aws

Amazon MSK API Reference

# Amazon Managed Streaming for Apache Kafka



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

# Amazon Managed Streaming for Apache Kafka: Amazon MSK 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

**API Reference ..... 1**
**Operations ..... 2**
**Resources ..... 13**
Apache Kafka Versions ..... 14
URI ..... 14
HTTP methods ..... 14
Schemas ..... 16
Properties ..... 17
See also ..... 18
BootstrapBrokers ..... 19
URI ..... 19
HTTP methods ..... 19
Schemas ..... 21
Properties ..... 22
See also ..... 26
Broker Count ..... 26
URI ..... 26
HTTP methods ..... 26
Schemas ..... 29
Properties ..... 30
See also ..... 31
Broker Storage ..... 32
URI ..... 32
HTTP methods ..... 32
Schemas ..... 34
Properties ..... 35
See also ..... 38
Broker Type ..... 38
URI ..... 38
HTTP methods ..... 38
Schemas ..... 40
Properties ..... 41
See also ..... 42
Brokers ..... 43

URI ... 43
HTTP methods ... 43
Schemas ... 47
Properties ... 48
See also ... 53
Cluster ... 54
URI ... 54
HTTP methods ... 54
Schemas ... 58
Properties ... 62
See also ... 85
Cluster Apache Kafka Version ... 86
URI ... 86
HTTP methods ... 87
Schemas ... 89
Properties ... 89
See also ... 91
Cluster Configuration ... 92
URI ... 92
HTTP methods ... 92
Schemas ... 94
Properties ... 95
See also ... 97
Cluster Connectivity ... 97
URI ... 97
HTTP methods ... 98
Schemas ... 100
Properties ... 101
See also ... 106
Cluster Operation ... 106
URI ... 106
HTTP methods ... 106
Schemas ... 108
Properties ... 115
See also ... 139
Cluster Operations ... 140

URI ... 140
HTTP methods ... 140
Schemas ... 143
Properties ... 149
See also ... 174
Cluster Security ... 174
URI ... 174
HTTP methods ... 175
Schemas ... 177
Properties ... 178
See also ... 184
Clusters ... 184
URI ... 184
HTTP methods ... 185
Schemas ... 189
Properties ... 196
See also ... 224
Clusters clusterArn Channels ... 225
URI ... 225
HTTP methods ... 225
Schemas ... 229
Properties ... 232
See also ... 249
Clusters clusterArn Channels channelArn ... 250
URI ... 250
HTTP methods ... 250
Schemas ... 255
Properties ... 258
See also ... 276
Clusters clusterArn Client-vpc-connection ... 277
URI ... 277
HTTP methods ... 277
Schemas ... 279
Properties ... 280
See also ... 281
Clusters clusterArn Client-vpc-connections ... 281

URI ... 281
HTTP methods ... 281
Schemas ... 284
Properties ... 285
See also ... 287
Clusters clusterArn Policy ... 288
URI ... 288
HTTP methods ... 288
Schemas ... 292
Properties ... 293
See also ... 295
Clusters clusterArn Rebalancing ... 296
URI ... 296
HTTP methods ... 297
Schemas ... 299
Properties ... 299
See also ... 301
Compatible Apache Kafka Versions ... 302
URI ... 302
HTTP methods ... 302
Schemas ... 304
Properties ... 304
See also ... 306
Configuration ... 306
URI ... 306
HTTP methods ... 306
Schemas ... 311
Properties ... 313
See also ... 317
Configuration Revision ... 318
URI ... 318
HTTP methods ... 318
Schemas ... 321
Properties ... 321
See also ... 323
Configuration Revisions ... 323

URI ... 323
HTTP methods ... 324
Schemas ... 326
Properties ... 327
See also ... 329
Configurations ... 329
URI ... 329
HTTP methods ... 329
Schemas ... 333
Properties ... 335
See also ... 340
Monitoring Properties ... 341
URI ... 341
HTTP methods ... 341
Schemas ... 344
Properties ... 346
See also ... 352
Reboot Broker ... 353
URI ... 353
HTTP methods ... 353
Schemas ... 355
Properties ... 356
See also ... 357
Scram Secrets ... 358
URI ... 358
HTTP methods ... 358
Schemas ... 364
Properties ... 365
See also ... 368
Tags ... 370
URI ... 370
HTTP methods ... 370
Schemas ... 376
Properties ... 376
See also ... 377
Topic ... 379

URI ... 379
HTTP methods ... 379
Schemas ... 384
Properties ... 385
See also ... 389
Topic Partitions ... 390
URI ... 390
HTTP methods ... 390
Schemas ... 393
Properties ... 394
See also ... 396
Topics ... 396
URI ... 396
HTTP methods ... 396
Schemas ... 401
Properties ... 402
See also ... 406
UpdateStorage ... 407
URI ... 407
HTTP methods ... 407
Schemas ... 409
Properties ... 410
See also ... 412
Vpc-connection ... 413
URI ... 413
HTTP methods ... 413
Schemas ... 414
Properties ... 415
See also ... 419
Vpc-connection arn ... 420
URI ... 420
HTTP methods ... 420
Schemas ... 423
Properties ... 424
See also ... 427
Vpc-connections ... 428

URI ..... 428
HTTP methods ..... 428
Schemas ..... 430
Properties ..... 431
See also ..... 434
**Document History ..... 435**
**AWS Glossary ..... 436**

# Amazon Managed Streaming for Apache Kafka API Reference

Amazon Managed Streaming for Apache Kafka (Amazon MSK) is a fully managed service that makes it easy for you to build and run applications that use Apache Kafka to process streaming data.

Amazon MSK provides the control-plane operations and lets you use Apache Kafka data-plane operations, such as producing and consuming data. It runs open-source versions of Apache Kafka, so existing applications, tooling, and plugins from partners and the Apache Kafka community are supported without requiring changes to application code.

For more information about Amazon MSK, see the *Amazon MSK Developer Guide*.

**MSK Replicator API**

For information about MSK Replicator APIs, see the Amazon MSK Replicator API Reference.

# Operations

The Amazon Managed Streaming for Apache Kafka REST API includes the following operations.

- BatchAssociateScramSecret

  Associates a list of SCRAM secrets with a cluster. SCRAM secrets are stored in the Amazon Secrets Manager service, and are used to authenticate clients using sign-in credentials. You can associate up to 10 secrets with a cluster at a time.
- BatchDisassociateScramSecret

  Disassociates a list of SCRAM secrets from a cluster. SCRAM secrets are stored in the Amazon Secrets Manager service, and are used to authenticate clients using sign-in credentials. You can disassociate up to 10 secrets from a cluster at a time.
- CreateChannel

  Creates a channel using the specified properties in the request on the cluster specified by the Amazon Resource Name (ARN) in the request.
- CreateCluster

  Creates a new MSK cluster. The following Python 3.6 examples shows how you can create a cluster that's distributed over two Availability Zones. Before you run this Python script, replace the example subnet and security-group IDs with the IDs of your subnets and security group. When you create an MSK cluster, its brokers get evenly distributed over a number of Availability Zones that's equal to the number of subnets that you specify in the `BrokerNodeGroupInfo` parameter. In this example, you can add a third subnet to get a cluster that's distributed over three Availability Zones.

```
import boto3

client = boto3.client('kafka')

response = client.create_cluster(
    BrokerNodeGroupInfo={
        'BrokerAZDistribution': 'DEFAULT',
        'ClientSubnets': [
            'subnet-012345678901fedcba',
            'subnet-9876543210abcdef01'
        ],
```

```
        'InstanceType': 'kafka.m5.large',
        'SecurityGroups': [
            'sg-012345abcdef789789'
        ]
    },
    ClusterName='SalesCluster',
    EncryptionInfo={
        'EncryptionInTransit': {
            'ClientBroker': 'TLS_PLAINTEXT',
            'InCluster': True
        }
    },
    EnhancedMonitoring='PER_TOPIC_PER_BROKER',
    KafkaVersion='2.2.1',
    NumberOfBrokerNodes=2
)

print(response)
```

- CreateConfiguration

  Creates a new MSK configuration. To see an example of how to use this operation, first save the following text to a file and name the file `config-file.txt`.

  ```
  auto.create.topics.enable = true

  zookeeper.connection.timeout.ms = 1000

  log.roll.ms = 604800000
  ```

  Now run the following Python 3.6 script in the folder where you saved `config-file.txt`. This script uses the properties specified in `config-file.txt` to create a configuration named `SalesClusterConfiguration`. This configuration can work with Apache Kafka versions 1.1.1 and 2.1.0.

  ```
  import boto3

  client = boto3.client('kafka')

  config_file = open('config-file.txt', 'r')

  server_properties = config_file.read()
  ```

```
response = client.create_configuration(
    Name='SalesClusterConfiguration',
    Description='The configuration to use on all sales clusters.',
    KafkaVersions=['1.1.1', '2.1.0'],
    ServerProperties=server_properties
)

print(response)
```

- CreateTopic

  Creates a topic using the specified properties in the request on the cluster specified by the Amazon Resource Name (ARN) in the request.

- CreateVpcConnection

  Create remote VPC connection.

- DeleteChannel

  Deletes the channel specified by the channelArn in the request from the cluster specified by the Amazon Resource Name (ARN) in the request.

- DeleteCluster

  Deletes the MSK cluster specified by the Amazon Resource Name (ARN) in the request, and all its revisions.

- DeleteClusterPolicy

  Delete cluster policy.

- DeleteConfiguration

  Deletes a cluster configuration and all its revisions.

- DeleteTopic

  Deletes the topic specified by the topicName in the request from the cluster specified by the Amazon Resource Name (ARN) in the request.

- DeleteVpcConnection

  Delete remote VPC connection.

- DescribeChannel

Returns details for a channel resource.

- DescribeCluster

  Returns a description of the MSK cluster whose Amazon Resource Name (ARN) is specified in the request. The following is a Python 3.6 example of how to use this operation. Before you run this Python script, replace the example cluster Amazon Resource Name (ARN) with the ARN of the cluster you want to describe. If you don't know the ARN of the cluster, you can use the `ListClusters` operation to list all the clusters and see their ARNs and full descriptions.

```
import boto3

client = boto3.client('kafka')

response = client.describe_cluster(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-abcd-cafe-abab-9876543210ab-4'
)

print(response)
```

  Note that the response to this operation only includes the `ZookeeperConnectStringTls` node in clusters created with Apache Kafka version 2.5.1 and later.

- DescribeClusterOperation

  Returns a description of the cluster operation specified by the Amazon Resource Name (ARN).

- DescribeConfiguration

  Returns a description of this MSK configuration.

- DescribeConfigurationRevision

  Returns a description of this revision of the configuration.

- DescribeTopic

  Returns details for a topic on a cluster.

  This API response reflects data that updates approximately every minute. For the most current topic state after making changes, allow approximately one minute before querying.

- DescribeTopicPartitions

Returns all partition information for a topic on a cluster.

This API response reflects data that updates approximately every minute. For the most current topic state after making changes, allow approximately one minute before querying.

- DescribeVpcConnection

  Describes Remote VPC Connection.

- GetBootstrapBrokers

  A list of brokers that a client can use to bootstrap. This list doesn't necessarily include all of the brokers in the cluster. The following Python 3.6 example shows how you can use the Amazon Resource Name (ARN) of a cluster to get its bootstrap brokers. If you don't know the ARN of your cluster, you can use the `ListClusters` operation to get the ARNs of all the clusters in this account and Region.

```
import boto3

client = boto3.client('kafka')

response = client.get_bootstrap_brokers(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4',
)

print(response['BootstrapBrokerString'])
```

- GetClusterPolicy

  Get cluster policy.

- GetCompatibleKafkaVersions

  Returns a list of the Apache Kafka versions to which you can update this cluster.

- ListChannels

  Returns all channels in a cluster.

- ListClientVpcConnections

  List client VPC connections.

- ListClusterOperations

Returns a list of all the operations that have been performed on the specified MSK cluster.

- ListClusters

  Returns a list of all the MSK clusters.

- ListConfigurationRevisions

  Returns a list of all the revisions of an MSK configuration.

- ListConfigurations

  Returns a list of all the MSK configurations.

- ListKafkaVersions

  Returns the Apache Kafka version objects.

- ListNodes

  Returns a list of the broker nodes in the cluster. The following Python 3.6 example first lists one node of a cluster. Because the cluster has more nodes, the response contains a token that the script then uses to list the remaining nodes.

```
import boto3

client = boto3.client('kafka')

list_nodes_response = client.list_nodes(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4',
    MaxResults=1
)

print('\n')
print('Here is the first node in the list:')
print('\n')
print(list_nodes_response['NodeInfoList'])

next_token = list_nodes_response['NextToken']

list_nodes_response = client.list_nodes(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4',
    NextToken=next_token
```

```
)

print('\n')
print('Here are the remaining nodes in the list:')
print('\n')
print(list_nodes_response['NodeInfoList'])
```

- ListScramSecrets

  Returns a list of SCRAM secrets associated with the cluster. SCRAM secrets are stored in the Amazon Secrets Manager service, and are used to authenticate clients using sign-in credentials.

- ListTagsForResource

  Returns a list of the tags associated with the specified resource.

- ListTopics

  Returns all topics in a cluster.

  This API response reflects data that updates approximately every minute. For the most current topic state after making changes, allow approximately one minute before querying.

- ListVpcConnections

  Lists all VPC connections.

- PutClusterPolicy

  Create or update cluster policy.

- RebootBroker

  Reboots a broker. In a given cluster, you can reboot one broker at a time.

  To reboot a broker, wait for the cluster status to be ACTIVE. This operation returns an error if you invoke it while the cluster status is HEALING. You must wait for the status to change from HEALING to ACTIVE before you reboot the broker.

- RejectClientVpcConnection

  Reject client VPC connection.

- TagResource

  Adds tags to the specified MSK resource.

- UntagResource

Removes the tags associated with the keys that are provided in the query.

- UpdateBrokerCount

  Updates the number of broker nodes in the cluster. You can use this operation to increase or decrease the number of brokers in an existing cluster.

  The following Python 3.6 example shows how you can increase the number of brokers in a cluster to 6 brokers. The update operation returns immediately, with a response that includes the Amazon Resource Name (ARN) that Amazon MSK assigns to this cluster operation. You can use that ARN to check the state of the operation. When the state changes from `PENDING` to `UPDATE_COMPLETE`, the operation is complete.

```
import boto3
import time

client = boto3.client('kafka')

update_broker_count_response = client.update_broker_count(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4',
    CurrentVersion='K12V3IB1VIZHHY',
    TargetNumberOfBrokerNodes=6
)

operation_arn = update_broker_count_response['ClusterOperationArn']

print(operation_arn)

describe_cluster_operation_response =
 client.describe_cluster_operation(ClusterOperationArn=operation_arn)
operation_state = describe_cluster_operation_response['ClusterOperationInfo']
['OperationState']
print(operation_state)

expanded = False

while not expanded:
    print('Sleeping for 15 seconds before checking to see if the cluster update is
 done...')
    time.sleep(15)
```

```
    describe_cluster_operation_response =
 client.describe_cluster_operation(ClusterOperationArn=operation_arn)
    operation_state = describe_cluster_operation_response['ClusterOperationInfo']
['OperationState']
    if 'UPDATE_COMPLETE' == operation_state:
        expanded = True
        print('The cluster has 6 brokers now.')
```

- UpdateBrokerStorage

  Updates the EBS storage associated with Amazon MSK brokers. You can increase the amount of EBS storage per broker. You can't decrease the storage. To increase storage, wait for the cluster to be in the ACTIVE state. Storage volumes remain available during this scaling-up operation.

- UpdateBrokerType

  For information about this operation, see Updating the broker type in the developer guide.

- UpdateChannel

  Updates the channel specified by the channelArn in the request from the cluster specified by the Amazon Resource Name (ARN) in the request.

- UpdateClusterConfiguration

  Updates the cluster with the configuration that is specified in the request body. Before you invoke this operation, ensure that the number of partitions per broker on your MSK cluster is under the limits described in Number of partitions per broker. You can't update the configuration of an MSK cluster that exceeds these limits.

- UpdateClusterKafkaVersion

  Updates the cluster to the specified Apache Kafka version. Before you invoke this operation, ensure that the number of partitions per broker on your MSK cluster is under the limits described in Number of partitions per broker. You can't update the Apache Kafka version for an MSK cluster that exceeds these limits.

- UpdateConfiguration

  Creates a new revision of the cluster configuration. The configuration must be in the `ACTIVE` state.

- UpdateConnectivity

  Updates the connectivity setting for the cluster.

- UpdateMonitoring

  Updates the monitoring settings for the cluster. You can use this operation to specify which Apache Kafka metrics you want Amazon MSK to send to Amazon CloudWatch. You can also specify settings for open monitoring with Prometheus. The following Python 3.6 example enables open monitoring with the Node Exporter. It also sets enhanced monitoring to PER_BROKER. For more information about monitoring, see Monitoring.

```
import boto3
import time

client = boto3.client('kafka')

update_monitoring_response = client.update_monitoring(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4',
    CurrentVersion='K12V3IB1VIZHHY',
    EnhancedMonitoring='PER_BROKER',
    OpenMonitoring={"Prometheus":{"JmxExporter":
{"EnabledInBroker":False},"NodeExporter":{"EnabledInBroker":True}}}
)

operation_arn = update_monitoring_response['ClusterOperationArn']
print('The ARN of the update operation is ' + operation_arn)

describe_cluster_operation_response =
 client.describe_cluster_operation(ClusterOperationArn=operation_arn)

operation_state = describe_cluster_operation_response['ClusterOperationInfo']
['OperationState']
print('The status of the update operation is ' + operation_state)

updated = False

while not updated:
    print('Sleeping for 15 seconds before checking to see if the monitoring update is
 done...')
    time.sleep(15)
    describe_cluster_operation_response =
 client.describe_cluster_operation(ClusterOperationArn=operation_arn)
    operation_state = describe_cluster_operation_response['ClusterOperationInfo']
['OperationState']
    if 'UPDATE_COMPLETE' == operation_state:
```

```
        updated = True
        print('You have successfully updated the monitoring settings.')
```

- UpdateRebalancing

  Use this resource to update the intelligent rebalancing status of an Amazon MSK Provisioned cluster with Express brokers.

- UpdateSecurity
- UpdateStorage
- UpdateTopic

  Updates topic partition or topic configs.

# Resources

The Amazon Managed Streaming for Apache Kafka REST API includes the following resources.

**Topics**


- Apache Kafka Versions
- BootstrapBrokers
- Broker Count
- Broker Storage
- Broker Type
- Brokers
- Cluster
- Cluster Apache Kafka Version
- Cluster Configuration
- Cluster Connectivity
- Cluster Operation
- Cluster Operations
- Cluster Security
- Clusters
- Clusters clusterArn Channels
- Clusters clusterArn Channels channelArn
- Clusters clusterArn Client-vpc-connection
- Clusters clusterArn Client-vpc-connections
- Clusters clusterArn Policy
- Clusters clusterArn Rebalancing
- Compatible Apache Kafka Versions
- Configuration
- Configuration Revision
- Configuration Revisions
- Configurations
- Monitoring Properties

- Reboot Broker
- Scram Secrets
- Tags
- Topic
- Topic Partitions
- Topics
- UpdateStorage
- Vpc-connection
- Vpc-connection arn
- Vpc-connections

# Apache Kafka Versions

Objects that represent Apache Kafka versions.

## URI

`/v1/kafka-versions`

## HTTP methods

### GET

**Operation ID:** `ListKafkaVersions`

Returns the Apache Kafka version objects.

**Query parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| `nextToken` | String | False | The paginated results marker. When the result of the operation is truncated, the call |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | returns `NextToken` in the response. To get the next batch, provide this token in your next request. |
| `maxResults` | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes a `NextToken` parameter. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | `ListKafkaVersionsResponse` | 200 response |
| 400 | `Error` | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | `Error` | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | `Error` | Access forbidden. Check your credentials and then retry your request. |

| Status code | Response model | Description |
| --- | --- | --- |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning the correct headers.

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Response bodies

#### ListKafkaVersionsResponse schema

```
{
  "nextToken": "string",
  "kafkaVersions": [
    {
      "version": "string",
      "status": enum
```

```
    }
  ]
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## KafkaVersion

Information about an Apache Kafka version.

**version**

The Apache Kafka version.

**Type**: string
**Required**: False

**status**

The status of the Apache Kafka version.

**Type**: KafkaVersionStatus
**Required**: False

## KafkaVersionStatus

The status of an Apache Kafka version.

```
ACTIVE
DEPRECATED
```

## ListKafkaVersionsResponse

Response for ListKafkaVersions.

**nextToken**

Paginated results marker.

**Type**: string
**Required**: False

**kafkaVersions**

An array of Apache Kafka version objects.

**Type**: Array of type KafkaVersion
**Required**: False

# See also

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

## ListKafkaVersions

- 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

# BootstrapBrokers

A list of brokers that a client application can use to bootstrap.

## URI

`/v1/clusters/`*`clusterArn`*`/bootstrap-brokers`

## HTTP methods

### GET

**Operation ID:** `GetBootstrapBrokers`

A list of brokers that a client can use to bootstrap. This list doesn't necessarily include all of the brokers in the cluster. The following Python 3.6 example shows how you can use the Amazon Resource Name (ARN) of a cluster to get its bootstrap brokers. If you don't know the ARN of your cluster, you can use the `ListClusters` operation to get the ARNs of all the clusters in this account and Region.

```
import boto3

client = boto3.client('kafka')

response = client.get_bootstrap_brokers(
```

```
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4',
)

print(response['BootstrapBrokerString'])
```

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | GetBootstrapBrokersResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |

| Status code | Response model | Description |
| --- | --- | --- |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning the correct headers.

#### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Response bodies

#### GetBootstrapBrokersResponse schema

```
{
```

```
  "bootstrapBrokerStringPublicSaslIam": "string",
  "bootstrapBrokerStringIpv6": "string",
  "bootstrapBrokerStringPublicSaslScram": "string",
  "bootstrapBrokerString": "string",
  "bootstrapBrokerStringTls": "string",
  "bootstrapBrokerStringVpcConnectivitySaslIam": "string",
  "bootstrapBrokerStringPublicTls": "string",
  "bootstrapBrokerStringVpcConnectivityTls": "string",
  "bootstrapBrokerStringVpcConnectivitySaslScram": "string",
  "bootstrapBrokerStringSaslIam": "string",
  "bootstrapBrokerStringTlsIpv6": "string",
  "bootstrapBrokerStringSaslScramIpv6": "string",
  "bootstrapBrokerStringSaslScram": "string",
  "bootstrapBrokerStringSaslIamIpv6": "string"
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string

**Required**: False

## GetBootstrapBrokersResponse

Returns a string containing one or more hostname:port pairs.

### bootstrapBrokerStringPublicSaslIam

A string that is one or more pairs of DNS names (or IP addresses) and SASL IAM ports for public access.

**Type**: string
**Required**: False

### bootstrapBrokerStringIpv6

A string that contains one or more host name (or IP) and port pairs for IPV6 connectivity.

**Type**: string
**Required**: False

### bootstrapBrokerStringPublicSaslScram

A string that is one or more pairs of DNS names (or IP addresses) and SASL IAM ports for public access.

**Type**: string
**Required**: False

### bootstrapBrokerString

A string containing one or more hostname:port pairs.

**Type**: string
**Required**: False

### bootstrapBrokerStringTls

A string containing one or more DNS names (or IP) and TLS port pairs. The following is an example.

```
{
    "BootstrapBrokerStringTls": "b-3.exampleClusterName.abcde.c2.kafka.us-
east-1.amazonaws.com:9094,b-1.exampleClusterName.abcde.c2.kafka.us-
east-1.amazonaws.com:9094,b-2.exampleClusterName.abcde.c2.kafka.us-
east-1.amazonaws.com:9094"
}
```

**Type**: string
**Required**: False

### bootstrapBrokerStringVpcConnectivitySaslIam

A string containing one or more dns name (or IP) and SASL IAM port pairs for VPC connectivity.

**Type**: string
**Required**: False

### bootstrapBrokerStringPublicTls

A string that is one or more pairs of DNS names (or IP addresses) and SASL IAM ports for public access.

**Type**: string
**Required**: False

### bootstrapBrokerStringVpcConnectivityTls

A string containing one or more dns name (or IP) and Tls port pairs for VPC connectivity.

**Type**: string
**Required**: False

### bootstrapBrokerStringVpcConnectivitySaslScram

A string containing one or more dns name (or IP) and SASL SCRAM port pairs for VPC connectivity.

**Type**: string
**Required**: False

**bootstrapBrokerStringSaslIam**

A string containing one or more dns name (or IP) and SASL IAM port pairs.

**Type**: string
**Required**: False

**bootstrapBrokerStringTlsIpv6**

A string that contains one or more DNS name (or IP) and TLS port pairs for IPV6 connectivity.

**Type**: string
**Required**: False

**bootstrapBrokerStringSaslScramIpv6**

A string that contains one or more DNS name (or IP) and SASL SCRAM port pairs for IPV6 connectivity.

**Type**: string
**Required**: False

**bootstrapBrokerStringSaslScram**

A string containing one or more dns name (or IP) and SASL SCRAM port pairs.

**Type**: string
**Required**: False

**bootstrapBrokerStringSaslIamIpv6**

A string that contains one or more DNS name (or IP) and SASL IAM port for IPV6 connectivity.

**Type**: string
**Required**: False

## See also

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

### GetBootstrapBrokers

- 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

# Broker Count

The number of broker nodes in a cluster.

## URI

`/v1/clusters/clusterArn/nodes/count`

## HTTP methods

### PUT

**Operation ID:** `UpdateBrokerCount`

Updates the number of broker nodes in the cluster. You can use this operation to increase or decrease the number of brokers in an existing cluster.

The following Python 3.6 example shows how you can increase the number of brokers in a cluster to 6 brokers. The update operation returns immediately, with a response that includes the Amazon Resource Name (ARN) that Amazon MSK assigns to this cluster operation. You can use that ARN to check the state of the operation. When the state changes from `PENDING` to `UPDATE_COMPLETE`, the operation is complete.

```
import boto3
import time

client = boto3.client('kafka')

update_broker_count_response = client.update_broker_count(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4',
    CurrentVersion='K12V3IB1VIZHHY',
    TargetNumberOfBrokerNodes=6
)

operation_arn = update_broker_count_response['ClusterOperationArn']

print(operation_arn)

describe_cluster_operation_response =
 client.describe_cluster_operation(ClusterOperationArn=operation_arn)
operation_state = describe_cluster_operation_response['ClusterOperationInfo']
['OperationState']
print(operation_state)

expanded = False

while not expanded:
    print('Sleeping for 15 seconds before checking to see if the cluster update is
 done...')
    time.sleep(15)
    describe_cluster_operation_response =
 client.describe_cluster_operation(ClusterOperationArn=operation_arn)
    operation_state = describe_cluster_operation_response['ClusterOperationInfo']
['OperationState']
    if 'UPDATE_COMPLETE' == operation_state:
        expanded = True
        print('The cluster has 6 brokers now.')
```

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | UpdateBrokerCountResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying |

| Status code | Response model | Description |
| --- | --- | --- |
| | | your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning the correct headers.

#### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### PUT schema

```
{
  "targetNumberOfBrokerNodes": integer,
  "currentVersion": "string"
}
```

## Response bodies

### UpdateBrokerCountResponse schema

```
{
  "clusterArn": "string",
  "clusterOperationArn": "string"
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Error

Returns information about an error.

### message

The description of the error.

**Type**: string
**Required**: False

### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## UpdateBrokerCountRequest

Request body for UpdateBrokerCount. Accepts a targetNumberOfBrokerNodes value that is 2 or 3 less than the current broker node count for the cluster (for 2 and 3 AZ clusters respectively).

**targetNumberOfBrokerNodes**

The number of broker nodes that you want the cluster to have after this operation completes successfully.

**Type**: integer
**Required**: True

**currentVersion**

The current version of the cluster.

**Type**: string
**Required**: True

## UpdateBrokerCountResponse

Response body for UpdateBrokerCount.

**clusterArn**

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

**clusterOperationArn**

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

# See also

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

### UpdateBrokerCount

- 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

# Broker Storage

Represents the EBS storage associated with the brokers.

## URI

`/v1/clusters/clusterArn/nodes/storage`

## HTTP methods

### PUT

**Operation ID:** `UpdateBrokerStorage`

Updates the EBS storage associated with Amazon MSK brokers. You can increase the amount of EBS storage per broker. You can't decrease the storage. To increase storage, wait for the cluster to be in the ACTIVE state. Storage volumes remain available during this scaling-up operation.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | UpdateBrokerStorageResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | None | Default response for CORS method |

# Schemas

## Request bodies

### PUT schema

```
{
  "targetBrokerEBSVolumeInfo": [
    {
      "volumeSizeGB": integer,
      "provisionedThroughput": {
        "volumeThroughput": integer,
        "enabled": boolean
      },
      "kafkaBrokerNodeId": "string"
    }
  ],
  "currentVersion": "string"
}
```

## Response bodies

### UpdateBrokerStorageResponse schema

```
{
  "clusterArn": "string",
  "clusterOperationArn": "string"
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## BrokerEBSVolumeInfo

Specifies the EBS volume upgrade information. The broker identifier must be set to the keyword ALL. This means the changes apply to all the brokers in the cluster.

