



API Reference

Amazon Kinesis Data Streams Service



API Version 2013-12-02

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

Amazon Kinesis Data Streams Service: API Reference

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

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

Table of Contents

Welcome	1
Actions	2
AddTagsToStream	4
Request Syntax	4
Request Parameters	4
Response Elements	5
Errors	6
Examples	7
See Also	8
CreateChannel	9
Request Syntax	9
Request Parameters	11
Response Syntax	13
Response Elements	14
Errors	15
Examples	17
See Also	20
CreateStream	21
Request Syntax	22
Request Parameters	22
Response Elements	24
Errors	24
Examples	25
See Also	26
DecreaseStreamRetentionPeriod	27
Request Syntax	27
Request Parameters	27
Response Elements	28
Errors	28
Examples	29
See Also	30
DeleteChannel	32
Request Syntax	32
Request Parameters	32

Response Elements	32
Errors	32
See Also	34
DeleteResourcePolicy	35
Request Syntax	35
Request Parameters	35
Response Elements	36
Errors	36
See Also	37
DeleteStream	38
Request Syntax	38
Request Parameters	38
Response Elements	40
Errors	40
Examples	41
See Also	42
DeregisterStreamConsumer	43
Request Syntax	43
Request Parameters	43
Response Elements	44
Errors	44
See Also	45
DescribeAccountSettings	47
Response Syntax	47
Response Elements	47
Errors	47
See Also	48
DescribeChannel	49
Request Syntax	49
Request Parameters	49
Response Syntax	49
Response Elements	51
Errors	51
See Also	52
DescribeLimits	54
Response Syntax	54

Response Elements	54
Errors	55
Examples	56
See Also	56
DescribeStream	58
Request Syntax	58
Request Parameters	59
Response Syntax	60
Response Elements	61
Errors	61
Examples	62
See Also	64
DescribeStreamConsumer	66
Request Syntax	66
Request Parameters	66
Response Syntax	67
Response Elements	68
Errors	68
See Also	69
DescribeStreamSummary	70
Request Syntax	70
Request Parameters	70
Response Syntax	71
Response Elements	72
Errors	72
See Also	73
DisableEnhancedMonitoring	74
Request Syntax	74
Request Parameters	74
Response Syntax	76
Response Elements	76
Errors	77
Examples	78
See Also	79
EnableEnhancedMonitoring	81
Request Syntax	81

Request Parameters	81
Response Syntax	83
Response Elements	83
Errors	84
Examples	85
See Also	86
GetRecords	88
Request Syntax	89
Request Parameters	89
Response Syntax	91
Response Elements	91
Errors	92
Examples	95
See Also	96
GetResourcePolicy	97
Request Syntax	97
Request Parameters	97
Response Syntax	98
Response Elements	98
Errors	98
See Also	99
GetShardIterator	101
Request Syntax	102
Request Parameters	102
Response Syntax	104
Response Elements	105
Errors	105
Examples	106
See Also	107
IncreaseStreamRetentionPeriod	108
Request Syntax	108
Request Parameters	108
Response Elements	109
Errors	109
Examples	111
See Also	111

ListChannels	113
Request Syntax	113
Request Parameters	113
Response Syntax	114
Response Elements	114
Errors	115
See Also	116
ListShards	117
Request Syntax	117
Request Parameters	118
Response Syntax	121
Response Elements	121
Errors	122
Examples	124
See Also	125
ListStreamConsumers	127
Request Syntax	127
Request Parameters	127
Response Syntax	129
Response Elements	129
Errors	130
See Also	131
ListStreams	133
Request Syntax	133
Request Parameters	133
Response Syntax	134
Response Elements	134
Errors	135
Examples	136
See Also	137
ListTagsForResource	138
Request Syntax	138
Request Parameters	138
Response Syntax	139
Response Elements	139
Errors	139

Examples	140
See Also	141
ListTagsForStream	143
Request Syntax	143
Request Parameters	143
Response Syntax	144
Response Elements	145
Errors	145
Examples	146
See Also	147
MergeShards	149
Request Syntax	150
Request Parameters	150
Response Elements	151
Errors	151
Examples	153
See Also	154
PutRecord	155
Request Syntax	156
Request Parameters	156
Response Syntax	158
Response Elements	159
Errors	160
Examples	162
See Also	163
PutRecords	164
Request Syntax	165
Request Parameters	166
Response Syntax	167
Response Elements	167
Errors	168
Examples	171
See Also	173
PutResourcePolicy	175
Request Syntax	175
Request Parameters	175

Response Elements	176
Errors	176
See Also	177
RegisterStreamConsumer	179
Request Syntax	179
Request Parameters	179
Response Syntax	181
Response Elements	181
Errors	181
Examples	182
See Also	183
RemoveTagsFromStream	185
Request Syntax	185
Request Parameters	185
Response Elements	186
Errors	186
Examples	188
See Also	188
SplitShard	190
Request Syntax	191
Request Parameters	191
Response Elements	193
Errors	193
Examples	194
See Also	195
StartStreamEncryption	196
Request Syntax	196
Request Parameters	196
Response Elements	198
Errors	198
See Also	201
StopStreamEncryption	202
Request Syntax	202
Request Parameters	202
Response Elements	204
Errors	204

See Also	205
SubscribeToShard	207
Request Syntax	207
Request Parameters	208
Response Syntax	209
Response Elements	210
Errors	210
See Also	211
TagResource	213
Request Syntax	213
Request Parameters	213
Response Elements	214
Errors	214
Examples	215
See Also	216
UntagResource	217
Request Syntax	217
Request Parameters	217
Response Elements	218
Errors	218
Examples	219
See Also	220
UpdateAccountSettings	221
Request Syntax	221
Request Parameters	221
Response Syntax	221
Response Elements	222
Errors	222
See Also	223
UpdateChannel	224
Request Syntax	224
Request Parameters	224
Response Syntax	225
Response Elements	227
Errors	227
See Also	229

UpdateMaxRecordSize	230
Request Syntax	230
Request Parameters	230
Response Elements	231
Errors	231
See Also	232
UpdateShardCount	234
Request Syntax	235
Request Parameters	235
Response Syntax	236
Response Elements	237
Errors	238
Examples	239
See Also	240
UpdateStreamMode	241
Request Syntax	241
Request Parameters	241
Response Elements	242
Errors	242
See Also	243
UpdateStreamWarmThroughput	245
Request Syntax	246
Request Parameters	246
Response Syntax	247
Response Elements	247
Errors	248
See Also	249
Data Types	251
ChannelDescription	253
Contents	253
See Also	255
ChannelEncryptionConfiguration	257
Contents	257
See Also	257
ChannelLoggingConfiguration	258
Contents	258

See Also	258
ChannelLoggingUpdateInput	259
Contents	259
See Also	259
ChannelStreamConfiguration	260
Contents	260
See Also	260
ChannelStreamDescription	261
Contents	261
See Also	261
ChannelStreamIdentifier	263
Contents	263
See Also	263
ChannelSummary	264
Contents	264
See Also	266
ChildShard	267
Contents	267
See Also	267
CloudWatchLogs	269
Contents	269
See Also	269
CloudWatchLogsUpdateInput	271
Contents	271
See Also	271
Consumer	273
Contents	273
See Also	274
ConsumerDescription	275
Contents	275
See Also	276
DeadLetterQueueS3Configuration	277
Contents	277
See Also	278
EnhancedMetrics	279
Contents	279

See Also	279
HashKeyRange	281
Contents	281
See Also	281
MinimumThroughputBillingCommitmentInput	282
Contents	282
See Also	282
MinimumThroughputBillingCommitmentOutput	283
Contents	283
See Also	283
PartitionField	285
Contents	285
See Also	285
PartitionSpec	286
Contents	286
See Also	286
PutRecordsRequestEntry	287
Contents	287
See Also	288
PutRecordsResultEntry	289
Contents	289
See Also	290
Record	291
Contents	291
See Also	292
RecordConfiguration	293
Contents	293
See Also	293
S3DestinationConfiguration	295
Contents	295
See Also	295
S3DestinationDescription	297
Contents	297
See Also	297
S3DestinationUpdateInput	298
Contents	298

See Also	298
S3StorageConfiguration	299
Contents	299
See Also	300
S3TablesConfiguration	301
Contents	301
See Also	302
S3TablesDestinationConfiguration	303
Contents	303
See Also	303
S3TablesDestinationDescription	305
Contents	305
See Also	305
S3TablesDestinationUpdateInput	307
Contents	307
See Also	307
SequenceNumberRange	308
Contents	308
See Also	308
Shard	309
Contents	309
See Also	310
ShardFilter	311
Contents	311
See Also	312
StartingPosition	313
Contents	313
See Also	314
StreamDescription	315
Contents	315
See Also	318
StreamDescriptionSummary	319
Contents	319
See Also	323
StreamFilter	324
Contents	324

See Also	324
StreamModeDetails	325
Contents	325
See Also	325
StreamSummary	326
Contents	326
See Also	327
SubscribeToShardEvent	328
Contents	328
See Also	329
SubscribeToShardEventStream	330
Contents	330
See Also	332
Tag	333
Contents	333
See Also	333
WarmThroughputObject	334
Contents	334
See Also	334
Common Error Types	335

Welcome

Amazon Kinesis Data Streams is a managed service that scales elastically for real-time processing of streaming big data.

This document was last published on September 21, 2026.

Actions

The following actions are supported:

- [AddTagsToStream](#)
- [CreateChannel](#)
- [CreateStream](#)
- [DecreaseStreamRetentionPeriod](#)
- [DeleteChannel](#)
- [DeleteResourcePolicy](#)
- [DeleteStream](#)
- [DeregisterStreamConsumer](#)
- [DescribeAccountSettings](#)
- [DescribeChannel](#)
- [DescribeLimits](#)
- [DescribeStream](#)
- [DescribeStreamConsumer](#)
- [DescribeStreamSummary](#)
- [DisableEnhancedMonitoring](#)
- [EnableEnhancedMonitoring](#)
- [GetRecords](#)
- [GetResourcePolicy](#)
- [GetShardIterator](#)
- [IncreaseStreamRetentionPeriod](#)
- [ListChannels](#)
- [ListShards](#)
- [ListStreamConsumers](#)
- [ListStreams](#)
- [ListTagsForResource](#)
- [ListTagsForStream](#)
- [MergeShards](#)

- [PutRecord](#)
- [PutRecords](#)
- [PutResourcePolicy](#)
- [RegisterStreamConsumer](#)
- [RemoveTagsFromStream](#)
- [SplitShard](#)
- [StartStreamEncryption](#)
- [StopStreamEncryption](#)
- [SubscribeToShard](#)
- [TagResource](#)
- [UntagResource](#)
- [UpdateAccountSettings](#)
- [UpdateChannel](#)
- [UpdateMaxRecordSize](#)
- [UpdateShardCount](#)
- [UpdateStreamMode](#)
- [UpdateStreamWarmThroughput](#)

AddTagsToStream

Adds or updates tags for the specified Kinesis data stream. You can assign up to 50 tags to a data stream.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

If tags have already been assigned to the stream, `AddTagsToStream` overwrites any existing tags that correspond to the specified tag keys.

[AddTagsToStream](#) has a limit of five transactions per second per account.

Request Syntax

```
{
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string",
  "Tags": {
    "string" : "string"
  }
}
```

Request Parameters

The request accepts the following data in JSON format.

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Tags

A set of up to 50 key-value pairs to use to create the tags. A tag consists of a required key and an optional value. You can add up to 50 tags per resource.

Type: String to string map

Map Entries: Maximum number of 200 items.

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

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

Required: Yes

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To add tags to a stream

The following JSON example adds two tags to the specified stream.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.AddTagsToStream
{
  "StreamName": "exampleStreamName",
  "Tags": {
    "Project" : "myProject",
    "Environment" : "Production"
  }
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
```

See Also

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

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

CreateChannel

Creates a channel that delivers records from a Kinesis data stream to a destination. A channel reads records from the specified stream and writes them to streaming tables on Apache Iceberg (Amazon S3 Tables) or to a general purpose Amazon S3 bucket.

You must specify either `S3DestinationConfiguration` or `S3TablesDestinationConfiguration`, but not both.

To use this operation, you must have permission to pass the specified service execution IAM role to Amazon Kinesis Data Streams (the `iam:PassRole` permission on that role).

Creating a channel is an asynchronous operation. Upon receiving the request, Amazon Kinesis Data Streams returns immediately with the channel in the `CREATING` state. After provisioning is complete, Amazon Kinesis Data Streams sets the state to `ACTIVE`. You can use [DescribeChannel](#) to check the current state.

This operation is only supported for data streams with the on-demand capacity mode.

This operation has a call limit of 5 transactions per second (TPS) for each AWS account. Exceeding 5 TPS results in a `LimitExceededException`.

Request Syntax

```
{
  "ChannelName": "string",
  "EncryptionConfiguration": {
    "EncryptionType": "string",
    "KeyId": "string"
  },
  "LoggingConfiguration": {
    "CloudWatchLogs": {
      "Enabled": boolean,
      "LogGroupName": "string",
      "LogStreamName": "string"
    }
  },
  "S3DestinationConfiguration": {
    "DataFreshnessInSeconds": number,
    "DeadLetterQueueS3Configuration": {
      "BucketARN": "string",
      "ErrorOutputPrefix": "string",

```

```

    "ExpectedBucketOwner": "string"
  },
  "StorageConfiguration": {
    "BucketARN": "string",
    "CompressionType": "string",
    "ExpectedBucketOwner": "string",
    "OutputKeyTemplate": "string",
    "StorageClass": "string"
  }
},
"S3TablesDestinationConfiguration": {
  "DataFreshnessInSeconds": number,
  "DeadLetterQueueS3Configuration": {
    "BucketARN": "string",
    "ErrorOutputPrefix": "string",
    "ExpectedBucketOwner": "string"
  },
  "S3TablesConfigurationList": [
    {
      "CompressionType": "string",
      "Namespace": "string",
      "PartitionSpec": {
        "PartitionFields": [
          {
            "SourceName": "string",
            "Transform": "string"
          }
        ]
      },
      "TableBucketARN": "string",
      "TableName": "string"
    }
  ],
  "ServiceExecutionRoleARN": "string",
  "StreamConfigurationList": [
    {
      "RecordConfiguration": {
        "GSRSchemaARN": "string",
        "RecordFormatType": "string"
      },
      "StreamARN": "string"
    }
  ]
},

```

```
"Tags": {  
  "string" : "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

ChannelName

The name of the channel. The name is unique within your AWS account and AWS Region.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

EncryptionConfiguration

The server-side encryption configuration that uses an AWS KMS key to encrypt data delivered to the destination.

Type: [ChannelEncryptionConfiguration](#) object

Required: No

LoggingConfiguration

The Amazon CloudWatch Logs configuration for the channel.

Type: [ChannelLoggingConfiguration](#) object

Required: No

S3DestinationConfiguration

The configuration for delivery to a general purpose Amazon S3 bucket. Specify this parameter when `S3TablesDestinationConfiguration` is not specified.

Type: [S3DestinationConfiguration](#) object

Required: No

S3TablesDestinationConfiguration

The configuration for delivery to streaming tables on Apache Iceberg in Amazon S3 Tables. Specify this parameter when `S3DestinationConfiguration` is not specified.

Type: [S3TablesDestinationConfiguration](#) object

Required: No

ServiceExecutionRoleARN

The Amazon Resource Name (ARN) of the IAM role that Amazon Kinesis Data Streams assumes to write records to the destination.

Type: String

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

Pattern: `arn:aws[-a-z0-9]*:iam::\d{12}:role/[a-zA-Z_0-9+=,.\@\-_/]+`

Required: Yes

StreamConfigurationList

The source stream configuration for the channel. Currently, one stream is supported per channel.

Type: Array of [ChannelStreamConfiguration](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10000 items.

Required: Yes

Tags

A set of key-value pairs to assign to the channel. A tag consists of a required key and an optional value.

Type: String to string map

Map Entries: Maximum number of 200 items.

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

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

Required: No

Response Syntax

```
{
  "ChannelDescription": {
    "ChannelARN": "string",
    "ChannelCreationTimestamp": number,
    "ChannelId": "string",
    "ChannelName": "string",
    "ChannelStatus": "string",
    "ChannelStatusReason": "string",
    "EncryptionConfiguration": {
      "EncryptionType": "string",
      "KeyId": "string"
    },
    "LoggingConfiguration": {
      "CloudWatchLogs": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
      }
    },
    "S3DestinationConfiguration": {
      "DataFreshnessInSeconds": number,
      "DeadLetterQueueS3Configuration": {
        "BucketARN": "string",
        "ErrorOutputPrefix": "string",
        "ExpectedBucketOwner": "string"
      },
      "StorageConfiguration": {
        "BucketARN": "string",
        "CompressionType": "string",
        "ExpectedBucketOwner": "string",
        "OutputKeyTemplate": "string",
        "StorageClass": "string"
      }
    },
    "S3TablesDestinationConfiguration": {
```

```

    "DataFreshnessInSeconds": number,
    "DeadLetterQueueS3Configuration": {
      "BucketARN": "string",
      "ErrorOutputPrefix": "string",
      "ExpectedBucketOwner": "string"
    },
    "S3TablesConfigurationList": [
      {
        "CompressionType": "string",
        "Namespace": "string",
        "PartitionSpec": {
          "PartitionFields": [
            {
              "SourceName": "string",
              "Transform": "string"
            }
          ]
        },
        "TableBucketARN": "string",
        "TableName": "string"
      }
    ],
    "ServiceExecutionRoleARN": "string",
    "StreamConfigurationList": [
      {
        "RecordConfiguration": {
          "GSRSchemaARN": "string",
          "RecordFormatType": "string"
        },
        "StreamARN": "string",
        "StreamCreationTimestamp": number
      }
    ]
  }
}

```

Response Elements

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

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

ChannelDescription

The configuration and current status of the channel, including its ARN, destination configuration, and lifecycle state. Immediately after creation, the state is CREATING.

