



Amazon MSK Replicator API Reference

Amazon Managed Streaming for Apache Kafka



Amazon Managed Streaming for Apache Kafka: Amazon MSK Replicator API Reference

Copyright © 2024 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

What Is Amazon Managed Streaming for Apache Kafka?	1
Operations	2
Resources	3
V1 Replicators	3
URI	3
HTTP methods	3
Schemas	7
Properties	9
See also	20
V1 Replicators replicatorArn	21
URI	21
HTTP methods	21
Schemas	25
Properties	26
See also	35
V1 Replicators replicatorArn Replication-info	36
URI	36
HTTP methods	36
Schemas	38
Properties	40
See also	43
Document History	45
AWS Glossary	46

What Is Amazon Managed Streaming for Apache Kafka?

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.

This Amazon MSK Replicator API Reference describes Amazon MSK Replicator operations: `CreateReplicator`, `DeleteReplicator`, `DescribeReplicator`, `ListReplicators`, `UpdateReplicationInfo`.

Other Amazon MSK API references include:

- [Amazon Managed Streaming for Apache Kafka API Reference 1.0](#)
- [Amazon Managed Streaming for Apache Kafka API Reference 2.0](#)
- [Amazon MSK Connect API Reference](#)

For more information about Amazon MSK, see the [Amazon MSK Developer Guide](#).

Operations

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

- [CreateReplicator](#)

Creates the replicator.

- [DeleteReplicator](#)

Deletes a replicator.

- [DescribeReplicator](#)

Describes a replicator.

- [ListReplicators](#)

Lists the replicators.

- [UpdateReplicationInfo](#)

Updates replication info of a replicator.

Resources

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

Topics

- [V1 Replicators](#)
- [V1 Replicators replicatorArn](#)
- [V1 Replicators replicatorArn Replication-info](#)

V1 Replicators

URI

/replication/v1/replicators

HTTP methods

GET

Operation ID: ListReplicators

Lists the replicators.

Query parameters

Name	Type	Required	Description
replicatorNameFilter	String	False	Returns replicators starting with given name.
nextToken	String	False	If the response of ListReplicators is truncated, it returns a NextToken in the response. This NextToken should be sent in the subsequent

Name	Type	Required	Description
maxResults	String	False	The maximum number of results to return in the response. If there are more results, the response includes a NextToken parameter.

Responses

Status code	Response model	Description
200	ListReplicatorsResponse	HTTP Status Code 200: OK.
400	None	HTTP Status Code 400: Bad request due to incorrect input. Correct your request and then retry it.
401	None	HTTP Status Code 401: Unauthorized request. The provided credentials couldn't be validated.
403	None	HTTP Status Code 403: Access forbidden. Correct your credentials and then retry your request.
404	None	HTTP Status Code 404: Resource not found due to

Status code	Response model	Description
		incorrect input. Correct your request and then retry it.
429	None	HTTP Status Code 429: Limit exceeded. Resource limit reached.
500	None	HTTP Status Code 500: Unexpected internal server error. Retrying your request might resolve the issue.
503	None	HTTP Status Code 503: Service Unavailable. Retrying your request in some time might resolve the issue.

POST

Operation ID: CreateReplicator

Creates the replicator.

Responses

Status code	Response model	Description
200	CreateReplicatorResponse	HTTP Status Code 200: OK.
400	None	HTTP Status Code 400: Bad request due to incorrect input. Correct your request and then retry it.
401	None	HTTP Status Code 401: Unauthorized request. The

Status code	Response model	Description
403	None	provided credentials couldn't be validated. HTTP Status Code 403: Access forbidden. Correct your credentials and then retry your request.
404	None	HTTP Status Code 404: Resource not found due to incorrect input. Correct your request and then retry it.
409	None	HTTP Status Code 409: Conflict. This replicator name already exists. Retry your request with another name.
429	None	HTTP Status Code 429: Limit exceeded. Resource limit reached.
500	None	HTTP Status Code 500: Unexpected internal server error. Retrying your request might resolve the issue.
503	None	HTTP Status Code 503: Service Unavailable. Retrying your request in some time might resolve the issue.