### volumeSizeGB

Size of the EBS volume to update.

**Type**: integer
**Required**: False

### provisionedThroughput

EBS volume provisioned throughput information.

**Type**: ProvisionedThroughput
**Required**: False

### kafkaBrokerNodeId

The ID of the broker to update. The only allowed value is ALL. This means that Amazon MSK applies the same storage update to all broker nodes.

**Type**: string
**Required**: True

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## ProvisionedThroughput

Contains information about provisioned throughput for EBS storage volumes attached to kafka broker nodes.

**volumeThroughput**

Throughput value of the EBS volumes for the data drive on each kafka broker node in MiB per second.

**Type**: integer
**Required**: False

**enabled**

Provisioned throughput is enabled or not.

**Type**: boolean
**Required**: False

## UpdateBrokerStorageRequest

Request object for UpdateBrokerStorage.

### targetBrokerEBSVolumeInfo

Describes the target volume size and the ID of the broker to apply the update to.

The value you specify for Target-Volume-in-GiB must be a whole number that is greater than 100 GiB.

The storage per broker after the update operation can't exceed 16384 GiB.

**Type**: Array of type BrokerEBSVolumeInfo
**Required**: True

### currentVersion

The version of the MSK cluster to update. Cluster versions aren't simple numbers. You can describe an MSK cluster to find its version. When this update operation is successful, it generates a new cluster version.

**Type**: string
**Required**: True

## UpdateBrokerStorageResponse

Response body for UpdateBrokerStorage.

### clusterArn

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

### clusterOperationArn

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string

**Required**: False

## See also

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

### UpdateBrokerStorage

- 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

# Broker Type

The type of brokers in the cluster. All of the brokers in a cluster are the same type.

## URI

`/v1/clusters/clusterArn/nodes/type`

## HTTP methods

### PUT

**Operation ID:** `UpdateBrokerType`

For information about this operation, see Updating the broker type in the developer guide.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | UpdateBrokerTypeResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying |

| Status code | Response model | Description |
| --- | --- | --- |
| | | your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning the correct headers.

#### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### PUT schema

```
{
  "targetInstanceType": "string",
  "currentVersion": "string"
}
```

### Response bodies

#### UpdateBrokerTypeResponse schema

```
{
  "clusterArn": "string",
  "clusterOperationArn": "string"
}
```

#### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

## Properties

### Error

Returns information about an error.

#### message

The description of the error.

**Type**: string
**Required**: False

#### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

### UpdateBrokerTypeRequest

Request body for UpdateBrokerType.

**targetInstanceType**

The type of Amazon EC2 instances to use for Kafka brokers.

**Type**: string
**Required**: True

**currentVersion**

Current cluster version.

**Type**: string
**Required**: True

## UpdateBrokerTypeResponse

Response body for UpdateBrokerType.

**clusterArn**

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

**clusterOperationArn**

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

# See also

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

## UpdateBrokerType

- 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

# Brokers

The broker nodes in the cluster.

## URI

`/v1/clusters/clusterArn/nodes`

## HTTP methods

### GET

**Operation ID:** `ListNodes`

Returns a list of the broker nodes in the cluster. The following Python 3.6 example first lists one node of a cluster. Because the cluster has more nodes, the response contains a token that the script then uses to list the remaining nodes.

```
import boto3

client = boto3.client('kafka')

list_nodes_response = client.list_nodes(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4',
    MaxResults=1
)
```

```
print('\n')
print('Here is the first node in the list:')
print('\n')
print(list_nodes_response['NodeInfoList'])

next_token = list_nodes_response['NextToken']

list_nodes_response = client.list_nodes(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4',
    NextToken=next_token
)

print('\n')
print('Here are the remaining nodes in the list:')
print('\n')
print(list_nodes_response['NodeInfoList'])
```

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Query parameters

| Name | Type | Required | Description |
|---|---|---|---|
| nextToken | String | False | The paginated results marker. When the result of the operation is truncated, the call returns NextToken in the response. To get the next batch, |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | provide this token in your next request. |
| maxResults | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes a NextToken parameter. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | ListNodesResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorize d. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. |

| Status code | Response model | Description |
| --- | --- | --- |
| | | Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

# Schemas

## Response bodies

### ListNodesResponse schema

```
{
  "nextToken": "string",
  "nodeInfoList": [
    {
      "zookeeperNodeInfo": {
        "zookeeperId": number,
        "clientVpcIpAddress": "string",
        "attachedENIId": "string",
        "endpoints": [
          "string"
        ],
        "zookeeperVersion": "string"
      },
      "instanceType": "string",
      "controllerNodeInfo": {
        "endpoints": [
          "string"
        ]
      },
      "nodeType": enum,
      "nodeARN": "string",
      "brokerNodeInfo": {
        "clientVpcIpAddress": "string",
        "attachedENIId": "string",
        "brokerId": number,
        "endpoints": [
          "string"
        ],
        "clientSubnet": "string",
        "currentBrokerSoftwareInfo": {
          "configurationRevision": integer,
          "kafkaVersion": "string",
          "configurationArn": "string"
        }
      },
      "addedToClusterTime": "string"
    }
```

```
  ]
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## BrokerNodeInfo

BrokerNodeInfo

### clientVpcIpAddress

The virtual private cloud (VPC) of the client.

**Type**: string
**Required**: False

### attachedENIId

The attached elastic network interface of the broker.

**Type**: string
**Required**: False

### brokerId

The ID of the broker.

**Type**: number
**Required**: False

### endpoints

Endpoints for accessing the broker.

**Type**: Array of type string

**Required**: False

#### clientSubnet

The client subnet to which this broker node belongs.

**Type**: string

**Required**: False

#### currentBrokerSoftwareInfo

Information about the version of software currently deployed on the brokers in the cluster.

**Type**: BrokerSoftwareInfo

**Required**: False

## BrokerSoftwareInfo

Information about the current software installed on the cluster.

#### configurationRevision

The revision of the configuration to use. This field isn't visible in this preview release.

**Type**: integer

**Required**: False

**Format**: int64

#### kafkaVersion

The version of Apache Kafka. You can use Amazon MSK to create clusters that use Apache Kafka versions 1.1.1 and 2.2.1. See Apache Kafka Versions.

**Type**: string

**Required**: False

**configurationArn**

The Amazon Resource Name (ARN) of the configuration used for the cluster. This field isn't visible in this preview release.

**Type**: string
**Required**: False

## ControllerNodeInfo

Controller Node Information.

**endpoints**

Endpoints for accessing the controller.

**Type**: Array of type string
**Required**: False

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## ListNodesResponse

Information about nodes in the cluster.

**nextToken**

The paginated results marker. When the result of a `ListNodes` operation is truncated, the call returns `NextToken` in the response. To get another batch of nodes, provide this token in your next request.

**Type**: string
**Required**: False

**nodeInfoList**

List containing a `NodeInfo` object. Doesn't contain entries for brokers that have been deleted.

**Type**: Array of type NodeInfo
**Required**: False

## NodeInfo

The node information object.

**zookeeperNodeInfo**

The ZookeeperNodeInfo.

**Type**: ZookeeperNodeInfo
**Required**: False

**instanceType**

The instance type.

**Type**: string
**Required**: False

**controllerNodeInfo**

The controller node information.

**Type**: ControllerNodeInfo

**Required**: False

### nodeType

The node type.

**Type**: NodeType
**Required**: False

### nodeARN

The Amazon Resource Name (ARN) of the node.

**Type**: string
**Required**: False

### brokerNodeInfo

The broker node info.

**Type**: BrokerNodeInfo
**Required**: False

### addedToClusterTime

The start time.

**Type**: string
**Required**: False

## NodeType

The broker or Apache ZooKeeper node.

BROKER

## ZookeeperNodeInfo

Apache ZooKeeper node information.

**zookeeperId**

The role-specific ID for Apache ZooKeeper.

**Type**: number
**Required**: False

**clientVpcIpAddress**

The virtual private cloud (VPC) IP address of the client.

**Type**: string
**Required**: False

**attachedENIId**

The attached elastic network interface of the broker.

**Type**: string
**Required**: False

**endpoints**

Endpoints for accessing the Apache ZooKeeper nodes.

**Type**: Array of type string
**Required**: False

**zookeeperVersion**

The version of Apache ZooKeeper.

**Type**: string
**Required**: False

## See also

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

### ListNodes

- 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

# Cluster

Represents an Amazon MSK cluster.

## URI

`/v1/clusters/`*`clusterArn`*

## HTTP methods

### GET

**Operation ID:** `DescribeCluster`

Returns a description of the MSK cluster whose Amazon Resource Name (ARN) is specified in the request. The following is a Python 3.6 example of how to use this operation. Before you run this Python script, replace the example cluster Amazon Resource Name (ARN) with the ARN of the cluster you want to describe. If you don't know the ARN of the cluster, you can use the `ListClusters` operation to list all the clusters and see their ARNs and full descriptions.

```
import boto3
```

```
client = boto3.client('kafka')

response = client.describe_cluster(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4'
)

print(response)
```

Note that the response to this operation only includes the `ZookeeperConnectStringTls` node in clusters created with Apache Kafka version 2.5.1 and later.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | DescribeClusterResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |

| Status code | Response model | Description |
| --- | --- | --- |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## DELETE

**Operation ID:** `DeleteCluster`

Deletes the MSK cluster specified by the Amazon Resource Name (ARN) in the request, and all its revisions.

### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Query parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| currentVersion | String | False | The current version of the MSK cluster. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | DeleteClusterResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | None | Default response for CORS method |

## Schemas

## Response bodies

### DescribeClusterResponse schema

```
{
  "clusterInfo": {
    "encryptionInfo": {
      "encryptionInTransit": {
        "inCluster": boolean,
        "clientBroker": enum
      },
      "encryptionAtRest": {
        "dataVolumeKMSKeyId": "string"
      }
    },
    "zookeeperConnectString": "string",
    "creationTime": "string",
    "customerActionStatus": enum,
    "zookeeperConnectStringTls": "string",
    "loggingInfo": {
      "brokerLogs": {
        "s3": {
```

```
          "bucket": "string",
          "prefix": "string",
          "enabled": boolean
        },
        "firehose": {
          "deliveryStream": "string",
          "enabled": boolean
        },
        "cloudWatchLogs": {
          "logGroup": "string",
          "enabled": boolean
        }
      },
      "authorizerLogs": {
        "s3": {
          "bucket": "string",
          "prefix": "string",
          "enabled": boolean
        },
        "firehose": {
          "deliveryStream": "string",
          "enabled": boolean
        },
        "cloudWatchLogs": {
          "logGroup": "string",
          "enabled": boolean
        }
      }
    },
    "currentVersion": "string",
    "tags": {
    },
    "numberOfBrokerNodes": integer,
    "clusterArn": "string",
    "activeOperationArn": "string",
    "enhancedMonitoring": enum,
    "storageMode": enum,
    "clusterName": "string",
    "stateInfo": {
      "code": "string",
      "message": "string"
    },
    "clientAuthentication": {
      "sasl": {
```

```
        "iam": {
          "enabled": boolean
        },
        "scram": {
          "enabled": boolean
        }
      },
      "unauthenticated": {
        "enabled": boolean
      },
      "tls": {
        "certificateAuthorityArnList": [
          "string"
        ],
        "enabled": boolean
      }
    },
    "state": enum,
    "brokerNodeGroupInfo": {
      "clientSubnets": [
        "string"
      ],
      "zoneIds": [
        "string"
      ],
      "instanceType": "string",
      "connectivityInfo": {
        "vpcConnectivity": {
          "clientAuthentication": {
            "sasl": {
              "iam": {
                "enabled": boolean
              },
              "scram": {
                "enabled": boolean
              }
            },
            "tls": {
              "enabled": boolean
            }
          }
        },
        "publicAccess": {
          "type": "string"
```

```
        },
        "networkType": "string"
      },
      "securityGroups": [
        "string"
      ],
      "brokerAZDistribution": enum,
      "storageInfo": {
        "ebsStorageInfo": {
          "provisionedThroughput": {
            "volumeThroughput": integer,
            "enabled": boolean
          },
          "volumeSize": integer
        }
      }
    },
    "openMonitoring": {
      "prometheus": {
        "nodeExporter": {
          "enabledInBroker": boolean
        },
        "jmxExporter": {
          "enabledInBroker": boolean
        }
      }
    },
    "rebalancing": {
      "status": enum
    },
    "currentBrokerSoftwareInfo": {
      "configurationRevision": integer,
      "kafkaVersion": "string",
      "configurationArn": "string"
    }
  }
}
```

### DeleteClusterResponse schema

```
{
  "clusterArn": "string",
  "state": enum
```

```
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## AuthorizerLogs

Authorizer Logs details for cluster.

### s3

S3 Log destination details.

**Type**: S3
**Required**: False

### firehose

Firehose Log destination details.

**Type**: Firehose
**Required**: False

### cloudWatchLogs

CloudWatch Log destination details.

**Type**: CloudWatchLogs
**Required**: False

## BrokerAZDistribution

This parameter is currently not in use.

DEFAULT

## BrokerLogs

The broker logs configuration for this MSK cluster.

**s3**

Details of the Amazon S3 destination for broker logs.

**Type**: S3
**Required**: False

**firehose**

Details of the Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: Firehose
**Required**: False

**cloudWatchLogs**

Details of the CloudWatch Logs destination for broker logs.

**Type**: CloudWatchLogs
**Required**: False

## BrokerNodeGroupInfo

Describes the setup to be used for the broker nodes in the cluster.

**clientSubnets**

The list of subnets to connect to in the client virtual private cloud (VPC). Amazon creates elastic network interfaces inside these subnets. Client applications use elastic network interfaces to produce and consume data.

If you use the US West (N. California) Region, specify exactly two subnets. For other Regions where Amazon MSK is available, you can specify either two or three subnets. The subnets that you specify

must be in distinct Availability Zones. When you create a cluster, Amazon MSK distributes the broker nodes evenly across the subnets that you specify.

Client subnets can't occupy the Availability Zone with ID `use1-az3`.

**Type**: Array of type string
**Required**: True

**zoneIds**

The zoneIds for brokers in customer account.

**Type**: Array of type string
**Required**: False

**instanceType**

The type of Amazon EC2 instances to use for brokers. The following instance types are allowed: kafka.m5.large, kafka.m5.xlarge, kafka.m5.2xlarge, kafka.m5.4xlarge, kafka.m5.8xlarge, kafka.m5.12xlarge, kafka.m5.16xlarge, and kafka.m5.24xlarge.

**Type**: string
**Required**: True
**MinLength**: 5
**MaxLength**: 32

**connectivityInfo**

Information about the cluster's connectivity setting.

**Type**: ConnectivityInfo
**Required**: False

**securityGroups**

The security groups to associate with the elastic network interfaces in order to specify who can connect to and communicate with the Amazon MSK cluster. If you don't specify a security group, Amazon MSK uses the default security group associated with the VPC. If you specify security

groups that were shared with you, you must ensure that you have permissions to them. Specifically, you need the `ec2:DescribeSecurityGroups` permission.

**Type**: Array of type string
**Required**: False

#### brokerAZDistribution

This parameter is currently not in use.

**Type**: BrokerAZDistribution
**Required**: False

#### storageInfo

Contains information about storage volumes attached to Amazon MSK broker nodes.

**Type**: StorageInfo
**Required**: False

### BrokerSoftwareInfo

Information about the current software installed on the cluster.

#### configurationRevision

The revision of the configuration to use. This field isn't visible in this preview release.

**Type**: integer
**Required**: False
**Format**: int64

#### kafkaVersion

The version of Apache Kafka. You can use Amazon MSK to create clusters that use Apache Kafka versions 1.1.1 and 2.2.1. See Apache Kafka Versions.

**Type**: string

**Required**: False

**configurationArn**

The Amazon Resource Name (ARN) of the configuration used for the cluster. This field isn't visible in this preview release.

**Type**: string
**Required**: False

## ClientAuthentication

Includes all client authentication information.

**sasl**

Details for client authentication using SASL. To turn on SASL, you must also turn on `EncryptionInTransit` by setting `inCluster` to true. You must set `clientBroker` to either `TLS` or `TLS_PLAINTEXT`. If you choose `TLS_PLAINTEXT`, then you must also set `unauthenticated` to true.

**Type**: Sasl
**Required**: False

**unauthenticated**

Details for ClientAuthentication using no authentication.

**Type**: Unauthenticated
**Required**: False

**tls**

Details for ClientAuthentication using TLS. To turn on TLS access control, you must also turn on `EncryptionInTransit` by setting `inCluster` to true and `clientBroker` to `TLS`.

**Type**: Tls
**Required**: False

## ClientBroker

Client-broker encryption in transit setting.

TLS
TLS_PLAINTEXT
PLAINTEXT

## CloudWatchLogs

Details of the CloudWatch Logs destination for broker logs.

### logGroup

The CloudWatch log group that is the destination for broker logs.

**Type**: string
**Required**: False

### enabled

Specifies whether broker logs get sent to the specified CloudWatch Logs destination.

**Type**: boolean
**Required**: True

## ClusterInfo

Returns information about a cluster.

### encryptionInfo

Includes all encryption-related information.

**Type**: EncryptionInfo
**Required**: False

### zookeeperConnectString

The connection string to use to connect to zookeeper cluster on plaintext port.

**Type**: string
**Required**: False

### creationTime

The time when the cluster was created.

**Type**: string
**Required**: False

### customerActionStatus

Determines if there is an action required from the customer.

**Type**: CustomerActionStatus
**Required**: False

### zookeeperConnectStringTls

The connection string to use to connect to the Apache ZooKeeper cluster on a TLS port.

**Type**: string
**Required**: False

### loggingInfo

You can configure your MSK cluster to send broker logs to different destination types. This is a container for the configuration details related to broker logs.

**Type**: LoggingInfo
**Required**: False

### currentVersion

The current version of the MSK cluster. Cluster versions aren't simple integers. You can obtain the current version by describing the cluster. An example version is `KTVPDKIKX0DER`.

**Type**: string

**Required**: False

**tags**

Tags attached to the cluster.

**Type**: object
**Required**: False

**numberOfBrokerNodes**

The number of broker nodes in the cluster.

**Type**: integer
**Required**: False

**clusterArn**

The Amazon Resource Name (ARN) that uniquely identifies the cluster.

**Type**: string
**Required**: False

**activeOperationArn**

Arn of active cluster operation.

**Type**: string
**Required**: False

**enhancedMonitoring**

Specifies which Apache Kafka metrics Amazon MSK gathers and sends to Amazon CloudWatch for this cluster. This property has three possible values: `DEFAULT`, `PER_BROKER`, and `PER_TOPIC_PER_BROKER`. For a list of the metrics associated with each of these three levels of monitoring, see Monitoring.

**Type**: EnhancedMonitoring

**Required**: False

### storageMode

This controls storage mode for supported storage tiers.

**Type**: StorageMode
**Required**: False

### clusterName

The name of the cluster.

**Type**: string
**Required**: False

### stateInfo

Includes information of the cluster state.

**Type**: StateInfo
**Required**: False

### clientAuthentication

Includes all client authentication information.

**Type**: ClientAuthentication
**Required**: False

### state

The state of the cluster. Amazon MSK automatically renews certificates on clusters every 13 months. It sets the state of the cluster to `MAINTENANCE` when it starts the certificate-update operation. It sets it back to `ACTIVE` when the update is done. While a cluster is in the `MAINTENANCE` state, you can continue to produce and consume data, but you can't perform any update operations on it. You can perform update operations on a cluster when it is in the `ACTIVE` state.

**Type**: ClusterState
**Required**: False

**brokerNodeGroupInfo**

Information about the broker nodes.

**Type**: BrokerNodeGroupInfo
**Required**: False

**openMonitoring**

Settings for open monitoring using Prometheus.

**Type**: OpenMonitoring
**Required**: False

**rebalancing**

Specifies if intelligent rebalancing is turned on for your MSK Provisioned cluster with Express brokers. For all new Express-based clusters that you create, intelligent rebalancing is turned on by default.

**Type**: Rebalancing
**Required**: False

**currentBrokerSoftwareInfo**

Information about the version of software currently deployed on the brokers in the cluster.

**Type**: BrokerSoftwareInfo
**Required**: False

## ClusterState

The state of the cluster. Amazon MSK automatically renews certificates on clusters every 13 months. It sets the state of the cluster to `MAINTENANCE` when it starts the certificate-update operation. It sets it back to `ACTIVE` when the update is done. While a cluster is in the

MAINTENANCE state, you can continue to produce and consume data, but you can't perform any update operations on it. You can perform update operations on a cluster when it is in the ACTIVE state.

ACTIVE
CREATING
UPDATING
DELETING
FAILED
MAINTENANCE
REBOOTING_BROKER
HEALING

## ConnectivityInfo

Broker access controls.

**vpcConnectivity**

VPC connection control settings for brokers

**Type**: VpcConnectivity
**Required**: False

**publicAccess**

Access control settings for the cluster's brokers.

**Type**: PublicAccess
**Required**: False

**networkType**

The network type of the cluster, which is IPv4 or DUAL. The DUAL network type uses both IPv4 and IPv6 addresses for your cluster and its resources. By default, a cluster uses the IPv4 network type.

**Type**: string
**Required**: False

**MinLength**: 4
**MaxLength**: 4

## CustomerActionStatus

A type of an action required from the customer.

`CRITICAL_ACTION_REQUIRED`
`ACTION_RECOMMENDED`
`NONE`

## DeleteClusterResponse

Returns information about the deleted cluster.

### clusterArn

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

### state

The state of the cluster. Amazon MSK automatically renews certificates on clusters every 13 months. It sets the state of the cluster to `MAINTENANCE` when it starts the certificate-update operation. It sets it back to `ACTIVE` when the update is done. While a cluster is in the `MAINTENANCE` state, you can continue to produce and consume data, but you can't perform any update operations on it. You can perform update operations on a cluster when it is in the `ACTIVE` state.

**Type**: ClusterState
**Required**: False

## DescribeClusterResponse

Returns information about a cluster.

**clusterInfo**

The cluster information.

**Type**: ClusterInfo
**Required**: False

## EBSStorageInfo

Contains information about the EBS storage volumes attached to the broker nodes.

**provisionedThroughput**

EBS volume provisioned throughput information.

**Type**: ProvisionedThroughput
**Required**: False

**volumeSize**

The size in GiB of the EBS volume for the data drive on each broker node.

**Type**: integer
**Required**: False
**Minimum**: 1
**Maximum**: 16384

## EncryptionAtRest

The data-volume encryption details. You can't update encryption at rest settings for existing clusters.

**dataVolumeKMSKeyId**

The Amazon Resource Name (ARN) of the Amazon KMS key for encrypting data at rest. If you don't specify a KMS key, MSK creates one for you and uses it.

**Type**: string
**Required**: True

## EncryptionInTransit

The settings for encrypting data in transit.

**inCluster**

When set to true, it indicates that data communication among the broker nodes of the cluster is encrypted. When set to false, the communication happens in plaintext.

The default value is true.

**Type**: boolean
**Required**: False

**clientBroker**

Indicates the encryption setting for data in transit between clients and brokers. You must set it to one of the following values.

`TLS` means that client-broker communication is enabled with TLS only.

`TLS_PLAINTEXT` means that client-broker communication is enabled for both TLS-encrypted, as well as plaintext data.

`PLAINTEXT` means that client-broker communication is enabled in plaintext only.

The default value is `TLS`.

**Type**: ClientBroker
**Required**: False

## EncryptionInfo

Includes encryption-related information, such as the Amazon KMS key used for encrypting data at rest and whether you want MSK to encrypt your data in transit.

**encryptionInTransit**

The details for encryption in transit.

**Type**: EncryptionInTransit

**Required**: False

#### encryptionAtRest

The data-volume encryption details.

**Type**: EncryptionAtRest
**Required**: False

## EnhancedMonitoring

Specifies which Apache Kafka metrics Amazon MSK gathers and sends to Amazon CloudWatch for this cluster. This property has three possible values: `DEFAULT`, `PER_BROKER`, and `PER_TOPIC_PER_BROKER`. For a list of the metrics associated with each of these three levels of monitoring, see Monitoring.

```
DEFAULT
PER_BROKER
PER_TOPIC_PER_BROKER
PER_TOPIC_PER_PARTITION
```

## Error

Returns information about an error.

#### message

The description of the error.

**Type**: string
**Required**: False

#### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## Firehose

Firehose details for BrokerLogs.

**deliveryStream**

The Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Kinesis Data Firehose delivery stream.

**Type**: boolean
**Required**: True

## IAM

Details for SASL/IAM client authentication.

**enabled**

SASL/IAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## JmxExporter

Indicates whether you want to enable or disable the JMX Exporter.

**enabledInBroker**

Indicates whether you want to enable or disable the JMX Exporter.

**Type**: boolean
**Required**: True

## LoggingInfo

You can configure your MSK cluster to send broker logs to different destination types. This is a container for the configuration details related to broker logs.

**brokerLogs**

You can configure your MSK cluster to send broker logs to different destination types. This configuration specifies the details of these destinations.

**Type**: BrokerLogs
**Required**: False

**authorizerLogs**

Authorizer Logs details.

**Type**: AuthorizerLogs
**Required**: False

## NodeExporter

Indicates whether you want to enable or disable the Node Exporter.

**enabledInBroker**

Indicates whether you want to enable or disable the Node Exporter.

**Type**: boolean
**Required**: True

## OpenMonitoring

JMX and Node monitoring for the MSK cluster.

**prometheus**

Prometheus exporter settings.

**Type**: Prometheus

**Required**: True

## Prometheus

Prometheus settings for open monitoring.

**nodeExporter**

Indicates whether you want to enable or disable the Node Exporter.

**Type**: NodeExporter
**Required**: False

**jmxExporter**

Indicates whether you want to enable or disable the JMX Exporter.

**Type**: JmxExporter
**Required**: False

## ProvisionedThroughput

Contains information about provisioned throughput for EBS storage volumes attached to kafka broker nodes.

**volumeThroughput**

Throughput value of the EBS volumes for the data drive on each kafka broker node in MiB per second.

**Type**: integer
**Required**: False

**enabled**

Provisioned throughput is enabled or not.

**Type**: boolean
**Required**: False

## PublicAccess

Broker access controls

**type**

DISABLED means that public access is turned off. SERVICE_PROVIDED_EIPS means that public access is turned on.

**Type**: string
**Required**: False

## Rebalancing

Specifies whether or not intelligent rebalancing is turned on for a newly created MSK Provisioned cluster with Express brokers. Intelligent rebalancing performs automatic partition balancing operations when you scale your clusters up or down.

By default, intelligent rebalancing is `ACTIVE` for all new Express-based clusters.

**status**

Intelligent rebalancing status. The default intelligent rebalancing status is `ACTIVE` for all new Express-based clusters.

**Type**: RebalancingStatus
**Required**: True

## RebalancingStatus

Intelligent rebalancing status. The default intelligent rebalancing status is `ACTIVE` for all new Express-based clusters.

`PAUSED`
`ACTIVE`

## S3

The details of the Amazon S3 destination for broker logs.

**bucket**

The name of the S3 bucket that is the destination for broker logs.

**Type**: string
**Required**: False

**prefix**

The S3 prefix that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Amazon S3 destination.

**Type**: boolean
**Required**: True

## Sasl

Details for client authentication using SASL. To turn on SASL, you must also turn on `EncryptionInTransit` by setting `inCluster` to true. You must set `clientBroker` to either `TLS` or `TLS_PLAINTEXT`. If you choose `TLS_PLAINTEXT`, then you must also set `unauthenticated` to true.

**iam**

Details for ClientAuthentication using IAM.

**Type**: IAM
**Required**: False

**scram**

Details for SASL/SCRAM client authentication.

**Type**: Scram

**Required**: False

## Scram

Details for SASL/SCRAM client authentication.

**enabled**

SASL/SCRAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## StateInfo

Includes information about the state of the cluster.

**code**

If the cluster is in an unusable state, this field contains the code that describes the issue.

**Type**: string
**Required**: False

**message**

If the cluster is in an unusable state, this field contains a message that describes the issue.

**Type**: string
**Required**: False

## StorageInfo

Contains information about storage volumes attached to Amazon MSK broker nodes.

**ebsStorageInfo**

EBS volume information.

**Type**: EBSStorageInfo
**Required**: False

## StorageMode

Controls storage mode for various supported storage tiers.

LOCAL
TIERED

## Tls

Details for client authentication using TLS.

**certificateAuthorityArnList**

List of AWS Private CA Amazon Resource Name (ARN)s.

**Type**: Array of type string
**Required**: False

**enabled**

TLS authentication is enabled or not.

**Type**: boolean
**Required**: False

## Unauthenticated

Details for allowing no client authentication.

**enabled**

Unauthenticated is enabled or not.

**Type**: boolean
**Required**: False

## VpcConnectivity

VPC connection control settings for brokers.

**clientAuthentication**

VPC connection control settings for brokers.

**Type**: VpcConnectivityClientAuthentication
**Required**: False

## VpcConnectivityClientAuthentication

Includes all client authentication information for VpcConnectivity.