Type: [ChannelDescription](#) object

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSAccessDeniedException

The ciphertext references a key that doesn't exist or that you don't have access to.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSDisabledException

The request was rejected because the specified customer master key (CMK) isn't enabled.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSInvalidStateException

The request was rejected because the state of the specified resource isn't valid for this request. For more information, see [How Key State Affects Use of a Customer Master Key](#) in the *AWS Key Management Service Developer Guide*.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSNotFoundException

The request was rejected because the specified entity or resource can't be found.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSOptInRequired

The AWS access key ID needs a subscription for the service.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSThrottlingException

The request was denied due to request throttling. For more information about throttling, see [Limits](#) in the *AWS Key Management Service Developer Guide*.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

Examples

To create a channel that delivers to streaming tables on Apache Iceberg

The following JSON example creates a channel that writes records from a Kinesis data stream to a streaming table on Apache Iceberg backed by Amazon S3 Tables.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
```

```
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.CreateChannel
{
  "ChannelName": "orders-iceberg-channel",
  "ServiceExecutionRoleARN": "arn:aws:iam::123456789012:role/KinesisDeliveryRole",
  "StreamConfigurationList": [
    {
      "StreamARN": "arn:aws:kinesis:us-east-1:123456789012:stream/orders-stream",
      "RecordConfiguration": {
        "RecordFormatType": "JSON",
        "GSRSchemaARN": "arn:aws:glue:us-east-1:123456789012:schema/my-registry/orders-
schema"
      }
    }
  ],
  "S3TablesDestinationConfiguration": {
    "DataFreshnessInSeconds": 300,
    "DeadLetterQueueS3Configuration": {
      "BucketARN": "arn:aws:s3::my-dlq-bucket",
      "ExpectedBucketOwner": "123456789012",
      "ErrorOutputPrefix": "errors/"
    },
    "S3TablesConfigurationList": [
      {
        "TableBucketARN": "arn:aws:s3tables:us-east-1:123456789012:bucket/analytics-
table-bucket",
        "Namespace": "analytics_db",
        "TableName": "orders_iceberg",
        "CompressionType": "ZSTD",
        "PartitionSpec": {
          "PartitionFields": [
            { "Transform": "TIME_HOUR", "SourceName": "event_timestamp" }
          ]
        }
      }
    ]
  }
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "ChannelDescription": {
    "ChannelName": "orders-iceberg-channel",
    "ChannelARN": "arn:aws:kinesis:us-east-1:123456789012:channel/orders-iceberg-channel",
    "ChannelId": "abcd1234efgh5678",
    "ChannelStatus": "CREATING",
    "ChannelCreationTimestamp": 1755500000,
    "ServiceExecutionRoleARN": "arn:aws:iam::123456789012:role/KinesisDeliveryRole",
    "StreamConfigurationList": [
      {
        "StreamARN": "arn:aws:kinesis:us-east-1:123456789012:stream/orders-stream",
        "StreamCreationTimestamp": 1755400000,
        "RecordConfiguration": {
          "RecordFormatType": "JSON",
          "GSRSchemaARN": "arn:aws:glue:us-east-1:123456789012:schema/my-registry/orders-schema"
        }
      }
    ],
    "S3TablesDestinationConfiguration": {
      "DataFreshnessInSeconds": 300,
      "DeadLetterQueueS3Configuration": {
        "BucketARN": "arn:aws:s3:::my-dlq-bucket",
        "ExpectedBucketOwner": "123456789012",
        "ErrorOutputPrefix": "errors/"
      }
    },
    "S3TablesConfigurationList": [
      {
        "TableBucketARN": "arn:aws:s3tables:us-east-1:123456789012:bucket/analytics-table-bucket",
        "Namespace": "analytics_db",
        "TableName": "orders_iceberg",
        "CompressionType": "ZSTD",
        "PartitionSpec": {
          "PartitionFields": [
            { "Transform": "TIME_HOUR", "SourceName": "event_timestamp" }
          ]
        }
      }
    ]
  }
}
```

```
    ]  
  }  
}  
]  
}  
}  
}
```

See Also

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

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

CreateStream

Creates a Kinesis data stream. A stream captures and transports data records that are continuously emitted from different data sources or *producers*. Scale-out within a stream is explicitly supported by means of shards, which are uniquely identified groups of data records in a stream.

You can create your data stream using either on-demand or provisioned capacity mode. Data streams with an on-demand mode require no capacity planning and automatically scale to handle gigabytes of write and read throughput per minute. With the on-demand mode, Kinesis Data Streams automatically manages the shards in order to provide the necessary throughput.

If you'd still like to proactively scale your on-demand data stream's capacity, you can unlock the warm throughput feature for on-demand data streams by enabling `MinimumThroughputBillingCommitment` for your account. Once your account has `MinimumThroughputBillingCommitment` enabled, you can specify the warm throughput in MiB per second that your stream can support in writes.

For the data streams with a provisioned mode, you must specify the number of shards for the data stream. Each shard can support reads up to five transactions per second, up to a maximum data read total of 2 MiB per second. Each shard can support writes up to 1,000 records per second, up to a maximum data write total of 1 MiB per second. If the amount of data input increases or decreases, you can add or remove shards.

The stream name identifies the stream. The name is scoped to the AWS account used by the application. It is also scoped by AWS Region. That is, two streams in two different accounts can have the same name, and two streams in the same account, but in two different Regions, can have the same name.

`CreateStream` is an asynchronous operation. Upon receiving a `CreateStream` request, Kinesis Data Streams immediately returns and sets the stream status to `CREATING`. After the stream is created, Kinesis Data Streams sets the stream status to `ACTIVE`. You should perform read and write operations only on an `ACTIVE` stream.

You receive a `LimitExceededException` when making a `CreateStream` request when you try to do one of the following:

- Have more than five streams in the `CREATING` state at any point in time.
- Create more shards than are authorized for your account.

For the default shard or on-demand throughput limits for an AWS account, see [Amazon Kinesis Data Streams Limits](#) in the *Amazon Kinesis Data Streams Developer Guide*. To increase this limit, [contact AWS Support](#).

You can use [DescribeStreamSummary](#) to check the stream status, which is returned in `StreamStatus`.

[CreateStream](#) has a limit of five transactions per second per account.

You can add tags to the stream when making a `CreateStream` request by setting the `Tags` parameter. If you pass the `Tags` parameter, in addition to having the `kinesis:CreateStream` permission, you must also have the `kinesis:AddTagsToStream` permission for the stream that will be created. The `kinesis:TagResource` permission won't work to tag streams on creation. Tags will take effect from the `CREATING` status of the stream, but you can't make any updates to the tags until the stream is in `ACTIVE` state.

Request Syntax

```
{
  "MaxRecordSizeInKiB": number,
  "ShardCount": number,
  "StreamModeDetails": {
    "StreamMode": "string"
  },
  "StreamName": "string",
  "Tags": {
    "string" : "string"
  },
  "WarmThroughputMiBps": number
}
```

Request Parameters

The request accepts the following data in JSON format.

[MaxRecordSizeInKiB](#)

The maximum record size of a single record in kibibyte (KiB) that you can write to, and read from a stream.

Type: Integer

Valid Range: Minimum value of 1024. Maximum value of 10240.

Required: No

ShardCount

The number of shards that the stream will use. The throughput of the stream is a function of the number of shards; more shards are required for greater provisioned throughput.

Type: Integer

Valid Range: Minimum value of 1.

Required: No

StreamModeDetails

Indicates the capacity mode of the data stream. Currently, in Kinesis Data Streams, you can choose between an **on-demand** capacity mode and a **provisioned** capacity mode for your data streams.

Type: [StreamModeDetails](#) object

Required: No

StreamName

A name to identify the stream. The stream name is scoped to the AWS account used by the application that creates the stream. It is also scoped by AWS Region. That is, two streams in two different AWS accounts can have the same name. Two streams in the same AWS account but in two different Regions can also have the same name.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

Tags

A set of up to 50 key-value pairs to use to create the tags. A tag consists of a required key and an optional value.

Type: String to string map

Map Entries: Maximum number of 200 items.

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

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

Required: No

WarmThroughputMiBps

The target warm throughput in MB/s that the stream should be scaled to handle. This represents the throughput capacity that will be immediately available for write operations.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

Response Elements

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

Errors

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

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

Examples

To create a stream

The following JSON example creates a stream with three shards.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.CreateStream
{
  "StreamName": "exampleStreamName",
```

```
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
```

See Also

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

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

DecreaseStreamRetentionPeriod

Decreases the Kinesis data stream's retention period, which is the length of time data records are accessible after they are added to the stream. The minimum value of a stream's retention period is 24 hours.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

This operation may result in lost data. For example, if the stream's retention period is 48 hours and is decreased to 24 hours, any data already in the stream that is older than 24 hours is inaccessible.

Request Syntax

```
{
  "RetentionPeriodHours": number,
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

RetentionPeriodHours

The new retention period of the stream, in hours. Must be less than the current retention period.

Type: Integer

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream to modify.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To decrease stream retention period

The following JSON example decreases a stream's retention period.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.DecreaseStreamRetentionPeriod
{
  "RetentionPeriodInHours": "24",
  "StreamName": "examplestream"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
```

See Also

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

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

- [AWS SDK for Ruby V3](#)

DeleteChannel

Deletes the specified channel. Deleting a channel stops delivery from the source stream to the destination. Data already delivered to the destination is not deleted.

A stream cannot be deleted while it has active channels. Use [ListChannels](#) with a stream filter to find the channels attached to a stream before deleting it.

This operation has a call limit of 5 transactions per second (TPS) for each AWS account. Exceeding 5 TPS results in a `LimitExceededException`.

Request Syntax

```
{
  "ChannelARN": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

[ChannelARN](#)

The Amazon Resource Name (ARN) of the channel to delete.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:channel/\S+`

Required: Yes

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

See Also

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

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

DeleteResourcePolicy

Delete a policy for the specified data stream or consumer. Request patterns can be one of the following:

- Data stream pattern: `arn:aws.*:kinesis:.*:\d{12}:.*stream/\S+`
- Consumer pattern: `^(arn):aws.*:kinesis:.*:\d{12}:.*stream\[a-zA-Z0-9_.-]+\[/consumer\[a-zA-Z0-9_.-]+\:[0-9]+`

Request Syntax

```
{  
  "ResourceARN": "string",  
  "StreamId": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

[ResourceARN](#)

The Amazon Resource Name (ARN) of the data stream or consumer.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:.*(stream|channel)/\S+`

Required: Yes

[StreamId](#)

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-zA-Z0-9]{20}-[a-zA-Z0-9]{3}`

Required: No

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

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

DeleteStream

Deletes a Kinesis data stream and all its shards and data. You must shut down any applications that are operating on the stream before you delete the stream. If an application attempts to operate on a deleted stream, it receives the exception `ResourceNotFoundException`.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

If the stream is in the `ACTIVE` state, you can delete it. After a `DeleteStream` request, the specified stream is in the `DELETING` state until Kinesis Data Streams completes the deletion.

Note: Kinesis Data Streams might continue to accept data read and write operations, such as [PutRecord](#), [PutRecords](#), and [GetRecords](#), on a stream in the `DELETING` state until the stream deletion is complete.

When you delete a stream, any shards in that stream are also deleted, and any tags are dissociated from the stream.

You can use the [DescribeStreamSummary](#) operation to check the state of the stream, which is returned in `StreamStatus`.

[DeleteStream](#) has a limit of five transactions per second per account.

Request Syntax

```
{
  "EnforceConsumerDeletion": boolean,
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

EnforceConsumerDeletion

If this parameter is unset (`null`) or if you set it to `false`, and the stream has registered consumers, the call to `DeleteStream` fails with a `ResourceInUseException`.

Type: Boolean

Required: No

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream to delete.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To delete a stream

The following JSON example deletes the specified stream.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.DeleteStream
{
  "StreamName": "exampleStreamName"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
```

See Also

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

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

DeregisterStreamConsumer

To deregister a consumer, provide its ARN. Alternatively, you can provide the ARN of the data stream and the name you gave the consumer when you registered it. You may also provide all three parameters, as long as they don't conflict with each other. If you don't know the name or ARN of the consumer that you want to deregister, you can use the [ListStreamConsumers](#) operation to get a list of the descriptions of all the consumers that are currently registered with a given data stream. The description of a consumer contains its name and ARN.

This operation has a limit of five transactions per second per stream.

Request Syntax

```
{
  "ConsumerARN": "string",
  "ConsumerName": "string",
  "StreamARN": "string",
  "StreamId": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

ConsumerARN

The ARN returned by Kinesis Data Streams when you registered the consumer. If you don't know the ARN of the consumer that you want to deregister, you can use the `ListStreamConsumers` operation to get a list of the descriptions of all the consumers that are currently registered with a given data stream. The description of a consumer contains its ARN.

Type: String

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

Pattern: `^(arn):aws.*:kinesis:.*:\d{12}:.*stream\[a-zA-Z0-9_.-]+\[/consumer\[a-zA-Z0-9_.-]+\[:[0-9]+\]`

Required: No

ConsumerName

The name that you gave to the consumer.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: No

StreamARN

The ARN of the Kinesis data stream that the consumer is registered with. For more information, see [Amazon Resource Names \(ARNs\) and AWS Service Namespaces](#).

Type: String

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

Pattern: arn:aws.*:kinesis:.*:\d{12}:stream/\S+

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: [a-z0-9]{20}-[a-z0-9]{3}

Required: No

Response Elements

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

Errors

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

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

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

DescribeAccountSettings

Describes the account-level settings for Amazon Kinesis Data Streams. This operation returns information about the minimum throughput billing commitments and other account-level configurations.

This API has a call limit of 5 transactions per second (TPS) for each AWS account. TPS over 5 will initiate the `LimitExceededException`.

Response Syntax

```
{
  "MinimumThroughputBillingCommitment": {
    "EarliestAllowedEndAt": number,
    "EndedAt": number,
    "StartedAt": number,
    "Status": "string"
  }
}
```

Response Elements

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

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

[MinimumThroughputBillingCommitment](#)

The current configuration of the minimum throughput billing commitment for your AWS account.

Type: [MinimumThroughputBillingCommitmentOutput](#) object

Errors

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

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

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

DescribeChannel

Describes the specified channel, including its configuration and current status.

Use this operation to verify that a channel reached the ACTIVE state after creation, or to diagnose a channel in the FAILED state by reading the ChannelStatusReason.

This operation has a call limit of 5 transactions per second (TPS) for each AWS account. Exceeding 5 TPS results in a LimitExceededException.

Request Syntax

```
{
  "ChannelARN": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

ChannelARN

The Amazon Resource Name (ARN) of the channel to describe.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:channel/\S+`

Required: Yes

Response Syntax

```
{
  "ChannelDescription": {
    "ChannelARN": "string",
    "ChannelCreationTimestamp": number,
    "ChannelId": "string",
    "ChannelName": "string",
    "ChannelStatus": "string",
    "ChannelStatusReason": "string",
  }
}
```

```

    "EncryptionConfiguration": {
      "EncryptionType": "string",
      "KeyId": "string"
    },
    "LoggingConfiguration": {
      "CloudWatchLogs": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
      }
    },
    "S3DestinationConfiguration": {
      "DataFreshnessInSeconds": number,
      "DeadLetterQueueS3Configuration": {
        "BucketARN": "string",
        "ErrorOutputPrefix": "string",
        "ExpectedBucketOwner": "string"
      },
      "StorageConfiguration": {
        "BucketARN": "string",
        "CompressionType": "string",
        "ExpectedBucketOwner": "string",
        "OutputKeyTemplate": "string",
        "StorageClass": "string"
      }
    },
    "S3TablesDestinationConfiguration": {
      "DataFreshnessInSeconds": number,
      "DeadLetterQueueS3Configuration": {
        "BucketARN": "string",
        "ErrorOutputPrefix": "string",
        "ExpectedBucketOwner": "string"
      },
      "S3TablesConfigurationList": [
        {
          "CompressionType": "string",
          "Namespace": "string",
          "PartitionSpec": {
            "PartitionFields": [
              {
                "SourceName": "string",
                "Transform": "string"
              }
            ]
          }
        }
      ]
    }
  }

```

```
    },
    "TableBucketARN": "string",
    "TableName": "string"
  }
],
},
"ServiceExecutionRoleARN": "string",
"StreamConfigurationList": [
  {
    "RecordConfiguration": {
      "GSRSchemaARN": "string",
      "RecordFormatType": "string"
    },
    "StreamARN": "string",
    "StreamCreationTimestamp": number
  }
]
}
```

Response Elements

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

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

ChannelDescription

The configuration and current status of the channel, including its ARN, source stream, destination configuration, and lifecycle state.

Type: [ChannelDescription](#) object

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

See Also

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

- [AWS Command Line Interface V2](#)

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

DescribeLimits

Describes the shard limits and usage for the account.

If you update your account limits, the old limits might be returned for a few minutes.

This operation has a limit of one transaction per second per account.

Response Syntax

```
{
  "ChannelCount": number,
  "ChannelCountLimit": number,
  "OnDemandStreamCount": number,
  "OnDemandStreamCountLimit": number,
  "OpenShardCount": number,
  "ShardLimit": number
}
```

Response Elements

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

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

ChannelCount

The number of channels in the account.

Type: Integer

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

ChannelCountLimit

The maximum number of channels allowed in the account.

Type: Integer

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

OnDemandStreamCount

Indicates the number of data streams with the on-demand capacity mode.

Type: Integer

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

OnDemandStreamCountLimit

The maximum number of data streams with the on-demand capacity mode.

Type: Integer

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

OpenShardCount

The number of open shards.

Type: Integer

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

ShardLimit

The maximum number of shards.

Type: Integer

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

Errors

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

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To display the shard limits for the account

The following example displays the shard limits for the account.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.DescribeLimits
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "OpenShardCount": 20,
  "ShardLimit": 70
}
```

See Also

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

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

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

DescribeStream

Describes the specified Kinesis data stream.

Note

This API has been revised. It's highly recommended that you use the [DescribeStreamSummary](#) API to get a summarized description of the specified Kinesis data stream and the [ListShards](#) API to list the shards in a specified data stream and obtain information about each shard.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

The information returned includes the stream name, Amazon Resource Name (ARN), creation time, enhanced metric configuration, and shard map. The shard map is an array of shard objects. For each shard object, there is the hash key and sequence number ranges that the shard spans, and the IDs of any earlier shards that played a role in creating the shard. Every record ingested in the stream is identified by a sequence number, which is assigned when the record is put into the stream.

You can limit the number of shards returned by each call. For more information, see [Retrieving Shards from a Stream](#) in the *Amazon Kinesis Data Streams Developer Guide*.

There are no guarantees about the chronological order shards returned. To process shards in chronological order, use the ID of the parent shard to track the lineage to the oldest shard.

This operation has a limit of 10 transactions per second per account.

Request Syntax

```
{  
  "ExclusiveStartShardId": "string",  
}
```

```
"Limit": number,  
"StreamARN": "string",  
"StreamId": "string",  
"StreamName": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

ExclusiveStartShardId

The shard ID of the shard to start with.

Specify this parameter to indicate that you want to describe the stream starting with the shard whose ID immediately follows `ExclusiveStartShardId`.

If you don't specify this parameter, the default behavior for `DescribeStream` is to describe the stream starting with the first shard in the stream.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Limit

The maximum number of shards to return in a single call. The default value is 100. If you specify a value greater than 100, at most 100 results are returned.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10000.

Required: No

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream to describe.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Syntax

```
{
  "StreamDescription": {
    "EncryptionType": "string",
    "EnhancedMonitoring": [
      {
        "ShardLevelMetrics": [ "string" ]
      }
    ],
    "HasMoreShards": boolean,
    "KeyId": "string",
```

```
"RetentionPeriodHours": number,
"Shards": [
  {
    "AdjacentParentShardId": "string",
    "HashKeyRange": {
      "EndingHashKey": "string",
      "StartingHashKey": "string"
    },
    "ParentShardId": "string",
    "SequenceNumberRange": {
      "EndingSequenceNumber": "string",
      "StartingSequenceNumber": "string"
    },
    "ShardId": "string"
  },
  ...
],
"StreamARN": "string",
"StreamCreationTimestamp": number,
"StreamModeDetails": {
  "StreamMode": "string"
},
"StreamName": "string",
"StreamStatus": "string"
}
```

Response Elements

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

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

StreamDescription

The current status of the stream, the stream Amazon Resource Name (ARN), an array of shard objects that comprise the stream, and whether there are more shards available.

Type: [StreamDescription](#) object

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To get information about a stream

The following JSON example describes the specified stream.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.DescribeStream
{
  "StreamName": "exampleStreamName"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "StreamDescription": {
    "EnhancedMonitoring": [
      {
        "ShardLevelMetrics": [
          "IncomingBytes",
          "OutgoingRecords"
        ]
      }
    ],
    "HasMoreShards": false,
    "RetentionPeriodHours": 24,
    "StreamCreationTimestamp": 1.468346745E9,
    "Shards": [
      {
        "HashKeyRange": {
          "EndingHashKey": "113427455640312821154458202477256070484",
          "StartingHashKey": "0"
        },
        "SequenceNumberRange": {
          "EndingSequenceNumber": "21269319989741826081360214168359141376",

```

```
    "StartingSequenceNumber": "21267647932558653966460912964485513216"
  },
  "ShardId": "shardId-000000000000"
},
{
  "HashKeyRange": {
    "EndingHashKey": "226854911280625642308916404954512140969",
    "StartingHashKey": "113427455640312821154458202477256070485"
  },
  "SequenceNumberRange": {
    "StartingSequenceNumber": "21267647932558653966460912964485513217"
  },
  "ShardId": "shardId-000000000001"
},
{
  "HashKeyRange": {
    "EndingHashKey": "340282366920938463463374607431768211455",
    "StartingHashKey": "226854911280625642308916404954512140970"
  },
  "SequenceNumberRange": {
    "StartingSequenceNumber": "21267647932558653966460912964485513218"
  },
  "ShardId": "shardId-000000000002"
}
],
"StreamARN": "arn:aws:kinesis:us-east-1:111122223333:exampleStreamName",
"StreamName": "exampleStreamName",
"StreamStatus": "ACTIVE"
}
}
```

See Also

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

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

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

DescribeStreamConsumer

To get the description of a registered consumer, provide the ARN of the consumer. Alternatively, you can provide the ARN of the data stream and the name you gave the consumer when you registered it. You may also provide all three parameters, as long as they don't conflict with each other. If you don't know the name or ARN of the consumer that you want to describe, you can use the [ListStreamConsumers](#) operation to get a list of the descriptions of all the consumers that are currently registered with a given data stream.

This operation has a limit of 20 transactions per second per stream.

Note

When making a cross-account call with `DescribeStreamConsumer`, make sure to provide the ARN of the consumer.

Request Syntax

```
{
  "ConsumerARN": "string",
  "ConsumerName": "string",
  "StreamARN": "string",
  "StreamId": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

ConsumerARN

The ARN returned by Kinesis Data Streams when you registered the consumer.

Type: String

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

Pattern: `^(arn):aws.*:kinesis:.*:\d{12}:.*stream[/][a-zA-Z0-9_.-]+[/]consumer[/][a-zA-Z0-9_.-]+:[0-9]+`

Required: No

ConsumerName

The name that you gave to the consumer.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

StreamARN

The ARN of the Kinesis data stream that the consumer is registered with. For more information, see [Amazon Resource Names \(ARNs\) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

Response Syntax

```
{
  "ConsumerDescription": {
    "ConsumerARN": "string",
```

```
"ConsumerCreationTimestamp": number,  
"ConsumerName": "string",  
"ConsumerStatus": "string",  
"StreamARN": "string"  
}  
}
```

Response Elements

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

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

ConsumerDescription

An object that represents the details of the consumer.

Type: [ConsumerDescription](#) object

Errors

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

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.
message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

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

DescribeStreamSummary

Provides a summarized description of the specified Kinesis data stream without the shard list.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

The information returned includes the stream name, Amazon Resource Name (ARN), status, record retention period, approximate creation time, monitoring, encryption details, and open shard count.