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

```
{
  "replicatorName": "string",
  "serviceExecutionRoleArn": "string",
  "replicationInfoList": [
    {
      "consumerGroupReplication": {
        "consumerGroupsToExclude": [
          "string"
        ],
        "detectAndCopyNewConsumerGroups": boolean,
        "consumerGroupsToReplicate": [
          "string"
        ],
        "synchroniseConsumerGroupOffsets": boolean
      },
      "targetCompressionType": enum,
      "topicReplication": {
        "copyAccessControlListsForTopics": boolean,
        "detectAndCopyNewTopics": boolean,
        "copyTopicConfigurations": boolean,
        "startingPosition": {
          "type": enum
        }
      },
      "topicsToReplicate": [
        "string"
      ],
      "topicsToExclude": [
        "string"
      ]
    }
  ]
}
```

```

    },
    "sourceKafkaClusterArn": "string",
    "targetKafkaClusterArn": "string"
  }
],
"description": "string",
"kafkaClusters": [
  {
    "amazonMskCluster": {
      "mskClusterArn": "string"
    },
    "vpcConfig": {
      "securityGroupIds": [
        "string"
      ],
      "subnetIds": [
        "string"
      ]
    }
  }
],
"tags": {
}
}

```

Response bodies

ListReplicatorsResponse schema

```

{
  "nextToken": "string",
  "replicators": [
    {
      "replicatorArn": "string",
      "replicatorName": "string",
      "creationTime": "string",
      "kafkaClustersSummary": [
        {
          "kafkaClusterAlias": "string",
          "amazonMskCluster": {
            "mskClusterArn": "string"
          }
        }
      ]
    }
  ]
}

```

```
    ],
    "replicatorState": enum,
    "isReplicatorReference": boolean,
    "replicationInfoSummaryList": [
      {
        "sourceKafkaClusterAlias": "string",
        "targetKafkaClusterAlias": "string"
      }
    ],
    "replicatorResourceArn": "string",
    "currentVersion": "string"
  }
]
}
```

CreateReplicatorResponse schema

```
{
  "replicatorArn": "string",
  "replicatorName": "string",
  "replicatorState": enum
}
```

Properties

AmazonMskCluster

Details of an Amazon MSK Cluster.

mskClusterArn

The Amazon Resource Name (ARN) of an Amazon MSK cluster.

Type: string

Required: True

ConsumerGroupReplication

Details about consumer group replication.

consumerGroupsToExclude

List of regular expression patterns indicating the consumer groups that should not be replicated.

Type: Array of type string

Required: False

MaxLength: 256

detectAndCopyNewConsumerGroups

Enables synchronization of consumer groups to target cluster.

Type: boolean

Required: False

consumerGroupsToReplicate

List of regular expression patterns indicating the consumer groups to copy.

Type: Array of type string

Required: True

MaxLength: 256

synchroniseConsumerGroupOffsets

Enables synchronization of consumer group offsets to target cluster. The translated offsets will be written to topic `__consumer_offsets`.

Type: boolean

Required: False

CreateReplicatorRequest

Request body for replicator.

replicatorName

The name of the replicator. Alpha-numeric characters with '-' are allowed.

Type: string

Required: True

Pattern: `^[0-9A-Za-z][0-9A-Za-z-]{0,}.$`

MinLength: 1

MaxLength: 128

serviceExecutionRoleArn

The ARN of the IAM role used by the replicator to access resources in the customer's account (e.g source and target clusters)

Type: string

Required: True

replicationInfoList

A list of replication configurations, where each configuration targets a given source cluster to target cluster replication flow.

Type: Array of type [ReplicationInfo](#)

Required: True

description

A summary description of the replicator.

Type: string

Required: False

MaxLength: 1024

kafkaClusters

Kafka Clusters to use in setting up sources / targets for replication.

Type: Array of type [KafkaCluster](#)

Required: True

tags

List of tags to attach to created Replicator.