**sasl**

Details for VpcConnectivity ClientAuthentication using SASL.

**Type**: VpcConnectivitySasl
**Required**: False

**tls**

Details for VpcConnectivity ClientAuthentication using TLS.

**Type**: VpcConnectivityTls
**Required**: False

## VpcConnectivityIAM

Details for SASL/IAM client authentication for VpcConnectivity.

**enabled**

SASL/IAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## VpcConnectivitySasl

Details for client authentication using SASL for VpcConnectivity.

**iam**

Details for ClientAuthentication using IAM for VpcConnectivity.

**Type**: VpcConnectivityIAM
**Required**: False

**scram**

Details for SASL/SCRAM client authentication for VpcConnectivity.

**Type**: VpcConnectivityScram
**Required**: False

## VpcConnectivityScram

Details for SASL/SCRAM client authentication for vpcConnectivity.

**enabled**

SASL/SCRAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## VpcConnectivityTls

Details for client authentication using TLS for vpcConnectivity.

**enabled**

TLS authentication is enabled or not.

**Type**: boolean
**Required**: False

# See also

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

### DescribeCluster

- 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

### DeleteCluster

- 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

## Cluster Apache Kafka Version

The Apache Kafka version that is on the cluster.

### URI

`/v1/clusters/clusterArn/version`

# HTTP methods

## PUT

**Operation ID:** `UpdateClusterKafkaVersion`

Updates the cluster to the specified Apache Kafka version. Before you invoke this operation, ensure that the number of partitions per broker on your MSK cluster is under the limits described in Number of partitions per broker. You can't update the Apache Kafka version for an MSK cluster that exceeds these limits.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | UpdateClusterKafkaVersionResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |

| Status code | Response model | Description |
| --- | --- | --- |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

# Schemas

## Request bodies

### PUT schema

```
{
  "targetKafkaVersion": "string",
  "configurationInfo": {
    "arn": "string",
    "revision": integer
  },
  "currentVersion": "string"
}
```

## Response bodies

### UpdateClusterKafkaVersionResponse schema

```
{
  "clusterArn": "string",
  "clusterOperationArn": "string"
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## ConfigurationInfo

Specifies the configuration to use for the brokers.

### arn

ARN of the configuration to use.

**Type**: string
**Required**: True

**revision**

The revision of the configuration to use.

**Type**: integer
**Required**: True
**Format**: int64
**Minimum**: 1

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## UpdateClusterKafkaVersionRequest

Request body for UpdateClusterKafkaVersion.

**targetKafkaVersion**

Target Apache Kafka version.

**Type**: string
**Required**: True

**configurationInfo**

The custom configuration that should be applied on the new version of cluster.

**Type**: ConfigurationInfo
**Required**: False

**currentVersion**

Current cluster version.

**Type**: string
**Required**: True

## UpdateClusterKafkaVersionResponse

Response body for UpdateClusterKafkaVersion.

**clusterArn**

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

**clusterOperationArn**

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

# See also

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

## UpdateClusterKafkaVersion

- 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

# Cluster Configuration

Represents the configuration of a specific cluster.

## URI

`/v1/clusters/clusterArn/configuration`

## HTTP methods

### PUT

**Operation ID:** `UpdateClusterConfiguration`

Updates the cluster with the configuration that is specified in the request body. Before you invoke this operation, ensure that the number of partitions per broker on your MSK cluster is under the limits described in Number of partitions per broker. You can't update the configuration of an MSK cluster that exceeds these limits.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | UpdateClusterConfigurationResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | None | Default response for CORS method |

# Schemas

## Request bodies

### PUT schema

```
{
  "configurationInfo": {
    "arn": "string",
    "revision": integer
  },
  "currentVersion": "string"
}
```

## Response bodies

### UpdateClusterConfigurationResponse schema

```
{
```

```
  "clusterArn": "string",
  "clusterOperationArn": "string"
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## ConfigurationInfo

Specifies the configuration to use for the brokers.

**arn**

ARN of the configuration to use.

**Type**: string
**Required**: True

**revision**

The revision of the configuration to use.

**Type**: integer
**Required**: True
**Format**: int64
**Minimum**: 1

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

#### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## UpdateClusterConfigurationRequest

Request body for UpdateClusterConfiguration.

#### configurationInfo

Represents the configuration that you want MSK to use for the cluster.

**Type**: ConfigurationInfo
**Required**: True

#### currentVersion

The version of the cluster that you want to update.

**Type**: string
**Required**: True

## UpdateClusterConfigurationResponse

Response body for UpdateClusterConfiguration.

#### clusterArn

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

**clusterOperationArn**

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string

**Required**: False

## See also

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

### UpdateClusterConfiguration

- 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

# Cluster Connectivity

Use this resource to update the connectivity setting for an MSK cluster.

## URI

`/v1/clusters/clusterArn/connectivity`

# HTTP methods

## PUT

**Operation ID:** `UpdateConnectivity`

Updates the connectivity setting for the cluster.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | UpdateConnectivity Response | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorize d. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. |

| Status code | Response model | Description |
| --- | --- | --- |
| | | Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

# Schemas

## Request bodies

### PUT schema

```
{
  "connectivityInfo": {
    "vpcConnectivity": {
      "clientAuthentication": {
        "sasl": {
          "iam": {
            "enabled": boolean
          },
          "scram": {
            "enabled": boolean
          }
        },
        "tls": {
          "enabled": boolean
        }
      }
    },
    "publicAccess": {
      "type": "string"
    },
    "networkType": "string"
  },
  "zookeeperAccess": {
    "enabled": boolean
  },
  "currentVersion": "string"
}
```

## Response bodies

### UpdateConnectivityResponse schema

```
{
  "clusterArn": "string",
  "clusterOperationArn": "string"
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## ConnectivityInfo

Broker access controls.

**vpcConnectivity**

VPC connection control settings for brokers

**Type**: VpcConnectivity
**Required**: False

**publicAccess**

Access control settings for the cluster's brokers.

**Type**: PublicAccess
**Required**: False

**networkType**

The network type of the cluster, which is IPv4 or DUAL. The DUAL network type uses both IPv4 and IPv6 addresses for your cluster and its resources. By default, a cluster uses the IPv4 network type.

**Type**: string
**Required**: False
**MinLength**: 4
**MaxLength**: 4

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## PublicAccess

Broker access controls

**type**

DISABLED means that public access is turned off. SERVICE_PROVIDED_EIPS means that public access is turned on.

**Type**: string
**Required**: False

## UpdateConnectivityRequest

Request body for UpdateConnectivity.

**connectivityInfo**

The target connectivity setting for the cluster.

**Type**: ConnectivityInfo
**Required**: False

**zookeeperAccess**

The target ZooKeeper access setting for the cluster.

**Type**: ZookeeperAccess
**Required**: False

**currentVersion**

The current version of the cluster.

**Type**: string
**Required**: True

## UpdateConnectivityResponse

Response body for UpdateConnectivity.

**clusterArn**

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

**clusterOperationArn**

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

## VpcConnectivity

VPC connection control settings for brokers.

**clientAuthentication**

VPC connection control settings for brokers.

**Type**: VpcConnectivityClientAuthentication
**Required**: False

## VpcConnectivityClientAuthentication

Includes all client authentication information for VpcConnectivity.

**sasl**

Details for VpcConnectivity ClientAuthentication using SASL.

**Type**: VpcConnectivitySasl
**Required**: False

**tls**

Details for VpcConnectivity ClientAuthentication using TLS.

**Type**: VpcConnectivityTls
**Required**: False

## VpcConnectivityIAM

Details for SASL/IAM client authentication for VpcConnectivity.

**enabled**

SASL/IAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## VpcConnectivitySasl

Details for client authentication using SASL for VpcConnectivity.

**iam**

Details for ClientAuthentication using IAM for VpcConnectivity.

**Type**: VpcConnectivityIAM
**Required**: False

**scram**

Details for SASL/SCRAM client authentication for VpcConnectivity.

**Type**: VpcConnectivityScram
**Required**: False

## VpcConnectivityScram

Details for SASL/SCRAM client authentication for vpcConnectivity.

**enabled**

SASL/SCRAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## VpcConnectivityTls

Details for client authentication using TLS for vpcConnectivity.

**enabled**

TLS authentication is enabled or not.

**Type**: boolean
**Required**: False

## ZookeeperAccess

ZooKeeper access control settings for the cluster.

**enabled**

Specifies whether direct Apache ZooKeeper client access is enabled or disabled for the cluster.

**Type**: boolean
**Required**: False

## See also

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

### UpdateConnectivity

- 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

# Cluster Operation

Represents an operation that was performed on an MSK cluster.

## URI

`/v1/operations/clusterOperationArn`

## HTTP methods

### GET

**Operation ID:** `DescribeClusterOperation`

Returns a description of the cluster operation specified by the Amazon Resource Name (ARN).

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterOperationArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the MSK cluster operation. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | DescribeClusterOperationResponse | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying |

| Status code | Response model | Description |
| --- | --- | --- |
| | | your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning the correct headers.

#### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterOperationArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the MSK cluster operation. |

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Response bodies

#### DescribeClusterOperationResponse schema

```
{
  "clusterOperationInfo": {
    "clusterArn": "string",
    "creationTime": "string",
    "clientRequestId": "string",
```

```
    "operationState": "string",
    "sourceClusterInfo": {
      "encryptionInfo": {
        "encryptionInTransit": {
          "inCluster": boolean,
          "clientBroker": enum
        },
        "encryptionAtRest": {
          "dataVolumeKMSKeyId": "string"
        }
      },
      "configurationInfo": {
        "arn": "string",
        "revision": integer
      },
      "brokerCountUpdateInfo": {
        "createdBrokerIds": [
          number
        ],
        "deletedBrokerIds": [
          number
        ]
      },
      "instanceType": "string",
      "zookeeperAccess": {
        "enabled": boolean
      },
      "loggingInfo": {
        "brokerLogs": {
          "s3": {
            "bucket": "string",
            "prefix": "string",
            "enabled": boolean
          },
          "firehose": {
            "deliveryStream": "string",
            "enabled": boolean
          },
          "cloudWatchLogs": {
            "logGroup": "string",
            "enabled": boolean
          }
        },
        "authorizerLogs": {
```

```
          "s3": {
            "bucket": "string",
            "prefix": "string",
            "enabled": boolean
          },
          "firehose": {
            "deliveryStream": "string",
            "enabled": boolean
          },
          "cloudWatchLogs": {
            "logGroup": "string",
            "enabled": boolean
          }
        }
      },
      "brokerEBSVolumeInfo": [
        {
          "volumeSizeGB": integer,
          "provisionedThroughput": {
            "volumeThroughput": integer,
            "enabled": boolean
          },
          "kafkaBrokerNodeId": "string"
        }
      ],
      "numberOfBrokerNodes": integer,
      "enhancedMonitoring": enum,
      "storageMode": enum,
      "kafkaVersion": "string",
      "connectivityInfo": {
        "vpcConnectivity": {
          "clientAuthentication": {
            "sasl": {
              "iam": {
                "enabled": boolean
              },
              "scram": {
                "enabled": boolean
              }
            },
            "tls": {
              "enabled": boolean
            }
          }
```

```
        },
        "publicAccess": {
          "type": "string"
        },
        "networkType": "string"
      },
      "clientAuthentication": {
        "sasl": {
          "iam": {
            "enabled": boolean
          },
          "scram": {
            "enabled": boolean
          }
        },
        "unauthenticated": {
          "enabled": boolean
        },
        "tls": {
          "certificateAuthorityArnList": [
            "string"
          ],
          "enabled": boolean
        }
      },
      "openMonitoring": {
        "prometheus": {
          "nodeExporter": {
            "enabledInBroker": boolean
          },
          "jmxExporter": {
            "enabledInBroker": boolean
          }
        }
      },
      "rebalancing": {
        "status": enum
      }
    },
    "errorInfo": {
      "errorString": "string",
      "errorCode": "string"
    },
    "vpcConnectionInfo": {
```

```
      "owner": "string",
      "vpcConnectionArn": "string",
      "creationTime": "string",
      "userIdentity": {
        "principalId": "string",
        "type": enum
      }
    },
    "operationType": "string",
    "endTime": "string",
    "operationSteps": [
      {
        "stepName": "string",
        "stepInfo": {
          "stepStatus": "string"
        }
      }
    ],
    "operationArn": "string",
    "targetClusterInfo": {
      "encryptionInfo": {
        "encryptionInTransit": {
          "inCluster": boolean,
          "clientBroker": enum
        },
        "encryptionAtRest": {
          "dataVolumeKMSKeyId": "string"
        }
      },
      "configurationInfo": {
        "arn": "string",
        "revision": integer
      },
      "brokerCountUpdateInfo": {
        "createdBrokerIds": [
          number
        ],
        "deletedBrokerIds": [
          number
        ]
      },
      "instanceType": "string",
      "zookeeperAccess": {
        "enabled": boolean
```

```
    },
    "loggingInfo": {
      "brokerLogs": {
        "s3": {
          "bucket": "string",
          "prefix": "string",
          "enabled": boolean
        },
        "firehose": {
          "deliveryStream": "string",
          "enabled": boolean
        },
        "cloudWatchLogs": {
          "logGroup": "string",
          "enabled": boolean
        }
      },
      "authorizerLogs": {
        "s3": {
          "bucket": "string",
          "prefix": "string",
          "enabled": boolean
        },
        "firehose": {
          "deliveryStream": "string",
          "enabled": boolean
        },
        "cloudWatchLogs": {
          "logGroup": "string",
          "enabled": boolean
        }
      }
    },
    "brokerEBSVolumeInfo": [
      {
        "volumeSizeGB": integer,
        "provisionedThroughput": {
          "volumeThroughput": integer,
          "enabled": boolean
        },
        "kafkaBrokerNodeId": "string"
      }
    ],
    "numberOfBrokerNodes": integer,
```

```
      "enhancedMonitoring": enum,
      "storageMode": enum,
      "kafkaVersion": "string",
      "connectivityInfo": {
        "vpcConnectivity": {
          "clientAuthentication": {
            "sasl": {
              "iam": {
                "enabled": boolean
              },
              "scram": {
                "enabled": boolean
              }
            },
            "tls": {
              "enabled": boolean
            }
          }
        },
        "publicAccess": {
          "type": "string"
        },
        "networkType": "string"
      },
      "clientAuthentication": {
        "sasl": {
          "iam": {
            "enabled": boolean
          },
          "scram": {
            "enabled": boolean
          }
        },
        "unauthenticated": {
          "enabled": boolean
        },
        "tls": {
          "certificateAuthorityArnList": [
            "string"
          ],
          "enabled": boolean
        }
      },
      "openMonitoring": {
```

```
        "prometheus": {
          "nodeExporter": {
            "enabledInBroker": boolean
          },
          "jmxExporter": {
            "enabledInBroker": boolean
          }
        }
      },
      "rebalancing": {
        "status": enum
      }
    }
  }
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## AuthorizerLogs

Authorizer Logs details for cluster.

**s3**

S3 Log destination details.

**Type**: S3
**Required**: False

**firehose**

Firehose Log destination details.

**Type**: Firehose

**Required**: False

#### cloudWatchLogs

CloudWatch Log destination details.

**Type**: CloudWatchLogs
**Required**: False

## BrokerCountUpdateInfo

Contains the list of broker ids being changed during a broker count update.

#### createdBrokerIds

List of Kafka Broker IDs being created. If operation is INCREASE_BROKER_COUNT, the list contains numeric ids of brokers added by the operation.

**Type**: Array of type number
**Required**: False

#### deletedBrokerIds

List of Kafka Broker IDs being deleted. If operation is DECREASE_BROKER_COUNT, the list contains numeric ids of brokers removed by the operation.

**Type**: Array of type number
**Required**: False

## BrokerEBSVolumeInfo

Specifies the EBS volume upgrade information. The broker identifier must be set to the keyword ALL. This means the changes apply to all the brokers in the cluster.

#### volumeSizeGB

Size of the EBS volume to update.

**Type**: integer

**Required**: False

#### provisionedThroughput

EBS volume provisioned throughput information.

**Type**: ProvisionedThroughput
**Required**: False

#### kafkaBrokerNodeId

The ID of the broker to update. The only allowed value is ALL. This means that Amazon MSK applies the same storage update to all broker nodes.

**Type**: string
**Required**: True

### BrokerLogs

The broker logs configuration for this MSK cluster.

#### s3

Details of the Amazon S3 destination for broker logs.

**Type**: S3
**Required**: False

#### firehose

Details of the Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: Firehose
**Required**: False

#### cloudWatchLogs

Details of the CloudWatch Logs destination for broker logs.

**Type**: CloudWatchLogs
**Required**: False

## ClientAuthentication

Includes all client authentication information.

**sasl**

Details for client authentication using SASL. To turn on SASL, you must also turn on `EncryptionInTransit` by setting `inCluster` to true. You must set `clientBroker` to either `TLS` or `TLS_PLAINTEXT`. If you choose `TLS_PLAINTEXT`, then you must also set `unauthenticated` to true.

**Type**: Sasl
**Required**: False

**unauthenticated**

Details for ClientAuthentication using no authentication.

**Type**: Unauthenticated
**Required**: False

**tls**

Details for ClientAuthentication using TLS. To turn on TLS access control, you must also turn on `EncryptionInTransit` by setting `inCluster` to true and `clientBroker` to `TLS`.

**Type**: Tls
**Required**: False

## ClientBroker

Client-broker encryption in transit setting.

`TLS`
`TLS_PLAINTEXT`
`PLAINTEXT`

## CloudWatchLogs

Details of the CloudWatch Logs destination for broker logs.

### logGroup

The CloudWatch log group that is the destination for broker logs.

**Type**: string
**Required**: False

### enabled

Specifies whether broker logs get sent to the specified CloudWatch Logs destination.

**Type**: boolean
**Required**: True

## ClusterOperationInfo

Returns information about a cluster operation.

### clusterArn

ARN of the cluster.

**Type**: string
**Required**: False

### creationTime

The time that the operation was created.

**Type**: string
**Required**: False

### clientRequestId

The ID of the API request that triggered this operation.

**Type**: string
**Required**: False

### operationState

State of the cluster operation.

**Type**: string
**Required**: False

### sourceClusterInfo

Information about cluster attributes before a cluster is updated.

**Type**: MutableClusterInfo
**Required**: False

### errorInfo

Describes the error if the operation fails.

**Type**: ErrorInfo
**Required**: False

### vpcConnectionInfo

Description of the VPC connection for CreateVpcConnection and DeleteVpcConnection operations.

**Type**: VpcConnectionInfo
**Required**: False

### operationType

Type of the cluster operation.

**Type**: string
**Required**: False

**endTime**

The time at which the operation finished.

**Type**: string
**Required**: False

**operationSteps**

Steps completed during the operation.

**Type**: Array of type ClusterOperationStep
**Required**: False

**operationArn**

ARN of the cluster operation.

**Type**: string
**Required**: False

**targetClusterInfo**

Information about cluster attributes after a cluster is updated.

**Type**: MutableClusterInfo
**Required**: False

## ClusterOperationStep

Step taken during a cluster operation.

**stepName**

The name of the step.

**Type**: string
**Required**: False

**stepInfo**

Information about the step and its status.

**Type**: ClusterOperationStepInfo
**Required**: False

## ClusterOperationStepInfo

Information about a step in an operation.

**stepStatus**

The step's current status.

**Type**: string
**Required**: False

## ConfigurationInfo

Specifies the configuration to use for the brokers.

**arn**

ARN of the configuration to use.

**Type**: string
**Required**: True

**revision**

The revision of the configuration to use.

**Type**: integer
**Required**: True
**Format**: int64
**Minimum**: 1

## ConnectivityInfo

Broker access controls.

### vpcConnectivity

VPC connection control settings for brokers

**Type**: VpcConnectivity
**Required**: False

### publicAccess

Access control settings for the cluster's brokers.

**Type**: PublicAccess
**Required**: False

### networkType

The network type of the cluster, which is IPv4 or DUAL. The DUAL network type uses both IPv4 and IPv6 addresses for your cluster and its resources. By default, a cluster uses the IPv4 network type.

**Type**: string
**Required**: False
**MinLength**: 4
**MaxLength**: 4

## DescribeClusterOperationResponse

Information about a cluster operation.

### clusterOperationInfo

Cluster operation information

**Type**: ClusterOperationInfo
**Required**: False

## EncryptionAtRest

The data-volume encryption details. You can't update encryption at rest settings for existing clusters.

**dataVolumeKMSKeyId**

The Amazon Resource Name (ARN) of the Amazon KMS key for encrypting data at rest. If you don't specify a KMS key, MSK creates one for you and uses it.

**Type**: string
**Required**: True

## EncryptionInTransit

The settings for encrypting data in transit.

**inCluster**

When set to true, it indicates that data communication among the broker nodes of the cluster is encrypted. When set to false, the communication happens in plaintext.

The default value is true.

**Type**: boolean
**Required**: False

**clientBroker**

Indicates the encryption setting for data in transit between clients and brokers. You must set it to one of the following values.

`TLS` means that client-broker communication is enabled with TLS only.

`TLS_PLAINTEXT` means that client-broker communication is enabled for both TLS-encrypted, as well as plaintext data.

`PLAINTEXT` means that client-broker communication is enabled in plaintext only.

The default value is `TLS`.

**Type**: ClientBroker
**Required**: False

## EncryptionInfo

Includes encryption-related information, such as the Amazon KMS key used for encrypting data at rest and whether you want MSK to encrypt your data in transit.

**encryptionInTransit**

The details for encryption in transit.

**Type**: EncryptionInTransit
**Required**: False

**encryptionAtRest**

The data-volume encryption details.

**Type**: EncryptionAtRest
**Required**: False

## EnhancedMonitoring

Specifies which Apache Kafka metrics Amazon MSK gathers and sends to Amazon CloudWatch for this cluster. This property has three possible values: `DEFAULT`, `PER_BROKER`, and `PER_TOPIC_PER_BROKER`. For a list of the metrics associated with each of these three levels of monitoring, see Monitoring.

`DEFAULT`
`PER_BROKER`
`PER_TOPIC_PER_BROKER`
`PER_TOPIC_PER_PARTITION`

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## ErrorInfo

Returns information about an error state of the cluster.

**errorString**

An optional field to provide more details about the error.

**Type**: string
**Required**: False

**errorCode**

A number describing the error programmatically.

**Type**: string
**Required**: False

## Firehose

Firehose details for BrokerLogs.

**deliveryStream**

The Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: string

**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Kinesis Data Firehose delivery stream.

**Type**: boolean
**Required**: True

## IAM

Details for SASL/IAM client authentication.

**enabled**

SASL/IAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## JmxExporter

Indicates whether you want to enable or disable the JMX Exporter.

**enabledInBroker**

Indicates whether you want to enable or disable the JMX Exporter.

**Type**: boolean
**Required**: True

## LoggingInfo

You can configure your MSK cluster to send broker logs to different destination types. This is a container for the configuration details related to broker logs.

**brokerLogs**

You can configure your MSK cluster to send broker logs to different destination types. This configuration specifies the details of these destinations.

**Type**: BrokerLogs
**Required**: False

**authorizerLogs**

Authorizer Logs details.

**Type**: AuthorizerLogs
**Required**: False

## MutableClusterInfo

Information about cluster attributes that can be updated via update APIs.

**encryptionInfo**

Encryption details.

**Type**: EncryptionInfo
**Required**: False

**configurationInfo**

Information about the changes in the configuration of the brokers.

**Type**: ConfigurationInfo
**Required**: False

**brokerCountUpdateInfo**

Describes brokers being changed during a broker count update.

**Type**: BrokerCountUpdateInfo
**Required**: False

**instanceType**

The broker type.

**Type**: string
**Required**: False

### zookeeperAccess

The ZooKeeper access setting for the cluster.

**Type**: ZookeeperAccess
**Required**: False

### loggingInfo

LoggingInfo details.

**Type**: LoggingInfo
**Required**: False

### brokerEBSVolumeInfo

Specifies the size of the EBS volume and the ID of the associated broker.

**Type**: Array of type BrokerEBSVolumeInfo
**Required**: False

### numberOfBrokerNodes

The number of broker nodes in the cluster.

**Type**: integer
**Required**: False

### enhancedMonitoring

The monitoring level.

**Type**: EnhancedMonitoring
**Required**: False

**storageMode**

This controls storage mode for supported storage tiers.

**Type**: StorageMode
**Required**: False

**kafkaVersion**

The Apache Kafka version.

**Type**: string
**Required**: False

**connectivityInfo**

Defines the connectivity setting of the cluster.

**Type**: ConnectivityInfo
**Required**: False

**clientAuthentication**

Client Authentication details.

**Type**: ClientAuthentication
**Required**: False

**openMonitoring**

Open monitoring details.

**Type**: OpenMonitoring
**Required**: False

**rebalancing**

Specifies if intelligent rebalancing is turned on for your cluster. The default intelligent rebalancing status is `ACTIVE` for all new MSK Provisioned clusters that you create with Express brokers.

**Type**: Rebalancing
**Required**: False

## NodeExporter

Indicates whether you want to enable or disable the Node Exporter.

**enabledInBroker**

Indicates whether you want to enable or disable the Node Exporter.

**Type**: boolean
**Required**: True

## OpenMonitoring

JMX and Node monitoring for the MSK cluster.

**prometheus**

Prometheus exporter settings.

**Type**: Prometheus
**Required**: True

## Prometheus

Prometheus settings for open monitoring.

**nodeExporter**

Indicates whether you want to enable or disable the Node Exporter.

**Type**: NodeExporter
**Required**: False

**jmxExporter**

Indicates whether you want to enable or disable the JMX Exporter.

**Type**: JmxExporter

**Required**: False

## ProvisionedThroughput

Contains information about provisioned throughput for EBS storage volumes attached to kafka broker nodes.

**volumeThroughput**

Throughput value of the EBS volumes for the data drive on each kafka broker node in MiB per second.

**Type**: integer
**Required**: False

**enabled**

Provisioned throughput is enabled or not.

**Type**: boolean
**Required**: False

## PublicAccess

Broker access controls

**type**

DISABLED means that public access is turned off. SERVICE_PROVIDED_EIPS means that public access is turned on.

**Type**: string
**Required**: False

## Rebalancing

Specifies whether or not intelligent rebalancing is turned on for a newly created MSK Provisioned cluster with Express brokers. Intelligent rebalancing performs automatic partition balancing operations when you scale your clusters up or down.

By default, intelligent rebalancing is `ACTIVE` for all new Express-based clusters.

**status**

Intelligent rebalancing status. The default intelligent rebalancing status is `ACTIVE` for all new Express-based clusters.

**Type**: RebalancingStatus
**Required**: True

## RebalancingStatus

Intelligent rebalancing status. The default intelligent rebalancing status is `ACTIVE` for all new Express-based clusters.

`PAUSED`
`ACTIVE`

## S3

The details of the Amazon S3 destination for broker logs.

**bucket**

The name of the S3 bucket that is the destination for broker logs.

**Type**: string
**Required**: False

**prefix**

The S3 prefix that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Amazon S3 destination.

**Type**: boolean
**Required**: True

## Sasl

Details for client authentication using SASL. To turn on SASL, you must also turn on `EncryptionInTransit` by setting `inCluster` to true. You must set `clientBroker` to either `TLS` or `TLS_PLAINTEXT`. If you choose `TLS_PLAINTEXT`, then you must also set `unauthenticated` to true.

**iam**

Details for ClientAuthentication using IAM.

**Type**: IAM
**Required**: False

**scram**

Details for SASL/SCRAM client authentication.

**Type**: Scram
**Required**: False

## Scram

Details for SASL/SCRAM client authentication.

**enabled**

SASL/SCRAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## StorageMode

Controls storage mode for various supported storage tiers.

LOCAL
TIERED

## Tls

Details for client authentication using TLS.

**certificateAuthorityArnList**

List of AWS Private CA Amazon Resource Name (ARN)s.

**Type**: Array of type string
**Required**: False

**enabled**

TLS authentication is enabled or not.

**Type**: boolean
**Required**: False

## Unauthenticated

Details for allowing no client authentication.

**enabled**

Unauthenticated is enabled or not.

**Type**: boolean
**Required**: False

## UserIdentity

Description of the requester that calls the API operation.

**principalId**

A unique identifier for the requester that calls the API operation.

**Type**: string
**Required**: False

#### type

The identity type of the requester that calls the API operation.

**Type**: UserIdentityType
**Required**: False

## UserIdentityType

The identity type of the requester that calls the API operation.

```
AWSAccount
AWSService
```

## VpcConnectionInfo

Description of the VPC connection.

#### owner

The owner of the VPC Connection.

**Type**: string
**Required**: False

#### vpcConnectionArn

The Amazon Resource Name (ARN) of the VPC connection.

**Type**: string
**Required**: False

#### creationTime

The time when Amazon MSK creates the VPC Connnection.

**Type**: string
**Required**: False

**userIdentity**

Description of the requester that calls the API operation.

**Type**: UserIdentity
**Required**: False

## VpcConnectivity

VPC connection control settings for brokers.

**clientAuthentication**

VPC connection control settings for brokers.

