aws

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 November 11, 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.

**InternalServerException**

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

**InternalServerException**

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.

**InternalServerException**

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.

**InternalServerException**

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 optimal 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.

**InternalServerException**

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 optimal 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.

### InternalServerException

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.

**InternalServerException**

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,
      "vCPUHour": 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,
      "vCPUHour": 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.

**InternalServerException**

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.

**InternalServerException**

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.

**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

# 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.

**InternalServerException**

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=n
 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.

**InternalServerException**

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.

**InternalServerException**

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.

**InternalServerException**

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

**InternalServerException**

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

**InternalServerException**

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.

**InternalServerException**

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\/[0-9a-zA-Z]+)?`

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.

**InternalServerException**

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.

**InternalServerException**

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}):\/applications\/[0-9a-zA-Z]+(\/jobruns\/[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.

**InternalServerException**

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.

**InternalServerException**

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]+`

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]+`

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: `([ -ÿ]+)`

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

**vCPUHour**

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 [arameter 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\r\n\t]*.*`

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

**vCPUHour**

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