Type: object

Required: False

CreateReplicatorResponse

Returns information about the created replicator.

replicatorArn

The Amazon Resource Name (ARN) of the replicator.

Type: string

Required: False

replicatorName

Name of the replicator provided by the customer.

Type: string

Required: False

replicatorState

State of the replicator.

Type: [ReplicatorState](#)

Required: False

KafkaCluster

Information about Kafka Cluster to be used as source / target for replication.

amazonMskCluster

Details of an Amazon MSK Cluster.

Type: [AmazonMskCluster](#)

Required: True

vpcConfig

Details of an Amazon VPC which has network connectivity to the Apache Kafka cluster.

Type: [KafkaClusterClientVpcConfig](#)

Required: True

KafkaClusterClientVpcConfig

Details of an Amazon VPC which has network connectivity to the Apache Kafka cluster.

securityGroupIds

The security groups to attach to the ENIs for the broker nodes.

Type: Array of type string

Required: False

subnetIds

The list of subnets in the client VPC to connect to.

Type: Array of type string

Required: True

KafkaClusterSummary

Summarized information about Kafka Cluster used as source / target for replication.

kafkaClusterAlias

The alias of the Kafka cluster. Used to prefix names of replicated topics.

Type: string

Required: False

amazonMskCluster

Details of an Amazon MSK Cluster.

Type: [AmazonMskCluster](#)

Required: False

ListReplicatorsResponse

The response contains an array containing replicator information and a NextToken if the response is truncated.

nextToken

If the response of ListReplicators is truncated, it returns a NextToken in the response. This NextToken should be sent in the subsequent request to ListReplicators.

Type: string

Required: False

replicators

List containing information of each of the replicators in the account.

Type: Array of type [ReplicatorSummary](#)

Required: False

ReplicationInfo

Specifies configuration for replication between a source and target Kafka cluster.

consumerGroupReplication

Configuration relating to consumer group replication.

Type: [ConsumerGroupReplication](#)

Required: True

targetCompressionType

The compression type to use when producing records to target cluster.

Type: [TargetCompressionType](#)

Required: True

topicReplication

Configuration relating to topic replication.

Type: [TopicReplication](#)

Required: True

sourceKafkaClusterArn

The ARN of the source Kafka cluster.

Type: string

Required: True

targetKafkaClusterArn

The ARN of the target Kafka cluster.

Type: string

Required: True

ReplicationInfoSummary

Summarized information of replication between clusters.

sourceKafkaClusterAlias

The alias of the source Kafka cluster.

Type: string

Required: False

targetKafkaClusterAlias

The alias of the target Kafka cluster.

Type: string

Required: False

ReplicationStartingPosition

Specifies the position in the topics to start replicating from.

type

The type of replication starting position.

Type: [ReplicationStartingPositionType](#)

Required: False

ReplicationStartingPositionType

The type of replication starting position.

LATEST

EARLIEST

ReplicatorState

State of a replicator.

RUNNING

CREATING

UPDATING

DELETING

FAILED

ReplicatorSummary

Information about a replicator.

replicatorArn

The Amazon Resource Name (ARN) of the replicator.

Type: string

Required: False

replicatorName

The name of the replicator.

Type: string

Required: False

creationTime

The time the replicator was created.

Type: string

Required: False

kafkaClustersSummary

Kafka Clusters used in setting up sources / targets for replication.

Type: Array of type [KafkaClusterSummary](#)

Required: False

replicatorState

State of the replicator.

Type: [ReplicatorState](#)

Required: False

isReplicatorReference

Whether this resource is a replicator reference.

Type: boolean

Required: False

replicationInfoSummaryList

A list of summarized information of replications between clusters.

Type: Array of type [ReplicationInfoSummary](#)

Required: False

replicatorResourceArn

The Amazon Resource Name (ARN) of the replicator resource in the region where the replicator was created.

Type: string

Required: False

currentVersion

The current version of the replicator.

Type: string

Required: False

TargetCompressionType

The type of compression to use producing records to the target cluster.