**Type**: VpcConnectivityClientAuthentication
**Required**: False

## VpcConnectivityClientAuthentication

Includes all client authentication information for VpcConnectivity.

**sasl**

Details for VpcConnectivity ClientAuthentication using SASL.

**Type**: VpcConnectivitySasl
**Required**: False

**tls**

Details for VpcConnectivity ClientAuthentication using TLS.

**Type**: VpcConnectivityTls
**Required**: False

## VpcConnectivityIAM

Details for SASL/IAM client authentication for VpcConnectivity.

**enabled**

SASL/IAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## VpcConnectivitySasl

Details for client authentication using SASL for VpcConnectivity.

**iam**

Details for ClientAuthentication using IAM for VpcConnectivity.

**Type**: VpcConnectivityIAM
**Required**: False

**scram**

Details for SASL/SCRAM client authentication for VpcConnectivity.

**Type**: VpcConnectivityScram
**Required**: False

## VpcConnectivityScram

Details for SASL/SCRAM client authentication for vpcConnectivity.

**enabled**

SASL/SCRAM authentication is enabled or not.

**Type**: boolean
**Required**: False

### VpcConnectivityTls

Details for client authentication using TLS for vpcConnectivity.

**enabled**

TLS authentication is enabled or not.

**Type**: boolean
**Required**: False

### ZookeeperAccess

ZooKeeper access control settings for the cluster.

**enabled**

Specifies whether direct Apache ZooKeeper client access is enabled or disabled for the cluster.

**Type**: boolean
**Required**: False

## See also

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

### DescribeClusterOperation

- 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

# Cluster Operations

Represents the operations that have been performed on an MSK cluster.

## URI

`/v1/clusters/clusterArn/operations`

## HTTP methods

### GET

**Operation ID:** `ListClusterOperations`

Returns a list of all the operations that have been performed on the specified MSK cluster.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Query parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| nextToken | String | False | The paginated results marker. When the result of the operation is truncated, the call |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | returns `NextToken` in the response. To get the next batch, provide this token in your next request. |
| `maxResults` | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes a `NextToken` parameter. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | `ListClusterOperationsResponse` | Successful response. |
| 400 | `Error` | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | `Error` | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | `Error` | Access forbidden. Check your credentials and then retry your request. |

| Status code | Response model | Description |
| --- | --- | --- |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

# Schemas

## Response bodies

### ListClusterOperationsResponse schema

```
{
  "nextToken": "string",
  "clusterOperationInfoList": [
    {
      "clusterArn": "string",
      "creationTime": "string",
      "clientRequestId": "string",
      "operationState": "string",
      "sourceClusterInfo": {
        "encryptionInfo": {
          "encryptionInTransit": {
            "inCluster": boolean,
            "clientBroker": enum
          },
          "encryptionAtRest": {
            "dataVolumeKMSKeyId": "string"
          }
        },
        "configurationInfo": {
          "arn": "string",
          "revision": integer
        },
        "brokerCountUpdateInfo": {
          "createdBrokerIds": [
            number
          ],
          "deletedBrokerIds": [
            number
          ]
        },
        "instanceType": "string",
        "zookeeperAccess": {
          "enabled": boolean
        },
        "loggingInfo": {
          "brokerLogs": {
            "s3": {
```

```
            "bucket": "string",
            "prefix": "string",
            "enabled": boolean
          },
          "firehose": {
            "deliveryStream": "string",
            "enabled": boolean
          },
          "cloudWatchLogs": {
            "logGroup": "string",
            "enabled": boolean
          }
        },
        "authorizerLogs": {
          "s3": {
            "bucket": "string",
            "prefix": "string",
            "enabled": boolean
          },
          "firehose": {
            "deliveryStream": "string",
            "enabled": boolean
          },
          "cloudWatchLogs": {
            "logGroup": "string",
            "enabled": boolean
          }
        }
      },
      "brokerEBSVolumeInfo": [
        {
          "volumeSizeGB": integer,
          "provisionedThroughput": {
            "volumeThroughput": integer,
            "enabled": boolean
          },
          "kafkaBrokerNodeId": "string"
        }
      ],
      "numberOfBrokerNodes": integer,
      "enhancedMonitoring": enum,
      "storageMode": enum,
      "kafkaVersion": "string",
      "connectivityInfo": {
```

```
          "vpcConnectivity": {
            "clientAuthentication": {
              "sasl": {
                "iam": {
                  "enabled": boolean
                },
                "scram": {
                  "enabled": boolean
                }
              },
              "tls": {
                "enabled": boolean
              }
            }
          },
          "publicAccess": {
            "type": "string"
          },
          "networkType": "string"
        },
        "clientAuthentication": {
          "sasl": {
            "iam": {
              "enabled": boolean
            },
            "scram": {
              "enabled": boolean
            }
          },
          "unauthenticated": {
            "enabled": boolean
          },
          "tls": {
            "certificateAuthorityArnList": [
              "string"
            ],
            "enabled": boolean
          }
        },
        "openMonitoring": {
          "prometheus": {
            "nodeExporter": {
              "enabledInBroker": boolean
            },
```

```
            "jmxExporter": {
              "enabledInBroker": boolean
            }
          }
        },
        "rebalancing": {
          "status": enum
        }
      },
      "errorInfo": {
        "errorString": "string",
        "errorCode": "string"
      },
      "vpcConnectionInfo": {
        "owner": "string",
        "vpcConnectionArn": "string",
        "creationTime": "string",
        "userIdentity": {
          "principalId": "string",
          "type": enum
        }
      },
      "operationType": "string",
      "endTime": "string",
      "operationSteps": [
        {
          "stepName": "string",
          "stepInfo": {
            "stepStatus": "string"
          }
        }
      ],
      "operationArn": "string",
      "targetClusterInfo": {
        "encryptionInfo": {
          "encryptionInTransit": {
            "inCluster": boolean,
            "clientBroker": enum
          },
          "encryptionAtRest": {
            "dataVolumeKMSKeyId": "string"
          }
        },
        "configurationInfo": {
```

```
      "arn": "string",
      "revision": integer
    },
    "brokerCountUpdateInfo": {
      "createdBrokerIds": [
        number
      ],
      "deletedBrokerIds": [
        number
      ]
    },
    "instanceType": "string",
    "zookeeperAccess": {
      "enabled": boolean
    },
    "loggingInfo": {
      "brokerLogs": {
        "s3": {
          "bucket": "string",
          "prefix": "string",
          "enabled": boolean
        },
        "firehose": {
          "deliveryStream": "string",
          "enabled": boolean
        },
        "cloudWatchLogs": {
          "logGroup": "string",
          "enabled": boolean
        }
      },
      "authorizerLogs": {
        "s3": {
          "bucket": "string",
          "prefix": "string",
          "enabled": boolean
        },
        "firehose": {
          "deliveryStream": "string",
          "enabled": boolean
        },
        "cloudWatchLogs": {
          "logGroup": "string",
          "enabled": boolean
```

```
            }
          }
        },
        "brokerEBSVolumeInfo": [
          {
            "volumeSizeGB": integer,
            "provisionedThroughput": {
              "volumeThroughput": integer,
              "enabled": boolean
            },
            "kafkaBrokerNodeId": "string"
          }
        ],
        "numberOfBrokerNodes": integer,
        "enhancedMonitoring": enum,
        "storageMode": enum,
        "kafkaVersion": "string",
        "connectivityInfo": {
          "vpcConnectivity": {
            "clientAuthentication": {
              "sasl": {
                "iam": {
                  "enabled": boolean
                },
                "scram": {
                  "enabled": boolean
                }
              },
              "tls": {
                "enabled": boolean
              }
            }
          },
          "publicAccess": {
            "type": "string"
          },
          "networkType": "string"
        },
        "clientAuthentication": {
          "sasl": {
            "iam": {
              "enabled": boolean
            },
            "scram": {
```

```
            "enabled": boolean
          }
        },
        "unauthenticated": {
          "enabled": boolean
        },
        "tls": {
          "certificateAuthorityArnList": [
            "string"
          ],
          "enabled": boolean
        }
      },
      "openMonitoring": {
        "prometheus": {
          "nodeExporter": {
            "enabledInBroker": boolean
          },
          "jmxExporter": {
            "enabledInBroker": boolean
          }
        }
      },
      "rebalancing": {
        "status": enum
      }
    }
  }
]
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## AuthorizerLogs

Authorizer Logs details for cluster.

**s3**

S3 Log destination details.

**Type**: S3
**Required**: False

**firehose**

Firehose Log destination details.

**Type**: Firehose
**Required**: False

**cloudWatchLogs**

CloudWatch Log destination details.

**Type**: CloudWatchLogs
**Required**: False

## BrokerCountUpdateInfo

Contains the list of broker ids being changed during a broker count update.

**createdBrokerIds**

List of Kafka Broker IDs being created. If operation is INCREASE_BROKER_COUNT, the list contains numeric ids of brokers added by the operation.

**Type**: Array of type number
**Required**: False

**deletedBrokerIds**

List of Kafka Broker IDs being deleted. If operation is DECREASE_BROKER_COUNT, the list contains numeric ids of brokers removed by the operation.

**Type**: Array of type number

**Required**: False

## BrokerEBSVolumeInfo

Specifies the EBS volume upgrade information. The broker identifier must be set to the keyword ALL. This means the changes apply to all the brokers in the cluster.

**volumeSizeGB**

Size of the EBS volume to update.

**Type**: integer
**Required**: False

**provisionedThroughput**

EBS volume provisioned throughput information.

**Type**: ProvisionedThroughput
**Required**: False

**kafkaBrokerNodeId**

The ID of the broker to update. The only allowed value is ALL. This means that Amazon MSK applies the same storage update to all broker nodes.

**Type**: string
**Required**: True

## BrokerLogs

The broker logs configuration for this MSK cluster.

**s3**

Details of the Amazon S3 destination for broker logs.

**Type**: S3
**Required**: False

**firehose**

Details of the Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: Firehose
**Required**: False

**cloudWatchLogs**

Details of the CloudWatch Logs destination for broker logs.

**Type**: CloudWatchLogs
**Required**: False

## ClientAuthentication

Includes all client authentication information.

**sasl**

Details for client authentication using SASL. To turn on SASL, you must also turn on `EncryptionInTransit` by setting `inCluster` to true. You must set `clientBroker` to either `TLS` or `TLS_PLAINTEXT`. If you choose `TLS_PLAINTEXT`, then you must also set `unauthenticated` to true.

**Type**: Sasl
**Required**: False

**unauthenticated**

Details for ClientAuthentication using no authentication.

**Type**: Unauthenticated
**Required**: False

**tls**

Details for ClientAuthentication using TLS. To turn on TLS access control, you must also turn on `EncryptionInTransit` by setting `inCluster` to true and `clientBroker` to `TLS`.

**Type**: Tls
**Required**: False

## ClientBroker

Client-broker encryption in transit setting.

TLS
TLS_PLAINTEXT
PLAINTEXT

## CloudWatchLogs

Details of the CloudWatch Logs destination for broker logs.

### logGroup

The CloudWatch log group that is the destination for broker logs.

**Type**: string
**Required**: False

### enabled

Specifies whether broker logs get sent to the specified CloudWatch Logs destination.

**Type**: boolean
**Required**: True

## ClusterOperationInfo

Returns information about a cluster operation.

### clusterArn

ARN of the cluster.

**Type**: string
**Required**: False

**creationTime**

The time that the operation was created.

**Type**: string
**Required**: False

**clientRequestId**

The ID of the API request that triggered this operation.

**Type**: string
**Required**: False

**operationState**

State of the cluster operation.

**Type**: string
**Required**: False

**sourceClusterInfo**

Information about cluster attributes before a cluster is updated.

**Type**: MutableClusterInfo
**Required**: False

**errorInfo**

Describes the error if the operation fails.

**Type**: ErrorInfo
**Required**: False

**vpcConnectionInfo**

Description of the VPC connection for CreateVpcConnection and DeleteVpcConnection operations.

**Type**: VpcConnectionInfo
**Required**: False

**operationType**

Type of the cluster operation.

**Type**: string
**Required**: False

**endTime**

The time at which the operation finished.

**Type**: string
**Required**: False

**operationSteps**

Steps completed during the operation.

**Type**: Array of type ClusterOperationStep
**Required**: False

**operationArn**

ARN of the cluster operation.

**Type**: string
**Required**: False

**targetClusterInfo**

Information about cluster attributes after a cluster is updated.

**Type**: MutableClusterInfo
**Required**: False

## ClusterOperationStep

Step taken during a cluster operation.

**stepName**

The name of the step.

**Type**: string
**Required**: False

**stepInfo**

Information about the step and its status.

**Type**: ClusterOperationStepInfo
**Required**: False

## ClusterOperationStepInfo

Information about a step in an operation.

**stepStatus**

The step's current status.

**Type**: string
**Required**: False

## ConfigurationInfo

Specifies the configuration to use for the brokers.

**arn**

ARN of the configuration to use.

**Type**: string
**Required**: True

**revision**

The revision of the configuration to use.

**Type**: integer
**Required**: True
**Format**: int64
**Minimum**: 1

## ConnectivityInfo

Broker access controls.

**vpcConnectivity**

VPC connection control settings for brokers

**Type**: VpcConnectivity
**Required**: False

**publicAccess**

Access control settings for the cluster's brokers.

**Type**: PublicAccess
**Required**: False

**networkType**

The network type of the cluster, which is IPv4 or DUAL. The DUAL network type uses both IPv4 and IPv6 addresses for your cluster and its resources. By default, a cluster uses the IPv4 network type.

**Type**: string
**Required**: False
**MinLength**: 4
**MaxLength**: 4

## EncryptionAtRest

The data-volume encryption details. You can't update encryption at rest settings for existing clusters.

**dataVolumeKMSKeyId**

The Amazon Resource Name (ARN) of the Amazon KMS key for encrypting data at rest. If you don't specify a KMS key, MSK creates one for you and uses it.

**Type**: string
**Required**: True

## EncryptionInTransit

The settings for encrypting data in transit.

**inCluster**

When set to true, it indicates that data communication among the broker nodes of the cluster is encrypted. When set to false, the communication happens in plaintext.

The default value is true.

**Type**: boolean
**Required**: False

**clientBroker**

Indicates the encryption setting for data in transit between clients and brokers. You must set it to one of the following values.

`TLS` means that client-broker communication is enabled with TLS only.

`TLS_PLAINTEXT` means that client-broker communication is enabled for both TLS-encrypted, as well as plaintext data.

`PLAINTEXT` means that client-broker communication is enabled in plaintext only.

The default value is `TLS`.

**Type**: ClientBroker
**Required**: False

## EncryptionInfo

Includes encryption-related information, such as the Amazon KMS key used for encrypting data at rest and whether you want MSK to encrypt your data in transit.

**encryptionInTransit**

The details for encryption in transit.

**Type**: EncryptionInTransit
**Required**: False

**encryptionAtRest**

The data-volume encryption details.

**Type**: EncryptionAtRest
**Required**: False

## EnhancedMonitoring

Specifies which Apache Kafka metrics Amazon MSK gathers and sends to Amazon CloudWatch for this cluster. This property has three possible values: `DEFAULT`, `PER_BROKER`, and `PER_TOPIC_PER_BROKER`. For a list of the metrics associated with each of these three levels of monitoring, see Monitoring.

`DEFAULT`
`PER_BROKER`
`PER_TOPIC_PER_BROKER`
`PER_TOPIC_PER_PARTITION`

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## ErrorInfo

Returns information about an error state of the cluster.

**errorString**

An optional field to provide more details about the error.

**Type**: string
**Required**: False

**errorCode**

A number describing the error programmatically.

**Type**: string
**Required**: False

## Firehose

Firehose details for BrokerLogs.

**deliveryStream**

The Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Kinesis Data Firehose delivery stream.

**Type**: boolean
**Required**: True

## IAM

Details for SASL/IAM client authentication.

**enabled**

SASL/IAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## JmxExporter

Indicates whether you want to enable or disable the JMX Exporter.

**enabledInBroker**

Indicates whether you want to enable or disable the JMX Exporter.

**Type**: boolean
**Required**: True

## ListClusterOperationsResponse

The response contains an array containing cluster operation information and a next token if the response is truncated. Displays operations of the type DECREASE_BROKER_COUNT in results list (after a successful broker removal operation).

**nextToken**

If the response of ListClusterOperations is truncated, Amazon MSK returns a NextToken in the response. Send this NextToken in your subsequent request to ListClusterOperations.

**Type**: string
**Required**: False

**clusterOperationInfoList**

An array of cluster operation information objects.

**Type**: Array of type ClusterOperationInfo
**Required**: False

## LoggingInfo

You can configure your MSK cluster to send broker logs to different destination types. This is a container for the configuration details related to broker logs.

**brokerLogs**

You can configure your MSK cluster to send broker logs to different destination types. This configuration specifies the details of these destinations.

**Type**: BrokerLogs
**Required**: False

**authorizerLogs**

Authorizer Logs details.

**Type**: AuthorizerLogs
**Required**: False

## MutableClusterInfo

Information about cluster attributes that can be updated via update APIs.

**encryptionInfo**

Encryption details.

**Type**: EncryptionInfo
**Required**: False

**configurationInfo**

Information about the changes in the configuration of the brokers.

**Type**: ConfigurationInfo
**Required**: False

**brokerCountUpdateInfo**

Describes brokers being changed during a broker count update.

**Type**: BrokerCountUpdateInfo
**Required**: False

**instanceType**

The broker type.

**Type**: string
**Required**: False

**zookeeperAccess**

The ZooKeeper access setting for the cluster.

**Type**: ZookeeperAccess
**Required**: False

**loggingInfo**

LoggingInfo details.

**Type**: LoggingInfo
**Required**: False

**brokerEBSVolumeInfo**

Specifies the size of the EBS volume and the ID of the associated broker.

**Type**: Array of type BrokerEBSVolumeInfo
**Required**: False

**numberOfBrokerNodes**

The number of broker nodes in the cluster.

**Type**: integer
**Required**: False

**enhancedMonitoring**

The monitoring level.

**Type**: EnhancedMonitoring
**Required**: False

**storageMode**

This controls storage mode for supported storage tiers.

**Type**: StorageMode
**Required**: False

**kafkaVersion**

The Apache Kafka version.

**Type**: string
**Required**: False

**connectivityInfo**

Defines the connectivity setting of the cluster.

**Type**: ConnectivityInfo
**Required**: False

**clientAuthentication**

Client Authentication details.

**Type**: ClientAuthentication
**Required**: False

**openMonitoring**

Open monitoring details.

**Type**: OpenMonitoring
**Required**: False

**rebalancing**

Specifies if intelligent rebalancing is turned on for your cluster. The default intelligent rebalancing status is `ACTIVE` for all new MSK Provisioned clusters that you create with Express brokers.

**Type**: Rebalancing
**Required**: False

## NodeExporter

Indicates whether you want to enable or disable the Node Exporter.

**enabledInBroker**

Indicates whether you want to enable or disable the Node Exporter.

**Type**: boolean
**Required**: True

## OpenMonitoring

JMX and Node monitoring for the MSK cluster.

**prometheus**

Prometheus exporter settings.

**Type**: Prometheus
**Required**: True

## Prometheus

Prometheus settings for open monitoring.

**nodeExporter**

Indicates whether you want to enable or disable the Node Exporter.

**Type**: NodeExporter
**Required**: False

**jmxExporter**

Indicates whether you want to enable or disable the JMX Exporter.

**Type**: JmxExporter
**Required**: False

## ProvisionedThroughput

Contains information about provisioned throughput for EBS storage volumes attached to kafka broker nodes.

**volumeThroughput**

Throughput value of the EBS volumes for the data drive on each kafka broker node in MiB per second.

**Type**: integer
**Required**: False

**enabled**

Provisioned throughput is enabled or not.

**Type**: boolean
**Required**: False

## PublicAccess

Broker access controls

**type**

DISABLED means that public access is turned off. SERVICE_PROVIDED_EIPS means that public access is turned on.

**Type**: string
**Required**: False

## Rebalancing

Specifies whether or not intelligent rebalancing is turned on for a newly created MSK Provisioned cluster with Express brokers. Intelligent rebalancing performs automatic partition balancing operations when you scale your clusters up or down.

By default, intelligent rebalancing is `ACTIVE` for all new Express-based clusters.

**status**

Intelligent rebalancing status. The default intelligent rebalancing status is `ACTIVE` for all new Express-based clusters.

**Type**: RebalancingStatus
**Required**: True

## RebalancingStatus

Intelligent rebalancing status. The default intelligent rebalancing status is `ACTIVE` for all new Express-based clusters.

PAUSED
ACTIVE

## S3

The details of the Amazon S3 destination for broker logs.

**bucket**

The name of the S3 bucket that is the destination for broker logs.

**Type**: string
**Required**: False

**prefix**

The S3 prefix that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Amazon S3 destination.

**Type**: boolean
**Required**: True

## Sasl

Details for client authentication using SASL. To turn on SASL, you must also turn on `EncryptionInTransit` by setting `inCluster` to true. You must set `clientBroker` to either `TLS` or `TLS_PLAINTEXT`. If you choose `TLS_PLAINTEXT`, then you must also set `unauthenticated` to true.

**iam**

Details for ClientAuthentication using IAM.

**Type**: IAM
**Required**: False

**scram**

Details for SASL/SCRAM client authentication.

**Type**: Scram
**Required**: False

## Scram

Details for SASL/SCRAM client authentication.

**enabled**

SASL/SCRAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## StorageMode

Controls storage mode for various supported storage tiers.

```
LOCAL
TIERED
```

## Tls

Details for client authentication using TLS.

**certificateAuthorityArnList**

List of AWS Private CA Amazon Resource Name (ARN)s.

**Type**: Array of type string
**Required**: False

**enabled**

TLS authentication is enabled or not.

**Type**: boolean
**Required**: False

## Unauthenticated

Details for allowing no client authentication.

**enabled**

Unauthenticated is enabled or not.

**Type**: boolean
**Required**: False

## UserIdentity

Description of the requester that calls the API operation.

**principalId**

A unique identifier for the requester that calls the API operation.

**Type**: string
**Required**: False

**type**

The identity type of the requester that calls the API operation.

**Type**: UserIdentityType
**Required**: False

## UserIdentityType

The identity type of the requester that calls the API operation.

AWSAccount
AWSService

## VpcConnectionInfo

Description of the VPC connection.

**owner**

The owner of the VPC Connection.

**Type**: string
**Required**: False

**vpcConnectionArn**

The Amazon Resource Name (ARN) of the VPC connection.

**Type**: string
**Required**: False

**creationTime**

The time when Amazon MSK creates the VPC Connnection.

**Type**: string
**Required**: False

**userIdentity**

Description of the requester that calls the API operation.

**Type**: UserIdentity
**Required**: False

## VpcConnectivity

VPC connection control settings for brokers.

**clientAuthentication**

VPC connection control settings for brokers.

**Type**: VpcConnectivityClientAuthentication
**Required**: False

## VpcConnectivityClientAuthentication

Includes all client authentication information for VpcConnectivity.

**sasl**

Details for VpcConnectivity ClientAuthentication using SASL.

**Type**: VpcConnectivitySasl
**Required**: False

**tls**

Details for VpcConnectivity ClientAuthentication using TLS.

**Type**: VpcConnectivityTls
**Required**: False

## VpcConnectivityIAM

Details for SASL/IAM client authentication for VpcConnectivity.

**enabled**

SASL/IAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## VpcConnectivitySasl

Details for client authentication using SASL for VpcConnectivity.

**iam**

Details for ClientAuthentication using IAM for VpcConnectivity.

**Type**: VpcConnectivityIAM
**Required**: False

**scram**

Details for SASL/SCRAM client authentication for VpcConnectivity.

**Type**: VpcConnectivityScram
**Required**: False

## VpcConnectivityScram

Details for SASL/SCRAM client authentication for vpcConnectivity.

**enabled**

SASL/SCRAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## VpcConnectivityTls

Details for client authentication using TLS for vpcConnectivity.

**enabled**

TLS authentication is enabled or not.

**Type**: boolean
**Required**: False

## ZookeeperAccess

ZooKeeper access control settings for the cluster.

**enabled**

Specifies whether direct Apache ZooKeeper client access is enabled or disabled for the cluster.

**Type**: boolean

**Required**: False

## See also

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

### ListClusterOperations

- 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

# Cluster Security

Use this resource to update the security settings of a cluster.

## URI

`/v1/clusters/clusterArn/security`

# HTTP methods

## OPTIONS

Enable CORS by returning the correct headers.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | None | Default response for CORS method |

## PATCH

**Operation ID:** `UpdateSecurity`

Updates security settings of the specified cluster.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | UpdateSecurityResponse | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

# Schemas

## Request bodies

### PATCH schema

```
{
  "encryptionInfo": {
    "encryptionInTransit": {
      "inCluster": boolean,
      "clientBroker": enum
    },
    "encryptionAtRest": {
      "dataVolumeKMSKeyId": "string"
    }
  },
  "clientAuthentication": {
    "sasl": {
      "iam": {
        "enabled": boolean
      },
      "scram": {
        "enabled": boolean
      }
    },
    "unauthenticated": {
      "enabled": boolean
    },
    "tls": {
      "certificateAuthorityArnList": [
        "string"
      ],
      "enabled": boolean
    }
  },
  "currentVersion": "string"
}
```

## Response bodies

### UpdateSecurityResponse schema

```
{
```

```
  "clusterArn": "string",
  "clusterOperationArn": "string"
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## ClientAuthentication

Includes all client authentication information.

**sasl**

Details for client authentication using SASL. To turn on SASL, you must also turn on `EncryptionInTransit` by setting `inCluster` to true. You must set `clientBroker` to either `TLS` or `TLS_PLAINTEXT`. If you choose `TLS_PLAINTEXT`, then you must also set `unauthenticated` to true.

**Type**: Sasl
**Required**: False

**unauthenticated**

Details for ClientAuthentication using no authentication.

**Type**: Unauthenticated
**Required**: False

**tls**

Details for ClientAuthentication using TLS. To turn on TLS access control, you must also turn on `EncryptionInTransit` by setting `inCluster` to true and `clientBroker` to `TLS`.

**Type**: Tls

**Required**: False

## ClientBroker

Client-broker encryption in transit setting.

TLS

TLS_PLAINTEXT

PLAINTEXT

## EncryptionAtRest

The data-volume encryption details. You can't update encryption at rest settings for existing clusters.

### dataVolumeKMSKeyId

The Amazon Resource Name (ARN) of the Amazon KMS key for encrypting data at rest. If you don't specify a KMS key, MSK creates one for you and uses it.

**Type**: string

**Required**: True

## EncryptionInTransit

The settings for encrypting data in transit.

### inCluster

When set to true, it indicates that data communication among the broker nodes of the cluster is encrypted. When set to false, the communication happens in plaintext.

The default value is true.

**Type**: boolean

**Required**: False

**clientBroker**

Indicates the encryption setting for data in transit between clients and brokers. You must set it to one of the following values.

`TLS` means that client-broker communication is enabled with TLS only.

`TLS_PLAINTEXT` means that client-broker communication is enabled for both TLS-encrypted, as well as plaintext data.

`PLAINTEXT` means that client-broker communication is enabled in plaintext only.

The default value is `TLS`.

**Type**: ClientBroker
**Required**: False

## EncryptionInfo

Includes encryption-related information, such as the Amazon KMS key used for encrypting data at rest and whether you want MSK to encrypt your data in transit.

**encryptionInTransit**

The details for encryption in transit.

**Type**: EncryptionInTransit
**Required**: False

**encryptionAtRest**

The data-volume encryption details.

**Type**: EncryptionAtRest
**Required**: False

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## IAM

Details for SASL/IAM client authentication.

**enabled**

SASL/IAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## Sasl

Details for client authentication using SASL. To turn on SASL, you must also turn on `EncryptionInTransit` by setting `inCluster` to true. You must set `clientBroker` to either `TLS` or `TLS_PLAINTEXT`. If you choose `TLS_PLAINTEXT`, then you must also set `unauthenticated` to true.

**iam**

Details for ClientAuthentication using IAM.

**Type**: IAM
**Required**: False

**scram**

Details for SASL/SCRAM client authentication.

**Type**: Scram
**Required**: False

## Scram

Details for SASL/SCRAM client authentication.

**enabled**

SASL/SCRAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## Tls

Details for client authentication using TLS.

**certificateAuthorityArnList**

List of AWS Private CA Amazon Resource Name (ARN)s.

**Type**: Array of type string
**Required**: False

**enabled**

TLS authentication is enabled or not.

**Type**: boolean
**Required**: False

## Unauthenticated

Details for allowing no client authentication.

**enabled**

Unauthenticated is enabled or not.

**Type**: boolean
**Required**: False

## UpdateSecurityRequest

Request body for UpdateSecurity.

**encryptionInfo**

The encryption info details.

**Type**: EncryptionInfo
**Required**: False

**clientAuthentication**

The client authentication info details.

**Type**: ClientAuthentication
**Required**: False

**currentVersion**

Current cluster version.

**Type**: string
**Required**: True

## UpdateSecurityResponse

Response body for UpdateSecurity.