[DescribeStreamSummary](#) has a limit of 20 transactions per second per account.

Request Syntax

```
{  
  "StreamARN": "string",  
  "StreamId": "string",  
  "StreamName": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

[StreamARN](#)

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream to describe.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Syntax

```
{
  "StreamDescriptionSummary": {
    "ChannelCount": number,
    "ConsumerCount": number,
    "EncryptionType": "string",
    "EnhancedMonitoring": [
      {
        "ShardLevelMetrics": [ "string" ]
      }
    ],
    "KeyId": "string",
    "MaxRecordSizeInKiB": number,
    "OpenShardCount": number,
    "RetentionPeriodHours": number,
    "StreamARN": "string",
    "StreamCreationTimestamp": number,
    "StreamId": "string",
    "StreamModeDetails": {
```

```
    "StreamMode": "string",
  },
  "StreamName": "string",
  "StreamStatus": "string",
  "WarmThroughput": {
    "CurrentMiBps": number,
    "TargetMiBps": number
  }
}
```

Response Elements

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

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

[StreamDescriptionSummary](#)

A [StreamDescriptionSummary](#) containing information about the stream.

Type: [StreamDescriptionSummary](#) object

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

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

DisableEnhancedMonitoring

Disables enhanced monitoring.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

Request Syntax

```
{  
  "ShardLevelMetrics": [ "string" ],  
  "StreamARN": "string",  
  "StreamId": "string",  
  "StreamName": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

ShardLevelMetrics

List of shard-level metrics to disable.

The following are the valid shard-level metrics. The value "ALL" disables every metric.

- IncomingBytes
- IncomingRecords
- OutgoingBytes
- OutgoingRecords
- WriteProvisionedThroughputExceeded
- ReadProvisionedThroughputExceeded
- IteratorAgeMilliseconds
- ALL

For more information, see [Monitoring the Amazon Kinesis Data Streams Service with Amazon CloudWatch](#) in the *Amazon Kinesis Data Streams Developer Guide*.

Type: Array of strings

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

Valid Values: IncomingBytes | IncomingRecords | OutgoingBytes
| OutgoingRecords | WriteProvisionedThroughputExceeded |
ReadProvisionedThroughputExceeded | IteratorAgeMilliseconds | ALL

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the Kinesis data stream for which to disable enhanced monitoring.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: No

Response Syntax

```
{
  "CurrentShardLevelMetrics": [ "string" ],
  "DesiredShardLevelMetrics": [ "string" ],
  "StreamARN": "string",
  "StreamName": "string"
}
```

Response Elements

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

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

[CurrentShardLevelMetrics](#)

Represents the current state of the metrics that are in the enhanced state before the operation.

Type: Array of strings

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

Valid Values: IncomingBytes | IncomingRecords | OutgoingBytes
| OutgoingRecords | WriteProvisionedThroughputExceeded |
ReadProvisionedThroughputExceeded | IteratorAgeMilliseconds | ALL

[DesiredShardLevelMetrics](#)

Represents the list of all the metrics that would be in the enhanced state after the operation.

Type: Array of strings

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

Valid Values: IncomingBytes | IncomingRecords | OutgoingBytes
| OutgoingRecords | WriteProvisionedThroughputExceeded |
ReadProvisionedThroughputExceeded | IteratorAgeMilliseconds | ALL

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

StreamName

The name of the Kinesis data stream.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To disable enhanced monitoring

The following JSON example disables enhanced monitoring for IncomingBytes and OutgoingRecords shard level metrics.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.DisableEnhancedMonitoring
```

```
{
  "ShardLevelMetrics": [
    "IncomingBytes", "OutgoingRecords"
  ],
  "StreamName": "exampleStreamName"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "StreamName": "exampleStreamName",
  "CurrentShardLevelMetrics": [
    "IncomingBytes",
    "OutgoingRecords"
  ],
  "DesiredShardLevelMetrics": []
}
```

See Also

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

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

EnableEnhancedMonitoring

Enables enhanced Kinesis data stream monitoring for shard-level metrics.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

Request Syntax

```
{
  "ShardLevelMetrics": [ "string" ],
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

ShardLevelMetrics

List of shard-level metrics to enable.

The following are the valid shard-level metrics. The value "ALL" enables every metric.

- IncomingBytes
- IncomingRecords
- OutgoingBytes
- OutgoingRecords
- WriteProvisionedThroughputExceeded
- ReadProvisionedThroughputExceeded
- IteratorAgeMilliseconds
- ALL

For more information, see [Monitoring the Amazon Kinesis Data Streams Service with Amazon CloudWatch](#) in the *Amazon Kinesis Data Streams Developer Guide*.

Type: Array of strings

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

Valid Values: IncomingBytes | IncomingRecords | OutgoingBytes
| OutgoingRecords | WriteProvisionedThroughputExceeded |
ReadProvisionedThroughputExceeded | IteratorAgeMilliseconds | ALL

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream for which to enable enhanced monitoring.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: No

Response Syntax

```
{  
  "CurrentShardLevelMetrics": [ "string" ],  
  "DesiredShardLevelMetrics": [ "string" ],  
  "StreamARN": "string",  
  "StreamName": "string"  
}
```

Response Elements

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

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

[CurrentShardLevelMetrics](#)

Represents the current state of the metrics that are in the enhanced state before the operation.

Type: Array of strings

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

Valid Values: IncomingBytes | IncomingRecords | OutgoingBytes
| OutgoingRecords | WriteProvisionedThroughputExceeded |
ReadProvisionedThroughputExceeded | IteratorAgeMilliseconds | ALL

[DesiredShardLevelMetrics](#)

Represents the list of all the metrics that would be in the enhanced state after the operation.

Type: Array of strings

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

Valid Values: IncomingBytes | IncomingRecords | OutgoingBytes
| OutgoingRecords | WriteProvisionedThroughputExceeded |
ReadProvisionedThroughputExceeded | IteratorAgeMilliseconds | ALL

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

StreamName

The name of the Kinesis data stream.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To enable enhanced monitoring

The following JSON example enables enhanced monitoring for IncomingBytes and OutgoingRecords shard level metrics.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.EnableEnhancedMonitoring
```

```
{
  "ShardLevelMetrics": [
    "IncomingBytes", "OutgoingRecords"
  ],
  "StreamName": "exampleStreamName"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "StreamName": "exampleStreamName",
  "CurrentShardLevelMetrics": [],
  "DesiredShardLevelMetrics": [
    "IncomingBytes",
    "OutgoingRecords"
  ]
}
```

See Also

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

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

GetRecords

Gets data records from a Kinesis data stream's shard.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

Specify a shard iterator using the `ShardIterator` parameter. The shard iterator specifies the position in the shard from which you want to start reading data records sequentially. If there are no records available in the portion of the shard that the iterator points to, [GetRecords](#) returns an empty list. It might take multiple calls to get to a portion of the shard that contains records.

You can scale by provisioning multiple shards per stream while considering service limits (for more information, see [Amazon Kinesis Data Streams Limits](#) in the *Amazon Kinesis Data Streams Developer Guide*). Your application should have one thread per shard, each reading continuously from its stream. To read from a stream continually, call [GetRecords](#) in a loop. Use [GetShardIterator](#) to get the shard iterator to specify in the first [GetRecords](#) call. [GetRecords](#) returns a new shard iterator in `NextShardIterator`. Specify the shard iterator returned in `NextShardIterator` in subsequent calls to [GetRecords](#). If the shard has been closed, the shard iterator can't return more data and [GetRecords](#) returns `null` in `NextShardIterator`. You can terminate the loop when the shard is closed, or when the shard iterator reaches the record with the sequence number or other attribute that marks it as the last record to process.

Each data record can be up to 1 MiB in size by default. Amazon Kinesis Data Streams supports large records up to 10 MiB in size, but the average throughput for your stream cannot exceed 1 MiB per second. For more information about how large records are handled, see [Large records](#). Each shard can read up to 2 MiB per second. You can ensure that your calls don't exceed the maximum supported size or throughput by using the `Limit` parameter to specify the maximum number of records that [GetRecords](#) can return. Consider your average record size when determining this limit. The maximum number of records that can be returned per call is 10,000.

The size of the data returned by [GetRecords](#) varies depending on the utilization of the shard. It is recommended that consumer applications retrieve records via the `GetRecords` command using the 5 TPS limit to remain caught up. Retrieving records less frequently can

lead to consumer applications falling behind. The maximum size of data that [GetRecords](#) can return is 10 MiB. If a call returns this amount of data, subsequent calls made within the next 5 seconds throw `ProvisionedThroughputExceededException`. If there is insufficient provisioned throughput on the stream, subsequent calls made within the next 1 second throw `ProvisionedThroughputExceededException`. [GetRecords](#) doesn't return any data when it throws an exception. For this reason, we recommend that you wait 1 second between calls to [GetRecords](#). However, it's possible that the application will get exceptions for longer than 1 second.

To detect whether the application is falling behind in processing, you can use the `MillisBehindLatest` response attribute. You can also monitor the stream using CloudWatch metrics and other mechanisms (see [Monitoring](#) in the *Amazon Kinesis Data Streams Developer Guide*).

Each Amazon Kinesis record includes a value, `ApproximateArrivalTimestamp`, that is set when a stream successfully receives and stores a record. This is commonly referred to as a server-side time stamp, whereas a client-side time stamp is set when a data producer creates or sends the record to a stream (a data producer is any data source putting data records into a stream, for example with [PutRecords](#)). The time stamp has millisecond precision. There are no guarantees about the time stamp accuracy, or that the time stamp is always increasing. For example, records in a shard or across a stream might have time stamps that are out of order.

This operation has a limit of five transactions per second per shard.

Request Syntax

```
{
  "DryRun": boolean,
  "Limit": number,
  "ShardIterator": "string",
  "StreamARN": "string",
  "StreamId": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

[DryRun](#)

Checks if your request will succeed. `DryRun` is an optional parameter.

Type: Boolean

Required: No

Limit

The maximum number of records to return. Specify a value of up to 10,000. If you specify a value that is greater than 10,000, [GetRecords](#) throws `InvalidArgumentException`. The default value is 10,000.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10000.

Required: No

ShardIterator

The position in the shard from which you want to start sequentially reading data records. A shard iterator specifies this position using the sequence number of a data record in the shard.

Type: String

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

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

Response Syntax

```
{
  "ChildShards": [
    {
      "HashKeyRange": {
        "EndingHashKey": "string",
        "StartingHashKey": "string"
      },
      "ParentShards": [ "string" ],
      "ShardId": "string"
    }
  ],
  "MillisBehindLatest": number,
  "NextShardIterator": "string",
  "Records": [
    {
      "ApproximateArrivalTimestamp": number,
      "Data": blob,
      "EncryptionType": "string",
      "PartitionKey": "string",
      "SequenceNumber": "string"
    }
  ]
}
```

Response Elements

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

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

ChildShards

The list of the current shard's child shards, returned in the `GetRecords` API's response only when the end of the current shard is reached.

Type: Array of [ChildShard](#) objects

MillisBehindLatest

The number of milliseconds the [GetRecords](#) response is from the tip of the stream, indicating how far behind current time the consumer is. A value of zero indicates that record processing is caught up, and there are no new records to process at this moment.

Type: Long

Valid Range: Minimum value of 0.

NextShardIterator

The next position in the shard from which to start sequentially reading data records. If set to `null`, the shard has been closed and the requested iterator does not return any more data.

Type: String

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

Records

The data records retrieved from the shard.

Type: Array of [Record](#) objects

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

DryRunOperationException

The request was rejected because the DryRun parameter was specified.

HTTP Status Code: 400

ExpiredIteratorException

The provided iterator exceeds the maximum age allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

InternalFailureException

The processing of the request failed because of an unknown error, exception, or failure.

HTTP Status Code: 500

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSAccessDeniedException

The ciphertext references a key that doesn't exist or that you don't have access to.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSDisabledException

The request was rejected because the specified customer master key (CMK) isn't enabled.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSInvalidStateException

The request was rejected because the state of the specified resource isn't valid for this request. For more information, see [How Key State Affects Use of a Customer Master Key](#) in the *AWS Key Management Service Developer Guide*.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSNotFoundException

The request was rejected because the specified entity or resource can't be found.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSOptInRequired

The AWS access key ID needs a subscription for the service.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSThrottlingException

The request was denied due to request throttling. For more information about throttling, see [Limits](#) in the *AWS Key Management Service Developer Guide*.

message

A message that provides information about the error.

HTTP Status Code: 400

ProvisionedThroughputExceededException

The request rate for the stream is too high, or the requested data is too large for the available throughput. Reduce the frequency or size of your requests. For more information, see [Streams Limits](#) in the *Amazon Kinesis Data Streams Developer Guide*, and [Error Retries and Exponential Backoff in AWS](#) in the *AWS General Reference*.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To get data from the shards in a stream

The following JSON example gets data from the shards in a stream.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.GetRecords
{
  "ShardIterator": "AAAAAAAAAAETyAYzd665+8e0X7JTsASDM/
Hr2rSwc0X2qz93iuA3udrjTH+ikQvpQk/1ZcMMLzRdAesqwBGPnsthzU0/CB1M/
U8/8oEqGwX3pKw0XyeDNRAAZyXBo3MqkQtCpXhr942BRTjvWKhfz70mCb2Ncfr8Tl2cBktooi6kJhr
+djN5WYkB38Rr3akRgCl9qaU4dY=",
  "Limit": 25
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
```

```
Date: <Date>
{
  "MillisBehindLatest": 2100,
  "NextShardIterator": "AAAAAAAAAAHsW8zCWf9164uy8Epue6WS3w6wmj4a4USt+CNvMd6uXQ
+HL5vAJMznqqC0DLKsIjuoiTi1BpT6nW0LN2M2D56zM5H8anHm30Gbri9ua+qaGgj+3XTyvbhpERfrezgLBhPB/
rIcVpykJbaSj5tmcXYRmFnqZBEyHwtZYFmh6hvWVFkIwLuMZLMrpWhG5r5hzkE=",
  "Records": [
    {
      "Data": "XzxkYXRhPl8w",
      "PartitionKey": "partitionKey",
      "ApproximateArrivalTimestamp": 1.441215410867E9,
      "SequenceNumber": "21269319989652663814458848515492872193"
    }
  ]
}
```

See Also

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

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

GetResourcePolicy

Returns a policy attached to the specified data stream or consumer. Request patterns can be one of the following:

- Data stream pattern: `arn:aws.*:kinesis:.*:\d{12}:.*stream/\S+`
- Consumer pattern: `^(arn):aws.*:kinesis:.*:\d{12}:.*stream\[a-zA-Z0-9_.-]+\consumer\[a-zA-Z0-9_.-]+:[0-9]+`

Request Syntax

```
{  
  "ResourceARN": "string",  
  "StreamId": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

ResourceARN

The Amazon Resource Name (ARN) of the data stream or consumer.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:.*(stream|channel)/\S+`

Required: Yes

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

Response Syntax

```
{  
  "Policy": "string"  
}
```

Response Elements

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

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

[Policy](#)

Details of the resource policy. This is formatted as a JSON string.

Type: String

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

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

GetShardIterator

Gets an Amazon Kinesis shard iterator. A shard iterator expires 5 minutes after it is returned to the requester.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

A shard iterator specifies the shard position from which to start reading data records sequentially. The position is specified using the sequence number of a data record in a shard. A sequence number is the identifier associated with every record ingested in the stream, and is assigned when a record is put into the stream. Each stream has one or more shards.

You must specify the shard iterator type. For example, you can set the `ShardIteratorType` parameter to read exactly from the position denoted by a specific sequence number by using the `AT_SEQUENCE_NUMBER` shard iterator type. Alternatively, the parameter can read right after the sequence number by using the `AFTER_SEQUENCE_NUMBER` shard iterator type, using sequence numbers returned by earlier calls to [PutRecord](#), [PutRecords](#), [GetRecords](#), or [DescribeStream](#). In the request, you can specify the shard iterator type `AT_TIMESTAMP` to read records from an arbitrary point in time, `TRIM_HORIZON` to cause `ShardIterator` to point to the last untrimmed record in the shard in the system (the oldest data record in the shard), or `LATEST` so that you always read the most recent data in the shard.

When you read repeatedly from a stream, use a [GetShardIterator](#) request to get the first shard iterator for use in your first [GetRecords](#) request and for subsequent reads use the shard iterator returned by the [GetRecords](#) request in `NextShardIterator`. A new shard iterator is returned by every [GetRecords](#) request in `NextShardIterator`, which you use in the `ShardIterator` parameter of the next [GetRecords](#) request.

If a [GetShardIterator](#) request is made too often, you receive a `ProvisionedThroughputExceededException`. For more information about throughput limits, see [GetRecords](#), and [Streams Limits](#) in the *Amazon Kinesis Data Streams Developer Guide*.

If the shard is closed, [GetShardIterator](#) returns a valid iterator for the last sequence number of the shard. A shard can be closed as a result of using [SplitShard](#) or [MergeShards](#).

[GetShardIterator](#) has a limit of five transactions per second per account per open shard.

Request Syntax

```
{  
  "DryRun": boolean,  
  "ShardId": "string",  
  "ShardIteratorType": "string",  
  "StartingSequenceNumber": "string",  
  "StreamARN": "string",  
  "StreamId": "string",  
  "StreamName": "string",  
  "Timestamp": number  
}
```

Request Parameters

The request accepts the following data in JSON format.

[DryRun](#)

Checks if your request will succeed. DryRun is an optional parameter.

Type: Boolean

Required: No

[ShardId](#)

The shard ID of the Kinesis Data Streams shard to get the iterator for.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

[ShardIteratorType](#)

Determines how the shard iterator is used to start reading data records from the shard.

The following are the valid Amazon Kinesis shard iterator types:

- **AT_SEQUENCE_NUMBER** - Start reading from the position denoted by a specific sequence number, provided in the value `StartingSequenceNumber`.
- **AFTER_SEQUENCE_NUMBER** - Start reading right after the position denoted by a specific sequence number, provided in the value `StartingSequenceNumber`.
- **AT_TIMESTAMP** - Start reading from the position denoted by a specific time stamp, provided in the value `Timestamp`.
- **TRIM_HORIZON** - Start reading at the last untrimmed record in the shard in the system, which is the oldest data record in the shard.
- **LATEST** - Start reading just after the most recent record in the shard, so that you always read the most recent data in the shard.

Type: String

Valid Values: `AT_SEQUENCE_NUMBER` | `AFTER_SEQUENCE_NUMBER` | `TRIM_HORIZON` | `LATEST` | `AT_TIMESTAMP`

Required: Yes

StartingSequenceNumber

The sequence number of the data record in the shard from which to start reading. Used with shard iterator type `AT_SEQUENCE_NUMBER` and `AFTER_SEQUENCE_NUMBER`.

Type: String

Pattern: `^(0|([1-9]\d{0,128}))$`

Required: No

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the Amazon Kinesis data stream.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Timestamp

The time stamp of the data record from which to start reading. Used with shard iterator type `AT_TIMESTAMP`. A time stamp is the Unix epoch date with precision in milliseconds. For example, `2016-04-04T19:58:46.480-00:00` or `1459799926.480`. If a record with this exact time stamp does not exist, the iterator returned is for the next (later) record. If the time stamp is older than the current trim horizon, the iterator returned is for the oldest untrimmed data record (`TRIM_HORIZON`).

Type: Timestamp

Required: No

Response Syntax

```
{
  "ShardIterator": "string"
}
```

Response Elements

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

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

ShardIterator

The position in the shard from which to start reading data records sequentially. A shard iterator specifies this position using the sequence number of a data record in a shard.

Type: String

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

DryRunOperationException

The request was rejected because the DryRun parameter was specified.

HTTP Status Code: 400

InternalFailureException

The processing of the request failed because of an unknown error, exception, or failure.

HTTP Status Code: 500

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

ProvisionedThroughputExceededException

The request rate for the stream is too high, or the requested data is too large for the available throughput. Reduce the frequency or size of your requests. For more information, see [Streams Limits](#) in the *Amazon Kinesis Data Streams Developer Guide*, and [Error Retries and Exponential Backoff in AWS](#) in the *AWS General Reference*.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To get a shard iterator

The following JSON example gets the specified shard iterator.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.GetShardIterator
{
  "StreamName": "exampleStreamName",
```

```
"ShardId": "shardId-000000000001",  
"ShardIteratorType": "LATEST"  
}
```

Sample Response

```
HTTP/1.1 200 OK  
x-amzn-RequestId: <RequestId>  
Content-Type: application/x-amz-json-1.1  
Content-Length: <PayloadSizeBytes>  
Date: <Date>  
{  
  "ShardIterator": "AAAAAAAAAAETyYAYzd665+8e0X7JTASDM/  
Hr2rSwc0X2qz93iuA3udrjTH+ikQvpQk/1ZcMMLzRdAesqwBGPnsthzU0/CB1M/  
U8/8oEqGwX3pKw0XyeDNRAAZyXBo3MqkQtCpXhr942BRTjvWKhFz70mCb2Ncfr8Tl2cBktooi6kJhr  
+djN5WYkB38Rr3akRgCl9qaU4dY="  
}
```

See Also

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

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

IncreaseStreamRetentionPeriod

Increases the Kinesis data stream's retention period, which is the length of time data records are accessible after they are added to the stream. The maximum value of a stream's retention period is 8760 hours (365 days).

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

If you choose a longer stream retention period, this operation increases the time period during which records that have not yet expired are accessible. However, it does not make previous, expired data (older than the stream's previous retention period) accessible after the operation has been called. For example, if a stream's retention period is set to 24 hours and is increased to 168 hours, any data that is older than 24 hours remains inaccessible to consumer applications.

Request Syntax

```
{
  "RetentionPeriodHours": number,
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

RetentionPeriodHours

The new retention period of the stream, in hours. Must be more than the current retention period.

