTypeSpecification

The specifications for a worker type.

Contents

imageConfiguration

The image configuration for a worker type.

Type: [ImageConfiguration](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

WorkerTypeSpecificationInput

The specifications for a worker type.

Contents

imageConfiguration

The image configuration for a worker type.

Type: [ImageConfigurationInput](#) object

Required: No

See Also

For more information about using this API in one of the language-specific AWS SDKs, see the following:

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

Common Parameters

The following list contains the parameters that all actions use for signing Signature Version 4 requests with a query string. Any action-specific parameters are listed in the topic for that action. For more information about Signature Version 4, see [Signing AWS API requests](#) in the *IAM User Guide*.

X-Amz-Algorithm

The hash algorithm that you used to create the request signature.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Valid Values: AWS4-HMAC-SHA256

Required: Conditional

X-Amz-Credential

The credential scope value, which is a string that includes your access key, the date, the region you are targeting, the service you are requesting, and a termination string ("aws4_request"). The value is expressed in the following format: *access_key/YYYYMMDD/region/service/aws4_request*.

For more information, see [Create a signed AWS API request](#) in the *IAM User Guide*.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

X-Amz-Date

The date that is used to create the signature. The format must be ISO 8601 basic format (YYYYMMDD'T'HHMMSS'Z'). For example, the following date time is a valid X-Amz-Date value: 20120325T120000Z.

Condition: X-Amz-Date is optional for all requests; it can be used to override the date used for signing requests. If the Date header is specified in the ISO 8601 basic format, X-Amz-Date is not required. When X-Amz-Date is used, it always overrides the value of the Date header. For more information, see [Elements of an AWS API request signature](#) in the *IAM User Guide*.

Type: string

Required: Conditional

X-Amz-Security-Token

The temporary security token that was obtained through a call to AWS Security Token Service (AWS STS). For a list of services that support temporary security credentials from AWS STS, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Condition: If you're using temporary security credentials from AWS STS, you must include the security token.

Type: string

Required: Conditional

X-Amz-Signature

Specifies the hex-encoded signature that was calculated from the string to sign and the derived signing key.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

X-Amz-SignedHeaders

Specifies all the HTTP headers that were included as part of the canonical request. For more information about specifying signed headers, see [Create a signed AWS API request](#) in the *IAM User Guide*.

Condition: Specify this parameter when you include authentication information in a query string instead of in the HTTP authorization header.

Type: string

Required: Conditional

Common Error Types

This section lists common error types that this AWS service may return. Not all services return all error types listed here. For errors specific to an API action for this service, see the topic for that API action.

AccessDeniedException

You don't have permission to perform this action. Verify that your IAM policy includes the required permissions.

HTTP Status Code: 403

ExpiredTokenException

The security token included in the request has expired. Request a new security token and try again.

HTTP Status Code: 403

IncompleteSignature

The request signature doesn't conform to AWS standards. Verify that you're using valid AWS credentials and that your request is properly formatted. If you're using an SDK, ensure it's up to date.

HTTP Status Code: 403

InternalFailure

The request can't be processed right now because of an internal server issue. Try again later. If the problem persists, contact AWS Support.

HTTP Status Code: 500

MalformedHttpRequestException

The request body can't be processed. This typically happens when the request body can't be decompressed using the specified content encoding algorithm. Verify that the content encoding header matches the compression format used.

HTTP Status Code: 400

NotAuthorized

You don't have permissions to perform this action. Verify that your IAM policy includes the required permissions.

HTTP Status Code: 401

OptInRequired

Your AWS account needs a subscription for this service. Verify that you've enabled the service in your account.

HTTP Status Code: 403

RequestAbortedException

The request was aborted before a response could be returned. This typically happens when the client closes the connection.

HTTP Status Code: 400

RequestEntityTooLargeException

The request entity is too large. Reduce the size of the request body and try again.

HTTP Status Code: 413

RequestTimeoutException

The request timed out. The server didn't receive the complete request within the expected time frame. Try again.

HTTP Status Code: 408

ServiceUnavailable

The service is temporarily unavailable. Try again later.

HTTP Status Code: 503

ThrottlingException

Your request rate is too high. The AWS SDKs automatically retry requests that receive this exception. Reduce the frequency of requests.

HTTP Status Code: 400

UnknownOperationException

The action or operation isn't recognized. Verify that the action name is spelled correctly and that it's supported by the API version you're using.

HTTP Status Code: 404

UnrecognizedClientException

The X.509 certificate or AWS access key ID you provided doesn't exist in our records. Verify that you're using valid credentials and that they haven't expired.

HTTP Status Code: 403

ValidationError

The input doesn't meet the required format or constraints. Check that all required parameters are included and that values are valid.

HTTP Status Code: 400