**clusterArn**

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

**clusterOperationArn**

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

## See also

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

### UpdateSecurity

- 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

# Clusters

Use this resource to create an Amazon MSK cluster and to get a list of existing clusters.

## URI

`/v1/clusters`

# HTTP methods

## GET

**Operation ID:** `ListClusters`

Returns a list of all the MSK clusters.

**Query parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| `nextToken` | String | False | The paginated results marker. When the result of the operation is truncated, the call returns `NextToken` in the response. To get the next batch, provide this token in your next request. |
| `clusterNameFilter` | String | False | Specify a prefix of the name of the clusters that you want to list. The service lists all the clusters whose names start with this prefix. |
| `maxResults` | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | a `NextToken` parameter. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | ListClustersResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## POST

**Operation ID:** `CreateCluster`

Creates a new MSK cluster. The following Python 3.6 examples shows how you can create a cluster that's distributed over two Availability Zones. Before you run this Python script, replace the example subnet and security-group IDs with the IDs of your subnets and security group. When you create an MSK cluster, its brokers get evenly distributed over a number of Availability Zones that's equal to the number of subnets that you specify in the `BrokerNodeGroupInfo` parameter. In this example, you can add a third subnet to get a cluster that's distributed over three Availability Zones.

```
import boto3

client = boto3.client('kafka')

response = client.create_cluster(
    BrokerNodeGroupInfo={
        'BrokerAZDistribution': 'DEFAULT',
        'ClientSubnets': [
            'subnet-012345678901fedcba',
            'subnet-9876543210abcdef01'
        ],
        'InstanceType': 'kafka.m5.large',
        'SecurityGroups': [
            'sg-012345abcdef789789'
        ]
    },
    ClusterName='SalesCluster',
    EncryptionInfo={
        'EncryptionInTransit': {
            'ClientBroker': 'TLS_PLAINTEXT',
            'InCluster': True
        }
    },
    EnhancedMonitoring='PER_TOPIC_PER_BROKER',
    KafkaVersion='2.2.1',
    NumberOfBrokerNodes=2
)

print(response)
```

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | CreateClusterResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 409 | Error | This cluster name already exists. Retry your request using another name. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning the correct headers.

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | The default response for a CORS method. |

## Schemas

### Request bodies

#### POST schema

```
{
  "encryptionInfo": {
    "encryptionInTransit": {
      "inCluster": boolean,
      "clientBroker": enum
    },
    "encryptionAtRest": {
      "dataVolumeKMSKeyId": "string"
    }
  },
  "configurationInfo": {
    "arn": "string",
    "revision": integer
  },
  "loggingInfo": {
    "brokerLogs": {
      "s3": {
        "bucket": "string",
        "prefix": "string",
        "enabled": boolean
      },
      "firehose": {
        "deliveryStream": "string",
        "enabled": boolean
      },
```

```
      "cloudWatchLogs": {
        "logGroup": "string",
        "enabled": boolean
      }
    },
    "authorizerLogs": {
      "s3": {
        "bucket": "string",
        "prefix": "string",
        "enabled": boolean
      },
      "firehose": {
        "deliveryStream": "string",
        "enabled": boolean
      },
      "cloudWatchLogs": {
        "logGroup": "string",
        "enabled": boolean
      }
    }
  },
  "tags": {
  },
  "numberOfBrokerNodes": integer,
  "enhancedMonitoring": enum,
  "storageMode": enum,
  "clusterName": "string",
  "kafkaVersion": "string",
  "clientAuthentication": {
    "sasl": {
      "iam": {
        "enabled": boolean
      },
      "scram": {
        "enabled": boolean
      }
    },
    "unauthenticated": {
      "enabled": boolean
    },
    "tls": {
      "certificateAuthorityArnList": [
        "string"
      ],
```

```
      "enabled": boolean
    }
  },
  "brokerNodeGroupInfo": {
    "clientSubnets": [
      "string"
    ],
    "zoneIds": [
      "string"
    ],
    "instanceType": "string",
    "connectivityInfo": {
      "vpcConnectivity": {
        "clientAuthentication": {
          "sasl": {
            "iam": {
              "enabled": boolean
            },
            "scram": {
              "enabled": boolean
            }
          },
          "tls": {
            "enabled": boolean
          }
        }
      },
      "publicAccess": {
        "type": "string"
      },
      "networkType": "string"
    },
    "securityGroups": [
      "string"
    ],
    "brokerAZDistribution": enum,
    "storageInfo": {
      "ebsStorageInfo": {
        "provisionedThroughput": {
          "volumeThroughput": integer,
          "enabled": boolean
        },
        "volumeSize": integer
      }
```

```
    }
  },
  "openMonitoring": {
    "prometheus": {
      "nodeExporter": {
        "enabledInBroker": boolean
      },
      "jmxExporter": {
        "enabledInBroker": boolean
      }
    }
  },
  "rebalancing": {
    "status": enum
  }
}
```

## Response bodies

### ListClustersResponse schema

```
{
  "nextToken": "string",
  "clusterInfoList": [
    {
      "encryptionInfo": {
        "encryptionInTransit": {
          "inCluster": boolean,
          "clientBroker": enum
        },
        "encryptionAtRest": {
          "dataVolumeKMSKeyId": "string"
        }
      },
      "zookeeperConnectString": "string",
      "creationTime": "string",
      "customerActionStatus": enum,
      "zookeeperConnectStringTls": "string",
      "loggingInfo": {
        "brokerLogs": {
          "s3": {
            "bucket": "string",
            "prefix": "string",
```

```
            "enabled": boolean
          },
          "firehose": {
            "deliveryStream": "string",
            "enabled": boolean
          },
          "cloudWatchLogs": {
            "logGroup": "string",
            "enabled": boolean
          }
        },
        "authorizerLogs": {
          "s3": {
            "bucket": "string",
            "prefix": "string",
            "enabled": boolean
          },
          "firehose": {
            "deliveryStream": "string",
            "enabled": boolean
          },
          "cloudWatchLogs": {
            "logGroup": "string",
            "enabled": boolean
          }
        }
      },
      "currentVersion": "string",
      "tags": {
      },
      "numberOfBrokerNodes": integer,
      "clusterArn": "string",
      "activeOperationArn": "string",
      "enhancedMonitoring": enum,
      "storageMode": enum,
      "clusterName": "string",
      "stateInfo": {
        "code": "string",
        "message": "string"
      },
      "clientAuthentication": {
        "sasl": {
          "iam": {
            "enabled": boolean
```

```
        },
        "scram": {
          "enabled": boolean
        }
      },
      "unauthenticated": {
        "enabled": boolean
      },
      "tls": {
        "certificateAuthorityArnList": [
          "string"
        ],
        "enabled": boolean
      }
    },
    "state": enum,
    "brokerNodeGroupInfo": {
      "clientSubnets": [
        "string"
      ],
      "zoneIds": [
        "string"
      ],
      "instanceType": "string",
      "connectivityInfo": {
        "vpcConnectivity": {
          "clientAuthentication": {
            "sasl": {
              "iam": {
                "enabled": boolean
              },
              "scram": {
                "enabled": boolean
              }
            },
            "tls": {
              "enabled": boolean
            }
          }
        },
        "publicAccess": {
          "type": "string"
        },
        "networkType": "string"
```

```
        },
        "securityGroups": [
          "string"
        ],
        "brokerAZDistribution": enum,
        "storageInfo": {
          "ebsStorageInfo": {
            "provisionedThroughput": {
              "volumeThroughput": integer,
              "enabled": boolean
            },
            "volumeSize": integer
          }
        }
      },
      "openMonitoring": {
        "prometheus": {
          "nodeExporter": {
            "enabledInBroker": boolean
          },
          "jmxExporter": {
            "enabledInBroker": boolean
          }
        }
      },
      "rebalancing": {
        "status": enum
      },
      "currentBrokerSoftwareInfo": {
        "configurationRevision": integer,
        "kafkaVersion": "string",
        "configurationArn": "string"
      }
    }
  ]
}
```

## CreateClusterResponse schema

```
{
  "clusterArn": "string",
  "clusterName": "string",
  "state": enum
```

```
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## AuthorizerLogs

Authorizer Logs details for cluster.

### s3

S3 Log destination details.

**Type**: S3
**Required**: False

### firehose

Firehose Log destination details.

**Type**: Firehose
**Required**: False

### cloudWatchLogs

CloudWatch Log destination details.

**Type**: CloudWatchLogs
**Required**: False

## BrokerAZDistribution

This parameter is currently not in use.

DEFAULT

## BrokerLogs

The broker logs configuration for this MSK cluster.

**s3**

Details of the Amazon S3 destination for broker logs.

**Type**: S3
**Required**: False

**firehose**

Details of the Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: Firehose
**Required**: False

**cloudWatchLogs**

Details of the CloudWatch Logs destination for broker logs.

**Type**: CloudWatchLogs
**Required**: False

## BrokerNodeGroupInfo

Describes the setup to be used for the broker nodes in the cluster.

**clientSubnets**

The list of subnets to connect to in the client virtual private cloud (VPC). Amazon creates elastic network interfaces inside these subnets. Client applications use elastic network interfaces to produce and consume data.

If you use the US West (N. California) Region, specify exactly two subnets. For other Regions where Amazon MSK is available, you can specify either two or three subnets. The subnets that you specify

must be in distinct Availability Zones. When you create a cluster, Amazon MSK distributes the broker nodes evenly across the subnets that you specify.

Client subnets can't occupy the Availability Zone with ID `use1-az3`.

**Type**: Array of type string
**Required**: True

**zoneIds**

The zoneIds for brokers in customer account.

**Type**: Array of type string
**Required**: False

**instanceType**

The type of Amazon EC2 instances to use for brokers. The following instance types are allowed: kafka.m5.large, kafka.m5.xlarge, kafka.m5.2xlarge, kafka.m5.4xlarge, kafka.m5.8xlarge, kafka.m5.12xlarge, kafka.m5.16xlarge, and kafka.m5.24xlarge.

**Type**: string
**Required**: True
**MinLength**: 5
**MaxLength**: 32

**connectivityInfo**

Information about the cluster's connectivity setting.

**Type**: ConnectivityInfo
**Required**: False

**securityGroups**

The security groups to associate with the elastic network interfaces in order to specify who can connect to and communicate with the Amazon MSK cluster. If you don't specify a security group, Amazon MSK uses the default security group associated with the VPC. If you specify security

groups that were shared with you, you must ensure that you have permissions to them. Specifically, you need the `ec2:DescribeSecurityGroups` permission.

**Type**: Array of type string
**Required**: False

#### brokerAZDistribution

This parameter is currently not in use.

**Type**: BrokerAZDistribution
**Required**: False

#### storageInfo

Contains information about storage volumes attached to Amazon MSK broker nodes.

**Type**: StorageInfo
**Required**: False

## BrokerSoftwareInfo

Information about the current software installed on the cluster.

#### configurationRevision

The revision of the configuration to use. This field isn't visible in this preview release.

**Type**: integer
**Required**: False
**Format**: int64

#### kafkaVersion

The version of Apache Kafka. You can use Amazon MSK to create clusters that use Apache Kafka versions 1.1.1 and 2.2.1. See Apache Kafka Versions.

**Type**: string

**Required**: False

**configurationArn**

The Amazon Resource Name (ARN) of the configuration used for the cluster. This field isn't visible in this preview release.

**Type**: string
**Required**: False

## ClientAuthentication

Includes all client authentication information.

**sasl**

Details for client authentication using SASL. To turn on SASL, you must also turn on `EncryptionInTransit` by setting `inCluster` to true. You must set `clientBroker` to either `TLS` or `TLS_PLAINTEXT`. If you choose `TLS_PLAINTEXT`, then you must also set `unauthenticated` to true.

**Type**: Sasl
**Required**: False

**unauthenticated**

Details for ClientAuthentication using no authentication.

**Type**: Unauthenticated
**Required**: False

**tls**

Details for ClientAuthentication using TLS. To turn on TLS access control, you must also turn on `EncryptionInTransit` by setting `inCluster` to true and `clientBroker` to `TLS`.

**Type**: Tls
**Required**: False

## ClientBroker

Client-broker encryption in transit setting.

TLS
TLS_PLAINTEXT
PLAINTEXT

## CloudWatchLogs

Details of the CloudWatch Logs destination for broker logs.

### logGroup

The CloudWatch log group that is the destination for broker logs.

**Type**: string
**Required**: False

### enabled

Specifies whether broker logs get sent to the specified CloudWatch Logs destination.

**Type**: boolean
**Required**: True

## ClusterInfo

Returns information about a cluster.

### encryptionInfo

Includes all encryption-related information.

**Type**: EncryptionInfo
**Required**: False

### zookeeperConnectString

The connection string to use to connect to zookeeper cluster on plaintext port.

**Type**: string
**Required**: False

### creationTime

The time when the cluster was created.

**Type**: string
**Required**: False

### customerActionStatus

Determines if there is an action required from the customer.

**Type**: CustomerActionStatus
**Required**: False

### zookeeperConnectStringTls

The connection string to use to connect to the Apache ZooKeeper cluster on a TLS port.

**Type**: string
**Required**: False

### loggingInfo

You can configure your MSK cluster to send broker logs to different destination types. This is a container for the configuration details related to broker logs.

**Type**: LoggingInfo
**Required**: False

### currentVersion

The current version of the MSK cluster. Cluster versions aren't simple integers. You can obtain the current version by describing the cluster. An example version is `KTVPDKIKX0DER`.

**Type**: string

**Required**: False

**tags**

Tags attached to the cluster.

**Type**: object
**Required**: False

**numberOfBrokerNodes**

The number of broker nodes in the cluster.

**Type**: integer
**Required**: False

**clusterArn**

The Amazon Resource Name (ARN) that uniquely identifies the cluster.

**Type**: string
**Required**: False

**activeOperationArn**

Arn of active cluster operation.

**Type**: string
**Required**: False

**enhancedMonitoring**

Specifies which Apache Kafka metrics Amazon MSK gathers and sends to Amazon CloudWatch for this cluster. This property has three possible values: `DEFAULT`, `PER_BROKER`, and `PER_TOPIC_PER_BROKER`. For a list of the metrics associated with each of these three levels of monitoring, see Monitoring.

**Type**: EnhancedMonitoring
**Required**: False

**storageMode**

This controls storage mode for supported storage tiers.

**Type**: StorageMode
**Required**: False

**clusterName**

The name of the cluster.

**Type**: string
**Required**: False

**stateInfo**

Includes information of the cluster state.

**Type**: StateInfo
**Required**: False

**clientAuthentication**

Includes all client authentication information.

**Type**: ClientAuthentication
**Required**: False

**state**

The state of the cluster. Amazon MSK automatically renews certificates on clusters every 13 months. It sets the state of the cluster to `MAINTENANCE` when it starts the certificate-update operation. It sets it back to `ACTIVE` when the update is done. While a cluster is in the `MAINTENANCE` state, you can continue to produce and consume data, but you can't perform any update operations on it. You can perform update operations on a cluster when it is in the `ACTIVE` state.

**Type**: ClusterState
**Required**: False

**brokerNodeGroupInfo**

Information about the broker nodes.

**Type**: BrokerNodeGroupInfo
**Required**: False

**openMonitoring**

Settings for open monitoring using Prometheus.

**Type**: OpenMonitoring
**Required**: False

**rebalancing**

Specifies if intelligent rebalancing is turned on for your MSK Provisioned cluster with Express brokers. For all new Express-based clusters that you create, intelligent rebalancing is turned on by default.

**Type**: Rebalancing
**Required**: False

**currentBrokerSoftwareInfo**

Information about the version of software currently deployed on the brokers in the cluster.

**Type**: BrokerSoftwareInfo
**Required**: False

## ClusterState

The state of the cluster. Amazon MSK automatically renews certificates on clusters every 13 months. It sets the state of the cluster to `MAINTENANCE` when it starts the certificate-update operation. It sets it back to `ACTIVE` when the update is done. While a cluster is in the `MAINTENANCE` state, you can continue to produce and consume data, but you can't perform any update operations on it. You can perform update operations on a cluster when it is in the `ACTIVE` state.

```
ACTIVE
CREATING
UPDATING
DELETING
FAILED
MAINTENANCE
REBOOTING_BROKER
HEALING
```

## ConfigurationInfo

Specifies the configuration to use for the brokers.

**arn**

ARN of the configuration to use.

**Type**: string
**Required**: True

**revision**

The revision of the configuration to use.

**Type**: integer
**Required**: True
**Format**: int64
**Minimum**: 1

## ConnectivityInfo

Broker access controls.

**vpcConnectivity**

VPC connection control settings for brokers

**Type**: VpcConnectivity
**Required**: False

**publicAccess**

Access control settings for the cluster's brokers.

**Type**: PublicAccess
**Required**: False

**networkType**

The network type of the cluster, which is IPv4 or DUAL. The DUAL network type uses both IPv4 and IPv6 addresses for your cluster and its resources. By default, a cluster uses the IPv4 network type.

**Type**: string
**Required**: False
**MinLength**: 4
**MaxLength**: 4

## CreateClusterRequest

Creates a cluster.

**encryptionInfo**

Includes all encryption-related information.

**Type**: EncryptionInfo
**Required**: False

**configurationInfo**

Represents the configuration that you want MSK to use for the cluster.

**Type**: ConfigurationInfo
**Required**: False

**loggingInfo**

Logging Info details.

**Type**: LoggingInfo

**Required**: False

### tags

Create tags when creating the cluster.

**Type**: object
**Required**: False

### numberOfBrokerNodes

The number of broker nodes in the cluster.

**Type**: integer
**Required**: True

### enhancedMonitoring

Specifies the level of monitoring for the MSK cluster. The possible values are `DEFAULT`, `PER_BROKER`, and `PER_TOPIC_PER_BROKER`.

**Type**: EnhancedMonitoring
**Required**: False

### storageMode

This controls storage mode for supported storage tiers.

**Type**: StorageMode
**Required**: False

### clusterName

The name of the cluster.

**Type**: string
**Required**: True
**MinLength**: 1
**MaxLength**: 64

**kafkaVersion**

The version of Apache Kafka. You can use Amazon MSK to create clusters that use Apache Kafka versions 1.1.1 and 2.2.1. See Apache Kafka Versions.

**Type**: string
**Required**: True
**MinLength**: 1
**MaxLength**: 128

**clientAuthentication**

Includes all client authentication related information.

**Type**: ClientAuthentication
**Required**: False

**brokerNodeGroupInfo**

Information about the broker nodes in the cluster.

**Type**: BrokerNodeGroupInfo
**Required**: True

**openMonitoring**

The settings for open monitoring.

**Type**: OpenMonitoringInfo
**Required**: False

**rebalancing**

Specifies if intelligent rebalancing should be turned on for a new MSK Provisioned cluster with Express brokers. For all new Express-based clusters that you create, intelligent rebalancing is turned on by default.

**Type**: Rebalancing

**Required**: False

## CreateClusterResponse

Returns information about the created cluster.

### clusterArn

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

### clusterName

The name of the MSK cluster.

**Type**: string
**Required**: False

### state

The state of the cluster. Amazon MSK automatically renews certificates on clusters every 13 months. It sets the state of the cluster to `MAINTENANCE` when it starts the certificate-update operation. It sets it back to `ACTIVE` when the update is done. While a cluster is in the `MAINTENANCE` state, you can continue to produce and consume data, but you can't perform any update operations on it. You can perform update operations on a cluster when it is in the `ACTIVE` state.

**Type**: ClusterState
**Required**: False

## CustomerActionStatus

A type of an action required from the customer.

```
CRITICAL_ACTION_REQUIRED
ACTION_RECOMMENDED
NONE
```

## EBSStorageInfo

Contains information about the EBS storage volumes attached to the broker nodes.

**provisionedThroughput**

EBS volume provisioned throughput information.

**Type**: ProvisionedThroughput
**Required**: False

**volumeSize**

The size in GiB of the EBS volume for the data drive on each broker node.

**Type**: integer
**Required**: False
**Minimum**: 1
**Maximum**: 16384

## EncryptionAtRest

The data-volume encryption details. You can't update encryption at rest settings for existing clusters.

**dataVolumeKMSKeyId**

The Amazon Resource Name (ARN) of the Amazon KMS key for encrypting data at rest. If you don't specify a KMS key, MSK creates one for you and uses it.

**Type**: string
**Required**: True

## EncryptionInTransit

The settings for encrypting data in transit.

**inCluster**

When set to true, it indicates that data communication among the broker nodes of the cluster is encrypted. When set to false, the communication happens in plaintext.

The default value is true.

**Type**: boolean
**Required**: False

**clientBroker**

Indicates the encryption setting for data in transit between clients and brokers. You must set it to one of the following values.

`TLS` means that client-broker communication is enabled with TLS only.

`TLS_PLAINTEXT` means that client-broker communication is enabled for both TLS-encrypted, as well as plaintext data.

`PLAINTEXT` means that client-broker communication is enabled in plaintext only.

The default value is `TLS`.

**Type**: ClientBroker
**Required**: False

## EncryptionInfo

Includes encryption-related information, such as the Amazon KMS key used for encrypting data at rest and whether you want MSK to encrypt your data in transit.

**encryptionInTransit**

The details for encryption in transit.

**Type**: EncryptionInTransit
**Required**: False

**encryptionAtRest**

The data-volume encryption details.

**Type**: EncryptionAtRest
**Required**: False

## EnhancedMonitoring

Specifies which Apache Kafka metrics Amazon MSK gathers and sends to Amazon CloudWatch for this cluster. This property has three possible values: DEFAULT, PER_BROKER, and PER_TOPIC_PER_BROKER. For a list of the metrics associated with each of these three levels of monitoring, see Monitoring.

DEFAULT
PER_BROKER
PER_TOPIC_PER_BROKER
PER_TOPIC_PER_PARTITION

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## Firehose

Firehose details for BrokerLogs.

**deliveryStream**

The Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Kinesis Data Firehose delivery stream.

**Type**: boolean
**Required**: True

## IAM

Details for SASL/IAM client authentication.

**enabled**

SASL/IAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## JmxExporter

Indicates whether you want to enable or disable the JMX Exporter.

**enabledInBroker**

Indicates whether you want to enable or disable the JMX Exporter.

**Type**: boolean
**Required**: True

## JmxExporterInfo

JMX Exporter details.

**enabledInBroker**

JMX Exporter being enabled in broker.

**Type**: boolean
**Required**: True

## ListClustersResponse

The response contains an array containing cluster information and a next token if the response is truncated.

### nextToken

The paginated results marker. When the result of a `ListClusters` operation is truncated, the call returns `NextToken` in the response. To get another batch of clusters, provide this token in your next request.

**Type**: string
**Required**: False

### clusterInfoList

Information on each of the MSK clusters in the response.

**Type**: Array of type ClusterInfo
**Required**: False

## LoggingInfo

You can configure your MSK cluster to send broker logs to different destination types. This is a container for the configuration details related to broker logs.

### brokerLogs

You can configure your MSK cluster to send broker logs to different destination types. This configuration specifies the details of these destinations.

**Type**: BrokerLogs
**Required**: False

### authorizerLogs

Authorizer Logs details.

**Type**: AuthorizerLogs
**Required**: False

## NodeExporter

Indicates whether you want to enable or disable the Node Exporter.

**enabledInBroker**

Indicates whether you want to enable or disable the Node Exporter.

**Type**: boolean
**Required**: True

## NodeExporterInfo

Node Exporter details.

**enabledInBroker**

Node Exporter being enabled in broker.

**Type**: boolean
**Required**: True

## OpenMonitoring

JMX and Node monitoring for the MSK cluster.

**prometheus**

Prometheus exporter settings.

**Type**: Prometheus
**Required**: True

## OpenMonitoringInfo

JMX and Node monitoring for cluster.

**prometheus**

Prometheus details.

**Type**: PrometheusInfo
**Required**: True

## Prometheus

Prometheus settings for open monitoring.

### nodeExporter

Indicates whether you want to enable or disable the Node Exporter.

**Type**: NodeExporter
**Required**: False

### jmxExporter

Indicates whether you want to enable or disable the JMX Exporter.

**Type**: JmxExporter
**Required**: False

## PrometheusInfo

Prometheus details.

### nodeExporter

Node Exporter details.

**Type**: NodeExporterInfo
**Required**: False

### jmxExporter

JMX Exporter details.

**Type**: JmxExporterInfo
**Required**: False

## ProvisionedThroughput

Contains information about provisioned throughput for EBS storage volumes attached to kafka broker nodes.

**volumeThroughput**

Throughput value of the EBS volumes for the data drive on each kafka broker node in MiB per second.

**Type**: integer
**Required**: False

**enabled**

Provisioned throughput is enabled or not.

**Type**: boolean
**Required**: False

## PublicAccess

Broker access controls

**type**

DISABLED means that public access is turned off. SERVICE_PROVIDED_EIPS means that public access is turned on.

**Type**: string
**Required**: False

## Rebalancing

Specifies whether or not intelligent rebalancing is turned on for a newly created MSK Provisioned cluster with Express brokers. Intelligent rebalancing performs automatic partition balancing operations when you scale your clusters up or down.

By default, intelligent rebalancing is `ACTIVE` for all new Express-based clusters.

**status**

Intelligent rebalancing status. The default intelligent rebalancing status is `ACTIVE` for all new Express-based clusters.

**Type**: RebalancingStatus
**Required**: True

## RebalancingStatus

Intelligent rebalancing status. The default intelligent rebalancing status is `ACTIVE` for all new Express-based clusters.

`PAUSED`
`ACTIVE`

## S3

The details of the Amazon S3 destination for broker logs.

**bucket**

The name of the S3 bucket that is the destination for broker logs.

**Type**: string
**Required**: False

**prefix**

The S3 prefix that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Amazon S3 destination.

**Type**: boolean
**Required**: True

## Sasl

Details for client authentication using SASL. To turn on SASL, you must also turn on `EncryptionInTransit` by setting `inCluster` to true. You must set `clientBroker` to either `TLS` or `TLS_PLAINTEXT`. If you choose `TLS_PLAINTEXT`, then you must also set `unauthenticated` to true.

**iam**

Details for ClientAuthentication using IAM.

**Type**: IAM
**Required**: False

**scram**

Details for SASL/SCRAM client authentication.

**Type**: Scram
**Required**: False

## Scram

Details for SASL/SCRAM client authentication.

**enabled**

SASL/SCRAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## StateInfo

Includes information about the state of the cluster.

**code**

If the cluster is in an unusable state, this field contains the code that describes the issue.

**Type**: string
**Required**: False

**message**

If the cluster is in an unusable state, this field contains a message that describes the issue.

**Type**: string
**Required**: False

## StorageInfo

Contains information about storage volumes attached to Amazon MSK broker nodes.

**ebsStorageInfo**

EBS volume information.

**Type**: EBSStorageInfo
**Required**: False

## StorageMode

Controls storage mode for various supported storage tiers.