Type: Integer

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream to modify.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To increase stream retention period

The following JSON example increases a stream's retention period.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.IncreaseStreamRetentionPeriod
{
  "RetentionPeriodInHours": "96",
  "StreamName": "examplestream"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
```

See Also

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

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

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

ListChannels

Lists the channels in your account. You can filter the results by source stream. The results are paginated. Use the `NextToken` value returned in the response to retrieve additional results.

Use this operation to find channels before deleting a stream, or to audit the channels configured in an AWS Region.

This operation has a call limit of 5 transactions per second (TPS) for each AWS account. Exceeding 5 TPS results in a `LimitExceededException`.

Request Syntax

```
{
  "MaxResults": number,
  "NextToken": "string",
  "StreamFilter": [
    {
      "StreamARN": "string",
      "StreamCreationTimestamp": number
    }
  ]
}
```

Request Parameters

The request accepts the following data in JSON format.

MaxResults

The maximum number of channels to return in a single call. The default value is 100. If you specify a value greater than 100, at most 100 results are returned.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10000.

Required: No

NextToken

The pagination token returned by a previous call. Specify this token to retrieve the next page of results.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1048576.

Required: No

StreamFilter

Filters the results to channels associated with the specified streams.

Type: Array of [StreamFilter](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10000 items.

Required: No

Response Syntax

```
{
  "ChannelSummaries": [
    {
      "ChannelARN": "string",
      "ChannelCreationTimestamp": number,
      "ChannelDestinationType": "string",
      "ChannelId": "string",
      "ChannelName": "string",
      "ChannelStatus": "string",
      "ChannelStatusReason": "string",
      "Streams": [
        {
          "StreamARN": "string",
          "StreamCreationTimestamp": number
        }
      ]
    }
  ],
  "NextToken": "string"
}
```

Response Elements

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

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

ChannelSummaries

A list of channel summaries.

Type: Array of [ChannelSummary](#) objects

NextToken

The pagination token to use in a subsequent call to retrieve the next page of results. This value is null when there are no more results to return.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1048576.

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

ExpiredNextTokenException

The pagination token passed to the operation is expired.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

See Also

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

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

ListShards

Lists the shards in a stream and provides information about each shard. This operation has a limit of 1000 transactions per second per data stream.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

This action does not list expired shards. For information about expired shards, see [Data Routing, Data Persistence, and Shard State after a Reshard](#).

Important

This API is a new operation that is used by the Amazon Kinesis Client Library (KCL). If you have a fine-grained IAM policy that only allows specific operations, you must update your policy to allow calls to this API. For more information, see [Controlling Access to Amazon Kinesis Data Streams Resources Using IAM](#).

Request Syntax

```
{
  "ExclusiveStartShardId": "string",
  "MaxResults": number,
  "NextToken": "string",
  "ShardFilter": {
    "ShardId": "string",
    "Timestamp": number,
    "Type": "string"
  },
  "StreamARN": "string",
  "StreamCreationTimestamp": number,
  "StreamId": "string",
  "StreamName": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

ExclusiveStartShardId

Specify this parameter to indicate that you want to list the shards starting with the shard whose ID immediately follows `ExclusiveStartShardId`.

If you don't specify this parameter, the default behavior is for `ListShards` to list the shards starting with the first one in the stream.

You cannot specify this parameter if you specify `NextToken`.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

MaxResults

The maximum number of shards to return in a single call to `ListShards`. The maximum number of shards to return in a single call. The default value is 1000. If you specify a value greater than 1000, at most 1000 results are returned.

When the number of shards to be listed is greater than the value of `MaxResults`, the response contains a `NextToken` value that you can use in a subsequent call to `ListShards` to list the next set of shards.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10000.

Required: No

NextToken

When the number of shards in the data stream is greater than the default value for the `MaxResults` parameter, or if you explicitly specify a value for `MaxResults` that is less than

the number of shards in the data stream, the response includes a pagination token named `NextToken`. You can specify this `NextToken` value in a subsequent call to `ListShards` to list the next set of shards.

Don't specify `StreamName` or `StreamCreationTimestamp` if you specify `NextToken` because the latter unambiguously identifies the stream.

You can optionally specify a value for the `MaxResults` parameter when you specify `NextToken`. If you specify a `MaxResults` value that is less than the number of shards that the operation returns if you don't specify `MaxResults`, the response will contain a new `NextToken` value. You can use the new `NextToken` value in a subsequent call to the `ListShards` operation.

Important

Tokens expire after 300 seconds. When you obtain a value for `NextToken` in the response to a call to `ListShards`, you have 300 seconds to use that value. If you specify an expired token in a call to `ListShards`, you get `ExpiredNextTokenException`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1048576.

Required: No

ShardFilter

Enables you to filter out the response of the `ListShards` API. You can only specify one filter at a time.

If you use the `ShardFilter` parameter when invoking the `ListShards` API, the `Type` is the required property and must be specified. If you specify the `AT_TRIM_HORIZON`, `FROM_TRIM_HORIZON`, or `AT_LATEST` types, you do not need to specify either the `ShardId` or the `Timestamp` optional properties.

If you specify the `AFTER_SHARD_ID` type, you must also provide the value for the optional `ShardId` property. The `ShardId` property is identical in functionality to the `ExclusiveStartShardId` parameter of the `ListShards` API. When `ShardId` property

is specified, the response includes the shards starting with the shard whose ID immediately follows the `ShardId` that you provided.

If you specify the `AT_TIMESTAMP` or `FROM_TIMESTAMP_ID` type, you must also provide the value for the optional `Timestamp` property. If you specify the `AT_TIMESTAMP` type, then all shards that were open at the provided timestamp are returned. If you specify the `FROM_TIMESTAMP` type, then all shards starting from the provided timestamp to `TIP` are returned.

Type: [ShardFilter](#) object

Required: No

[StreamARN](#)

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

[StreamCreationTimestamp](#)

Specify this input parameter to distinguish data streams that have the same name. For example, if you create a data stream and then delete it, and you later create another data stream with the same name, you can use this input parameter to specify which of the two streams you want to list the shards for.

You cannot specify this parameter if you specify the `NextToken` parameter.

Type: Timestamp

Required: No

[StreamId](#)

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the data stream whose shards you want to list.

You cannot specify this parameter if you specify the `NextToken` parameter.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Syntax

```
{
  "NextToken": "string",
  "Shards": [
    {
      "AdjacentParentShardId": "string",
      "HashKeyRange": {
        "EndingHashKey": "string",
        "StartingHashKey": "string"
      },
      "ParentShardId": "string",
      "SequenceNumberRange": {
        "EndingSequenceNumber": "string",
        "StartingSequenceNumber": "string"
      },
      "ShardId": "string"
    }
  ]
}
```

Response Elements

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

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

NextToken

When the number of shards in the data stream is greater than the default value for the `MaxResults` parameter, or if you explicitly specify a value for `MaxResults` that is less than the number of shards in the data stream, the response includes a pagination token named `NextToken`. You can specify this `NextToken` value in a subsequent call to `ListShards` to list the next set of shards. For more information about the use of this pagination token when calling the `ListShards` operation, see [ListShards:NextToken](#).

Important

Tokens expire after 300 seconds. When you obtain a value for `NextToken` in the response to a call to `ListShards`, you have 300 seconds to use that value. If you specify an expired token in a call to `ListShards`, you get `ExpiredNextTokenException`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1048576.

Shards

An array of JSON objects. Each object represents one shard and specifies the IDs of the shard, the shard's parent, and the shard that's adjacent to the shard's parent. Each object also contains the starting and ending hash keys and the starting and ending sequence numbers for the shard.

Type: Array of [Shard](#) objects

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

ExpiredNextTokenException

The pagination token passed to the operation is expired.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To list the shards of a stream

The following JSON example lists three of the shards of the stream that is specified using `StreamName`. If the stream has more than three shards, the response includes a pagination token named `NextToken`. You can use `NextToken` in your input to a subsequent call to `ListShards` to get a listing of the next set of shards. When there are no more shards to be listed, the response doesn't contain `NextToken`.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.ListShards
{
  "StreamName": "exampleStreamName",
  "MaxResults": 3
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "NextToken": "AAAAAAAAAAGK9EEG0sJqVhCUS2JsgigQ5dcpB4q9PYswrH2oK44Skbjtm+WR0xA7/hrAFFsohevH1/0yPnbzKBS1byPyCZuVcokYtQe/b1m4c0SCI7jctPT0oUTLRdwSRirKm9dp9YC/EL+kZH0vYAUnztVGs0APEFC3ECf/bVC927bDZBbRRzy/440HfWmrCLcbcWqehRh5D14WnL3yLsumhiHDkyuxSlkBepauvMnNLtT01RtmQ5Q5reoujfq2gzeCS0t",
  "Shards": [
    {
      "ShardId": "shardId-000000000001",
      "HashKeyRange": {
```

```

        "EndingHashKey": "68056473384187692692674921486353642280",
        "StartingHashKey": "34028236692093846346337460743176821145"
    },
    "SequenceNumberRange": {
        "StartingSequenceNumber":
"4957984403772733356165064238440708846556371693205002258"
    }
},
{
    "ShardId": "shardId-000000000002",
    "HashKeyRange": {
        "EndingHashKey": "102084710076281539039012382229530463436",
        "StartingHashKey": "68056473384187692692674921486353642281"
    },
    "SequenceNumberRange": {
        "StartingSequenceNumber":
"49579844037749634101363594861582244564829020124710982690"
    }
},
{
    "ShardId": "shardId-000000000003",
    "HashKeyRange": {
        "EndingHashKey": "136112946768375385385349842972707284581",
        "StartingHashKey": "102084710076281539039012382229530463437"
    },
    "SequenceNumberRange": {
        "StartingSequenceNumber":
"49579844037771934846562125484723780283101668556216963122"
    }
}
]
}

```

See Also

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

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

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

ListStreamConsumers

Lists the consumers registered to receive data from a stream using enhanced fan-out, and provides information about each consumer.

This operation has a limit of 5 transactions per second per stream.

Request Syntax

```
{
  "MaxResults": number,
  "NextToken": "string",
  "StreamARN": "string",
  "StreamCreationTimestamp": number,
  "StreamId": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

MaxResults

The maximum number of consumers that you want a single call of `ListStreamConsumers` to return. The default value is 100. If you specify a value greater than 100, at most 100 results are returned.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10000.

Required: No

NextToken

When the number of consumers that are registered with the data stream is greater than the default value for the `MaxResults` parameter, or if you explicitly specify a value for `MaxResults` that is less than the number of consumers that are registered with the data stream, the response includes a pagination token named `NextToken`. You can specify this `NextToken` value in a subsequent call to `ListStreamConsumers` to list the next set of registered consumers.

Don't specify `StreamName` or `StreamCreationTimestamp` if you specify `NextToken` because the latter unambiguously identifies the stream.

You can optionally specify a value for the `MaxResults` parameter when you specify `NextToken`. If you specify a `MaxResults` value that is less than the number of consumers that the operation returns if you don't specify `MaxResults`, the response will contain a new `NextToken` value. You can use the new `NextToken` value in a subsequent call to the `ListStreamConsumers` operation to list the next set of consumers.

 Important

Tokens expire after 300 seconds. When you obtain a value for `NextToken` in the response to a call to `ListStreamConsumers`, you have 300 seconds to use that value. If you specify an expired token in a call to `ListStreamConsumers`, you get `ExpiredNextTokenException`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1048576.

Required: No

StreamARN

The ARN of the Kinesis data stream for which you want to list the registered consumers. For more information, see [Amazon Resource Names \(ARNs\) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

StreamCreationTimestamp

Specify this input parameter to distinguish data streams that have the same name. For example, if you create a data stream and then delete it, and you later create another data stream with the same name, you can use this input parameter to specify which of the two streams you want to list the consumers for.

You can't specify this parameter if you specify the `NextToken` parameter.

Type: Timestamp

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

Response Syntax

```
{
  "Consumers": [
    {
      "ConsumerARN": "string",
      "ConsumerCreationTimestamp": number,
      "ConsumerName": "string",
      "ConsumerStatus": "string"
    }
  ],
  "NextToken": "string"
}
```

Response Elements

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

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

Consumers

An array of JSON objects. Each object represents one registered consumer.

Type: Array of [Consumer](#) objects

NextToken

When the number of consumers that are registered with the data stream is greater than the default value for the `MaxResults` parameter, or if you explicitly specify a value for `MaxResults` that is less than the number of registered consumers, the response includes a pagination token named `NextToken`. You can specify this `NextToken` value in a subsequent call to `ListStreamConsumers` to list the next set of registered consumers. For more information about the use of this pagination token when calling the `ListStreamConsumers` operation, see [ListStreamConsumers:NextToken](#).

Important

Tokens expire after 300 seconds. When you obtain a value for `NextToken` in the response to a call to `ListStreamConsumers`, you have 300 seconds to use that value. If you specify an expired token in a call to `ListStreamConsumers`, you get `ExpiredNextTokenException`.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1048576.

Errors

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

ExpiredNextTokenException

The pagination token passed to the operation is expired.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

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

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

ListStreams

Lists your Kinesis data streams.

The number of streams may be too large to return from a single call to `ListStreams`. You can limit the number of returned streams using the `Limit` parameter. If you do not specify a value for the `Limit` parameter, Kinesis Data Streams uses the default limit, which is currently 100.

You can detect if there are more streams available to list by using the `HasMoreStreams` flag from the returned output. If there are more streams available, you can request more streams by using the name of the last stream returned by the `ListStreams` request in the `ExclusiveStartStreamName` parameter in a subsequent request to `ListStreams`. The group of stream names returned by the subsequent request is then added to the list. You can continue this process until all the stream names have been collected in the list.

[ListStreams](#) has a limit of five transactions per second per account.

Request Syntax

```
{  
  "ExclusiveStartStreamName": "string",  
  "Limit": number,  
  "NextToken": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

[ExclusiveStartStreamName](#)

The name of the stream to start the list with.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Limit

The maximum number of streams to list. The default value is 100. If you specify a value greater than 100, at most 100 results are returned.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 10000.

Required: No

NextToken

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1048576.

Required: No

Response Syntax

```
{
  "HasMoreStreams": boolean,
  "NextToken": "string",
  "StreamNames": [ "string" ],
  "StreamSummaries": [
    {
      "StreamARN": "string",
      "StreamCreationTimestamp": number,
      "StreamModeDetails": {
        "StreamMode": "string"
      },
      "StreamName": "string",
      "StreamStatus": "string"
    }
  ]
}
```

Response Elements

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

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

HasMoreStreams

If set to `true`, there are more streams available to list.

Type: Boolean

NextToken

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1048576.

StreamNames

The names of the streams that are associated with the AWS account making the `ListStreams` request.

Type: Array of strings

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

Pattern: `[a-zA-Z0-9_.-]+`

StreamSummaries

Type: Array of [StreamSummary](#) objects

Errors

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

ExpiredNextTokenException

The pagination token passed to the operation is expired.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To list your streams

The following JSON example lists your streams, starting with the specified stream.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.ListStreams
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
```

```
"HasMoreStreams": false,
"StreamNames": [
  "exampleStreamName"
]
```

See Also

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

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

ListTagsForResource

List all tags added to the specified Kinesis resource. Each tag is a label consisting of a user-defined key and value. Tags can help you manage, identify, organize, search for, and filter resources.

For more information about tagging Kinesis resources, see [Tag your Amazon Kinesis Data Streams resources](#).

Request Syntax

```
{
  "ResourceARN": "string",
  "StreamId": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

ResourceARN

The Amazon Resource Name (ARN) of the Kinesis resource for which to list tags.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:.*(stream|channel)/\S+`

Required: Yes

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

Response Syntax

```
{
  "Tags": [
    {
      "Key": "string",
      "Value": "string"
    }
  ]
}
```

Response Elements

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

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

Tags

An array of tags associated with the specified Kinesis resource.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To list all tags added to a resource

The following JSON example lists all the tags for a specific consumer.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
```

```
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.ListTagsForResource
{
  "ResourceARN": "arn:aws:kinesis:us-east-1:123456789012:/stream/myStream/consumer/
myConsumer:174353308"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "Tags" : [
    {
      "Key": "Project",
      "Value": "myProject"
    },
    {
      "Key": "Environment",
      "Value": "Production"
    }
  ]
}
```

See Also

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

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

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

ListTagsForStream

Lists the tags for the specified Kinesis data stream. This operation has a limit of five transactions per second per account.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

Request Syntax

```
{
  "ExclusiveStartTagKey": "string",
  "Limit": number,
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

ExclusiveStartTagKey

The key to use as the starting point for the list of tags. If this parameter is set, `ListTagsForStream` gets all tags that occur after `ExclusiveStartTagKey`.

Type: String

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

Required: No

Limit

The number of tags to return. If this number is less than the total number of tags associated with the stream, `HasMoreTags` is set to `true`. To list additional tags, set `ExclusiveStartTagKey` to the last key in the response.

Type: Integer

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

Required: No

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Syntax

```
{  
  "HasMoreTags": boolean,}
```

```
"Tags": [  
  {  
    "Key": "string",  
    "Value": "string"  
  }  
]
```

Response Elements

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

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

HasMoreTags

If set to `true`, more tags are available. To request additional tags, set `ExclusiveStartTagKey` to the key of the last tag returned.

Type: Boolean

Tags

A list of tags associated with `StreamName`, starting with the first tag after `ExclusiveStartTagKey` and up to the specified `Limit`.

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 200 items.

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To list the tags for a stream

The following JSON example lists the tags for the specified stream.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.ListTagsForStream
```

```
{
  "StreamName": "exampleStreamName"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "HasMoreTags": "false",
  "Tags" : [
    {
      "Key": "Project",
      "Value": "myProject"
    },
    {
      "Key": "Environment",
      "Value": "Production"
    }
  ]
}
```

See Also

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

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

- [AWS SDK for Ruby V3](#)

MergeShards

MergeShards merges two adjacent shards in a Kinesis data stream and combines them into a single shard to reduce the stream's capacity to ingest and transport data. This API is only supported for the data streams with the provisioned capacity mode. Two shards are considered adjacent if the union of the hash key ranges for the two shards form a contiguous set with no gaps. For example, if you have two shards, one with a hash key range of 276...381 and the other with a hash key range of 382...454, then you could merge these two shards into a single shard that would have a hash key range of 276...454. After the merge, the single child shard receives data for all hash key values covered by the two parent shards.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

MergeShards is called when there is a need to reduce the overall capacity of a stream because of excess capacity that is not being used. You must specify the shard to be merged and the adjacent shard for a stream. For more information about merging shards, see [Merge Two Shards](#) in the *Amazon Kinesis Data Streams Developer Guide*.

If the stream is in the `ACTIVE` state, you can call MergeShards. If a stream is in the `CREATING`, `UPDATING`, or `DELETING` state, MergeShards returns a `ResourceInUseException`. If the specified stream does not exist, MergeShards returns a `ResourceNotFoundException`.

You can use [DescribeStreamSummary](#) to check the state of the stream, which is returned in `StreamStatus`.

MergeShards is an asynchronous operation. Upon receiving a MergeShards request, Amazon Kinesis Data Streams immediately returns a response and sets the `StreamStatus` to `UPDATING`. After the operation is completed, Kinesis Data Streams sets the `StreamStatus` to `ACTIVE`. Read and write operations continue to work while the stream is in the `UPDATING` state.

You use [DescribeStreamSummary](#) and the [ListShards](#) APIs to determine the shard IDs that are specified in the MergeShards request.

If you try to operate on too many streams in parallel using [CreateStream](#), [DeleteStream](#), MergeShards, or [SplitShard](#), you receive a `LimitExceededException`.

MergeShards has a limit of five transactions per second per account.