NONE

GZIP

SNAPPY

LZ4

ZSTD

TopicReplication

Details about topic replication.

copyAccessControlListsForTopics

Whether to periodically configure remote topic ACLs to match their corresponding upstream topics.

Type: boolean

Required: False

detectAndCopyNewTopics

Whether to periodically check for new topics and partitions.

Type: boolean

Required: False

copyTopicConfigurations

Whether to periodically configure remote topics to match their corresponding upstream topics.

Type: boolean

Required: False

startingPosition

Specifies the position in the topics to start replicating from.

Type: [ReplicationStartingPosition](#)

Required: False

topicsToReplicate

List of regular expression patterns indicating the topics to copy.

Type: Array of type string

Required: True

MaxLength: 249

topicsToExclude

List of regular expression patterns indicating the topics that should not be replicated.

Type: Array of type string

Required: False

MaxLength: 249

See also

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

ListReplicators

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

CreateReplicator

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

V1 Replicators replicatorArn

URI

/replication/v1/replicators/*replicatorArn*

HTTP methods

GET

Operation ID: DescribeReplicator

Describes a replicator.

Path parameters

Name	Type	Required	Description
<i>replicatorArn</i>	String	True	The Amazon Resource Name (ARN) of the replicator to be described.

Responses

Status code	Response model	Description
200	DescribeReplicator Response	HTTP Status Code 200: OK.
400	None	HTTP Status Code 400: Bad request due to incorrect input. Correct your request and then retry it.
401	None	HTTP Status Code 401: Unauthorized request. The provided credentials couldn't be validated.

Status code	Response model	Description
403	None	HTTP Status Code 403: Access forbidden. Correct your credentials and then retry your request.
404	None	HTTP Status Code 404: Resource not found due to incorrect input. Correct your request and then retry it.
429	None	HTTP Status Code 429: Limit exceeded. Resource limit reached.
500	None	HTTP Status Code 500: Unexpected internal server error. Retrying your request might resolve the issue.
503	None	HTTP Status Code 503: Service Unavailable. Retrying your request in some time might resolve the issue.

DELETE

Operation ID: DeleteReplicator

Deletes a replicator.

Path parameters

Name	Type	Required	Description
<i>replicatorArn</i>	String	True	The Amazon Resource Name (ARN) of the

Name	Type	Required	Description
			replicator to be described.

Query parameters

Name	Type	Required	Description
currentVersion	String	False	The current version of the replicator.

Responses

Status code	Response model	Description
200	DeleteReplicatorResponse	HTTP Status Code 200: OK.
400	None	HTTP Status Code 400: Bad request due to incorrect input. Correct your request and then retry it.
401	None	HTTP Status Code 401: Unauthorized request. The provided credentials couldn't be validated.
403	None	HTTP Status Code 403: Access forbidden. Correct your credentials and then retry your request.
404	None	HTTP Status Code 404: Resource not found due to incorrect input. Correct your request and then retry it.

Status code	Response model	Description
429	None	HTTP Status Code 429: Limit exceeded. Resource limit reached.
500	None	HTTP Status Code 500: Unexpected internal server error. Retrying your request might resolve the issue.
503	None	HTTP Status Code 503: Service Unavailable. Retrying your request in some time might resolve the issue.

OPTIONS

Enable CORS by returning correct headers

Path parameters

Name	Type	Required	Description
<i>replicatorArn</i>	String	True	The Amazon Resource Name (ARN) of the replicator to be described.