```
LOCAL
TIERED
```

## Tls

Details for client authentication using TLS.

**certificateAuthorityArnList**

List of AWS Private CA Amazon Resource Name (ARN)s.

**Type**: Array of type string
**Required**: False

**enabled**

TLS authentication is enabled or not.

**Type**: boolean
**Required**: False

## Unauthenticated

Details for allowing no client authentication.

**enabled**

Unauthenticated is enabled or not.

**Type**: boolean
**Required**: False

## VpcConnectivity

VPC connection control settings for brokers.

**clientAuthentication**

VPC connection control settings for brokers.

**Type**: VpcConnectivityClientAuthentication
**Required**: False

## VpcConnectivityClientAuthentication

Includes all client authentication information for VpcConnectivity.

**sasl**

Details for VpcConnectivity ClientAuthentication using SASL.

**Type**: VpcConnectivitySasl
**Required**: False

**tls**

Details for VpcConnectivity ClientAuthentication using TLS.

**Type**: VpcConnectivityTls
**Required**: False

## VpcConnectivityIAM

Details for SASL/IAM client authentication for VpcConnectivity.

**enabled**

SASL/IAM authentication is enabled or not.

**Type**: boolean
**Required**: False

## VpcConnectivitySasl

Details for client authentication using SASL for VpcConnectivity.

**iam**

Details for ClientAuthentication using IAM for VpcConnectivity.

**Type**: VpcConnectivityIAM
**Required**: False

**scram**

Details for SASL/SCRAM client authentication for VpcConnectivity.

**Type**: VpcConnectivityScram
**Required**: False

## VpcConnectivityScram

Details for SASL/SCRAM client authentication for vpcConnectivity.

**enabled**

SASL/SCRAM authentication is enabled or not.

**Type**: boolean
**Required**: False

### VpcConnectivityTls

Details for client authentication using TLS for vpcConnectivity.

**enabled**

TLS authentication is enabled or not.

**Type**: boolean
**Required**: False

## See also

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

### ListClusters

- 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

### CreateCluster

- 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

# Clusters clusterArn Channels

## URI

`/v1/clusters/clusterArn/channels`

## HTTP methods

### GET

**Operation ID:** `ListChannels`

Returns all channels in a cluster.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Query parameters

| Name | Type | Required | Description |
|---|---|---|---|
| topicNameFilter | String | False | Filters results to channels whose topic name matches the specified value. |
| nextToken | String | False | The paginated results marker. When the result of the operation is truncated, the call returns NextToken in the response. To get the next batch, provide this token in your next request. |
| maxResults | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes a NextToken parameter. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | ListChannelsResponse | Successful response. |

| Status code | Response model | Description |
| --- | --- | --- |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorize d. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## POST

**Operation ID:** `CreateChannel`

Creates a channel using the specified properties in the request on the cluster specified by the Amazon Resource Name (ARN) in the request.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | CreateChannelResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 409 | Error | This cluster name already exists. Retry your request using another name. |
| 429 | Error | 429 response |

| Status code | Response model | Description |
| --- | --- | --- |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning correct headers

#### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### POST schema

```
{
  "s3DestinationConfiguration": {
    "dataFreshnessInSeconds": integer,
```

```
    "serviceExecutionRoleArn": "string",
    "deadLetterQueueS3": {
      "bucketArn": "string",
      "expectedBucketOwner": "string",
      "errorOutputPrefix": "string"
    },
    "storage": {
      "storageClass": enum,
      "bucketArn": "string",
      "outputKeyTemplate": "string",
      "outputPrefix": "string",
      "expectedBucketOwner": "string",
      "compressionType": enum
    }
  },
  "topicConfigurationList": [
    {
      "recordConverter": {
        "valueConverter": enum
      },
      "recordSchema": {
        "gsrArn": "string"
      },
      "topicArn": "string"
    }
  ],
  "encryptionConfiguration": {
    "kmsKeyArn": "string"
  },
  "channelName": "string",
  "loggingInfo": {
    "s3": {
      "bucket": "string",
      "prefix": "string",
      "enabled": boolean
    },
    "firehose": {
      "deliveryStream": "string",
      "enabled": boolean
    },
    "cloudWatchLogs": {
      "logGroup": "string",
      "enabled": boolean
    }
```

```
  },
  "icebergDestinationConfiguration": {
    "dataFreshnessInSeconds": integer,
    "serviceExecutionRoleArn": "string",
    "catalog": {
      "catalogArn": "string",
      "warehouseLocation": "string"
    },
    "schemaEvolution": {
      "enableSchemaEvolution": boolean
    },
    "deadLetterQueueS3": {
      "bucketArn": "string",
      "expectedBucketOwner": "string",
      "errorOutputPrefix": "string"
    },
    "compressionType": enum,
    "destinationTableList": [
      {
        "destinationTableName": "string",
        "destinationDatabaseName": "string",
        "partitionSpec": {
          "partitionStrategy": enum,
          "sourceList": [
            {
              "sourceName": "string"
            }
          ]
        }
      }
    ],
    "tableCreation": {
      "enableTableCreation": boolean
    },
    "appendOnly": boolean
  },
  "tags": {
  }
}
```

## Response bodies

### ListChannelsResponse schema

```
{
  "channels": [
    {
      "clusterOperationArn": "string",
      "channelArn": "string",
      "creationTime": "string",
      "destinationType": enum,
      "channelName": "string",
      "status": enum
    }
  ],
  "nextToken": "string"
}
```

### CreateChannelResponse schema

```
{
  "clusterOperationArn": "string",
  "channelArn": "string"
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Catalog

Configuration of the AWS Glue Data Catalog and S3 Tables warehouse used by the Apache Iceberg destination.

**catalogArn**

The Amazon Resource Name (ARN) of the federated AWS Glue Data Catalog that projects the S3 Tables bucket. If omitted, MSK derives the catalog ARN from warehouseLocation.

**Type**: string
**Required**: False

**warehouseLocation**

The Amazon Resource Name (ARN) of the S3 Tables bucket that backs the Apache Iceberg warehouse.

**Type**: string
**Required**: False

## ChannelDestinationType

The type of destination configured for the channel.

ICEBERG
S3

## ChannelInfo

Summary information about a channel returned by ListChannels.

**clusterOperationArn**

The Amazon Resource Name (ARN) of the in-flight cluster operation. Returned only while the channel is in CREATING, UPDATING, or DELETING.

**Type**: string
**Required**: False

**channelArn**

The Amazon Resource Name (ARN) that uniquely identifies the channel.

**Type**: string
**Required**: True

### creationTime

The time when the channel was created.

**Type**: string
**Required**: True

### destinationType

The type of destination configured for the channel.

**Type**: ChannelDestinationType
**Required**: True

### channelName

The name of the channel.

**Type**: string
**Required**: True

### status

The current lifecycle state of the channel.

**Type**: ChannelStatus
**Required**: True

## ChannelLoggingInfo

Configuration for the destinations to which the channel publishes operational logs.

### s3

Configuration for delivering logs to an Amazon S3 bucket.

**Type**: S3
**Required**: False

#### firehose

Configuration for delivering logs to an Amazon Data Firehose delivery stream.

**Type**: Firehose
**Required**: False

#### cloudWatchLogs

Configuration for delivering logs to Amazon CloudWatch Logs.

**Type**: CloudWatchLogs
**Required**: False

## ChannelStatus

The lifecycle state of a channel.

```
CREATING
ACTIVE
UPDATING
DELETING
FAILED
SUSPENDING
SUSPENDED
```

## CloudWatchLogs

Details of the CloudWatch Logs destination for broker logs.

#### logGroup

The CloudWatch log group that is the destination for broker logs.

**Type**: string

**Required**: False

#### enabled

Specifies whether broker logs get sent to the specified CloudWatch Logs destination.

**Type**: boolean
**Required**: True

## CreateChannelRequest

Creates a Channel that streams records from an Amazon MSK Express cluster topic to Amazon S3 or Apache Iceberg.

#### s3DestinationConfiguration

The Amazon S3 destination for the channel. Mutually exclusive with icebergDestinationConfiguration.

**Type**: S3DestinationConfiguration
**Required**: False

#### topicConfigurationList

The list of topic configurations for the channel. Currently exactly one topic must be specified.

**Type**: Array of type TopicConfiguration
**Required**: True

#### encryptionConfiguration

The encryption configuration applied to the channel.

**Type**: EncryptionConfiguration
**Required**: False

#### channelName

The name of the channel. Must be unique within the cluster.

**Type**: string
**Required**: True

#### loggingInfo

The destinations to which the channel publishes operational logs.

**Type**: ChannelLoggingInfo
**Required**: False

#### icebergDestinationConfiguration

The Apache Iceberg destination for the channel. Mutually exclusive with s3DestinationConfiguration.

**Type**: IcebergDestinationConfiguration
**Required**: False

#### tags

The tags attached to the channel.

**Type**: object
**Required**: False

## CreateChannelResponse

Returns the channel ARN and the cluster-operation ARN that tracks the asynchronous create.

#### clusterOperationArn

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

#### channelArn

The Amazon Resource Name (ARN) that uniquely identifies the channel.

**Type**: string
**Required**: True

## DeadLetterQueueS3

Configuration of the Amazon S3 bucket where records that fail to deliver are stored.

### bucketArn

The Amazon Resource Name (ARN) of the dead-letter Amazon S3 bucket.

**Type**: string
**Required**: True

### expectedBucketOwner

Optional 12-digit AWS account ID expected to own the dead-letter Amazon S3 bucket.

**Type**: string
**Required**: False

### errorOutputPrefix

An optional prefix prepended to every dead-letter Amazon S3 object key.

**Type**: string
**Required**: False

## DestinationTable

Configuration of an Apache Iceberg destination table.

### destinationTableName

The name of the destination Apache Iceberg table.

**Type**: string
**Required**: False

**destinationDatabaseName**

The name of the destination namespace (database) in the AWS Glue Data Catalog.

**Type**: string
**Required**: False

**partitionSpec**

The partition specification for the destination table.

**Type**: PartitionSpec
**Required**: False

## EncryptionConfiguration

The AWS KMS encryption configuration applied to data at rest.

**kmsKeyArn**

The Amazon Resource Name (ARN) of the AWS KMS key used to encrypt the data.

**Type**: string
**Required**: True

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## Firehose

Firehose details for BrokerLogs.

**deliveryStream**

The Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Kinesis Data Firehose delivery stream.

**Type**: boolean
**Required**: True

## IcebergCompressionType

Compression codec for Iceberg table data files. Defaults to ZSTD.

ZSTD
SNAPPY

## IcebergDestinationConfiguration

Configuration of an Apache Iceberg destination for a channel.

**dataFreshnessInSeconds**

The maximum time, in seconds, that records buffer in MSK before being flushed to the destination. Allowed range: 300 to 900. Default: 600.

**Type**: integer

**Required**: False

**serviceExecutionRoleArn**

The Amazon Resource Name (ARN) of the IAM role that MSK assumes to access the destination table, the AWS Glue Data Catalog, and the dead-letter Amazon S3 bucket.

**Type**: string
**Required**: True

**catalog**

The AWS Glue Data Catalog and S3 Tables warehouse used by the destination.

**Type**: Catalog
**Required**: False

**schemaEvolution**

Configuration controlling whether the destination table's schema is evolved to match incoming records.

**Type**: SchemaEvolution
**Required**: True

**deadLetterQueueS3**

The Amazon S3 bucket and prefix where MSK writes records that fail to deliver.

**Type**: DeadLetterQueueS3
**Required**: True

**compressionType**

The compression codec for Iceberg table data files. Defaults to ZSTD.

**Type**: IcebergCompressionType
**Required**: False

**destinationTableList**

The destination Iceberg tables. Currently exactly one table must be specified.

**Type**: Array of type DestinationTable
**Required**: True

**tableCreation**

Configuration controlling whether MSK creates the destination table if it does not already exist.

**Type**: TableCreation
**Required**: True

**appendOnly**

Whether the destination is append-only. Must be true; updates and deletes are not supported.

**Type**: boolean
**Required**: True

## ListChannelsResponse

Returns the list of channels in the cluster.

**channels**

The list of channels in the cluster.

**Type**: Array of type ChannelInfo
**Required**: False

**nextToken**

If the response from ListChannels is truncated, this token is included. Send it as the nextToken parameter on a subsequent ListChannels call to retrieve the next page.

**Type**: string
**Required**: False

## PartitionSource

A source column used by an Apache Iceberg destination table's partition specification.

**sourceName**

The name of the source column.

**Type**: string
**Required**: False

## PartitionSpec

Partition specification for an Apache Iceberg destination table.

**partitionStrategy**

The partitioning strategy applied to records written to the table.

**Type**: PartitionStrategy
**Required**: True

**sourceList**

The source columns used by the partitioning strategy. For TIME_HOUR, must contain exactly one source column whose value is a timestamp.

**Type**: Array of type PartitionSource
**Required**: False

## PartitionStrategy

The partitioning strategy used to partition records in the destination Apache Iceberg table.

`TIME_HOUR`

## RecordConverter

Configuration that controls how Apache Kafka record values are deserialized for the destination.

**valueConverter**

The deserialization format applied to Apache Kafka record values.

**Type**: ValueConverter
**Required**: True

## RecordSchema

Schema configuration that controls how Apache Kafka record values are validated.

**gsrArn**

The Amazon Resource Name (ARN) of the AWS Glue Schema Registry schema (not registry) used to validate records for the destination Apache Iceberg table.

**Type**: string
**Required**: True

## S3

The details of the Amazon S3 destination for broker logs.

**bucket**

The name of the S3 bucket that is the destination for broker logs.

**Type**: string
**Required**: False

**prefix**

The S3 prefix that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Amazon S3 destination.

**Type**: boolean

**Required**: True

## S3CompressionType

The compression codec applied to delivered Amazon S3 objects.

`NONE`

`GZIP`

`ZSTD`

## S3DestinationConfiguration

Configuration of an Amazon S3 destination for a channel.

### dataFreshnessInSeconds

The maximum time, in seconds, that records buffer in MSK before being flushed to the destination. Allowed range: 300 to 900. Default: 600.

**Type**: integer

**Required**: False

### serviceExecutionRoleArn

The Amazon Resource Name (ARN) of the IAM role that MSK assumes to write to the destination Amazon S3 bucket and the dead-letter bucket.

**Type**: string

**Required**: True

### deadLetterQueueS3

The Amazon S3 bucket and prefix where MSK writes records that fail to deliver.

**Type**: DeadLetterQueueS3

**Required**: True

**storage**

The Amazon S3 bucket, prefix, and storage class for delivered records.

**Type**: S3Storage
**Required**: True

## S3Storage

Storage configuration for an Amazon S3 destination bucket.

**storageClass**

The Amazon S3 storage class for delivered objects.

**Type**: S3StorageClass
**Required**: True

**bucketArn**

The Amazon Resource Name (ARN) of the destination Amazon S3 bucket.

**Type**: string
**Required**: True

**outputKeyTemplate**

An optional template that controls the Amazon S3 object key for each delivered record. Supports the placeholders !{partition-id}, !{sequence-number}, and !{kafka-offset}.

**Type**: string
**Required**: False

**outputPrefix**

An optional prefix prepended to every Amazon S3 object key written by the channel.

**Type**: string
**Required**: False

**expectedBucketOwner**

Optional 12-digit AWS account ID expected to own the Amazon S3 bucket.

**Type**: string
**Required**: False

**compressionType**

The compression codec applied to delivered Amazon S3 objects.

**Type**: S3CompressionType
**Required**: True

## S3StorageClass

The Amazon S3 storage class applied to delivered objects.

```
STANDARD
INTELLIGENT_TIERING
GLACIER_IR
```

## SchemaEvolution

Configuration controlling whether the Apache Iceberg destination table's schema is evolved as incoming records change.

**enableSchemaEvolution**

Whether to allow MSK to evolve the destination table's schema. Must be false for the current release.

**Type**: boolean
**Required**: False

## TableCreation

Configuration controlling whether MSK creates the destination Apache Iceberg table if it does not already exist.

**enableTableCreation**

Whether MSK creates the destination table on the customer's behalf. Must be true for the current release.

**Type**: boolean
**Required**: False

## TopicConfiguration

Configuration of an Apache Kafka topic that feeds a channel.

**recordConverter**

Configuration that controls how Apache Kafka record values are deserialized for the destination.

**Type**: RecordConverter
**Required**: True

**recordSchema**

The schema used to validate records when the value converter requires one (for example, JSON_SCHEMA_GSR).

**Type**: RecordSchema
**Required**: False

**topicArn**

The Amazon Resource Name (ARN) that uniquely identifies the topic.

**Type**: string
**Required**: True

## ValueConverter

The deserialization format applied to Apache Kafka record values.

BYTE_ARRAY
JSON

JSON_SCHEMA_GSR
STRING

## See also

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

### ListChannels

- 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

### CreateChannel

- 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

# Clusters clusterArn Channels channelArn

## URI

`/v1/clusters/clusterArn/channels/channelArn`

## HTTP methods

### GET

**Operation ID:** `DescribeChannel`

Returns details for a channel resource.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |
| *channelArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the channel. |

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | DescribeChannelResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |

| Status code | Response model | Description |
| --- | --- | --- |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## PUT

**Operation ID:** `UpdateChannel`

Updates the channel specified by the channelArn in the request from the cluster specified by the Amazon Resource Name (ARN) in the request.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *channelArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the channel. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | UpdateChannelResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |

| Status code | Response model | Description |
| --- | --- | --- |
| 503 | Error | 503 response |

## DELETE

**Operation ID:** `DeleteChannel`

Deletes the channel specified by the channelArn in the request from the cluster specified by the Amazon Resource Name (ARN) in the request.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |
| *channelArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the channel. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | DeleteChannelResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |

| Status code | Response model | Description |
| --- | --- | --- |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning correct headers

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |
| *channelArn* | String | True | The Amazon Resource Name (ARN) that |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | uniquely identifies the channel. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### PUT schema

```
{
  "icebergDestinationUpdate": {
    "dataFreshnessInSeconds": integer
  },
  "s3DestinationUpdate": {
    "dataFreshnessInSeconds": integer
  }
}
```

### Response bodies

#### DescribeChannelResponse schema

```
{
  "s3DestinationConfiguration": {
    "dataFreshnessInSeconds": integer,
    "serviceExecutionRoleArn": "string",
    "deadLetterQueueS3": {
      "bucketArn": "string",
      "expectedBucketOwner": "string",
      "errorOutputPrefix": "string"
    },
```

```
    "storage": {
      "storageClass": enum,
      "bucketArn": "string",
      "outputKeyTemplate": "string",
      "outputPrefix": "string",
      "expectedBucketOwner": "string",
      "compressionType": enum
    }
  },
  "clusterOperationArn": "string",
  "creationTime": "string",
  "loggingInfo": {
    "s3": {
      "bucket": "string",
      "prefix": "string",
      "enabled": boolean
    },
    "firehose": {
      "deliveryStream": "string",
      "enabled": boolean
    },
    "cloudWatchLogs": {
      "logGroup": "string",
      "enabled": boolean
    }
  },
  "icebergDestinationConfiguration": {
    "dataFreshnessInSeconds": integer,
    "serviceExecutionRoleArn": "string",
    "catalog": {
      "catalogArn": "string",
      "warehouseLocation": "string"
    },
    "schemaEvolution": {
      "enableSchemaEvolution": boolean
    },
    "deadLetterQueueS3": {
      "bucketArn": "string",
      "expectedBucketOwner": "string",
      "errorOutputPrefix": "string"
    },
    "compressionType": enum,
    "destinationTableList": [
      {
```

```
        "destinationTableName": "string",
        "destinationDatabaseName": "string",
        "partitionSpec": {
          "partitionStrategy": enum,
          "sourceList": [
            {
              "sourceName": "string"
            }
          ]
        }
      }
    ],
    "tableCreation": {
      "enableTableCreation": boolean
    },
    "appendOnly": boolean
  },
  "tags": {
  },
  "channelArn": "string",
  "topicConfigurationList": [
    {
      "recordConverter": {
        "valueConverter": enum
      },
      "recordSchema": {
        "gsrArn": "string"
      },
      "topicArn": "string"
    }
  ],
  "stateInfo": {
    "code": "string",
    "message": "string"
  },
  "encryptionConfiguration": {
    "kmsKeyArn": "string"
  },
  "destinationType": enum,
  "channelName": "string",
  "status": enum
}
```

### UpdateChannelResponse schema

```
{
  "clusterOperationArn": "string",
  "channelArn": "string"
}
```

### DeleteChannelResponse schema

```
{
  "clusterOperationArn": "string",
  "channelArn": "string"
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Catalog

Configuration of the AWS Glue Data Catalog and S3 Tables warehouse used by the Apache Iceberg destination.

### catalogArn

The Amazon Resource Name (ARN) of the federated AWS Glue Data Catalog that projects the S3 Tables bucket. If omitted, MSK derives the catalog ARN from warehouseLocation.

**Type**: string
**Required**: False

### warehouseLocation

The Amazon Resource Name (ARN) of the S3 Tables bucket that backs the Apache Iceberg warehouse.

**Type**: string
**Required**: False

## ChannelDestinationType

The type of destination configured for the channel.

ICEBERG
S3

## ChannelLoggingInfo

Configuration for the destinations to which the channel publishes operational logs.

**s3**

Configuration for delivering logs to an Amazon S3 bucket.

**Type**: S3
**Required**: False

**firehose**

Configuration for delivering logs to an Amazon Data Firehose delivery stream.

**Type**: Firehose
**Required**: False

**cloudWatchLogs**

Configuration for delivering logs to Amazon CloudWatch Logs.

**Type**: CloudWatchLogs
**Required**: False

## ChannelStateInfo

Additional context for the current channel state, populated when the channel is in FAILED.

**code**

A short, machine-readable code identifying the failure cause.

**Type**: string
**Required**: False

**message**

A human-readable message describing the failure.

**Type**: string
**Required**: False

## ChannelStatus

The lifecycle state of a channel.

```
CREATING
ACTIVE
UPDATING
DELETING
FAILED
SUSPENDING
SUSPENDED
```

## CloudWatchLogs

Details of the CloudWatch Logs destination for broker logs.

**logGroup**

The CloudWatch log group that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified CloudWatch Logs destination.

**Type**: boolean
**Required**: True

## DeadLetterQueueS3

Configuration of the Amazon S3 bucket where records that fail to deliver are stored.

**bucketArn**

The Amazon Resource Name (ARN) of the dead-letter Amazon S3 bucket.

**Type**: string
**Required**: True

**expectedBucketOwner**

Optional 12-digit AWS account ID expected to own the dead-letter Amazon S3 bucket.

**Type**: string
**Required**: False

**errorOutputPrefix**

An optional prefix prepended to every dead-letter Amazon S3 object key.

**Type**: string
**Required**: False

## DeleteChannelResponse

Returns the channel ARN and the cluster-operation ARN that tracks the asynchronous delete.

**clusterOperationArn**

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

#### channelArn

The Amazon Resource Name (ARN) that uniquely identifies the channel.

**Type**: string
**Required**: True

## DescribeChannelResponse

Contains the current configuration and state of a channel.

#### s3DestinationConfiguration

The Amazon S3 destination for the channel, if configured.

**Type**: S3DestinationConfiguration
**Required**: False

#### clusterOperationArn

The Amazon Resource Name (ARN) of the in-flight cluster operation. Returned only while the channel is in CREATING, UPDATING, or DELETING.

**Type**: string
**Required**: False

#### creationTime

The time when the channel was created.

**Type**: string
**Required**: True

#### loggingInfo

The destinations to which the channel publishes operational logs.

**Type**: ChannelLoggingInfo
**Required**: False

### icebergDestinationConfiguration

The Apache Iceberg destination for the channel, if configured.

**Type**: IcebergDestinationConfiguration
**Required**: False

### tags

The tags attached to the channel.

**Type**: object
**Required**: False

### channelArn

The Amazon Resource Name (ARN) that uniquely identifies the channel.

**Type**: string
**Required**: True

### topicConfigurationList

The list of topic configurations for the channel.

**Type**: Array of type TopicConfiguration
**Required**: True

### stateInfo

Additional context for the current channel state, populated when the channel is in FAILED.

**Type**: ChannelStateInfo
**Required**: False

**encryptionConfiguration**

The encryption configuration applied to the channel.

**Type**: EncryptionConfiguration
**Required**: False

**destinationType**

The type of destination configured for the channel.

**Type**: ChannelDestinationType
**Required**: True

**channelName**

The name of the channel.

**Type**: string
**Required**: True

**status**

The current lifecycle state of the channel.

**Type**: ChannelStatus
**Required**: True

## DestinationTable

Configuration of an Apache Iceberg destination table.

**destinationTableName**

The name of the destination Apache Iceberg table.

**Type**: string
**Required**: False

**destinationDatabaseName**

The name of the destination namespace (database) in the AWS Glue Data Catalog.

**Type**: string
**Required**: False

**partitionSpec**

The partition specification for the destination table.

**Type**: PartitionSpec
**Required**: False

## EncryptionConfiguration

The AWS KMS encryption configuration applied to data at rest.

**kmsKeyArn**

The Amazon Resource Name (ARN) of the AWS KMS key used to encrypt the data.

**Type**: string
**Required**: True

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## Firehose

Firehose details for BrokerLogs.

**deliveryStream**

The Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Kinesis Data Firehose delivery stream.

**Type**: boolean
**Required**: True

## IcebergCompressionType

Compression codec for Iceberg table data files. Defaults to ZSTD.

ZSTD
SNAPPY

## IcebergDestinationConfiguration

Configuration of an Apache Iceberg destination for a channel.

**dataFreshnessInSeconds**

The maximum time, in seconds, that records buffer in MSK before being flushed to the destination. Allowed range: 300 to 900. Default: 600.

**Type**: integer

**Required**: False

**serviceExecutionRoleArn**

The Amazon Resource Name (ARN) of the IAM role that MSK assumes to access the destination table, the AWS Glue Data Catalog, and the dead-letter Amazon S3 bucket.

**Type**: string
**Required**: True

**catalog**

The AWS Glue Data Catalog and S3 Tables warehouse used by the destination.

**Type**: Catalog
**Required**: False

**schemaEvolution**

Configuration controlling whether the destination table's schema is evolved to match incoming records.

**Type**: SchemaEvolution
**Required**: True

**deadLetterQueueS3**

The Amazon S3 bucket and prefix where MSK writes records that fail to deliver.

**Type**: DeadLetterQueueS3
**Required**: True

**compressionType**

The compression codec for Iceberg table data files. Defaults to ZSTD.

**Type**: IcebergCompressionType
**Required**: False

**destinationTableList**

The destination Iceberg tables. Currently exactly one table must be specified.

**Type**: Array of type DestinationTable
**Required**: True

**tableCreation**

Configuration controlling whether MSK creates the destination table if it does not already exist.

**Type**: TableCreation
**Required**: True

**appendOnly**

Whether the destination is append-only. Must be true; updates and deletes are not supported.

**Type**: boolean
**Required**: True

## IcebergDestinationUpdate

Update payload for an Apache Iceberg destination.

**dataFreshnessInSeconds**

The maximum time, in seconds, that records buffer in MSK before being flushed to the destination. Allowed range: 300 to 900.

**Type**: integer
**Required**: True

## PartitionSource

A source column used by an Apache Iceberg destination table's partition specification.

**sourceName**

The name of the source column.

**Type**: string
**Required**: False

## PartitionSpec

Partition specification for an Apache Iceberg destination table.

### partitionStrategy

The partitioning strategy applied to records written to the table.

**Type**: PartitionStrategy
**Required**: True

### sourceList

The source columns used by the partitioning strategy. For TIME_HOUR, must contain exactly one source column whose value is a timestamp.

**Type**: Array of type PartitionSource
**Required**: False

## PartitionStrategy

The partitioning strategy used to partition records in the destination Apache Iceberg table.

`TIME_HOUR`

## RecordConverter

Configuration that controls how Apache Kafka record values are deserialized for the destination.

### valueConverter

The deserialization format applied to Apache Kafka record values.

**Type**: ValueConverter
**Required**: True

## RecordSchema

Schema configuration that controls how Apache Kafka record values are validated.

**gsrArn**

The Amazon Resource Name (ARN) of the AWS Glue Schema Registry schema (not registry) used to validate records for the destination Apache Iceberg table.

**Type**: string
**Required**: True

## S3

The details of the Amazon S3 destination for broker logs.

**bucket**

The name of the S3 bucket that is the destination for broker logs.

**Type**: string
**Required**: False

**prefix**

The S3 prefix that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Amazon S3 destination.

**Type**: boolean
**Required**: True

## S3CompressionType

The compression codec applied to delivered Amazon S3 objects.

NONE

GZIP

ZSTD

## S3DestinationConfiguration

Configuration of an Amazon S3 destination for a channel.

### dataFreshnessInSeconds

The maximum time, in seconds, that records buffer in MSK before being flushed to the destination. Allowed range: 300 to 900. Default: 600.

**Type**: integer

**Required**: False

### serviceExecutionRoleArn

The Amazon Resource Name (ARN) of the IAM role that MSK assumes to write to the destination Amazon S3 bucket and the dead-letter bucket.

**Type**: string

**Required**: True

### deadLetterQueueS3

The Amazon S3 bucket and prefix where MSK writes records that fail to deliver.

**Type**: DeadLetterQueueS3

**Required**: True

### storage

The Amazon S3 bucket, prefix, and storage class for delivered records.

**Type**: S3Storage

**Required**: True

## S3DestinationUpdate

Update payload for an Amazon S3 destination.

### dataFreshnessInSeconds

The maximum time, in seconds, that records buffer in MSK before being flushed to the destination. Allowed range: 300 to 900.

**Type**: integer
**Required**: True

## S3Storage

Storage configuration for an Amazon S3 destination bucket.

### storageClass

The Amazon S3 storage class for delivered objects.

**Type**: S3StorageClass
**Required**: True

### bucketArn

The Amazon Resource Name (ARN) of the destination Amazon S3 bucket.

**Type**: string
**Required**: True

### outputKeyTemplate

An optional template that controls the Amazon S3 object key for each delivered record. Supports the placeholders !{partition-id}, !{sequence-number}, and !{kafka-offset}.

**Type**: string
**Required**: False

**outputPrefix**

An optional prefix prepended to every Amazon S3 object key written by the channel.

**Type**: string
**Required**: False

**expectedBucketOwner**

Optional 12-digit AWS account ID expected to own the Amazon S3 bucket.

**Type**: string
**Required**: False

**compressionType**

The compression codec applied to delivered Amazon S3 objects.

**Type**: S3CompressionType
**Required**: True

## S3StorageClass

The Amazon S3 storage class applied to delivered objects.