Request Syntax

```
{
  "AdjacentShardToMerge": "string",
  "ShardToMerge": "string",
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

[AdjacentShardToMerge](#)

The shard ID of the adjacent shard for the merge.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: Yes

[ShardToMerge](#)

The shard ID of the shard to combine with the adjacent shard for the merge.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream for the merge.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

Examples

To merge two adjacent shards

The following JSON example merges two adjacent shards.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.MergeShards
{
  "StreamName": "exampleStreamName",
  "ShardToMerge": "shardId-0000000000000",
  "AdjacentShardToMerge": "shardId-00000000000001"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
```

See Also

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

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

PutRecord

Writes a single data record into an Amazon Kinesis data stream. Call `PutRecord` to send data into the stream for real-time ingestion and subsequent processing, one record at a time. Each shard can support writes up to 1,000 records per second, up to a maximum data write total of 10 MiB per second.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

You must specify the name of the stream that captures, stores, and transports the data; a partition key; and the data blob itself.

The data blob can be any type of data; for example, a segment from a log file, geographic/location data, website clickstream data, and so on.

The partition key is used by Kinesis Data Streams to distribute data across shards. Kinesis Data Streams segregates the data records that belong to a stream into multiple shards, using the partition key associated with each data record to determine the shard to which a given data record belongs.

Partition keys are Unicode strings, with a maximum length limit of 256 characters for each key. An MD5 hash function is used to map partition keys to 128-bit integer values and to map associated data records to shards using the hash key ranges of the shards. You can override hashing the partition key to determine the shard by explicitly specifying a hash value using the `ExplicitHashKey` parameter. For more information, see [Adding Data to a Stream](#) in the *Amazon Kinesis Data Streams Developer Guide*.

`PutRecord` returns the shard ID of where the data record was placed and the sequence number that was assigned to the data record.

Sequence numbers increase over time and are specific to a shard within a stream, not across all shards within a stream. To guarantee strictly increasing ordering, write serially to a shard and use the `SequenceNumberForOrdering` parameter. For more information, see [Adding Data to a Stream](#) in the *Amazon Kinesis Data Streams Developer Guide*.

Important

After you write a record to a stream, you cannot modify that record or its order within the stream.

If a `PutRecord` request cannot be processed because of insufficient provisioned throughput on the shard involved in the request, `PutRecord` throws `ProvisionedThroughputExceededException`.

By default, data records are accessible for 24 hours from the time that they are added to a stream. You can use [IncreaseStreamRetentionPeriod](#) or [DecreaseStreamRetentionPeriod](#) to modify this retention period.

Request Syntax

```
{  
  "Data": blob,  
  "DryRun": boolean,  
  "ExplicitHashKey": "string",  
  "PartitionKey": "string",  
  "SequenceNumberForOrdering": "string",  
  "StreamARN": "string",  
  "StreamId": "string",  
  "StreamName": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

Data

The data blob to put into the record, which is base64-encoded when the blob is serialized. When the data blob (the payload before base64-encoding) is added to the partition key size, the total size must not exceed the maximum record size (10 MiB).

Type: Base64-encoded binary data object

Length Constraints: Minimum length of 0. Maximum length of 10485760.

Required: Yes

DryRun

Checks if your request will succeed. DryRun is an optional parameter.

Type: Boolean

Required: No

ExplicitHashKey

The hash value used to explicitly determine the shard the data record is assigned to by overriding the partition key hash.

Type: String

Pattern: `^(0|([1-9]\d{0,38}))$`

Required: No

PartitionKey

Determines which shard in the stream the data record is assigned to. Partition keys are Unicode strings with a maximum length limit of 256 characters for each key. Amazon Kinesis Data Streams uses the partition key as input to a hash function that maps the partition key and associated data to a specific shard. Specifically, an MD5 hash function is used to map partition keys to 128-bit integer values and to map associated data records to shards. As a result of this hashing mechanism, all data records with the same partition key map to the same shard within the stream.

Type: String

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

Required: Yes

SequenceNumberForOrdering

Guarantees strictly increasing sequence numbers, for puts from the same client and to the same partition key. Usage: set the SequenceNumberForOrdering of record n to the sequence number of record $n-1$ (as returned in the result when putting record $n-1$). If this parameter is not set, records are coarsely ordered based on arrival time.

Type: String

Pattern: `^(0|([1-9]\d{0,128}))$`

Required: No

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream to put the data record into.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Syntax

```
{
```

```
"EncryptionType": "string",  
"SequenceNumber": "string",  
"ShardId": "string"  
}
```

Response Elements

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

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

EncryptionType

The encryption type to use on the record. This parameter can be one of the following values:

- NONE: Do not encrypt the records in the stream.
- KMS: Use server-side encryption on the records in the stream using a customer-managed AWS KMS key.

Type: String

Valid Values: NONE | KMS

SequenceNumber

The sequence number identifier that was assigned to the put data record. The sequence number for the record is unique across all records in the stream. A sequence number is the identifier associated with every record put into the stream.

Type: String

Pattern: `^(0|([1-9]\d{0,128}))$`

ShardId

The shard ID of the shard where the data record was placed.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

DryRunOperationException

The request was rejected because the DryRun parameter was specified.

HTTP Status Code: 400

InternalFailureException

The processing of the request failed because of an unknown error, exception, or failure.

HTTP Status Code: 500

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSAccessDeniedException

The ciphertext references a key that doesn't exist or that you don't have access to.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSDisabledException

The request was rejected because the specified customer master key (CMK) isn't enabled.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSInvalidStateException

The request was rejected because the state of the specified resource isn't valid for this request. For more information, see [How Key State Affects Use of a Customer Master Key](#) in the *AWS Key Management Service Developer Guide*.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSNotFoundException

The request was rejected because the specified entity or resource can't be found.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSOptInRequired

The AWS access key ID needs a subscription for the service.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSThrottlingException

The request was denied due to request throttling. For more information about throttling, see [Limits](#) in the *AWS Key Management Service Developer Guide*.

message

A message that provides information about the error.

HTTP Status Code: 400

ProvisionedThroughputExceededException

The request rate for the stream is too high, or the requested data is too large for the available throughput. Reduce the frequency or size of your requests. For more information, see [Streams Limits](#) in the *Amazon Kinesis Data Streams Developer Guide*, and [Error Retries and Exponential Backoff in AWS](#) in the *AWS General Reference*.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To add data to a stream

The following JSON example adds data to the specified stream.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.PutRecord
{
  "StreamName": "exampleStreamName",
  "Data": "XzxkYXRhPl8x",
  "PartitionKey": "partitionKey"
```

```
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "SequenceNumber": "21269319989653637946712965403778482177",
  "ShardId": "shardId-0000000000001"
}
```

See Also

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

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

PutRecords

Writes multiple data records into a Kinesis data stream in a single call (also referred to as a `PutRecords` request). Use this operation to send data into the stream for data ingestion and processing.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

Each `PutRecords` request can support up to 500 records. Each record in the request can be as large as 10 MiB, up to a limit of 10 MiB for the entire request, including partition keys. Each shard can support writes up to 1,000 records per second, up to a maximum data write total of 1 MB per second.

You must specify the name of the stream that captures, stores, and transports the data; and an array of request `Records`, with each record in the array requiring a partition key and data blob. The record size limit applies to the total size of the partition key and data blob.

The data blob can be any type of data; for example, a segment from a log file, geographic/location data, website clickstream data, and so on.

The partition key is used by Kinesis Data Streams as input to a hash function that maps the partition key and associated data to a specific shard. An MD5 hash function is used to map partition keys to 128-bit integer values and to map associated data records to shards. As a result of this hashing mechanism, all data records with the same partition key map to the same shard within the stream. For more information, see [Adding Data to a Stream](#) in the *Amazon Kinesis Data Streams Developer Guide*.

Each record in the `Records` array may include an optional parameter, `ExplicitHashKey`, which overrides the partition key to shard mapping. This parameter allows a data producer to determine explicitly the shard where the record is stored. For more information, see [Adding Multiple Records with PutRecords](#) in the *Amazon Kinesis Data Streams Developer Guide*.

The `PutRecords` response includes an array of response `Records`. Each record in the response array directly correlates with a record in the request array using natural ordering, from the top to

the bottom of the request and response. The response `Records` array always includes the same number of records as the request array.

The response `Records` array includes both successfully and unsuccessfully processed records. Kinesis Data Streams attempts to process all records in each `PutRecords` request. A single record failure does not stop the processing of subsequent records. As a result, `PutRecords` doesn't guarantee the ordering of records. If you need to read records in the same order they are written to the stream, use [PutRecord](#) instead of `PutRecords`, and write to the same shard.

A successfully processed record includes `ShardId` and `SequenceNumber` values. The `ShardId` parameter identifies the shard in the stream where the record is stored. The `SequenceNumber` parameter is an identifier assigned to the put record, unique to all records in the stream.

An unsuccessfully processed record includes `ErrorCode` and `ErrorMessage` values. `ErrorCode` reflects the type of error and can be one of the following values: `ProvisionedThroughputExceededException` or `InternalFailure`. `ErrorMessage` provides more detailed information about the `ProvisionedThroughputExceededException` exception including the account ID, stream name, and shard ID of the record that was throttled. For more information about partially successful responses, see [Adding Multiple Records with PutRecords](#) in the *Amazon Kinesis Data Streams Developer Guide*.

Important

After you write a record to a stream, you cannot modify that record or its order within the stream.

By default, data records are accessible for 24 hours from the time that they are added to a stream. You can use [IncreaseStreamRetentionPeriod](#) or [DecreaseStreamRetentionPeriod](#) to modify this retention period.

Request Syntax

```
{
  "DryRun": boolean,
  "Records": [
    {
      "Data": blob,
      "ExplicitHashKey": "string",
      "PartitionKey": "string"
    }
  ]
}
```

```
    }  
  ],  
  "StreamARN": "string",  
  "StreamId": "string",  
  "StreamName": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

DryRun

Checks if your request will succeed. DryRun is an optional parameter.

Type: Boolean

Required: No

Records

The records associated with the request.

Type: Array of [PutRecordsRequestEntry](#) objects

Array Members: Minimum number of 1 item. Maximum number of 500 items.

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The stream name associated with the request.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Syntax

```
{
  "EncryptionType": "string",
  "FailedRecordCount": number,
  "Records": [
    {
      "ErrorCode": "string",
      "ErrorMessage": "string",
      "SequenceNumber": "string",
      "ShardId": "string"
    }
  ]
}
```

Response Elements

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

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

EncryptionType

The encryption type used on the records. This parameter can be one of the following values:

- **NONE:** Do not encrypt the records.
- **KMS:** Use server-side encryption on the records using a customer-managed AWS KMS key.

Type: String

Valid Values: NONE | KMS

FailedRecordCount

The number of unsuccessfully processed records in a PutRecords request.

Type: Integer

Valid Range: Minimum value of 1.

Records

An array of successfully and unsuccessfully processed record results. A record that is successfully added to a stream includes `SequenceNumber` and `ShardId` in the result. A record that fails to be added to a stream includes `ErrorCode` and `ErrorMessage` in the result.

Type: Array of [PutRecordsResultEntry](#) objects

Array Members: Minimum number of 1 item. Maximum number of 500 items.

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

DryRunOperationException

The request was rejected because the `DryRun` parameter was specified.

HTTP Status Code: 400

InternalFailureException

The processing of the request failed because of an unknown error, exception, or failure.

HTTP Status Code: 500

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSAccessDeniedException

The ciphertext references a key that doesn't exist or that you don't have access to.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSDisabledException

The request was rejected because the specified customer master key (CMK) isn't enabled.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSInvalidStateException

The request was rejected because the state of the specified resource isn't valid for this request. For more information, see [How Key State Affects Use of a Customer Master Key](#) in the *AWS Key Management Service Developer Guide*.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSNotFoundException

The request was rejected because the specified entity or resource can't be found.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSOptInRequired

The AWS access key ID needs a subscription for the service.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSThrottlingException

The request was denied due to request throttling. For more information about throttling, see [Limits](#) in the *AWS Key Management Service Developer Guide*.

message

A message that provides information about the error.

HTTP Status Code: 400

ProvisionedThroughputExceededException

The request rate for the stream is too high, or the requested data is too large for the available throughput. Reduce the frequency or size of your requests. For more information, see [Streams Limits](#) in the *Amazon Kinesis Data Streams Developer Guide*, and [Error Retries and Exponential Backoff in AWS](#) in the *AWS General Reference*.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To add data to a stream, with complete success

The following JSON example adds data to the specified stream with a successful response.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.PutRecords
{
  "Records": [
    {
      "Data": "XzxkYXRhPl8x",
      "PartitionKey": "partitionKey1"
    },
    {
      "Data": "f1PxFQo92Afh",
      "PartitionKey": "partitionKey2"
    },
    {
      "Data": "Gi4sEdd08HypA",
      "PartitionKey": "partitionKey3"
    }
  ],
  "StreamName": "exampleStreamName"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
```

```
{
  "FailedRecordCount": 0,
  "Records": [
    {
      "SequenceNumber":
"49543463076548007577105092703039560359975228518395019266",
      "ShardId": "shardId-000000000000"
    },
    {
      "SequenceNumber":
"49543463076570308322303623326179887152428262250726293522",
      "ShardId": "shardId-000000000001"
    },
    {
      "SequenceNumber":
"49543463076570308322303623326179887152428262250726293588",
      "ShardId": "shardId-000000000003"
    }
  ]
}
```

To add data to a stream, with a partially successful response

The following JSON example adds data to the specified stream with a partially successful response and contains failed records.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.PutRecords
{
  "Records": [
    {
      "Data": "XzxkYXRhPl8x",
      "PartitionKey": "partitionKey1"
    },
    {
```

```
        "Data": "f1PxFQo92Afh",
        "PartitionKey": "partitionKey2"
    },
    {
        "Data": "Gi4sEdd08HypA",
        "PartitionKey": "partitionKey3"
    }
],
"StreamName": "exampleStreamName"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
    "FailedRecordCount": 2,
    "Records": [
        {
            "SequenceNumber":
"49543463076548007577105092703039560359975228518395012686",
            "ShardId": "shardId-000000000000"
        },
        {
            "ErrorCode": "ProvisionedThroughputExceededException",
            "ErrorMessage": "Rate exceeded for shard shardId-000000000001 in stream
exampleStreamName under account 111111111111."
        },
        {
            "ErrorCode": "InternalFailure",
            "ErrorMessage": "Internal service failure."
        }
    ]
}
```

See Also

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

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

PutResourcePolicy

Attaches a resource-based policy to a data stream or registered consumer. If you are using an identity other than the root user of the AWS account that owns the resource, the calling identity must have the `PutResourcePolicy` permissions on the specified Kinesis Data Streams resource and belong to the owner's account in order to use this operation. If you don't have `PutResourcePolicy` permissions, Amazon Kinesis Data Streams returns a `403 Access Denied` error. If you receive a `ResourceNotFoundException`, check to see if you passed a valid stream or consumer resource.

Request patterns can be one of the following:

- Data stream pattern: `arn:aws.*:kinesis:.*:\d{12}:.*stream/\S+`
- Consumer pattern: `^(arn):aws.*:kinesis:.*:\d{12}:.*stream\[a-zA-Z0-9_.-]+\[/consumer\[a-zA-Z0-9_.-]+:[0-9]+`

For more information, see [Controlling Access to Amazon Kinesis Data Streams Resources Using IAM](#).

Request Syntax

```
{  
  "Policy": "string",  
  "ResourceARN": "string",  
  "StreamId": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

[Policy](#)

Details of the resource policy. It must include the identity of the principal and the actions allowed on this resource. This is formatted as a JSON string.

Type: String

Required: Yes

ResourceARN

The Amazon Resource Name (ARN) of the data stream or consumer.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:.*(stream|channel)/\S+`

Required: Yes

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

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

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

RegisterStreamConsumer

Registers a consumer with a Kinesis data stream. When you use this operation, the consumer you register can then call [SubscribeToShard](#) to receive data from the stream using enhanced fan-out, at a rate of up to 2 MiB per second for every shard you subscribe to. This rate is unaffected by the total number of consumers that read from the same stream.

You can add tags to the registered consumer when making a RegisterStreamConsumer request by setting the Tags parameter. If you pass the Tags parameter, in addition to having the `kinesis:RegisterStreamConsumer` permission, you must also have the `kinesis:TagResource` permission for the consumer that will be registered. Tags will take effect from the CREATING status of the consumer.

With On-demand Advantage streams, you can register up to 50 consumers per stream to use Enhanced Fan-out. With On-demand Standard and Provisioned streams, you can register up to 20 consumers per stream to use Enhanced Fan-out. A given consumer can only be registered with one stream at a time.

For an example of how to use this operation, see [Enhanced Fan-Out Using the Kinesis Data Streams API](#).

The use of this operation has a limit of five transactions per second per account. Also, only 5 consumers can be created simultaneously. In other words, you cannot have more than 5 consumers in a CREATING status at the same time. Registering a 6th consumer while there are 5 in a CREATING status results in a `LimitExceededException`.

Request Syntax

```
{
  "ConsumerName": "string",
  "StreamARN": "string",
  "StreamId": "string",
  "Tags": {
    "string" : "string"
  }
}
```

Request Parameters

The request accepts the following data in JSON format.

ConsumerName

For a given Kinesis data stream, each consumer must have a unique name. However, consumer names don't have to be unique across data streams.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: Yes

StreamARN

The ARN of the Kinesis data stream that you want to register the consumer with. For more info, see [Amazon Resource Names \(ARNs\) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

Tags

A set of up to 50 key-value pairs. A tag consists of a required key and an optional value.

Type: String to string map

Map Entries: Maximum number of 200 items.

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

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

Required: No

Response Syntax

```
{
  "Consumer": {
    "ConsumerARN": "string",
    "ConsumerCreationTimestamp": number,
    "ConsumerName": "string",
    "ConsumerStatus": "string"
  }
}
```

Response Elements

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

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

[Consumer](#)

An object that represents the details of the consumer you registered. When you register a consumer, it gets an ARN that is generated by Kinesis Data Streams.

Type: [Consumer](#) object

Errors

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

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To create a consumer and add tags to it

The following JSON example creates a consumer with the specified tags.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
```

```
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.RegisterStreamConsumer
{
  "StreamARN": "arn:aws:kinesis:us-east-1:123456789012:/stream/myStream",
  "ConsumerName" : "myConsumer",
  "Tags" : [
    {
      "Key": "Project",
      "Value": "myProject"
    },
    {
      "Key": "Environment",
      "Value": "Production"
    }
  ]
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "Consumer": {
    "ConsumerARN": "string",
    "ConsumerCreationTimestamp": number,
    "ConsumerName": "string",
    "ConsumerStatus": "string"
  }
}
```

See Also

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

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

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

RemoveTagsFromStream

Removes tags from the specified Kinesis data stream. Removed tags are deleted and cannot be recovered after this operation successfully completes.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

If you specify a tag that does not exist, it is ignored.

[RemoveTagsFromStream](#) has a limit of five transactions per second per account.

Request Syntax

```
{
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string",
  "TagKeys": [ "string" ]
}
```

Request Parameters

The request accepts the following data in JSON format.

[StreamARN](#)

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

TagKeys

A list of tag keys. Each corresponding tag is removed from the stream.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 50 items.

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

Required: Yes

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To remove tags from a stream

The following JSON example removes the specified tag from the specified stream.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.RemoveTagsFromStream
{
  "StreamName": "exampleStreamName",
  "TagKeys": ["Project", "Environment"]
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
```

See Also

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

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

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

SplitShard

Splits a shard into two new shards in the Kinesis data stream, to increase the stream's capacity to ingest and transport data. `SplitShard` is called when there is a need to increase the overall capacity of a stream because of an expected increase in the volume of data records being ingested. This API is only supported for the data streams with the provisioned capacity mode.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

You can also use `SplitShard` when a shard appears to be approaching its maximum utilization; for example, the producers sending data into the specific shard are suddenly sending more than previously anticipated. You can also call `SplitShard` to increase stream capacity, so that more Kinesis Data Streams applications can simultaneously read data from the stream for real-time processing.

You must specify the shard to be split and the new hash key, which is the position in the shard where the shard gets split in two. In many cases, the new hash key might be the average of the beginning and ending hash key, but it can be any hash key value in the range being mapped into the shard. For more information, see [Split a Shard](#) in the *Amazon Kinesis Data Streams Developer Guide*.

You can use [DescribeStreamSummary](#) and the [ListShards](#) APIs to determine the shard ID and hash key values for the `ShardToSplit` and `NewStartingHashKey` parameters that are specified in the `SplitShard` request.

`SplitShard` is an asynchronous operation. Upon receiving a `SplitShard` request, Kinesis Data Streams immediately returns a response and sets the stream status to `UPDATING`. After the operation is completed, Kinesis Data Streams sets the stream status to `ACTIVE`. Read and write operations continue to work while the stream is in the `UPDATING` state.

You can use [DescribeStreamSummary](#) to check the status of the stream, which is returned in `StreamStatus`. If the stream is in the `ACTIVE` state, you can call `SplitShard`.

If the specified stream does not exist, [DescribeStreamSummary](#) returns a `ResourceNotFoundException`. If you try to create more shards than are authorized for your account, you receive a `LimitExceededException`.