Responses

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

Schemas

Response bodies

DescribeReplicatorResponse schema

```
{
  "replicatorArn": "string",
  "creationTime": "string",
  "kafkaClusters": [
    {
      "kafkaClusterAlias": "string",
      "amazonMskCluster": {
        "mskClusterArn": "string"
      },
      "vpcConfig": {
        "securityGroupIds": [
          "string"
        ],
        "subnetIds": [
          "string"
        ]
      }
    }
  ],
  "currentVersion": "string",
  "tags": {
  },
  "replicatorDescription": "string",
  "replicatorName": "string",
  "serviceExecutionRoleArn": "string",
  "replicationInfoList": [
    {
      "consumerGroupReplication": {
        "consumerGroupsToExclude": [
          "string"
        ],
        "detectAndCopyNewConsumerGroups": boolean,
        "consumerGroupsToReplicate": [
          "string"
        ],
        "synchroniseConsumerGroupOffsets": boolean
      },
    }
  ],
}
```

```
"targetCompressionType": enum,
"sourceKafkaClusterAlias": "string",
"topicReplication": {
  "copyAccessControlListsForTopics": boolean,
  "detectAndCopyNewTopics": boolean,
  "copyTopicConfigurations": boolean,
  "startingPosition": {
    "type": enum
  },
  "topicsToReplicate": [
    "string"
  ],
  "topicsToExclude": [
    "string"
  ]
},
"targetKafkaClusterAlias": "string"
}
],
"stateInfo": {
  "code": "string",
  "message": "string"
},
"replicatorState": enum,
"isReplicatorReference": boolean,
"replicatorResourceArn": "string"
}
```

DeleteReplicatorResponse schema

```
{
  "replicatorArn": "string",
  "replicatorState": enum
}
```

Properties

AmazonMskCluster

Details of an Amazon MSK Cluster.

mskClusterArn

The Amazon Resource Name (ARN) of an Amazon MSK cluster.

Type: string

Required: True

ConsumerGroupReplication

Details about consumer group replication.

consumerGroupsToExclude

List of regular expression patterns indicating the consumer groups that should not be replicated.

Type: Array of type string

Required: False

MaxLength: 256

detectAndCopyNewConsumerGroups

Enables synchronization of consumer groups to target cluster.

Type: boolean

Required: False

consumerGroupsToReplicate

List of regular expression patterns indicating the consumer groups to copy.

Type: Array of type string

Required: True

MaxLength: 256

synchroniseConsumerGroupOffsets

Enables synchronization of consumer group offsets to target cluster. The translated offsets will be written to topic `__consumer_offsets`.

Type: boolean

Required: False

DeleteReplicatorResponse

Returns information about the deleted replicator.

replicatorArn

The Amazon Resource Name (ARN) of the replicator.

Type: string

Required: False

replicatorState

State of the replicator.

Type: [ReplicatorState](#)

Required: False

DescribeReplicatorResponse

Response body for DescribeReplicator.

replicatorArn

The Amazon Resource Name (ARN) of the replicator.

Type: string

Required: False

creationTime

The time when the replicator was created.

Type: string

Required: False

kafkaClusters

Kafka Clusters used in setting up sources / targets for replication.

Type: Array of type [KafkaClusterDescription](#)

Required: False

currentVersion

The current version number of the replicator.

Type: string

Required: False

tags

List of tags attached to the Replicator.

Type: object

Required: False

replicatorDescription

The description of the replicator.

Type: string

Required: False

replicatorName

The name of the replicator.

Type: string

Required: False

serviceExecutionRoleArn

The Amazon Resource Name (ARN) of the IAM role used by the replicator to access resources in the customer's account (e.g source and target clusters)

Type: string

Required: False

replicationInfoList

A list of replication configurations, where each configuration targets a given source cluster to target cluster replication flow.

Type: Array of type [ReplicationInfoDescription](#)

Required: False

stateInfo

Details about the state of the replicator.

Type: [ReplicationStateInfo](#)

Required: False

replicatorState

State of the replicator.

Type: [ReplicatorState](#)

Required: False

isReplicatorReference

Whether this resource is a replicator reference.

Type: boolean

Required: False

replicatorResourceArn

The Amazon Resource Name (ARN) of the replicator resource in the region where the replicator was created.

Type: string

Required: False

KafkaClusterClientVpcConfig

Details of an Amazon VPC which has network connectivity to the Apache Kafka cluster.

securityGroupIds

The security groups to attach to the ENIs for the broker nodes.

Type: Array of type string

Required: False

subnetIds

The list of subnets in the client VPC to connect to.

Type: Array of type string

Required: True

KafkaClusterDescription

Information about Kafka Cluster used as source / target for replication.

kafkaClusterAlias

The alias of the Kafka cluster. Used to prefix names of replicated topics.