```
STANDARD
INTELLIGENT_TIERING
GLACIER_IR
```

## SchemaEvolution

Configuration controlling whether the Apache Iceberg destination table's schema is evolved as incoming records change.

**enableSchemaEvolution**

Whether to allow MSK to evolve the destination table's schema. Must be false for the current release.

**Type**: boolean
**Required**: False

## TableCreation

Configuration controlling whether MSK creates the destination Apache Iceberg table if it does not already exist.

**enableTableCreation**

Whether MSK creates the destination table on the customer's behalf. Must be true for the current release.

**Type**: boolean
**Required**: False

## TopicConfiguration

Configuration of an Apache Kafka topic that feeds a channel.

**recordConverter**

Configuration that controls how Apache Kafka record values are deserialized for the destination.

**Type**: RecordConverter
**Required**: True

**recordSchema**

The schema used to validate records when the value converter requires one (for example, JSON_SCHEMA_GSR).

**Type**: RecordSchema
**Required**: False

**topicArn**

The Amazon Resource Name (ARN) that uniquely identifies the topic.

**Type**: string
**Required**: True

## UpdateChannelRequest

Updates an existing channel's destination configuration. You must update the same destination type the channel was created with; the destination type cannot be changed.

### icebergDestinationUpdate

Updates fields on an Apache Iceberg destination. Use only when the channel was created with an Iceberg destination.

**Type**: IcebergDestinationUpdate
**Required**: False

### s3DestinationUpdate

Updates fields on an Amazon S3 destination. Use only when the channel was created with an Amazon S3 destination.

**Type**: S3DestinationUpdate
**Required**: False

## UpdateChannelResponse

Returns the channel ARN and the cluster-operation ARN that tracks the asynchronous update.

### clusterOperationArn

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

### channelArn

The Amazon Resource Name (ARN) that uniquely identifies the channel.

**Type**: string

**Required**: True

### ValueConverter

The deserialization format applied to Apache Kafka record values.

`BYTE_ARRAY`
`JSON`
`JSON_SCHEMA_GSR`
`STRING`

## See also

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

### DescribeChannel

- 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

### UpdateChannel

- 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

### DeleteChannel

- 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

# Clusters clusterArn Client-vpc-connection

## URI

`/v1/clusters/`*`clusterArn`*`/client-vpc-connection`

## HTTP methods

### PUT

**Operation ID:** `RejectClientVpcConnection`

Reject client VPC connection.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | RejectClientVpcConnectionResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying |

| Status code | Response model | Description |
| --- | --- | --- |
| | | your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning correct headers.

#### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### PUT schema

```
{
  "vpcConnectionArn": "string"
}
```

## Response bodies

### RejectClientVpcConnectionResponse schema

```
{
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Error

Returns information about an error.

### message

The description of the error.

**Type**: string
**Required**: False

### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## RejectClientVpcConnectionRequest

Reject VPC Connection

### vpcConnectionArn

VPC Connection Amazon Resource Name (ARN).

**Type**: string
**Required**: True
**MinLength**: 1

### RejectClientVpcConnectionResponse

Blocks client connections connecting to the cluster

## See also

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

### RejectClientVpcConnection

- 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

# Clusters clusterArn Client-vpc-connections

## URI

`/v1/clusters/clusterArn/client-vpc-connections`

## HTTP methods

### GET

**Operation ID:** `ListClientVpcConnections`

List client VPC connections.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Query parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| nextToken | String | False | The paginated results marker. When the result of the operation is truncated, the call returns NextToken in the response. To get the next batch, provide this token in your next request. |
| maxResults | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes a NextToken parameter. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | ListClientVpcConnectionsResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning correct headers.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | None | Default response for CORS method |

## Schemas

### Response bodies

#### ListClientVpcConnectionsResponse schema

```
{
  "clientVpcConnections": [
    {
      "owner": "string",
      "vpcConnectionArn": "string",
      "creationTime": "string",
      "state": enum,
      "authentication": "string"
    }
  ],
  "nextToken": "string"
}
```

#### Error schema

```
{
  "message": "string",
```

```
  "invalidParameter": "string"
}
```

# Properties

## ClientVpcConnection

VPC Connection description

### owner

The Owner of the VPC Connection.

**Type**: string
**Required**: False

### vpcConnectionArn

The Amazon Resource Name (ARN) of the Remote VPC.

**Type**: string
**Required**: True

### creationTime

The time which the VPC Connnection is created.

**Type**: string
**Required**: False

### state

State of the Remote VPC Connection.

**Type**: VpcConnectionState
**Required**: False

### authentication

The type of private link authentication.

**Type**: string
**Required**: False

## Error

Returns information about an error.

### message

The description of the error.

**Type**: string
**Required**: False

### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## ListClientVpcConnectionsResponse

The response contains an array vpc connections.

### clientVpcConnections

An array of client vpc connections information objects.

**Type**: Array of type ClientVpcConnection
**Required**: False

### nextToken

If the response of ListClientVpcConnections is truncated, it returns a NextToken in the response. This Nexttoken should be sent in the subsequent request to ListClientVpcConnections.

**Type**: string

**Required**: False

### VpcConnectionState

State of the vpc connection

```
CREATING
AVAILABLE
INACTIVE
UPDATING
DEACTIVATING
DELETING
FAILED
REJECTED
REJECTING
```

## See also

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

### ListClientVpcConnections

- 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

# Clusters clusterArn Policy

## URI

`/v1/clusters/clusterArn/policy`

## HTTP methods

### GET

**Operation ID:** `GetClusterPolicy`

Get cluster policy.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | GetClusterPolicyResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |

| Status code | Response model | Description |
| --- | --- | --- |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## PUT

**Operation ID:** `PutClusterPolicy`

Create or update cluster policy.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | PutClusterPolicyResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## DELETE

**Operation ID:** `DeleteClusterPolicy`

Delete cluster policy.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | DeleteClusterPolicyResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying |

| Status code | Response model | Description |
| --- | --- | --- |
| | | your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning correct headers.

#### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### PUT schema

```
{
  "currentVersion": "string",
  "policy": "string"
}
```

## Response bodies

### GetClusterPolicyResponse schema

```
{
  "currentVersion": "string",
  "policy": "string"
}
```

### PutClusterPolicyResponse schema

```
{
  "currentVersion": "string"
}
```

### DeleteClusterPolicyResponse schema

```
{
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## DeleteClusterPolicyResponse

Delete resource policy for MSK cluster

## Error

Returns information about an error.

### message

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## GetClusterPolicyResponse

Returns resource policy for MSK cluster

**currentVersion**

Resource policy version

**Type**: string
**Required**: False

**policy**

Resource policy attached to the MSK cluster

**Type**: string
**Required**: False

## PutClusterPolicyRequest

Create or update resource policy for cluster

**currentVersion**

Current cluster policy version.

**Type**: string
**Required**: False
**MinLength**: 1

**policy**

Resource policy for cluster

**Type**: string
**Required**: True
**MinLength**: 1
**MaxLength**: 20480

### PutClusterPolicyResponse

Create or update cluster policy

**currentVersion**

Resource policy version

**Type**: string
**Required**: False

## See also

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

### GetClusterPolicy

- 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

### PutClusterPolicy

- 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

### DeleteClusterPolicy

- 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

# Clusters clusterArn Rebalancing

## URI

`/v1/clusters/clusterArn/rebalancing`

# HTTP methods

## PUT

**Operation ID:** `UpdateRebalancing`

Use this resource to update the intelligent rebalancing status of an Amazon MSK Provisioned cluster with Express brokers.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | UpdateRebalancingResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. |

| Status code | Response model | Description |
| --- | --- | --- |
| | | Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Use this resource to update the intelligent rebalancing status of an Amazon MSK Provisioned cluster with Express brokers.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

# Schemas

## Request bodies

### PUT schema

```
{
  "currentVersion": "string",
  "rebalancing": {
    "status": enum
  }
}
```

## Response bodies

### UpdateRebalancingResponse schema

```
{
  "clusterArn": "string",
  "clusterOperationArn": "string"
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Error

Returns information about an error.

### message

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## Rebalancing

Specifies whether or not intelligent rebalancing is turned on for a newly created MSK Provisioned cluster with Express brokers. Intelligent rebalancing performs automatic partition balancing operations when you scale your clusters up or down.

By default, intelligent rebalancing is `ACTIVE` for all new Express-based clusters.

**status**

Intelligent rebalancing status. The default intelligent rebalancing status is `ACTIVE` for all new Express-based clusters.

**Type**: RebalancingStatus
**Required**: True

## RebalancingStatus

Intelligent rebalancing status. The default intelligent rebalancing status is `ACTIVE` for all new Express-based clusters.

```
PAUSED
ACTIVE
```

## UpdateRebalancingRequest

Updates the intelligent rebalancing status for a new MSK Provisioned cluster with Express brokers.

**currentVersion**

The current version of the cluster.

**Type**: string

**Required**: True

**rebalancing**

Specifies if intelligent rebalancing should be turned on for your cluster. The default intelligent rebalancing status is `ACTIVE` for all new MSK Provisioned clusters that you create with Express brokers.

**Type**: Rebalancing
**Required**: True

## UpdateRebalancingResponse

Provides information about the intelligent rebalancing update for a cluster.

**clusterArn**

The Amazon Resource Name (ARN) of the cluster whose intelligent rebalancing status you've updated.

**Type**: string
**Required**: False

**clusterOperationArn**

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

# See also

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

## UpdateRebalancing

- 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

# Compatible Apache Kafka Versions

The Apache Kafka versions to which you can update the cluster.

## URI

`/v1/compatible-kafka-versions`

## HTTP methods

### GET

**Operation ID:** `GetCompatibleKafkaVersions`

Returns a list of the Apache Kafka versions to which you can update this cluster.

**Query parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| `clusterArn` | String | False | The Amazon Resource Name (ARN) of the cluster check. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | GetCompatibleKafkaVersionsResponse | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Response bodies

#### GetCompatibleKafkaVersionsResponse schema

```
{
  "compatibleKafkaVersions": [
    {
      "sourceVersion": "string",
      "targetVersions": [
        "string"
      ]
    }
  ]
}
```

#### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

## Properties

### CompatibleKafkaVersion

Contains source Apache Kafka versions and compatible target Apache Kafka versions.

#### sourceVersion

An Apache Kafka version.

**Type**: string
**Required**: False

#### targetVersions

A list of Apache Kafka versions.

**Type**: Array of type string
**Required**: False

## Error

Returns information about an error.

#### message

The description of the error.

**Type**: string
**Required**: False

#### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## GetCompatibleKafkaVersionsResponse

Response body for GetCompatibleKafkaVersions.

#### compatibleKafkaVersions

A list of CompatibleKafkaVersion objects.

**Type**: Array of type CompatibleKafkaVersion
**Required**: False

## See also

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

### GetCompatibleKafkaVersions

- 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

# Configuration

Represents an MSK configuration. Use this path to describe the configuration.

## URI

`/v1/configurations/arn`

## HTTP methods

### GET

**Operation ID:** `DescribeConfiguration`

Returns a description of this MSK configuration.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *arn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies an MSK configuration and all of its revisions . |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | DescribeConfigurationResponse | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |

| Status code | Response model | Description |
| --- | --- | --- |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## PUT

**Operation ID:** UpdateConfiguration

Creates a new revision of the cluster configuration. The configuration must be in the ACTIVE state.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *arn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies an MSK configuration and all of its revisions . |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | UpdateConfiguration Response | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |

| Status code | Response model | Description |
| --- | --- | --- |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 409 | Error | This cluster name already exists. Retry your request using another name. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## DELETE

**Operation ID:** `DeleteConfiguration`

Deletes a cluster configuration and all its revisions.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *arn* | String | True | The Amazon Resource Name (ARN) that |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | uniquely identifies an MSK configuration and all of its revisions . |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | DeleteConfiguration Response | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |

| Status code | Response model | Description |
| --- | --- | --- |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning the correct headers.

#### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *arn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies an MSK configuration and all of its revisions . |

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### PUT schema

```
{
  "description": "string",
  "serverProperties": "string"
}
```

## Response bodies

### DescribeConfigurationResponse schema

```
{
  "creationTime": "string",
  "name": "string",
  "description": "string",
  "kafkaVersions": [
    "string"
  ],
  "state": enum,
  "arn": "string",
  "latestRevision": {
    "creationTime": "string",
    "description": "string",
    "revision": integer
  }
}
```

### UpdateConfigurationResponse schema

```
{
  "arn": "string",
  "latestRevision": {
    "creationTime": "string",
    "description": "string",
    "revision": integer
  }
}
```

### DeleteConfigurationResponse schema

```
{
  "state": enum,
  "Arn": "string"
}
```

### Error schema

```
{
```

```
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## ConfigurationRevision

Describes a configuration revision.

### creationTime

The time when the configuration revision was created.

**Type**: string
**Required**: True

### description

The description of the configuration revision.

**Type**: string
**Required**: False

### revision

The revision number.

**Type**: integer
**Required**: True
**Format**: int64

## ConfigurationState

State of a kafka configuration

```
ACTIVE
DELETING
DELETE_FAILED
```

## DeleteConfigurationResponse

Returns information about the deleted configuration.

**state**

State of the configuration.

**Type**: ConfigurationState
**Required**: False

**Arn**

The Amazon Resource Name (ARN) of the configuration.

**Type**: string
**Required**: False

## DescribeConfigurationResponse

Response body for DescribeConfiguration.

**creationTime**

The time when the configuration was created.

**Type**: string
**Required**: True

**name**

The name of the configuration. Configuration names are strings that match the regex "^[0-9A-Za-z][0-9A-Za-z-]{0,}$".

**Type**: string
**Required**: True

**description**

The description of the configuration.

**Type**: string
**Required**: True

### kafkaVersions

The versions of Apache Kafka with which you can use this MSK configuration.

**Type**: Array of type string
**Required**: True

### state

State of the configuration.

**Type**: ConfigurationState
**Required**: False

### arn

The Amazon Resource Name (ARN) of the configuration.

**Type**: string
**Required**: True

### latestRevision

Latest revision of the configuration.

**Type**: ConfigurationRevision
**Required**: True

## Error

Returns information about an error.

### message

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## UpdateConfigurationRequest

Update an MSK configuration.

**description**

The description of the configuration.

**Type**: string
**Required**: False

**serverProperties**

Contents of the `server.properties` file. When using the API, you must ensure that the contents of the file are base64 encoded. When using the console, the SDK, or the CLI, the contents of `server.properties` can be in plaintext.

**Type**: string
**Required**: True

## UpdateConfigurationResponse

Response body for UpdateConfiguration

**arn**

The Amazon Resource Name (ARN) of the configuration.

**Type**: string

**Required**: False

#### latestRevision

Latest revision of the configuration.

**Type**: ConfigurationRevision
**Required**: False

## See also

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

### DescribeConfiguration

- 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

### UpdateConfiguration

- 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

### DeleteConfiguration

- 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

## Configuration Revision

Represents a specific revision of an MSK configuration.

### URI

`/v1/configurations/arn/revisions/revision`

### HTTP methods

### GET

**Operation ID:** `DescribeConfigurationRevision`

Returns a description of this revision of the configuration.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *revision* | String | True | A string that uniquely identifies a revision of an MSK configuration. |
| *arn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies an MSK configuration and all of its revisions. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | DescribeConfigurationRevisionResponse | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. |

| Status code | Response model | Description |
| --- | --- | --- |
| | | Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *revision* | String | True | A string that uniquely identifies a revision of an MSK configura tion. |
| *arn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies an MSK configuration and all of its revisions . |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Response bodies

#### DescribeConfigurationRevisionResponse schema

```
{
  "creationTime": "string",
  "description": "string",
  "serverProperties": "string",
  "arn": "string",
  "revision": integer
}
```

#### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

## Properties

### DescribeConfigurationRevisionResponse

Response body for DescribeConfigurationRevision.

#### creationTime

The time when the configuration was created.

**Type**: string
**Required**: True

**description**

The description of the configuration.

**Type**: string
**Required**: True

**serverProperties**

Contents of the `server.properties` file. When using the API, you must ensure that the contents of the file are base64 encoded. When using the console, the SDK, or the CLI, the contents of `server.properties` can be in plaintext.

**Type**: string
**Required**: True

**arn**

The Amazon Resource Name (ARN) of the configuration.

**Type**: string
**Required**: True

**revision**

The revision number.

**Type**: integer
**Required**: True
**Format**: int64

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## See also

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

### DescribeConfigurationRevision

- 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

# Configuration Revisions

Represents the revisions of an MSK configuration.

## URI

`/v1/configurations/arn/revisions`

# HTTP methods

## GET

**Operation ID:** `ListConfigurationRevisions`

Returns a list of all the revisions of an MSK configuration.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *arn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies an MSK configuration and all of its revisions . |

### Query parameters

| Name | Type | Required | Description |
|---|---|---|---|
| nextToken | String | False | The paginated results marker. When the result of the operation is truncated, the call returns NextToken in the response. To get the next batch, provide this token in your next request. |
| maxResults | String | False | The maximum number of results to return in the response (default |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | maximum 100 results per API call). If there are more results, the response includes a `NextToken` parameter. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | ListConfigurationRevisionsResponse | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying |

| Status code | Response model | Description |
| --- | --- | --- |
| | | your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning the correct headers.

#### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *arn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies an MSK configuration and all of its revisions . |

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Response bodies

#### ListConfigurationRevisionsResponse schema

```
{
  "nextToken": "string",
```

```
  "revisions": [
    {
      "creationTime": "string",
      "description": "string",
      "revision": integer
    }
  ]
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## ConfigurationRevision

Describes a configuration revision.

### creationTime

The time when the configuration revision was created.

**Type**: string
**Required**: True

### description

The description of the configuration revision.

**Type**: string
**Required**: False

### revision

The revision number.

**Type**: integer
**Required**: True
**Format**: int64

## Error

Returns information about an error.

### message

The description of the error.

**Type**: string
**Required**: False

### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## ListConfigurationRevisionsResponse

Information about revisions of an MSK configuration.

### nextToken

Paginated results marker.

**Type**: string
**Required**: False

### revisions

List of ConfigurationRevision objects.

**Type**: Array of type ConfigurationRevision

**Required**: False

## See also

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

### ListConfigurationRevisions

- 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

# Configurations

A collection of MSK configurations.

## URI

`/v1/configurations`

## HTTP methods

### GET

**Operation ID:** `ListConfigurations`

Returns a list of all the MSK configurations.

### Query parameters

| Name | Type | Required | Description |
|---|---|---|---|
| nextToken | String | False | The paginated results marker. When the result of the operation is truncated, the call returns NextToken in the response. To get the next batch, provide this token in your next request. |
| maxResults | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes a NextToken parameter. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | ListConfigurations Response | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |

| Status code | Response model | Description |
| --- | --- | --- |
| 401 | Error | The request is not authorize d. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## POST

**Operation ID:** `CreateConfiguration`

Creates a new MSK configuration. To see an example of how to use this operation, first save the following text to a file and name the file `config-file.txt`.

```
auto.create.topics.enable = true

zookeeper.connection.timeout.ms = 1000

log.roll.ms = 604800000
```

Now run the following Python 3.6 script in the folder where you saved `config-file.txt`. This script uses the properties specified in `config-file.txt` to create a configuration named

`SalesClusterConfiguration`. This configuration can work with Apache Kafka versions 1.1.1 and 2.1.0.

```
import boto3

client = boto3.client('kafka')

config_file = open('config-file.txt', 'r')

server_properties = config_file.read()

response = client.create_configuration(
    Name='SalesClusterConfiguration',
    Description='The configuration to use on all sales clusters.',
    KafkaVersions=['1.1.1', '2.1.0'],
    ServerProperties=server_properties
)

print(response)
```

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | CreateConfigurationResponse | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |

| Status code | Response model | Description |
| --- | --- | --- |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 409 | Error | This cluster name already exists. Retry your request using another name. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning the correct headers.

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### POST schema

```
{
```

```
  "name": "string",
  "description": "string",
  "kafkaVersions": [
    "string"
  ],
  "serverProperties": "string"
}
```

## Response bodies

### ListConfigurationsResponse schema

```
{
  "nextToken": "string",
  "configurations": [
    {
      "creationTime": "string",
      "name": "string",
      "description": "string",
      "kafkaVersions": [
        "string"
      ],
      "state": enum,
      "arn": "string",
      "latestRevision": {
        "creationTime": "string",
        "description": "string",
        "revision": integer
      }
    }
  ]
}
```

### CreateConfigurationResponse schema

```
{
  "creationTime": "string",
  "name": "string",
  "state": enum,
  "arn": "string",
  "latestRevision": {
    "creationTime": "string",
```

```
    "description": "string",
    "revision": integer
  }
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Configuration

Represents an MSK Configuration.

**creationTime**

The time when the configuration was created.

**Type**: string
**Required**: True

**name**

The name of the configuration. Configuration names are strings that match the regex "^[0-9A-Za-z][0-9A-Za-z-]{0,}$".

**Type**: string
**Required**: True

**description**

The description of the configuration.

**Type**: string
**Required**: True

**kafkaVersions**

An array of the versions of Apache Kafka with which you can use this MSK configuration. You can use this configuration for an MSK cluster only if the Apache Kafka version specified for the cluster appears in this array.

**Type**: Array of type string
**Required**: False

**state**

State of the configuration.

**Type**: ConfigurationState
**Required**: False

**arn**

The Amazon Resource Name (ARN) of the configuration.

**Type**: string
**Required**: True

**latestRevision**

Latest revision of the configuration.

**Type**: ConfigurationRevision
**Required**: True

## ConfigurationRevision

Describes a configuration revision.

**creationTime**

The time when the configuration revision was created.

**Type**: string
**Required**: True

**description**

The description of the configuration revision.

**Type**: string
**Required**: False

**revision**

The revision number.

**Type**: integer
**Required**: True
**Format**: int64

## ConfigurationState

State of a kafka configuration

```
ACTIVE
DELETING
DELETE_FAILED
```

## CreateConfigurationRequest

Request body for CreateConfiguration.

**name**

The name of the configuration. Configuration names are strings that match the regex "^[0-9A-Za-z][0-9A-Za-z-]{0,}$".

**Type**: string
**Required**: True

**description**

The description of the configuration.

**Type**: string

**Required**: False

**kafkaVersions**

The versions of Apache Kafka with which you can use this MSK configuration.

**Type**: Array of type string

**Required**: False

**serverProperties**

Contents of the `server.properties` file. When using the API, you must ensure that the contents of the file are base64 encoded. When using the console, the SDK, or the CLI, the contents of `server.properties` can be in plaintext.

**Type**: string

**Required**: True

## CreateConfigurationResponse

Response body for CreateConfiguration

**creationTime**

The time when the configuration was created.

**Type**: string

**Required**: False

**name**

The name of the configuration. Configuration names are strings that match the regex "^[0-9A-Za-z][0-9A-Za-z-]{0,}$".

**Type**: string

**Required**: False

**state**

State of the configuration.

**Type**: ConfigurationState
**Required**: False

**arn**

The Amazon Resource Name (ARN) of the configuration.

**Type**: string
**Required**: False

**latestRevision**

Latest revision of the configuration.

**Type**: ConfigurationRevision
**Required**: False

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

### ListConfigurationsResponse

The response contains an array of Configuration and a next token if the response is truncated.

**nextToken**

The paginated results marker. When the result of a `ListConfigurations` operation is truncated, the call returns `NextToken` in the response. To get another batch of configurations, provide this token in your next request.

**Type**: string
**Required**: False

**configurations**

An array of MSK configurations.

**Type**: Array of type Configuration
**Required**: False

## See also

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

### ListConfigurations

- 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

### CreateConfiguration

- 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

## Monitoring Properties

The options for monitoring an Amazon MSK cluster. You can specify which Apache Kafka metrics you want Amazon MSK to gather and send to Amazon CloudWatch. You can also configure open monitoring to gather metrics with Prometheus or Prometheus-compatible tools.

### URI

`/v1/clusters/`*`clusterArn`*`/monitoring`

### HTTP methods

#### PUT

**Operation ID:** `UpdateMonitoring`

Updates the monitoring settings for the cluster. You can use this operation to specify which Apache Kafka metrics you want Amazon MSK to send to Amazon CloudWatch. You can also specify settings for open monitoring with Prometheus. The following Python 3.6 example enables open monitoring with the Node Exporter. It also sets enhanced monitoring to PER_BROKER. For more information about monitoring, see Monitoring.

```
import boto3
```

```
import time

client = boto3.client('kafka')

update_monitoring_response = client.update_monitoring(
    ClusterArn='arn:aws:kafka:us-east-1:0123456789019:cluster/SalesCluster/abcd1234-
abcd-cafe-abab-9876543210ab-4',
    CurrentVersion='K12V3IB1VIZHHY',
    EnhancedMonitoring='PER_BROKER',
    OpenMonitoring={"Prometheus":{"JmxExporter":
{"EnabledInBroker":False},"NodeExporter":{"EnabledInBroker":True}}}
)

operation_arn = update_monitoring_response['ClusterOperationArn']
print('The ARN of the update operation is ' + operation_arn)

describe_cluster_operation_response =
 client.describe_cluster_operation(ClusterOperationArn=operation_arn)

operation_state = describe_cluster_operation_response['ClusterOperationInfo']
['OperationState']
print('The status of the update operation is ' + operation_state)

updated = False

while not updated:
    print('Sleeping for 15 seconds before checking to see if the monitoring update is
 done...')
    time.sleep(15)
    describe_cluster_operation_response =
 client.describe_cluster_operation(ClusterOperationArn=operation_arn)
    operation_state = describe_cluster_operation_response['ClusterOperationInfo']
['OperationState']
    if 'UPDATE_COMPLETE' == operation_state:
        updated = True
        print('You have successfully updated the monitoring settings.')
```

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | UpdateMonitoringRequest | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

# Schemas

## Request bodies

### PUT schema

```
{
  "enhancedMonitoring": enum,
  "loggingInfo": {
    "brokerLogs": {
      "s3": {
        "bucket": "string",
        "prefix": "string",
        "enabled": boolean
      },
      "firehose": {
        "deliveryStream": "string",
        "enabled": boolean
      },
```

```
      "cloudWatchLogs": {
        "logGroup": "string",
        "enabled": boolean
      }
    },
    "authorizerLogs": {
      "s3": {
        "bucket": "string",
        "prefix": "string",
        "enabled": boolean
      },
      "firehose": {
        "deliveryStream": "string",
        "enabled": boolean
      },
      "cloudWatchLogs": {
        "logGroup": "string",
        "enabled": boolean
      }
    }
  },
  "openMonitoring": {
    "prometheus": {
      "nodeExporter": {
        "enabledInBroker": boolean
      },
      "jmxExporter": {
        "enabledInBroker": boolean
      }
    }
  },
  "currentVersion": "string"
}
```

## Response bodies

### UpdateMonitoringRequest schema

```
{
  "clusterArn": "string",
  "clusterOperationArn": "string"
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## AuthorizerLogs

Authorizer Logs details for cluster.

**s3**

S3 Log destination details.

**Type**: S3
**Required**: False

**firehose**

Firehose Log destination details.

**Type**: Firehose
**Required**: False

**cloudWatchLogs**

CloudWatch Log destination details.

**Type**: CloudWatchLogs
**Required**: False

## BrokerLogs

The broker logs configuration for this MSK cluster.

**s3**

Details of the Amazon S3 destination for broker logs.

**Type**: S3
**Required**: False

**firehose**

Details of the Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: Firehose
**Required**: False

**cloudWatchLogs**

Details of the CloudWatch Logs destination for broker logs.

**Type**: CloudWatchLogs
**Required**: False

## CloudWatchLogs

Details of the CloudWatch Logs destination for broker logs.

**logGroup**

The CloudWatch log group that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified CloudWatch Logs destination.

**Type**: boolean
**Required**: True

## EnhancedMonitoring

Specifies which Apache Kafka metrics Amazon MSK gathers and sends to Amazon CloudWatch for this cluster. This property has three possible values: DEFAULT, PER_BROKER, and

PER_TOPIC_PER_BROKER. For a list of the metrics associated with each of these three levels of monitoring, see Monitoring.

DEFAULT
PER_BROKER
PER_TOPIC_PER_BROKER
PER_TOPIC_PER_PARTITION

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## Firehose

Firehose details for BrokerLogs.

**deliveryStream**

The Kinesis Data Firehose delivery stream that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Kinesis Data Firehose delivery stream.

**Type**: boolean

**Required**: True

## JmxExporterInfo

JMX Exporter details.

### enabledInBroker

JMX Exporter being enabled in broker.

**Type**: boolean

**Required**: True

## LoggingInfo

You can configure your MSK cluster to send broker logs to different destination types. This is a container for the configuration details related to broker logs.

### brokerLogs

You can configure your MSK cluster to send broker logs to different destination types. This configuration specifies the details of these destinations.

**Type**: BrokerLogs

**Required**: False

### authorizerLogs

Authorizer Logs details.

**Type**: AuthorizerLogs

**Required**: False

## NodeExporterInfo

Node Exporter details.

**enabledInBroker**

Node Exporter being enabled in broker.

**Type**: boolean
**Required**: True

## OpenMonitoringInfo

JMX and Node monitoring for cluster.

**prometheus**

Prometheus details.

**Type**: PrometheusInfo
**Required**: True

## PrometheusInfo

Prometheus details.

**nodeExporter**

Node Exporter details.

**Type**: NodeExporterInfo
**Required**: False

**jmxExporter**

JMX Exporter details.

**Type**: JmxExporterInfo
**Required**: False

## S3

The details of the Amazon S3 destination for broker logs.

**bucket**

The name of the S3 bucket that is the destination for broker logs.

**Type**: string
**Required**: False

**prefix**

The S3 prefix that is the destination for broker logs.

**Type**: string
**Required**: False

**enabled**

Specifies whether broker logs get sent to the specified Amazon S3 destination.

**Type**: boolean
**Required**: True

## UpdateMonitoringRequest

Request body for UpdateMonitoring.