For the default shard limit for an AWS account, see [Kinesis Data Streams Limits](#) in the *Amazon Kinesis Data Streams Developer Guide*. To increase this limit, [contact AWS Support](#).

If you try to operate on too many streams simultaneously using [CreateStream](#), [DeleteStream](#), [MergeShards](#), and/or [SplitShard](#), you receive a `LimitExceededException`.

`SplitShard` has a limit of five transactions per second per account.

Request Syntax

```
{
  "NewStartingHashKey": "string",
  "ShardToSplit": "string",
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

[NewStartingHashKey](#)

A hash key value for the starting hash key of one of the child shards created by the split. The hash key range for a given shard constitutes a set of ordered contiguous positive integers. The value for `NewStartingHashKey` must be in the range of hash keys being mapped into the shard. The `NewStartingHashKey` hash key value and all higher hash key values in hash key range are distributed to one of the child shards. All the lower hash key values in the range are distributed to the other child shard.

Type: String

Pattern: `^(0|([1-9]\d{0,38}))$`

Required: Yes

ShardToSplit

The shard ID of the shard to split.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: arn:aws.*:kinesis:.*:\d{12}:stream/\S+

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: [a-z0-9]{20}-[a-z0-9]{3}

Required: No

StreamName

The name of the stream for the shard split.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: No

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.
message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

Examples

To split a shard

The following JSON example splits the specified shard.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.SplitShard
{
  "StreamName": "exampleStreamName",
  "ShardToSplit": "shardId-000000000000",
  "NewStartingHashKey": "10"
}
```

Sample Response

```
HTTP/1.1 200 OK
```

```
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
```

See Also

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

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

StartStreamEncryption

Enables or updates server-side encryption using an AWS KMS key for a specified stream.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

Starting encryption is an asynchronous operation. Upon receiving the request, Kinesis Data Streams returns immediately and sets the status of the stream to `UPDATING`. After the update is complete, Kinesis Data Streams sets the status of the stream back to `ACTIVE`. Updating or applying encryption normally takes a few seconds to complete, but it can take minutes. You can continue to read and write data to your stream while its status is `UPDATING`. Once the status of the stream is `ACTIVE`, encryption begins for records written to the stream.

API Limits: You can successfully apply a new AWS KMS key for server-side encryption 25 times in a rolling 24-hour period.

Note: It can take up to 5 seconds after the stream is in an `ACTIVE` status before all records written to the stream are encrypted. After you enable encryption, you can verify that encryption is applied by inspecting the API response from `PutRecord` or `PutRecords`.

Request Syntax

```
{
  "EncryptionType": "string",
  "KeyId": "string",
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

EncryptionType

The encryption type to use. The only valid value is KMS.

Type: String

Valid Values: KMS

Required: Yes

KeyId

The GUID for the customer-managed AWS KMS key to use for encryption. This value can be a globally unique identifier, a fully specified Amazon Resource Name (ARN) to either an alias or a key, or an alias name prefixed by "alias/". You can also use a master key owned by Kinesis Data Streams by specifying the alias `aws/kinesis`.

- Key ARN example: `arn:aws:kms:us-east-1:123456789012:key/12345678-1234-1234-1234-123456789012`
- Alias ARN example: `arn:aws:kms:us-east-1:123456789012:alias/MyAliasName`
- Globally unique key ID example: `12345678-1234-1234-1234-123456789012`
- Alias name example: `alias/MyAliasName`
- Master key owned by Kinesis Data Streams: `alias/aws/kinesis`

Type: String

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

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream for which to start encrypting records.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSAccessDeniedException

The ciphertext references a key that doesn't exist or that you don't have access to.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSDisabledException

The request was rejected because the specified customer master key (CMK) isn't enabled.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSInvalidStateException

The request was rejected because the state of the specified resource isn't valid for this request. For more information, see [How Key State Affects Use of a Customer Master Key](#) in the *AWS Key Management Service Developer Guide*.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSNotFoundException

The request was rejected because the specified entity or resource can't be found.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSOptInRequired

The AWS access key ID needs a subscription for the service.

message

A message that provides information about the error.

HTTP Status Code: 400

KMSThrottlingException

The request was denied due to request throttling. For more information about throttling, see [Limits](#) in the *AWS Key Management Service Developer Guide*.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

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

StopStreamEncryption

Disables server-side encryption for a specified stream.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

Stopping encryption is an asynchronous operation. Upon receiving the request, Kinesis Data Streams returns immediately and sets the status of the stream to `UPDATING`. After the update is complete, Kinesis Data Streams sets the status of the stream back to `ACTIVE`. Stopping encryption normally takes a few seconds to complete, but it can take minutes. You can continue to read and write data to your stream while its status is `UPDATING`. Once the status of the stream is `ACTIVE`, records written to the stream are no longer encrypted by Kinesis Data Streams.

API Limits: You can successfully disable server-side encryption 25 times in a rolling 24-hour period.

Note: It can take up to 5 seconds after the stream is in an `ACTIVE` status before all records written to the stream are no longer subject to encryption. After you disabled encryption, you can verify that encryption is not applied by inspecting the API response from `PutRecord` or `PutRecords`.

Request Syntax

```
{
  "EncryptionType": "string",
  "KeyId": "string",
  "StreamARN": "string",
  "StreamId": "string",
  "StreamName": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

EncryptionType

The encryption type. The only valid value is KMS.

Type: String

Valid Values: KMS

Required: Yes

KeyId

The GUID for the customer-managed AWS KMS key to use for encryption. This value can be a globally unique identifier, a fully specified Amazon Resource Name (ARN) to either an alias or a key, or an alias name prefixed by "alias/". You can also use a master key owned by Kinesis Data Streams by specifying the alias `aws/kinesis`.

- Key ARN example: `arn:aws:kms:us-east-1:123456789012:key/12345678-1234-1234-1234-123456789012`
- Alias ARN example: `arn:aws:kms:us-east-1:123456789012:alias/MyAliasName`
- Globally unique key ID example: `12345678-1234-1234-1234-123456789012`
- Alias name example: `alias/MyAliasName`
- Master key owned by Kinesis Data Streams: `alias/aws/kinesis`

Type: String

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

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream on which to stop encrypting records.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

- [AWS Command Line Interface V2](#)
- [AWS SDK for .NET V4](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)

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

SubscribeToShard

This operation establishes an HTTP/2 connection between the consumer you specify in the `ConsumerARN` parameter and the shard you specify in the `ShardId` parameter. After the connection is successfully established, Kinesis Data Streams pushes records from the shard to the consumer over this connection. Before you call this operation, call [RegisterStreamConsumer](#) to register the consumer with Kinesis Data Streams.

When the `SubscribeToShard` call succeeds, your consumer starts receiving events of type [SubscribeToShardEvent](#) over the HTTP/2 connection for up to 5 minutes, after which time you need to call `SubscribeToShard` again to renew the subscription if you want to continue to receive records.

You can make one call to `SubscribeToShard` per second per registered consumer per shard. For example, if you have a 4000 shard stream and two registered stream consumers, you can make one `SubscribeToShard` request per second for each combination of shard and registered consumer, allowing you to subscribe both consumers to all 4000 shards in one second.

If you call `SubscribeToShard` again with the same `ConsumerARN` and `ShardId` within 5 seconds of a successful call, you'll get a `ResourceInUseException`. If you call `SubscribeToShard` 5 seconds or more after a successful call, the second call takes over the subscription and the previous connection expires or fails with a `ResourceInUseException`.

For an example of how to use this operation, see [Enhanced Fan-Out Using the Kinesis Data Streams API](#).

Request Syntax

```
{
  "ConsumerARN": "string",
  "DryRun": boolean,
  "ShardId": "string",
  "StartingPosition": {
    "SequenceNumber": "string",
    "Timestamp": number,
    "Type": "string"
  },
  "StreamId": "string"
}
```

Request Parameters

The request accepts the following data in JSON format.

ConsumerARN

For this parameter, use the value you obtained when you called [RegisterStreamConsumer](#).

Type: String

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

Pattern: `^(arn):aws.*:kinesis:.*:\d{12}:.*stream\[a-zA-Z0-9_.-]+\[/consumer\[a-zA-Z0-9_.-]+\:[0-9]+\`

Required: Yes

DryRun

Checks if your request will succeed. DryRun is an optional parameter.

Type: Boolean

Required: No

ShardId

The ID of the shard you want to subscribe to. To see a list of all the shards for a given stream, use [ListShards](#).

Type: String

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

Pattern: `\[a-zA-Z0-9_.-]+\`

Required: Yes

StartingPosition

The starting position in the data stream from which to start streaming.

Type: [StartingPosition](#) object

Required: Yes

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: [a-z0-9]{20}-[a-z0-9]{3}

Required: No

Response Syntax

```
{
  "EventStream": {
    "InternalFailureException": {
    },
    "KMSAccessDeniedException": {
    },
    "KMSThrottlingException": {
    },
    "KMSDisabledException": {
    },
    "KMSInvalidStateException": {
    },
    "KMSNotFoundException": {
    },
    "KMSShutdownException": {
    },
    "KMSOptInRequired": {
    },
    "ResourceInUseException": {
    },
    "ResourceNotFoundException": {
    },
    "SubscribeToShardEvent": {
      "ChildShards": [
        {
          "HashKeyRange": {
            "EndingHashKey": "string",
            "StartingHashKey": "string"
          },
          "ParentShards": [ "string" ],
          "ShardId": "string"
        }
      ]
    }
  }
}
```

```
    }
  ],
  "ContinuationSequenceNumber": "string",
  "MillisBehindLatest": number,
  "Records": [
    {
      "ApproximateArrivalTimestamp": number,
      "Data": blob,
      "EncryptionType": "string",
      "PartitionKey": "string",
      "SequenceNumber": "string"
    }
  ]
}
```

Response Elements

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

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

EventStream

The event stream that your consumer can use to read records from the shard.

Type: [SubscribeToShardEventStream](#) object

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

DryRunOperationException

The request was rejected because the DryRun parameter was specified.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

See Also

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

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

TagResource

Adds or updates tags for the specified Kinesis resource. Each tag is a label consisting of a user-defined key and value. Tags can help you manage, identify, organize, search for, and filter resources. You can assign up to 50 tags to a Kinesis resource.

Request Syntax

```
{
  "ResourceARN": "string",
  "StreamId": "string",
  "Tags": {
    "string" : "string"
  }
}
```

Request Parameters

The request accepts the following data in JSON format.

ResourceARN

The Amazon Resource Name (ARN) of the Kinesis resource to which to add tags.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:.*(stream|channel)/\S+`

Required: Yes

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

Tags

An array of tags to be added to the Kinesis resource. A tag consists of a required key and an optional value. You can add up to 50 tags per resource.

Tags may only contain Unicode letters, digits, white space, or these symbols: `_ . : / = + - @`.

Type: String to string map

Map Entries: Maximum number of 200 items.

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

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

Required: Yes

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To add tags to a resource

The following JSON example adds two tags to a specified consumer.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
```

```
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.TagResource
{
  "ResourceARN": "arn:aws:kinesis:us-east-1:123456789012:/stream/myStream/consumer/
myConsumer:174353308",
  "Tags": {
    "Project" : "myProject",
    "Environment" : "Production"
  }
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
```

See Also

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

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

UntagResource

Removes tags from the specified Kinesis resource. Removed tags are deleted and can't be recovered after this operation completes successfully.

Request Syntax

```
{  
  "ResourceARN": "string",  
  "StreamId": "string",  
  "TagKeys": [ "string" ]  
}
```

Request Parameters

The request accepts the following data in JSON format.

ResourceARN

The Amazon Resource Name (ARN) of the Kinesis resource from which to remove tags.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:.*(stream|channel)/\S+`

Required: Yes

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

TagKeys

A list of tag key-value pairs. Existing tags of the resource whose keys are members of this list will be removed from the Kinesis resource.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 50 items.

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

Required: Yes

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

Examples

To remove tags from a resource

The following JSON example removes the specified tag from a consumer.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.UntagResource
{
  "ResourceARN": "arn:aws:kinesis:us-east-1:123456789012:/stream/myStream/consumer/
myConsumer:174353308",
  "TagKeys": ["Project", "Environment"]
}
```

See Also

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

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

UpdateAccountSettings

Updates the account-level settings for Amazon Kinesis Data Streams.

Updating account settings is a synchronous operation. Upon receiving the request, Kinesis Data Streams will return immediately with your account's updated settings.

API limits

- Certain account configurations have minimum commitment windows. Attempting to update your settings prior to the end of the minimum commitment window might have certain restrictions.
- This API has a call limit of 5 transactions per second (TPS) for each AWS account. TPS over 5 will initiate the `LimitExceededException`.

Request Syntax

```
{  
  "MinimumThroughputBillingCommitment": {  
    "Status": "string"  
  }  
}
```

Request Parameters

The request accepts the following data in JSON format.

MinimumThroughputBillingCommitment

Specifies the minimum throughput billing commitment configuration for your account.

Type: [MinimumThroughputBillingCommitmentInput](#) object

Required: Yes

Response Syntax

```
{  
  "MinimumThroughputBillingCommitment": {  
    "EarliestAllowedEndAt": number,  
  }  
}
```

```
"EndedAt": number,  
"StartedAt": number,  
"Status": "string"  
}  
}
```

Response Elements

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

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

MinimumThroughputBillingCommitment

The updated configuration of the minimum throughput billing commitment for your account.

Type: [MinimumThroughputBillingCommitmentOutput](#) object

Errors

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

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

See Also

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

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

UpdateChannel

Updates the data freshness interval or the Amazon CloudWatch Logs configuration of an existing channel. You cannot change the destination, source stream, record format, schema, encryption configuration, or service execution role of an existing channel. To change any other setting, delete the channel and create a new one.

Updating a channel is an asynchronous operation. Upon receiving the request, Amazon Kinesis Data Streams sets the channel to the `UPDATING` state and returns immediately. After the change is applied, Amazon Kinesis Data Streams sets the channel back to the `ACTIVE` state.

This operation has a call limit of 5 transactions per second (TPS) for each AWS account. Exceeding 5 TPS results in a `LimitExceededException`.

Request Syntax

```
{
  "ChannelARN": "string",
  "LoggingConfiguration": {
    "CloudWatchLogs": {
      "Enabled": boolean,
      "LogGroupName": "string",
      "LogStreamName": "string"
    }
  },
  "S3DestinationConfiguration": {
    "DataFreshnessInSeconds": number
  },
  "S3TablesDestinationConfiguration": {
    "DataFreshnessInSeconds": number
  }
}
```

Request Parameters

The request accepts the following data in JSON format.

ChannelARN

The Amazon Resource Name (ARN) of the channel to update.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:channel/\S+`

Required: Yes

LoggingConfiguration

The updated Amazon CloudWatch Logs configuration for the channel.

Type: [ChannelLoggingUpdateInput](#) object

Required: No

S3DestinationConfiguration

The updated configuration for a general purpose Amazon S3 destination. Specify this parameter when the channel delivers to a general purpose Amazon S3 bucket. Only `DataFreshnessInSeconds` can be updated.

Type: [S3DestinationUpdateInput](#) object

Required: No

S3TablesDestinationConfiguration

The updated configuration for a streaming table destination. Specify this parameter when the channel delivers to streaming tables on Apache Iceberg in Amazon S3 Tables. Only `DataFreshnessInSeconds` can be updated.

Type: [S3TablesDestinationUpdateInput](#) object

Required: No

Response Syntax

```
{
  "ChannelDescription": {
    "ChannelARN": "string",
    "ChannelCreationTimestamp": number,
    "ChannelId": "string",
    "ChannelName": "string",
    "ChannelStatus": "string",
    "ChannelStatusReason": "string",
    "EncryptionConfiguration": {
```

```

    "EncryptionType": "string",
    "KeyId": "string"
  },
  "LoggingConfiguration": {
    "CloudWatchLogs": {
      "Enabled": boolean,
      "LogGroupName": "string",
      "LogStreamName": "string"
    }
  },
  "S3DestinationConfiguration": {
    "DataFreshnessInSeconds": number,
    "DeadLetterQueueS3Configuration": {
      "BucketARN": "string",
      "ErrorOutputPrefix": "string",
      "ExpectedBucketOwner": "string"
    },
    "StorageConfiguration": {
      "BucketARN": "string",
      "CompressionType": "string",
      "ExpectedBucketOwner": "string",
      "OutputKeyTemplate": "string",
      "StorageClass": "string"
    }
  },
  "S3TablesDestinationConfiguration": {
    "DataFreshnessInSeconds": number,
    "DeadLetterQueueS3Configuration": {
      "BucketARN": "string",
      "ErrorOutputPrefix": "string",
      "ExpectedBucketOwner": "string"
    },
    "S3TablesConfigurationList": [
      {
        "CompressionType": "string",
        "Namespace": "string",
        "PartitionSpec": {
          "PartitionFields": [
            {
              "SourceName": "string",
              "Transform": "string"
            }
          ]
        }
      }
    ]
  },

```

```
        "TableBucketARN": "string",
        "TableName": "string"
      }
    ],
    "ServiceExecutionRoleARN": "string",
    "StreamConfigurationList": [
      {
        "RecordConfiguration": {
          "GSRSchemaARN": "string",
          "RecordFormatType": "string"
        },
        "StreamARN": "string",
        "StreamCreationTimestamp": number
      }
    ]
  }
}
```

Response Elements

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

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

ChannelDescription

The configuration and current status of the channel after the update, including its ARN, destination configuration, and lifecycle state. Immediately after the request, the state is UPDATING.

Type: [ChannelDescription](#) object

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

See Also

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

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

UpdateMaxRecordSize

This allows you to update the MaxRecordSize of a single record that you can write to, and read from a stream. You can ingest and digest single records up to 10240 KiB.

Request Syntax

```
{  
  "MaxRecordSizeInKiB": number,  
  "StreamARN": "string",  
  "StreamId": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

MaxRecordSizeInKiB

The maximum record size of a single record in KiB that you can write to, and read from a stream. Specify a value between 1024 and 10240 KiB (1 to 10 MiB). If you specify a value that is out of this range, UpdateMaxRecordSize sends back a ValidationException message.

Type: Integer

Valid Range: Minimum value of 1024. Maximum value of 10240.

Required: Yes

StreamARN

The Amazon Resource Name (ARN) of the stream for the MaxRecordSize update.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

Response Elements

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

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

See Also

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

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

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

UpdateShardCount

Updates the shard count of the specified stream to the specified number of shards. This API is only supported for the data streams with the provisioned capacity mode.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

Updating the shard count is an asynchronous operation. Upon receiving the request, Kinesis Data Streams returns immediately and sets the status of the stream to `UPDATING`. After the update is complete, Kinesis Data Streams sets the status of the stream back to `ACTIVE`. Depending on the size of the stream, the scaling action could take a few minutes to complete. You can continue to read and write data to your stream while its status is `UPDATING`.

To update the shard count, Kinesis Data Streams performs splits or merges on individual shards. This can cause short-lived shards to be created, in addition to the final shards. These short-lived shards count towards your total shard limit for your account in the Region.

When using this operation, we recommend that you specify a target shard count that is a multiple of 25% (25%, 50%, 75%, 100%). You can specify any target value within your shard limit. However, if you specify a target that isn't a multiple of 25%, the scaling action might take longer to complete.

This operation has the following default limits. By default, you cannot do the following:

- Scale more than ten times per rolling 24-hour period per stream
- Scale up to more than double your current shard count for a stream
- Scale down below half your current shard count for a stream
- Scale up to more than 10000 shards in a stream
- Scale a stream with more than 10000 shards down unless the result is less than 10000 shards
- Scale up to more than the shard limit for your account
- Make over 10 TPS. TPS over 10 will trigger the `LimitExceededException`

For the default limits for an AWS account, see [Streams Limits](#) in the *Amazon Kinesis Data Streams Developer Guide*. To request an increase in the call rate limit, the shard limit for this API, or your overall shard limit, use the [limits form](#).

Request Syntax

```
{  
  "ScalingType": "string",  
  "StreamARN": "string",  
  "StreamId": "string",  
  "StreamName": "string",  
  "TargetShardCount": number  
}
```

Request Parameters

The request accepts the following data in JSON format.

ScalingType

The scaling type. Uniform scaling creates shards of equal size.

Type: String

Valid Values: UNIFORM_SCALING

Required: Yes

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

TargetShardCount

The new number of shards. This value has the following default limits. By default, you cannot do the following:

- Set this value to more than double your current shard count for a stream.
- Set this value below half your current shard count for a stream.
- Set this value to more than 10000 shards in a stream (the default limit for shard count per stream is 10000 per account per region), unless you request a limit increase.
- Scale a stream with more than 10000 shards down unless you set this value to less than 10000 shards.

Type: Integer

Valid Range: Minimum value of 1.

Required: Yes

Response Syntax

```
{  
  "CurrentShardCount": number,  
  "StreamARN": "string",
```

```
"StreamName": "string",  
"TargetShardCount": number  
}
```

Response Elements

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

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

CurrentShardCount

The current number of shards.

Type: Integer

Valid Range: Minimum value of 1.

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

StreamName

The name of the stream.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

TargetShardCount

The updated number of shards.

Type: Integer

Valid Range: Minimum value of 1.

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

Examples

To update the shard count of the specified stream from 2 to 4

This example illustrates one usage of UpdateShardCount.

Sample Request