Type: string

Required: False

amazonMskCluster

Details of an Amazon MSK Cluster.

Type: [AmazonMskCluster](#)

Required: False

vpcConfig

Details of an Amazon VPC which has network connectivity to the Apache Kafka cluster.

Type: [KafkaClusterClientVpcConfig](#)

Required: False

ReplicationInfoDescription

Specifies configuration for replication between a source and target Kafka cluster (sourceKafkaClusterAlias -> targetKafkaClusterAlias)

consumerGroupReplication

Configuration relating to consumer group replication.

Type: [ConsumerGroupReplication](#)

Required: False

targetCompressionType

The compression type to use when producing records to target cluster.

Type: [TargetCompressionType](#)

Required: False

sourceKafkaClusterAlias

The alias of the source Kafka cluster.

Type: string

Required: False

topicReplication

Configuration relating to topic replication.

Type: [TopicReplication](#)

Required: False

targetKafkaClusterAlias

The alias of the target Kafka cluster.

Type: string

Required: False

ReplicationStartingPosition

Specifies the position in the topics to start replicating from.

type

The type of replication starting position.

Type: [ReplicationStartingPositionType](#)

Required: False

ReplicationStartingPositionType

The type of replication starting position.

LATEST

EARLIEST

ReplicationStateInfo

Details about the state of a replicator.

code

Code that describes the state of the replicator.

Type: string

Required: False

message

Message that describes the state of the replicator.

Type: string

Required: False

ReplicatorState

State of a replicator.

RUNNING
CREATING
UPDATING
DELETING
FAILED

TargetCompressionType

The type of compression to use producing records to the target cluster.

NONE
GZIP
SNAPPY
LZ4
ZSTD

TopicReplication

Details about topic replication.

copyAccessControlListsForTopics

Whether to periodically configure remote topic ACLs to match their corresponding upstream topics.

Type: boolean
Required: False

detectAndCopyNewTopics

Whether to periodically check for new topics and partitions.

Type: boolean
Required: False

copyTopicConfigurations

Whether to periodically configure remote topics to match their corresponding upstream topics.

Type: boolean

Required: False

startingPosition

Specifies the position in the topics to start replicating from.

Type: [ReplicationStartingPosition](#)

Required: False

topicsToReplicate

List of regular expression patterns indicating the topics to copy.

Type: Array of type string

Required: True

MaxLength: 249

topicsToExclude

List of regular expression patterns indicating the topics that should not be replicated.

Type: Array of type string

Required: False

MaxLength: 249

See also

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

DescribeReplicator

- [AWS Command Line Interface](#)

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

DeleteReplicator

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

V1 Replicators replicatorArn Replication-info

URI

/replication/v1/replicators/*replicatorArn*/replication-info

HTTP methods

PUT

Operation ID: UpdateReplicationInfo

Updates replication info of a replicator.

Path parameters

Name	Type	Required	Description
<i>replicatorArn</i>	String	True	The Amazon Resource Name (ARN) of the replicator to be described.

Responses

Status code	Response model	Description
200	UpdateReplicationInfoResponse	HTTP Status Code 200: OK.
400	None	HTTP Status Code 400: Bad request due to incorrect input. Correct your request and then retry it.
401	None	HTTP Status Code 401: Unauthorized request. The provided credentials couldn't be validated.
403	None	HTTP Status Code 403: Access forbidden. Correct your credentials and then retry your request.
404	None	HTTP Status Code 404: Resource not found due to incorrect input. Correct your request and then retry it.
429	None	HTTP Status Code 429: Limit exceeded. Resource limit reached.

Status code	Response model	Description
500	None	HTTP Status Code 500: Unexpected internal server error. Retrying your request might resolve the issue.
503	None	HTTP Status Code 503: Service Unavailable. Retrying your request in some time might resolve the issue.

OPTIONS

Enable CORS by returning correct headers

Path parameters

Name	Type	Required	Description
<i>replicatorArn</i>	String	True	The Amazon Resource Name (ARN) of the replicator to be described.