**enhancedMonitoring**

Specifies which Apache Kafka metrics Amazon MSK gathers and sends to Amazon CloudWatch for this cluster.

**Type**: EnhancedMonitoring
**Required**: False

**loggingInfo**

LoggingInfo details.

**Type**: LoggingInfo
**Required**: False

**openMonitoring**

The settings for open monitoring.

**Type**: OpenMonitoringInfo
**Required**: False

**currentVersion**

The version of the MSK cluster to update. Cluster versions aren't simple numbers. You can describe an MSK cluster to find its version. When this update operation is successful, it generates a new cluster version.

**Type**: string
**Required**: True

## UpdateMonitoringResponse

Request body for UpdateMonitoring.

**clusterArn**

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

**clusterOperationArn**

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

# See also

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

### UpdateMonitoring

- 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

# Reboot Broker

## URI

`/v1/clusters/clusterArn/reboot-broker`

## HTTP methods

### PUT

**Operation ID:** `RebootBroker`

Reboots a broker. In a given cluster, you can reboot one broker at a time.

To reboot a broker, wait for the cluster status to be ACTIVE. This operation returns an error if you invoke it while the cluster status is HEALING. You must wait for the status to change from HEALING to ACTIVE before you reboot the broker.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | RebootBrokerResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning the correct headers.

#### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

#### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### PUT schema

```
{
  "brokerIds": [
    "string"
  ]
}
```

### Response bodies

#### RebootBrokerResponse schema

```
{
  "clusterArn": "string",
```

```
  "clusterOperationArn": "string"
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## RebootBrokerRequest

The request body for the RebootBroker action.

**brokerIds**

The list of broker IDs to be rebooted. Specify only one broker ID.

**Type**: Array of type string

**Required**: True

### RebootBrokerResponse

The response body for RebootBroker.

#### clusterArn

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

#### clusterOperationArn

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

## See also

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

### RebootBroker

- 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

# Scram Secrets

Represents a secret stored in the Amazon Secrets Manager that can be used to authenticate with a cluster using your sign-in credentials.

## URI

`/v1/clusters/clusterArn/scram-secrets`

## HTTP methods

### GET

**Operation ID:** `ListScramSecrets`

Returns a list of SCRAM secrets associated with the cluster. SCRAM secrets are stored in the Amazon Secrets Manager service, and are used to authenticate clients using sign-in credentials.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Query parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| nextToken | String | False | The paginated results marker. When the result of the operation is truncated, the call |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | returns `NextToken` in the response. To get the next batch, provide this token in your next request. |
| `maxResults` | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes a `NextToken` parameter. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | ListScramSecretsResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |

| Status code | Response model | Description |
| --- | --- | --- |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## POST

**Operation ID:** `BatchAssociateScramSecret`

Associates a list of SCRAM secrets with a cluster. SCRAM secrets are stored in the Amazon Secrets Manager service, and are used to authenticate clients using sign-in credentials. You can associate up to 10 secrets with a cluster at a time.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | BatchAssociateScramSecretResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | None | Default response for CORS method |

## PATCH

**Operation ID:** `BatchDisassociateScramSecret`

Disassociates a list of SCRAM secrets from a cluster. SCRAM secrets are stored in the Amazon Secrets Manager service, and are used to authenticate clients using sign-in credentials. You can disassociate up to 10 secrets from a cluster at a time.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | BatchDisassociateScramSecretResponse | 200 response |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

# Schemas

## Request bodies

### POST schema

```
{
  "secretArnList": [
    "string"
  ]
}
```

### PATCH schema

```
{
  "secretArnList": [
    "string"
  ]
}
```

## Response bodies

### ListScramSecretsResponse schema

```
{
  "secretArnList": [
    "string"
  ],
  "nextToken": "string"
}
```

### BatchAssociateScramSecretResponse schema

```
{
  "clusterArn": "string",
  "unprocessedScramSecrets": [
    {
      "secretArn": "string",
      "errorMessage": "string",
      "errorCode": "string"
```

```
    }
  ]
}
```

### BatchDisassociateScramSecretResponse schema

```
{
  "clusterArn": "string",
  "unprocessedScramSecrets": [
    {
      "secretArn": "string",
      "errorMessage": "string",
      "errorCode": "string"
    }
  ]
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

## Properties

### BatchAssociateScramSecretRequest

Request body for BatchAssociateScramSecret.

#### secretArnList

List of Amazon Secrets Manager secret Amazon Resource Name (ARN)s.

**Type**: Array of type string
**Required**: True

### BatchAssociateScramSecretResponse

Response body for BatchAssociateScramSecret.

**clusterArn**

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

**unprocessedScramSecrets**

List of errors when associating secrets to cluster.

**Type**: Array of type UnprocessedScramSecret
**Required**: False

## BatchDisassociateScramSecretRequest

Request body for BatchDisassociateScramSecret.

**secretArnList**

List of Amazon Secrets Manager secret Amazon Resource Name (ARN)s.

**Type**: Array of type string
**Required**: True

## BatchDisassociateScramSecretResponse

Response body for BatchDisassociateScramSecret.

**clusterArn**

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

**unprocessedScramSecrets**

List of errors when disassociating secrets to cluster.

**Type**: Array of type UnprocessedScramSecret
**Required**: False

## Error

Returns information about an error.

### message

The description of the error.

**Type**: string
**Required**: False

### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## ListScramSecretsResponse

Information about scram secrets associated to the cluster.

### secretArnList

The list of scram secrets associated with the cluster.

**Type**: Array of type string
**Required**: False

### nextToken

Paginated results marker.

**Type**: string
**Required**: False

### UnprocessedScramSecret

Error info for scram secret associate/disassociate failure.

#### secretArn

Amazon Secrets Manager secret Amazon Resource Name (ARN).

**Type**: string
**Required**: False

#### errorMessage

Error message for associate/disassociate failure.

**Type**: string
**Required**: False

#### errorCode

Error code for associate/disassociate failure.

**Type**: string
**Required**: False

## See also

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

### ListScramSecrets

- 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

## BatchAssociateScramSecret

- 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

## BatchDisassociateScramSecret

- 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

# Tags

Represents the set of tags for an MSK resource. A tag is a key-value pair that you define for the cluster. Using tags is a simple yet powerful way to manage Amazon resources and organize data, including billing data.

## URI

`/v1/tags/resourceArn`

## HTTP methods

### GET

**Operation ID:** `ListTagsForResource`

Returns a list of the tags associated with the specified resource.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *resourceArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the resource that's associated with the tags. |

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | ListTagsForResourceResponse | Success response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |

| Status code | Response model | Description |
| --- | --- | --- |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## POST

**Operation ID:** `TagResource`

Adds tags to the specified MSK resource.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *resourceArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the resource that's |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | associated with the tags. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 204 | None | 204 response |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorize d. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## DELETE

**Operation ID:** `UntagResource`

Removes the tags associated with the keys that are provided in the query.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *resourceArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the resource that's associated with the tags. |

### Query parameters

| Name | Type | Required | Description |
|---|---|---|---|
| tagKeys | String | True | Tag keys must be unique for a given cluster. In addition, the following restrictions apply:<br>• Each tag key must be unique. If you add a tag with a key that's already in use, your new tag overwrites the existing key-value pair.<br>• You can't start a tag key with `aws:` because this prefix |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | is reserved for use by Amazon. Amazon creates tags that begin with this prefix on your behalf, but you can't edit or delete them.<br>• Tag keys must be between 1 and 128 Unicode characters in length.<br>• Tag keys must consist of the following characters: Unicode letters, digits, white space, and the following special characters: `_ . / = + - @`. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 204 | None | 204 response |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |

| Status code | Response model | Description |
| --- | --- | --- |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *resourceArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the resource that's associated with the tags. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### POST schema

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

### Response bodies

#### ListTagsForResourceResponse schema

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

#### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

## Properties

### Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

### ListTagsForResourceResponse

Response of listing tags for a resource.

**tags**

The key-value pair for the resource tag.

**Type**: object
**Required**: True

### TagResourceRequest

Tag a resource.

**tags**

The key-value pair for the resource tag.

**Type**: object
**Required**: True

## See also

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

## ListTagsForResource

- 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

- 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

- 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

# Topic

## URI

`/v1/clusters/clusterArn/topics/topicName`

## HTTP methods

### GET

**Operation ID:** `DescribeTopic`

Returns details for a topic on a cluster.

This API response reflects data that updates approximately every minute. For the most current topic state after making changes, allow approximately one minute before querying.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |
| *topicName* | String | True | The name of the topic. |

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | DescribeTopicResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## PUT

**Operation ID:** `UpdateTopic`

Updates topic partition or topic configs.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |
| *topicName* | String | True | The name of the topic. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | UpdateTopicResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorize d. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying |

| Status code | Response model | Description |
| --- | --- | --- |
| | | your request might resolve the issue. |
| 503 | Error | 503 response |

## DELETE

**Operation ID:** `DeleteTopic`

Deletes the topic specified by the topicName in the request from the cluster specified by the Amazon Resource Name (ARN) in the request.

### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |
| *topicName* | String | True | The name of the topic. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | DeleteTopicResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |

| Status code | Response model | Description |
| --- | --- | --- |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning correct headers

### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |
| *topicName* | String | True | The name of the topic. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

# Schemas

## Request bodies

### PUT schema

```
{
  "partitionCount": integer,
  "configs": "string",
  "tags": {
  }
}
```

## Response bodies

### DescribeTopicResponse schema

```
{
  "replicationFactor": number,
  "partitionCount": number,
  "configs": "string",
  "topicName": "string",
  "topicArn": "string",
  "status": enum
}
```

### UpdateTopicResponse schema

```
{
  "topicName": "string",
  "topicArn": "string",
  "status": enum
}
```

### DeleteTopicResponse schema

```
{
  "topicName": "string",
  "topicArn": "string",
  "status": enum
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## DeleteTopicResponse

Returns information about the deleted topic.

### topicName

Name of the topic.

**Type**: string
**Required**: False

### topicArn

ARN of the topic.

**Type**: string
**Required**: False

### status

Status of the topic.

**Type**: TopicStatus

**Required**: False

## DescribeTopicResponse

The response contains metadata for a topic on a cluster.

### replicationFactor

Replication factor for a topic

**Type**: number
**Required**: False

### partitionCount

Partition count for a topic

**Type**: number
**Required**: False

### configs

Kafka configs for a topic.

**Type**: string
**Required**: False

### topicName

Name for a topic.

**Type**: string
**Required**: False

### topicArn

ARN of the topic.

**Type**: string

**Required**: False

#### status

Status of the topic.

**Type**: TopicStatus
**Required**: False

## Error

Returns information about an error.

#### message

The description of the error.

**Type**: string
**Required**: False

#### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## TopicStatus

Status of a kafka topic

```
CREATING
UPDATING
DELETING
ACTIVE
```

## UpdateTopicRequest

Request body for UpdateTopic.

**partitionCount**

How many partitions in the topic to be created.

**Type**: integer
**Required**: False

**configs**

Base64 encoded kafka configs.

**Type**: string
**Required**: False

**tags**

Tags attached to the topic.

**Type**: object
**Required**: False

## UpdateTopicResponse

Returns information about the updated topic.

**topicName**

Name of the topic.

**Type**: string
**Required**: False

**topicArn**

ARN of the topic.

**Type**: string
**Required**: False

**status**

Status of the topic.

**Type**: TopicStatus
**Required**: False

## See also

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

### DescribeTopic

- 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

### UpdateTopic

- 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

### DeleteTopic

- 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

## Topic Partitions

## URI

`/v1/clusters/clusterArn/topics/topicName/partitions`

## HTTP methods

### GET

**Operation ID:** `DescribeTopicPartitions`

Returns all partition information for a topic on a cluster.

This API response reflects data that updates approximately every minute. For the most current topic state after making changes, allow approximately one minute before querying.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |
| *topicName* | String | True | The name of the topic. |

### Query parameters

| Name | Type | Required | Description |
|---|---|---|---|
| nextToken | String | False | The paginated results marker. When the result of the operation is truncated, the call returns NextToken in the response. To get the next batch, provide this token in your next request. |
| maxResults | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes a NextToken parameter. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | DescribeTopicPartitionsResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning correct headers

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |
| *topicName* | String | True | The name of the topic. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | None | Default response for CORS method |

## Schemas

### Response bodies

#### DescribeTopicPartitionsResponse schema

```
{
  "partitions": [
    {
      "leader": number,
      "partition": number,
      "replicas": [
        number
      ],
      "isr": [
        number
      ]
    }
  ],
  "nextToken": "string"
```

```
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## DescribeTopicPartitionsResponse

The response contains information about partitions for a topic on a cluster.

**partitions**

List containing partition info.

**Type**: Array of type TopicPartitionInfo
**Required**: False

**nextToken**

If the response of DescribeTopicPartitions is truncated, it returns a NextToken in the response. This NextToken should be sent in the subsequent request to DescribeTopicPartitions.

**Type**: string
**Required**: False

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string

**Required**: False

### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## TopicPartitionInfo

Includes information about a partition.

### leader

The broker ID of the leader for this partition.

**Type**: number
**Required**: False

### partition

The partition number.

**Type**: number
**Required**: False

### replicas

The list of broker IDs that are replicas for this partition.

**Type**: Array of type number
**Required**: False

### isr

The list of in-sync replica broker IDs for this partition.

**Type**: Array of type number

**Required**: False

## See also

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

### DescribeTopicPartitions

- 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

# Topics

## URI

`/v1/clusters/`*`clusterArn`*`/topics`

## HTTP methods

### GET

**Operation ID:** `ListTopics`

Returns all topics in a cluster.

This API response reflects data that updates approximately every minute. For the most current topic state after making changes, allow approximately one minute before querying.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Query parameters

| Name | Type | Required | Description |
|---|---|---|---|
| topicNameFilter | String | False | Filters results to channels whose topic name matches the specified value. |
| nextToken | String | False | The paginated results marker. When the result of the operation is truncated, the call returns NextToken in the response. To get the next batch, provide this token in your next request. |
| maxResults | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | a NextToken parameter. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | ListTopicsResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorize d. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## POST

**Operation ID:** `CreateTopic`

Creates a topic using the specified properties in the request on the cluster specified by the Amazon Resource Name (ARN) in the request.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | CreateTopicResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorize d. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |

| Status code | Response model | Description |
| --- | --- | --- |
| 409 | Error | This cluster name already exists. Retry your request using another name. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning correct headers

### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

# Schemas

## Request bodies

### POST schema

```
{
  "partitionCount": integer,
  "configs": "string",
  "replicationFactor": integer,
  "topicName": "string",
  "tags": {
  }
}
```

## Response bodies

### ListTopicsResponse schema

```
{
  "topics": [
    {
      "partitionCount": number,
      "replicationFactor": number,
      "topicName": "string",
      "topicArn": "string",
      "outOfSyncReplicaCount": number
    }
  ],
  "nextToken": "string"
}
```

### CreateTopicResponse schema

```
{
  "topicName": "string",
  "topicArn": "string",
  "status": enum
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## CreateTopicRequest

Request body for CreateTopic.

### partitionCount

How many partitions in the topic to be created.

**Type**: integer
**Required**: True

### configs

Base64 encoded Kafka configs.

**Type**: string
**Required**: False

### replicationFactor

The replication factor of the topic.

**Type**: integer
**Required**: True

### topicName

Name of the topic.

**Type**: string
**Required**: True

**tags**

Tags attached to the topic.

**Type**: object
**Required**: False

## CreateTopicResponse

Returns information about the created topic.

**topicName**

Name of the topic.

**Type**: string
**Required**: False

**topicArn**

ARN of the topic.

**Type**: string
**Required**: False

**status**

Status of the topic.

**Type**: TopicStatus
**Required**: False

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

#### invalidParameter

The parameter that caused the error.

**Type**: string
**Required**: False

## ListTopicsResponse

The response contains an array of topics in a cluster.

#### topics

List containing topics info.

**Type**: Array of type TopicInfo
**Required**: False

#### nextToken

If the response of ListTopics is truncated, it returns a NextToken in the response. This NextToken should be sent in the subsequent request to ListTopics.

**Type**: string
**Required**: False

## TopicInfo

Includes identification info about the topic.

#### partitionCount

Partition count for a topic

**Type**: number

**Required**: False

### replicationFactor

Replication factor for a topic

**Type**: number
**Required**: False

### topicName

Name for a topic.

**Type**: string
**Required**: False

### topicArn

ARN of the topic.

**Type**: string
**Required**: False

### outOfSyncReplicaCount

Number of out of sync replica for topic

**Type**: number
**Required**: False

## TopicStatus

Status of a kafka topic

```
CREATING
UPDATING
DELETING
```

```
ACTIVE
```

## See also

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

### ListTopics

- 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

### CreateTopic

- 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

# UpdateStorage

Updates cluster broker volume size (or) sets cluster storage mode to TIERED.

## URI

`/v1/clusters/clusterArn/storage`

## HTTP methods

### PUT

**Operation ID:** `UpdateStorage`

Updates cluster broker volume size (or) sets cluster storage mode to TIERED.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | UpdateStorageResponse | Successful response. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |

| Status code | Response model | Description |
| --- | --- | --- |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning the correct headers.

### Path parameters

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *clusterArn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies the cluster. |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method |

## Schemas

### Request bodies

#### PUT schema

```
{
  "volumeSizeGB": integer,
  "storageMode": enum,
  "provisionedThroughput": {
    "volumeThroughput": integer,
    "enabled": boolean
  },
  "currentVersion": "string"
}
```

### Response bodies

#### UpdateStorageResponse schema

```
{
  "clusterArn": "string",
  "clusterOperationArn": "string"
}
```

#### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

## ProvisionedThroughput

Contains information about provisioned throughput for EBS storage volumes attached to kafka broker nodes.

**volumeThroughput**

Throughput value of the EBS volumes for the data drive on each kafka broker node in MiB per second.

**Type**: integer
**Required**: False

**enabled**

Provisioned throughput is enabled or not.

**Type**: boolean
**Required**: False

## StorageMode

Controls storage mode for various supported storage tiers.

LOCAL
TIERED

## UpdateStorageRequest

Request object for UpdateStorageApi.

### volumeSizeGB

size of the EBS volume to update.

**Type**: integer
**Required**: False

### storageMode

Controls storage mode for supported storage tiers.

**Type**: StorageMode
**Required**: False

### provisionedThroughput

EBS volume provisioned throughput information.

**Type**: ProvisionedThroughput
**Required**: False

### currentVersion

The version of cluster to update from. A successful operation will then generate a new version.

**Type**: string
**Required**: True

## UpdateStorageResponse

Response body for UpdateStorageResponse Api.

### clusterArn

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: False

### clusterOperationArn

The Amazon Resource Name (ARN) of the cluster operation.

**Type**: string
**Required**: False

# See also

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

## UpdateStorage

- 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

# Vpc-connection

## URI

`/v1/vpc-connection`

## HTTP methods

### POST

**Operation ID:** `CreateVpcConnection`

Create remote VPC connection.

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | CreateVpcConnectionResponse | HTTP Status Code 200: OK. |
| 400 | Error | The request isn't valid because the input is incorrect. Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |

| Status code | Response model | Description |
| --- | --- | --- |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning correct headers.

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method. |

## Schemas

### Request bodies

#### POST schema

```
{
  "clientSubnets": [
    "string"
  ],
  "targetClusterArn": "string",
  "vpcId": "string",
  "securityGroups": [
    "string"
  ],
  "tags": {
  },
  "authentication": "string"
}
```

## Response bodies

### CreateVpcConnectionResponse schema

```
{
  "vpcConnectionArn": "string",
  "clientSubnets": [
    "string"
  ],
  "creationTime": "string",
  "vpcId": "string",
  "securityGroups": [
    "string"
  ],
  "state": enum,
  "tags": {
  },
  "authentication": "string"
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## CreateVpcConnectionRequest

Creates a remote VPC connection for account client.

### clientSubnets

The list of subnets in the client VPC to connect to.

**Type**: Array of type string

**Required**: True

**targetClusterArn**

The Amazon Resource Name (ARN) of the cluster.

**Type**: string
**Required**: True

**vpcId**

The VPC id of the remote client.

**Type**: string
**Required**: True

**securityGroups**

The security groups to attach to the ENIs for the broker nodes.

**Type**: Array of type string
**Required**: False

**tags**

Create tags when creating the VPC connection.

**Type**: object
**Required**: False

**authentication**

The authentication type for the client VPC connection. Specify one of these auth type strings: `SASL_IAM`, `SASL_SCRAM`, or `TLS`.

**Type**: string
**Required**: True
**MinLength**: 3
**MaxLength**: 10

## CreateVpcConnectionResponse

Returns information about the created VPC connection.

### vpcConnectionArn

The Amazon Resource Name (ARN) of the remote VPC connection.

**Type**: string
**Required**: False

### clientSubnets

The list of subnets in the client VPC to connect to.

**Type**: Array of type string
**Required**: False

### creationTime

The time when the VPC connection was created.

**Type**: string
**Required**: False

### vpcId

The VPC id of the remote client.

**Type**: string
**Required**: False

### securityGroups

The security groups to attach to the ENIs for the broker nodes.

**Type**: Array of type string
**Required**: False

**state**

State of the VPC connection.

**Type**: VpcConnectionState
**Required**: False

**tags**

Tags attached to the VPC connection.

**Type**: object
**Required**: False

**authentication**

The type of authentication to be uses by remote clients.

**Type**: string
**Required**: False
**MinLength**: 3
**MaxLength**: 10

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string

**Required**: False

### VpcConnectionState

State of the vpc connection

```
CREATING
AVAILABLE
INACTIVE
UPDATING
DEACTIVATING
DELETING
FAILED
REJECTED
REJECTING
```

## See also

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

### CreateVpcConnection

- 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

# Vpc-connection arn

## URI

`/v1/vpc-connection/arn`

## HTTP methods

### GET

**Operation ID:** `DescribeVpcConnection`

Describes Remote VPC Connection.

**Path parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| *arn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies an MSK configuration and all of its revisions . |

**Responses**

| Status code | Response model | Description |
|---|---|---|
| 200 | DescribeVpcConnectionResponse | HTTP Status Code 200: OK. |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |

| Status code | Response model | Description |
| --- | --- | --- |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## DELETE

**Operation ID:** `DeleteVpcConnection`

Delete remote VPC connection.

**Path parameters**

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| *arn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies an MSK configuration |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | and all of its revisions . |

### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | DeleteVpcConnectionResponse | HTTP Status Code 200: OK. |
| 400 | Error | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| 401 | Error | The request is not authorized. The provided credentials couldn't be validated. |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

## OPTIONS

Enable CORS by returning correct headers.

### Path parameters

| Name | Type | Required | Description |
|---|---|---|---|
| *arn* | String | True | The Amazon Resource Name (ARN) that uniquely identifies an MSK configuration and all of its revisions . |

### Responses

| Status code | Response model | Description |
|---|---|---|
| 200 | None | Default response for CORS method. |

# Schemas

## Response bodies

### DescribeVpcConnectionResponse schema

```
{
  "vpcConnectionArn": "string",
  "creationTime": "string",
  "targetClusterArn": "string",
  "vpcId": "string",
  "subnets": [
    "string"
  ],
  "securityGroups": [
    "string"
  ],
```

```
  "state": enum,
  "tags": {
  },
  "authentication": "string"
}
```

### DeleteVpcConnectionResponse schema

```
{
  "vpcConnectionArn": "string",
  "state": enum
}
```

### Error schema

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## DeleteVpcConnectionResponse

Returns information about the deleted VPC connection.

### vpcConnectionArn

The Amazon Resource Name (ARN) of the Remote VPC.

**Type**: string
**Required**: False

### state

State of the Remote VPC Connection.

**Type**: VpcConnectionState
**Required**: False

## DescribeVpcConnectionResponse

Response body for DescribeVpcConnection.

### vpcConnectionArn

The Amazon Resource Name (ARN) of the Remote VPC.

**Type**: string
**Required**: True

### creationTime

The time when the configuration was created.

**Type**: string
**Required**: True

### targetClusterArn

The Amazon Resource Name (ARN) of the target cluster.

**Type**: string
**Required**: True

### vpcId

The description of the vpcId.

**Type**: string
**Required**: False

### subnets

The list of subnets in Remote VPC Connection.

**Type**: Array of type string
**Required**: False

**securityGroups**

The list of security groups in Remote VPC Connection.

**Type**: Array of type string
**Required**: False

**state**

State of the Remote VPC Connection.

**Type**: VpcConnectionState
**Required**: False

**tags**

Tags attached to the vpc connection.

**Type**: object
**Required**: False

**authentication**

The type of private link authentication.

**Type**: string
**Required**: False

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string
**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string
**Required**: False

### VpcConnectionState

State of the vpc connection

```
CREATING
AVAILABLE
INACTIVE
UPDATING
DEACTIVATING
DELETING
FAILED
REJECTED
REJECTING
```

## See also

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

### DescribeVpcConnection

- 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

### DeleteVpcConnection

- 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

## Vpc-connections

### URI

`/v1/vpc-connections`

### HTTP methods

#### GET

**Operation ID:** `ListVpcConnections`

Lists all VPC connections.

**Query parameters**

| Name | Type | Required | Description |
|---|---|---|---|
| `nextToken` | String | False | The paginated results marker. |

| Name | Type | Required | Description |
| --- | --- | --- | --- |
| | | | When the result of the operation is truncated, the call returns `NextToken` in the response. To get the next batch, provide this token in your next request. |
| `maxResults` | String | False | The maximum number of results to return in the response (default maximum 100 results per API call). If there are more results, the response includes a `NextToken` parameter. |

**Responses**

| Status code | Response model | Description |
| --- | --- | --- |
| `200` | `ListVpcConnections Response` | HTTP Status Code 200: OK. |
| `400` | `Error` | The request isn't valid because the input is incorrect . Correct your input and then submit it again. |
| `401` | `Error` | The request is not authorize d. The provided credentials couldn't be validated. |

| Status code | Response model | Description |
| --- | --- | --- |
| 403 | Error | Access forbidden. Check your credentials and then retry your request. |
| 404 | Error | The resource could not be found due to incorrect input. Correct the input, then retry the request. |
| 429 | Error | 429 response |
| 500 | Error | There was an unexpected internal server error. Retrying your request might resolve the issue. |
| 503 | Error | 503 response |

### OPTIONS

Enable CORS by returning correct headers.

#### Responses

| Status code | Response model | Description |
| --- | --- | --- |
| 200 | None | Default response for CORS method. |

## Schemas

### Response bodies

#### ListVpcConnectionsResponse schema

```
{
```

```
  "nextToken": "string",
  "vpcConnections": [
    {
      "vpcConnectionArn": "string",
      "creationTime": "string",
      "targetClusterArn": "string",
      "vpcId": "string",
      "state": enum,
      "authentication": "string"
    }
  ]
}
```

**Error schema**

```
{
  "message": "string",
  "invalidParameter": "string"
}
```

# Properties

## Error

Returns information about an error.

**message**

The description of the error.

**Type**: string

**Required**: False

**invalidParameter**

The parameter that caused the error.

**Type**: string

**Required**: False

## ListVpcConnectionsResponse

The response contains an array of vpcConnections and a next token if the response is truncated.

**nextToken**

If the response of ListVpcConnections is truncated, it returns a NextToken in the response. This NextToken should be sent in the subsequent request to ListVpcConnections.

**Type**: string
**Required**: False

**vpcConnections**

An array of VPC Connection.

**Type**: Array of type VpcConnection
**Required**: False

## VpcConnection

Vpc Connection description

**vpcConnectionArn**

The Amazon Resource Name (ARN) of the Remote Vpc.

**Type**: string
**Required**: True

**creationTime**

The time which the VPC Connnection is created.

**Type**: string
**Required**: False

**targetClusterArn**

The Amazon Resource Name (ARN) of the target cluster.

**Type**: string
**Required**: True

### vpcId

The description of the vpcId.

**Type**: string
**Required**: False

### state

State of the Remote VPC Connection.

**Type**: VpcConnectionState
**Required**: False

### authentication

The type of private link authentication.

**Type**: string
**Required**: False

## VpcConnectionState

State of the vpc connection

```
CREATING
AVAILABLE
INACTIVE
UPDATING
DEACTIVATING
DELETING
FAILED
REJECTED
REJECTING
```

## See also

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

### ListVpcConnections

- 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

# Document History for Amazon Managed Streaming for Apache Kafka API Reference

The following table describes the documentation for this release of the *Amazon Managed Streaming for Apache Kafka API Reference*.

- **API version:** 2019-05-30
- **Latest documentation update:** October 26, 2022.

| Change | Description | Date |
| --- | --- | --- |
| Amazon MSK GA release | This is the general-availability release of the Amazon MSK API Reference to support Amazon MSK Replicator. | October 17, 2023 |
| Amazon MSK GA release | This is the general-availability release of the Amazon MSK API Reference to support tiered storage. | October 26, 2022 |
| Amazon MSK GA release | This is the general-availability release of the Amazon MSK API Reference. | May 30, 2019 |
| Amazon MSK preview release | This is the preview release of the Amazon MSK API Reference. | November 29, 2018 |

# AWS Glossary

For the latest AWS terminology, see the AWS glossary in the *AWS Glossary Reference*.