```
POST / HTTP/1.1
Host: kinesis.<region>.<domain>
Content-Length: <PayloadSizeBytes>
User-Agent: <UserAgentString>
Content-Type: application/x-amz-json-1.1
Authorization: <AuthParams>
Connection: Keep-Alive
X-Amz-Date: <Date>
X-Amz-Target: Kinesis_20131202.UpdateShardCount
{
  "StreamName": "exampleStreamName",
  "TargetShardCount": 4,
  "ScalingType": "UNIFORM_SCALING"
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
```

```
Date: <Date>
{
  "TargetShardCount": 4,
  "StreamName": "exampleStreamName",
  "CurrentShardCount": 2
}
```

See Also

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

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

UpdateStreamMode

Updates the capacity mode of the data stream. Currently, in Kinesis Data Streams, you can choose between an **on-demand** capacity mode and a **provisioned** capacity mode for your data stream.

If you'd still like to proactively scale your on-demand data stream's capacity, you can unlock the warm throughput feature for on-demand data streams by enabling `MinimumThroughputBillingCommitment` for your account. Once your account has `MinimumThroughputBillingCommitment` enabled, you can specify the warm throughput in MiB per second that your stream can support in writes.

Request Syntax

```
{
  "StreamARN": "string",
  "StreamId": "string",
  "StreamModeDetails": {
    "StreamMode": "string"
  },
  "WarmThroughputMiBps": number
}
```

Request Parameters

The request accepts the following data in JSON format.

StreamARN

Specifies the ARN of the data stream whose capacity mode you want to update.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

[StreamModeDetails](#)

Specifies the capacity mode to which you want to set your data stream. Currently, in Kinesis Data Streams, you can choose between an **on-demand** capacity mode and a **provisioned** capacity mode for your data streams.

Type: [StreamModeDetails](#) object

Required: Yes

[WarmThroughputMiBps](#)

The target warm throughput in MB/s that the stream should be scaled to handle. This represents the throughput capacity that will be immediately available for write operations. This field is only valid when the stream mode is being updated to on-demand.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

Response Elements

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

Errors

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

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

See Also

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

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

UpdateStreamWarmThroughput

Updates the warm throughput configuration for the specified Amazon Kinesis Data Streams on-demand data stream. Updates the warm throughput configuration for the specified on-demand data stream. Use this operation to scale your stream to a specified throughput level before anticipated traffic spikes, or to release excess capacity after traffic has decreased.

Note

When invoking this API, you must use either the `StreamARN` or the `StreamName` parameter, or both. It is recommended that you use the `StreamARN` input parameter when you invoke this API.

Updating the warm throughput is an asynchronous operation. Upon receiving the request, Kinesis Data Streams returns immediately and sets the status of the stream to `UPDATING`. After the update is complete, Kinesis Data Streams sets the status of the stream back to `ACTIVE`. Depending on the size of the stream, the scaling action could take a few minutes to complete. You can continue to read and write data to your stream while its status is `UPDATING`.

This operation is only supported for data streams with the on-demand capacity mode in accounts that have `MinimumThroughputBillingCommitment` enabled. Provisioned capacity mode streams do not support warm throughput configuration.

To release excess capacity, call the API again and set the warm throughput to the same or a lower value.

This operation has the following default limits. By default, you cannot do the following:

- Scale to more than 10 GiBps for an on-demand stream.
- This API has a call limit of 5 transactions per second (TPS) for each AWS account. TPS over 5 will initiate the `LimitExceededException`.

For the default limits for an AWS account, see [Streams Limits](#) in the *Amazon Kinesis Data Streams Developer Guide*. To request an increase in the call rate limit, the shard limit for this API, or your overall shard limit, use the [limits form](#).

Request Syntax

```
{  
  "StreamARN": "string",  
  "StreamId": "string",  
  "StreamName": "string",  
  "WarmThroughputMiBps": number  
}
```

Request Parameters

The request accepts the following data in JSON format.

StreamARN

The ARN of the stream to be updated.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamName

The name of the stream to be updated.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: No

WarmThroughputMiBps

The target warm throughput in MB/s that the stream should be scaled to handle. This represents the throughput capacity that will be immediately available for write operations.

Type: Integer

Valid Range: Minimum value of 0.

Required: Yes

Response Syntax