Responses

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

Schemas

Request bodies

PUT schema

```
{
  "consumerGroupReplication": {
    "consumerGroupsToExclude": [
      "string"
    ],
    "detectAndCopyNewConsumerGroups": boolean,
    "consumerGroupsToReplicate": [
      "string"
    ],
    "synchroniseConsumerGroupOffsets": boolean
  },
  "topicReplication": {
    "copyAccessControlListsForTopics": boolean,
    "detectAndCopyNewTopics": boolean,
    "copyTopicConfigurations": boolean,
    "topicsToReplicate": [
      "string"
    ],
    "topicsToExclude": [
      "string"
    ]
  },
  "sourceKafkaClusterArn": "string",
  "targetKafkaClusterArn": "string",
  "currentVersion": "string"
}
```

Response bodies

UpdateReplicationInfoResponse schema

```
{
  "replicatorArn": "string",
  "replicatorState": enum
}
```

Properties

ConsumerGroupReplicationUpdate

Details about consumer group replication.

consumerGroupsToExclude

List of regular expression patterns indicating the consumer groups that should not be replicated.

Type: Array of type string

Required: True

MaxLength: 256

detectAndCopyNewConsumerGroups

Enables synchronization of consumer groups to target cluster.

Type: boolean

Required: True

consumerGroupsToReplicate

List of regular expression patterns indicating the consumer groups to copy.

Type: Array of type string

Required: True

MaxLength: 256

synchroniseConsumerGroupOffsets

Enables synchronization of consumer group offsets to target cluster. The translated offsets will be written to topic `__consumer_offsets`.

Type: boolean

Required: True

ReplicatorState

State of a replicator.

RUNNING
CREATING
UPDATING
DELETING
FAILED

TopicReplicationUpdate

Details for updating the topic replication of a replicator.

copyAccessControlListsForTopics

Whether to periodically configure remote topic ACLs to match their corresponding upstream topics.

Type: boolean

Required: True

detectAndCopyNewTopics

Whether to periodically check for new topics and partitions.

Type: boolean

Required: True

copyTopicConfigurations

Whether to periodically configure remote topics to match their corresponding upstream topics.

Type: boolean

Required: True

topicsToReplicate

List of regular expression patterns indicating the topics to copy.

Type: Array of type string

Required: True

MaxLength: 249

topicsToExclude

List of regular expression patterns indicating the topics that should not be replicated.

Type: Array of type string

Required: True

MaxLength: 249

UpdateReplicationInfoRequest

Parameters for updating replication information between source and target Kafka clusters of a replicator.

consumerGroupReplication

Updated consumer group replication information.

Type: [ConsumerGroupReplicationUpdate](#)

Required: False

topicReplication

Updated topic replication information.

Type: [TopicReplicationUpdate](#)

Required: False

sourceKafkaClusterArn

The ARN of the source Kafka cluster.

Type: string

Required: True

targetKafkaClusterArn

The ARN of the target Kafka cluster.

Type: string

Required: True

currentVersion

Current replicator version.

Type: string

Required: True

UpdateReplicationInfoResponse

Updated Replication information of a replicator.

replicatorArn

The Amazon Resource Name (ARN) of the replicator.

Type: string

Required: False

replicatorState

State of the replicator.

Type: [ReplicatorState](#)

Required: False

See also

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

UpdateReplicationInfo

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

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

Document History

The following table describes the documentation for this release of ServiceName.

- **API version: 1.0**
- **Latest documentation update:** April 10, 2024

The following table describes the important changes to the *Amazon MSK Replicator API Reference*. For notification about updates to this documentation, you can subscribe to an RSS feed.

Change	Description	Date
Added ReplicationStartingPosition	Amazon MSK supports replication from earliest position in your source cluster's topics in addition to latest position.	March 14, 2024
Added Amazon MSK Replicator APIs	Amazon MSK Replicator supports replicating existing data on your topics across Amazon Managed Streaming for Apache Kafka (Amazon MSK) clusters.	Oct 27, 2023

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.