```
{
  "StreamARN": "string",
  "StreamName": "string",
  "WarmThroughput": {
    "CurrentMiBps": number,
    "TargetMiBps": number
  }
}
```

Response Elements

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

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

StreamARN

The ARN of the stream that was updated.

Type: String

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

Pattern: arn:aws.*:kinesis:.*:\d{12}:stream/\S+

StreamName

The name of the stream that was updated.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

WarmThroughput

Specifies the updated warm throughput configuration for your data stream.

Type: [WarmThroughputObject](#) object

Errors

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

AccessDeniedException

Specifies that you do not have the permissions required to perform this operation.

HTTP Status Code: 400

InvalidArgumentException

A specified parameter exceeds its restrictions, is not supported, or can't be used. For more information, see the returned message.

message

A message that provides information about the error.

HTTP Status Code: 400

LimitExceededException

The requested resource exceeds the maximum number allowed, or the number of concurrent stream requests exceeds the maximum number allowed.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

message

A message that provides information about the error.

HTTP Status Code: 400

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

message

A message that provides information about the error.

HTTP Status Code: 400

ValidationException

Specifies that you tried to invoke this API for a data stream with the on-demand capacity mode. This API is only supported for data streams with the provisioned capacity mode.

HTTP Status Code: 400

See Also

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

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

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

Data Types

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

Note

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

The following data types are supported:

- [ChannelDescription](#)
- [ChannelEncryptionConfiguration](#)
- [ChannelLoggingConfiguration](#)
- [ChannelLoggingUpdateInput](#)
- [ChannelStreamConfiguration](#)
- [ChannelStreamDescription](#)
- [ChannelStreamIdentifier](#)
- [ChannelSummary](#)
- [ChildShard](#)
- [CloudWatchLogs](#)
- [CloudWatchLogsUpdateInput](#)
- [Consumer](#)
- [ConsumerDescription](#)
- [DeadLetterQueueS3Configuration](#)
- [EnhancedMetrics](#)
- [HashKeyRange](#)
- [MinimumThroughputBillingCommitmentInput](#)
- [MinimumThroughputBillingCommitmentOutput](#)
- [PartitionField](#)
- [PartitionSpec](#)

- [PutRecordsRequestEntry](#)
- [PutRecordsResultEntry](#)
- [Record](#)
- [RecordConfiguration](#)
- [S3DestinationConfiguration](#)
- [S3DestinationDescription](#)
- [S3DestinationUpdateInput](#)
- [S3StorageConfiguration](#)
- [S3TablesConfiguration](#)
- [S3TablesDestinationConfiguration](#)
- [S3TablesDestinationDescription](#)
- [S3TablesDestinationUpdateInput](#)
- [SequenceNumberRange](#)
- [Shard](#)
- [ShardFilter](#)
- [StartingPosition](#)
- [StreamDescription](#)
- [StreamDescriptionSummary](#)
- [StreamFilter](#)
- [StreamModeDetails](#)
- [StreamSummary](#)
- [SubscribeToShardEvent](#)
- [SubscribeToShardEventStream](#)
- [Tag](#)
- [WarmThroughputObject](#)

ChannelDescription

Describes the configuration and current status of a channel.

Contents

ChannelARN

The Amazon Resource Name (ARN) of the channel.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:channel/\S+`

Required: Yes

ChannelCreationTimestamp

The time at which the channel was created.

Type: Timestamp

Required: Yes

ChannelId

The unique identifier of the channel.

Type: String

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

Required: Yes

ChannelName

The name of the channel.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: Yes

ChannelStatus

The current status of the channel. Valid values:

- CREATING - The channel is being created.
- ACTIVE - The channel is ready to deliver records.
- UPDATING - The channel configuration is being updated.
- DELETING - The channel is being deleted.
- FAILED - See `ChannelStatusReason` for the failure cause.

Type: String

Valid Values: CREATING | ACTIVE | UPDATING | DELETING | FAILED

Required: Yes

LoggingConfiguration

The Amazon CloudWatch Logs configuration for the channel.

Type: [ChannelLoggingConfiguration](#) object

Required: Yes

ServiceExecutionRoleARN

The Amazon Resource Name (ARN) of the IAM role that Amazon Kinesis Data Streams assumes to write records to the destination.

Type: String

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

Pattern: `arn:aws[-a-z0-9]*:iam::\d{12}:role/[a-zA-Z_0-9+=,.\@-_/]+`

Required: Yes

StreamConfigurationList

The source stream configuration for the channel.

Type: Array of [ChannelStreamDescription](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10000 items.

Required: Yes

ChannelStatusReason

A message describing the reason for a FAILED status.

Type: String

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

Required: No

EncryptionConfiguration

The AWS KMS key configuration that Amazon Kinesis Data Streams uses to encrypt data delivered to the channel's destination.

Type: [ChannelEncryptionConfiguration](#) object

Required: No

S3DestinationConfiguration

The configuration for delivery to a general purpose Amazon S3 bucket. Present only when the channel destination is a general purpose Amazon S3 bucket.

Type: [S3DestinationDescription](#) object

Required: No

S3TablesDestinationConfiguration

The configuration for delivery to streaming tables on Apache Iceberg in Amazon S3 Tables. Present only when the channel destination is a streaming table.

Type: [S3TablesDestinationDescription](#) object

Required: No

See Also

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

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

ChannelEncryptionConfiguration

Specifies the AWS KMS key that Amazon Kinesis Data Streams uses to encrypt data delivered to the channel's destination.

Contents

EncryptionType

The encryption type. The only valid value is KMS.

Type: String

Valid Values: KMS

Required: Yes

KeyId

The identifier of the customer managed AWS KMS key. You cannot use the Amazon Kinesis Data Streams service key (aws/kinesis).

Type: String

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

Required: Yes

See Also

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

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

ChannelLoggingConfiguration

The Amazon CloudWatch Logs configuration for a channel.

Contents

CloudWatchLogs

The Amazon CloudWatch Logs settings for the channel.

Type: [CloudWatchLogs](#) object

Required: Yes

See Also

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

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

ChannelLoggingUpdateInput

The updated Amazon CloudWatch Logs configuration for a channel. Used in [UpdateChannel](#).

Contents

CloudWatchLogs

The updated Amazon CloudWatch Logs settings, including whether logging is enabled and the target log group and log stream.

Type: [CloudWatchLogsUpdateInput](#) object

Required: Yes

See Also

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

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

ChannelStreamConfiguration

Specifies the source stream and record configuration when creating a channel.

Contents

RecordConfiguration

The record format configuration for the source stream.

Type: [RecordConfiguration](#) object

Required: Yes

StreamARN

The Amazon Resource Name (ARN) of the source Kinesis data stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

See Also

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

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

ChannelStreamDescription

Describes the source stream of a channel.

Contents

RecordConfiguration

The record format configuration for the source stream.

Type: [RecordConfiguration](#) object

Required: Yes

StreamARN

The Amazon Resource Name (ARN) of the source Kinesis data stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

StreamCreationTimestamp

The time at which the source stream was created.

Type: Timestamp

Required: Yes

See Also

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

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

ChannelStreamIdentifier

Identifies a source stream associated with a channel.

Contents

StreamARN

The Amazon Resource Name (ARN) of the source Kinesis data stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

StreamCreationTimestamp

The time at which the source stream was created.

Type: Timestamp

Required: Yes

See Also

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

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

ChannelSummary

A summary of a channel, returned by [ListChannels](#).

Contents

ChannelARN

The Amazon Resource Name (ARN) of the channel.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:channel/\S+`

Required: Yes

ChannelCreationTimestamp

The time at which the channel was created.

Type: Timestamp

Required: Yes

ChannelDestinationType

The destination type of the channel. Valid values:

- S3 - Delivery to a general purpose Amazon S3 bucket.
- S3_TABLES - Delivery to streaming tables on Apache Iceberg.

Type: String

Valid Values: S3 | S3_TABLES

Required: Yes

ChannelId

The unique identifier of the channel.

Type: String

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

Required: Yes

ChannelName

The name of the channel.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

ChannelStatus

The current status of the channel. Valid values:

- CREATING - The channel is being created.
- ACTIVE - The channel is ready to deliver records.
- UPDATING - The channel configuration is being updated.
- DELETING - The channel is being deleted.
- FAILED - See `ChannelStatusReason` for the failure cause.

Type: String

Valid Values: CREATING | ACTIVE | UPDATING | DELETING | FAILED

Required: Yes

Streams

The source streams associated with the channel.

Type: Array of [ChannelStreamIdentifier](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10000 items.

Required: Yes

ChannelStatusReason

A message describing the reason for a FAILED status.

Type: String

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

Required: No

See Also

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

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

ChildShard

Output parameter of the GetRecords API. The existing child shard of the current shard.

Contents

HashKeyRange

The range of possible hash key values for the shard, which is a set of ordered contiguous positive integers.

Type: [HashKeyRange](#) object

Required: Yes

ParentShards

The current shard that is the parent of the existing child shard.

Type: Array of strings

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

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

ShardId

The shard ID of the existing child shard of the current shard.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

See Also

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

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

CloudWatchLogs

The Amazon CloudWatch Logs settings for channel logging.

Contents

Enabled

Specifies whether logging to Amazon CloudWatch Logs is enabled.

Type: Boolean

Required: Yes

LogGroupName

The name of the Amazon CloudWatch Logs log group. Defaults to `/aws/kinesis/{channelName}/{channelId}`.

Type: String

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

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

Required: No

LogStreamName

The name of the Amazon CloudWatch Logs log stream. Defaults to `DestinationDelivery`.

Type: String

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

Pattern: `^[^:]*`

Required: No

See Also

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

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

CloudWatchLogsUpdateInput

The updated Amazon CloudWatch Logs settings for a channel.

Contents

Enabled

Specifies whether logging to Amazon CloudWatch Logs is enabled.

Type: Boolean

Required: Yes

LogGroupName

The name of the Amazon CloudWatch Logs log group.

Type: String

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

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

Required: No

LogStreamName

The name of the Amazon CloudWatch Logs log stream.

Type: String

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

Pattern: `[^:]*`

Required: No

See Also

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

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

Consumer

An object that represents the details of the consumer you registered. This type of object is returned by [RegisterStreamConsumer](#).

Contents

ConsumerARN

When you register a consumer, Kinesis Data Streams generates an ARN for it. You need this ARN to be able to call [SubscribeToShard](#).

If you delete a consumer and then create a new one with the same name, it won't have the same ARN. That's because consumer ARNs contain the creation timestamp. This is important to keep in mind if you have IAM policies that reference consumer ARNs.

Type: String

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

Pattern: `^(arn):aws.*:kinesis:.*:\d{12}:.*stream[/][a-zA-Z0-9_.-]+[/]consumer[/][a-zA-Z0-9_.-]+:[0-9]+`

Required: Yes

ConsumerCreationTimestamp

Type: Timestamp

Required: Yes

ConsumerName

The name of the consumer is something you choose when you register the consumer.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: Yes

ConsumerStatus

A consumer can't read data while in the CREATING or DELETING states.

Type: String

Valid Values: CREATING | DELETING | ACTIVE

Required: Yes

See Also

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

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

ConsumerDescription

An object that represents the details of a registered consumer. This type of object is returned by [DescribeStreamConsumer](#).

Contents

ConsumerARN

When you register a consumer, Kinesis Data Streams generates an ARN for it. You need this ARN to be able to call [SubscribeToShard](#).

If you delete a consumer and then create a new one with the same name, it won't have the same ARN. That's because consumer ARNs contain the creation timestamp. This is important to keep in mind if you have IAM policies that reference consumer ARNs.

Type: String

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

Pattern: `^(arn):aws.*:kinesis:.*:\d{12}:.*stream[/][a-zA-Z0-9_.-]+[/]consumer[/][a-zA-Z0-9_.-]+:[0-9]+`

Required: Yes

ConsumerCreationTimestamp

Type: Timestamp

Required: Yes

ConsumerName

The name of the consumer is something you choose when you register the consumer.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: Yes

ConsumerStatus

A consumer can't read data while in the CREATING or DELETING states.

Type: String

Valid Values: CREATING | DELETING | ACTIVE

Required: Yes

StreamARN

The ARN of the stream with which you registered the consumer.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

See Also

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

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

DeadLetterQueueS3Configuration

The Amazon S3 dead-letter queue configuration for records that cannot be delivered.

Contents

BucketARN

The Amazon Resource Name (ARN) of the dead-letter queue Amazon S3 bucket.

Type: String

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

Pattern: `arn:aws[-a-z0-9]*:s3::[a-z0-9._-]{3,63}`

Required: Yes

ExpectedBucketOwner

The AWS account ID of the expected owner of the dead-letter queue bucket.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `\d{12}`

Required: Yes

ErrorOutputPrefix

The Amazon S3 key prefix for error records.

Type: String

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

Pattern: `[0-9A-Za-z!\-_'.*()\/]+`

Required: No

See Also

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

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

EnhancedMetrics

Represents enhanced metrics types.

Contents

ShardLevelMetrics

List of shard-level metrics.

The following are the valid shard-level metrics. The value "ALL" enhances every metric.

- IncomingBytes
- IncomingRecords
- OutgoingBytes
- OutgoingRecords
- WriteProvisionedThroughputExceeded
- ReadProvisionedThroughputExceeded
- IteratorAgeMilliseconds
- ALL

For more information, see [Monitoring the Amazon Kinesis Data Streams Service with Amazon CloudWatch](#) in the *Amazon Kinesis Data Streams Developer Guide*.

Type: Array of strings

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

Valid Values: IncomingBytes | IncomingRecords | OutgoingBytes
| OutgoingRecords | WriteProvisionedThroughputExceeded |
ReadProvisionedThroughputExceeded | IteratorAgeMilliseconds | ALL

Required: No

See Also

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

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

HashKeyRange

The range of possible hash key values for the shard, which is a set of ordered contiguous positive integers.

Contents

EndingHashKey

The ending hash key of the hash key range.

Type: String

Pattern: `^(0|([1-9]\d{0,38}))$`

Required: Yes

StartingHashKey

The starting hash key of the hash key range.

Type: String

Pattern: `^(0|([1-9]\d{0,38}))$`

Required: Yes

See Also

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

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

MinimumThroughputBillingCommitmentInput

Represents the request parameters for configuring minimum throughput billing commitment.

Note

- Minimum throughput billing commitments provide cost savings on on-demand data streams in exchange for committing to a minimum level of throughput usage.
- Commitments have a minimum duration of 24 hours that must be honored before they can be disabled.
- If you attempt to disable a commitment before the minimum commitment period ends, the commitment will be scheduled for automatic disable at the earliest allowed end time.
- You can cancel a pending disable by enabling the commitment again before the earliest allowed end time.

Contents

Status

The desired status of the minimum throughput billing commitment.

Type: String

Valid Values: ENABLED | DISABLED

Required: Yes

See Also

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

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

MinimumThroughputBillingCommitmentOutput

Represents the current status of minimum throughput billing commitment for an account.

Contents

Status

The current status of the minimum throughput billing commitment.

Type: String

Valid Values: ENABLED | DISABLED | ENABLED_UNTIL_EARLIEST_ALLOWED_END

Required: Yes

EarliestAllowedEndAt

The earliest timestamp when the commitment can be ended.

Type: Timestamp

Required: No

EndedAt

The timestamp when the commitment was ended.

Type: Timestamp

Required: No

StartedAt

The timestamp when the commitment was started.

Type: Timestamp

Required: No

See Also

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

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

PartitionField

Specifies a single partition field.

Contents

SourceName

The name of the source column used for partitioning. This column must be of the `timestampz` type.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: `[a-zA-Z0-9\.\_\-]+`

Required: Yes

Transform

The partition transform to apply. The only valid value is `TIME_HOUR`.

Type: String

Valid Values: `TIME_HOUR`

Required: Yes

See Also

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

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

PartitionSpec

Specifies how the destination table is partitioned.

Contents

PartitionFields

The list of partition fields.

Type: Array of [PartitionField](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10 items.

Required: Yes

See Also

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

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

PutRecordsRequestEntry

Represents the output for PutRecords.

Contents

Data

The data blob to put into the record, which is base64-encoded when the blob is serialized. When the data blob (the payload before base64-encoding) is added to the partition key size, the total size must not exceed the maximum record size (10 MiB).

Type: Base64-encoded binary data object

Length Constraints: Minimum length of 0. Maximum length of 10485760.

Required: Yes

PartitionKey

Determines which shard in the stream the data record is assigned to. Partition keys are Unicode strings with a maximum length limit of 256 characters for each key. Amazon Kinesis Data Streams uses the partition key as input to a hash function that maps the partition key and associated data to a specific shard. Specifically, an MD5 hash function is used to map partition keys to 128-bit integer values and to map associated data records to shards. As a result of this hashing mechanism, all data records with the same partition key map to the same shard within the stream.

Type: String

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

Required: Yes

ExplicitHashKey

The hash value used to determine explicitly the shard that the data record is assigned to by overriding the partition key hash.

Type: String

Pattern: `^(0|([1-9]\d{0,38}))$`

Required: No

See Also

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

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

PutRecordsResultEntry

Represents the result of an individual record from a PutRecords request. A record that is successfully added to a stream includes `SequenceNumber` and `ShardId` in the result. A record that fails to be added to the stream includes `ErrorCode` and `ErrorMessage` in the result.

Contents

ErrorCode

The error code for an individual record result. ErrorCodes can be either `ProvisionedThroughputExceededException` or `InternalFailure`.

Type: String

Required: No

ErrorMessage

The error message for an individual record result. An `ErrorCode` value of `ProvisionedThroughputExceededException` has an error message that includes the account ID, stream name, and shard ID. An `ErrorCode` value of `InternalFailure` has the error message "Internal Service Failure".

Type: String

Required: No

SequenceNumber

The sequence number for an individual record result.

Type: String

Pattern: `^(0|([1-9]\d{0,128}))$`

Required: No

ShardId

The shard ID for an individual record result.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: No

See Also

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

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

Record

The unit of data of the Kinesis data stream, which is composed of a sequence number, a partition key, and a data blob.

Contents

Data

The data blob. The data in the blob is both opaque and immutable to Kinesis Data Streams, which does not inspect, interpret, or change the data in the blob in any way. When the data blob (the payload before base64-encoding) is added to the partition key size, the total size must not exceed the maximum record size (10 MiB).

Type: Base64-encoded binary data object

Length Constraints: Minimum length of 0. Maximum length of 10485760.

Required: Yes

PartitionKey

Identifies which shard in the stream the data record is assigned to.

Type: String

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

Required: Yes

SequenceNumber

The unique identifier of the record within its shard.

Type: String

Pattern: `^(0|([1-9]\d{0,128}))$`

Required: Yes

ApproximateArrivalTimestamp

The approximate time that the record was inserted into the stream.

Type: Timestamp

Required: No

EncryptionType

The encryption type used on the record. This parameter can be one of the following values:

- NONE: Do not encrypt the records in the stream.
- KMS: Use server-side encryption on the records in the stream using a customer-managed AWS KMS key.

Type: String

Valid Values: NONE | KMS

Required: No

See Also

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

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

RecordConfiguration

Specifies the format of records read from the source stream.

Contents

RecordFormatType

The format of records on the source stream. Valid values:

- GSR_JSON - Supported only for streaming table (Amazon S3 Tables) destinations.
- JSON - Supported for both general purpose Amazon S3 and streaming table destinations.
- STRING - Supported only for general purpose Amazon S3 destinations.
- BYTE_ARRAY - Supported only for general purpose Amazon S3 destinations.

Type: String

Valid Values: GSR_JSON | JSON | STRING | BYTE_ARRAY

Required: Yes

GSRSchemaARN

The Amazon Resource Name (ARN) of the AWS Glue Schema Registry schema used to validate records. Required when the channel destination is a streaming table.

Type: String

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

Pattern: `arn:aws[-a-z0-9]*:glue:[-a-z0-9]+:\d{12}:schema/[-a-zA-Z0-9_$.]+/[-a-zA-Z0-9_$.]+`

Required: No

See Also

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

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

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

S3DestinationConfiguration

The configuration for delivery to a general purpose Amazon S3 bucket. Used in [CreateChannel](#).

Contents

StorageConfiguration

The Amazon S3 storage configuration for the channel.

Type: [S3StorageConfiguration](#) object

Required: Yes

DataFreshnessInSeconds

The maximum age, in seconds, of undelivered data before the channel delivers it to the destination. The default value is 300 seconds.

Type: Integer

Required: No

DeadLetterQueueS3Configuration

The dead-letter queue configuration for records that cannot be delivered. Optional for general purpose Amazon S3 destinations. If not specified, it defaults to the destination bucket with an error prefix.

Type: [DeadLetterQueueS3Configuration](#) object

Required: No

See Also

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

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

S3DestinationDescription

The configuration for delivery to a general purpose Amazon S3 bucket. Returned in [ChannelDescription](#).

Contents

DataFreshnessInSeconds

The maximum age, in seconds, of undelivered data.

Type: Integer

Required: Yes

DeadLetterQueueS3Configuration

The dead-letter queue configuration for records that cannot be delivered.

Type: [DeadLetterQueueS3Configuration](#) object

Required: Yes

StorageConfiguration

The Amazon S3 storage configuration for the channel.

Type: [S3StorageConfiguration](#) object

Required: Yes

See Also

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

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

S3DestinationUpdateInput

The updated configuration for a general purpose Amazon S3 destination. Used in [UpdateChannel](#). Only DataFreshnessInSeconds can be updated.

Contents

DataFreshnessInSeconds

The maximum age, in seconds, of undelivered data before the channel delivers it to the destination.

Type: Integer

Required: Yes

See Also

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

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

S3StorageConfiguration

The Amazon S3 storage settings for a general purpose Amazon S3 destination.

Contents

BucketARN

The Amazon Resource Name (ARN) of the destination Amazon S3 bucket.

Type: String

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

Pattern: `arn:aws[-a-z0-9]*:s3:::[a-z0-9._-]{3,63}`

Required: Yes

CompressionType

The compression applied to delivered objects. Valid values:

- NONE - No compression.
- GZIP - gzip compression.
- ZSTD - Zstandard compression.

Type: String

Valid Values: NONE | GZIP | ZSTD

Required: Yes

ExpectedBucketOwner

The AWS account ID of the expected owner of the destination bucket. This value helps prevent delivery to an unintended bucket if ownership changes.

Type: String

Length Constraints: Fixed length of 12.

Pattern: `\d{12}`

Required: Yes

OutputKeyTemplate

The template used to construct the Amazon S3 object key for delivered objects. If not specified, a default template is used.

Type: String

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

Pattern: `[0-9A-Za-z!\-_'.*()\|/=:{}]+`

Required: No

StorageClass

The Amazon S3 storage class for delivered objects. Valid values:

- STANDARD - The default storage class, for frequently accessed data.
- INTELLIGENT_TIERING - Automatically moves objects to the most cost-effective access tier based on usage patterns.
- GLACIER_IR - Low-cost storage for rarely accessed data that requires millisecond retrieval.

Type: String

Valid Values: STANDARD | INTELLIGENT_TIERING | GLACIER_IR

Required: No

See Also

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

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

S3TablesConfiguration

Specifies a destination streaming table on Apache Iceberg.

Contents

CompressionType

The compression applied to Parquet data files. Valid values:

- NONE - No compression.
- ZSTD - Zstandard compression.
- SNAPPY - Snappy compression.

Type: String

Valid Values: NONE | ZSTD | SNAPPY

Required: Yes

Namespace

The namespace (database) of the destination table.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Pattern: [0-9a-z_]+

Required: Yes

TableBucketARN

The Amazon Resource Name (ARN) of the Amazon S3 table bucket.

Type: String

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

Pattern: arn:aws[-a-z0-9]*:s3tables:[-a-z0-9]+:\d{12}:bucket/[-a-z0-9_-]{3,63}

Required: Yes

TableName

The name of the destination table. Amazon Kinesis Data Streams creates this table in the specified table bucket.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

PartitionSpec

The partitioning specification for the destination table.

Type: [PartitionSpec](#) object

Required: No

See Also

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

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

S3TablesDestinationConfiguration

The configuration for delivery to streaming tables on Apache Iceberg. Used in [CreateChannel](#).

Contents

DeadLetterQueueS3Configuration

The dead-letter queue configuration for records that cannot be delivered. Required for streaming table destinations.

Type: [DeadLetterQueueS3Configuration](#) object

Required: Yes

S3TablesConfigurationList

The list of streaming table configurations. Currently, one table is supported per channel.

Type: Array of [S3TablesConfiguration](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10000 items.

Required: Yes

DataFreshnessInSeconds

The maximum age, in seconds, of undelivered data before the channel delivers it to the destination. The default value is 300 seconds.

Type: Integer

Required: No

See Also

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

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

S3TablesDestinationDescription

The configuration for delivery to streaming tables on Apache Iceberg. Returned in [ChannelDescription](#).

Contents

DataFreshnessInSeconds

The maximum age, in seconds, of undelivered data.

Type: Integer

Required: Yes

DeadLetterQueueS3Configuration

The dead-letter queue configuration for records that cannot be delivered.

Type: [DeadLetterQueueS3Configuration](#) object

Required: Yes

S3TablesConfigurationList

The list of streaming table configurations.

Type: Array of [S3TablesConfiguration](#) objects

Array Members: Minimum number of 1 item. Maximum number of 10000 items.

Required: Yes

See Also

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

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

S3TablesDestinationUpdateInput

The updated configuration for a streaming table destination. Used in [UpdateChannel](#). Only `DataFreshnessInSeconds` can be updated.

Contents

DataFreshnessInSeconds

The maximum age, in seconds, of undelivered data before the channel delivers it to the destination.

Type: Integer

Required: Yes

See Also

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

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

SequenceNumberRange

The range of possible sequence numbers for the shard.

Contents

StartingSequenceNumber

The starting sequence number for the range.

Type: String

Pattern: `^(0|([1-9]\d{0,128}))$`

Required: Yes

EndingSequenceNumber

The ending sequence number for the range. Shards that are in the OPEN state have an ending sequence number of `null`.

Type: String

Pattern: `^(0|([1-9]\d{0,128}))$`

Required: No

See Also

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

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

Shard

A uniquely identified group of data records in a Kinesis data stream.

Contents

HashKeyRange

The range of possible hash key values for the shard, which is a set of ordered contiguous positive integers.

Type: [HashKeyRange](#) object

Required: Yes

SequenceNumberRange

The range of possible sequence numbers for the shard.

Type: [SequenceNumberRange](#) object

Required: Yes

ShardId

The unique identifier of the shard within the stream.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

AdjacentParentShardId

The shard ID of the shard adjacent to the shard's parent.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: No

ParentShardId

The shard ID of the shard's parent.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: No

See Also

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

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

ShardFilter

The request parameter used to filter out the response of the `ListShards` API.

Contents

Type

The shard type specified in the `ShardFilter` parameter. This is a required property of the `ShardFilter` parameter.

You can specify the following valid values:

- `AFTER_SHARD_ID` - the response includes all the shards, starting with the shard whose ID immediately follows the `ShardId` that you provided.
- `AT_TRIM_HORIZON` - the response includes all the shards that were open at `TRIM_HORIZON`.
- `FROM_TRIM_HORIZON` - (default), the response includes all the shards within the retention period of the data stream (trim to tip).
- `AT_LATEST` - the response includes only the currently open shards of the data stream.
- `AT_TIMESTAMP` - the response includes all shards whose start timestamp is less than or equal to the given timestamp and end timestamp is greater than or equal to the given timestamp or still open.
- `FROM_TIMESTAMP` - the response includes all closed shards whose end timestamp is greater than or equal to the given timestamp and also all open shards. Corrected to `TRIM_HORIZON` of the data stream if `FROM_TIMESTAMP` is less than the `TRIM_HORIZON` value.

Type: String

Valid Values: `AFTER_SHARD_ID` | `AT_TRIM_HORIZON` | `FROM_TRIM_HORIZON` | `AT_LATEST` | `AT_TIMESTAMP` | `FROM_TIMESTAMP`

Required: Yes

ShardId

The exclusive start `shardID` specified in the `ShardFilter` parameter. This property can only be used if the `AFTER_SHARD_ID` shard type is specified.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: No

Timestamp

The timestamps specified in the `ShardFilter` parameter. A timestamp is a Unix epoch date with precision in milliseconds. For example, `2016-04-04T19:58:46.480-00:00` or `1459799926.480`. This property can only be used if `FROM_TIMESTAMP` or `AT_TIMESTAMP` shard types are specified.

Type: Timestamp

Required: No

See Also

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

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

StartingPosition

The starting position in the data stream from which to start streaming.

Contents

Type

You can set the starting position to one of the following values:

AT_SEQUENCE_NUMBER: Start streaming from the position denoted by the sequence number specified in the `SequenceNumber` field.

AFTER_SEQUENCE_NUMBER: Start streaming right after the position denoted by the sequence number specified in the `SequenceNumber` field.

AT_TIMESTAMP: Start streaming from the position denoted by the time stamp specified in the `Timestamp` field.

TRIM_HORIZON: Start streaming at the last untrimmed record in the shard, which is the oldest data record in the shard.

LATEST: Start streaming just after the most recent record in the shard, so that you always read the most recent data in the shard.

Type: String

Valid Values: `AT_SEQUENCE_NUMBER` | `AFTER_SEQUENCE_NUMBER` | `TRIM_HORIZON` | `LATEST` | `AT_TIMESTAMP`

Required: Yes

SequenceNumber

The sequence number of the data record in the shard from which to start streaming. To specify a sequence number, set `StartingPosition` to `AT_SEQUENCE_NUMBER` or `AFTER_SEQUENCE_NUMBER`.

Type: String

Pattern: `^(0|([1-9]\d{0,128}))$`

Required: No

Timestamp

The time stamp of the data record from which to start reading. To specify a time stamp, set `StartingPosition` to Type `AT_TIMESTAMP`. A time stamp is the Unix epoch date with precision in milliseconds. For example, `2016-04-04T19:58:46.480-00:00` or `1459799926.480`. If a record with this exact time stamp does not exist, records will be streamed from the next (later) record. If the time stamp is older than the current trim horizon, records will be streamed from the oldest untrimmed data record (`TRIM_HORIZON`).

Type: Timestamp

Required: No

See Also

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

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

StreamDescription

Represents the output for [DescribeStream](#).

Contents

EnhancedMonitoring

Represents the current enhanced monitoring settings of the stream.

Type: Array of [EnhancedMetrics](#) objects

Required: Yes

HasMoreShards

If set to `true`, more shards in the stream are available to describe.

Type: Boolean

Required: Yes

RetentionPeriodHours

The current retention period, in hours. Minimum value of 24. Maximum value of 168.

Type: Integer

Required: Yes

Shards

The shards that comprise the stream.

Type: Array of [Shard](#) objects

Required: Yes

StreamARN

The Amazon Resource Name (ARN) for the stream being described.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

StreamCreationTimestamp

The approximate time that the stream was created.

Type: Timestamp

Required: Yes

StreamName

The name of the stream being described.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: Yes

StreamStatus

The current status of the stream being described. The stream status is one of the following states:

- **CREATING** - The stream is being created. Kinesis Data Streams immediately returns and sets `StreamStatus` to **CREATING**.
- **DELETING** - The stream is being deleted. The specified stream is in the **DELETING** state until Kinesis Data Streams completes the deletion.
- **ACTIVE** - The stream exists and is ready for read and write operations or deletion. You should perform read and write operations only on an **ACTIVE** stream.
- **UPDATING** - Shards in the stream are being merged or split. Read and write operations continue to work while the stream is in the **UPDATING** state.

Type: String

Valid Values: **CREATING** | **DELETING** | **ACTIVE** | **UPDATING**

Required: Yes

EncryptionType

The server-side encryption type used on the stream. This parameter can be one of the following values:

- **NONE**: Do not encrypt the records in the stream.
- **KMS**: Use server-side encryption on the records in the stream using a customer-managed AWS KMS key.

Type: String

Valid Values: NONE | KMS

Required: No

KeyId

The GUID for the customer-managed AWS KMS key to use for encryption. This value can be a globally unique identifier, a fully specified ARN to either an alias or a key, or an alias name prefixed by "alias/". You can also use a master key owned by Kinesis Data Streams by specifying the alias `aws/kinesis`.

- Key ARN example: `arn:aws:kms:us-east-1:123456789012:key/12345678-1234-1234-1234-123456789012`
- Alias ARN example: `arn:aws:kms:us-east-1:123456789012:alias/MyAliasName`
- Globally unique key ID example: `12345678-1234-1234-1234-123456789012`
- Alias name example: `alias/MyAliasName`
- Master key owned by Kinesis Data Streams: `alias/aws/kinesis`

Type: String

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

Required: No

StreamModeDetails

Specifies the capacity mode to which you want to set your data stream. Currently, in Kinesis Data Streams, you can choose between an **on-demand** capacity mode and a **provisioned** capacity mode for your data streams.

Type: [StreamModeDetails](#) object

Required: No

See Also

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

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

StreamDescriptionSummary

Represents the output for [DescribeStreamSummary](#)

Contents

EnhancedMonitoring

Represents the current enhanced monitoring settings of the stream.

Type: Array of [EnhancedMetrics](#) objects

Required: Yes

OpenShardCount

The number of open shards in the stream.

Type: Integer

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

Required: Yes

RetentionPeriodHours

The current retention period, in hours.

Type: Integer

Required: Yes

StreamARN

The Amazon Resource Name (ARN) for the stream being described.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

StreamCreationTimestamp

The approximate time that the stream was created.

Type: Timestamp

Required: Yes

StreamName

The name of the stream being described.

Type: String

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

Pattern: [a-zA-Z0-9_.-]+

Required: Yes

StreamStatus

The current status of the stream being described. The stream status is one of the following states:

- **CREATING** - The stream is being created. Kinesis Data Streams immediately returns and sets **StreamStatus** to **CREATING**.
- **DELETING** - The stream is being deleted. The specified stream is in the **DELETING** state until Kinesis Data Streams completes the deletion.
- **ACTIVE** - The stream exists and is ready for read and write operations or deletion. You should perform read and write operations only on an **ACTIVE** stream.
- **UPDATING** - Shards in the stream are being merged or split. Read and write operations continue to work while the stream is in the **UPDATING** state.

Type: String

Valid Values: CREATING | DELETING | ACTIVE | UPDATING

Required: Yes

ChannelCount

The number of channels associated with the stream.

Type: Integer

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

Required: No

ConsumerCount

The number of enhanced fan-out consumers registered with the stream.

Type: Integer

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

Required: No

EncryptionType

The encryption type used. This value is one of the following:

- KMS
- NONE

Type: String

Valid Values: NONE | KMS

Required: No

KeyId

The GUID for the customer-managed AWS KMS key to use for encryption. This value can be a globally unique identifier, a fully specified ARN to either an alias or a key, or an alias name prefixed by "alias/". You can also use a master key owned by Kinesis Data Streams by specifying the alias `aws/kinesis`.

- Key ARN example: `arn:aws:kms:us-east-1:123456789012:key/12345678-1234-1234-1234-123456789012`
- Alias ARN example: `arn:aws:kms:us-east-1:123456789012:alias/MyAliasName`
- Globally unique key ID example: `12345678-1234-1234-1234-123456789012`
- Alias name example: `alias/MyAliasName`
- Master key owned by Kinesis Data Streams: `alias/aws/kinesis`

Type: String

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

Required: No

MaxRecordSizeInKiB

The maximum record size of a single record in kibibyte (KiB) that you can write to, and read from a stream.

Type: Integer

Valid Range: Minimum value of 1024. Maximum value of 10240.

Required: No

StreamId

Not Implemented. Reserved for future use.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 24.

Pattern: `[a-z0-9]{20}-[a-z0-9]{3}`

Required: No

StreamModeDetails

Specifies the capacity mode to which you want to set your data stream. Currently, in Kinesis Data Streams, you can choose between an **on-demand** capacity mode and a **provisioned** capacity mode for your data streams.

Type: [StreamModeDetails](#) object

Required: No

WarmThroughput

The warm throughput in MB/s for the stream. This represents the throughput capacity that will be immediately available for write operations.

Type: [WarmThroughputObject](#) object

Required: No

See Also

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

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

StreamFilter

Filters [ListChannels](#) results by source stream.

Contents

StreamARN

The Amazon Resource Name (ARN) of the source stream to filter by.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

StreamCreationTimestamp

The creation timestamp of the source stream.

Type: Timestamp

Required: No

See Also

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

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

StreamModeDetails

Specifies the capacity mode to which you want to set your data stream. Currently, in Kinesis Data Streams, you can choose between an **on-demand** capacity mode and a **provisioned** capacity mode for your data streams.

Contents

StreamMode

Specifies the capacity mode to which you want to set your data stream. Currently, in Kinesis Data Streams, you can choose between an **on-demand** capacity mode and a **provisioned** capacity mode for your data streams.

Type: String

Valid Values: PROVISIONED | ON_DEMAND

Required: Yes

See Also

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

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

StreamSummary

The summary of a stream.

Contents

StreamARN

The ARN of the stream.

Type: String

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

Pattern: `arn:aws.*:kinesis:.*:\d{12}:stream/\S+`

Required: Yes

StreamName

The name of a stream.

Type: String

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

Pattern: `[a-zA-Z0-9_.-]+`

Required: Yes

StreamStatus

The status of the stream.

Type: String

Valid Values: CREATING | DELETING | ACTIVE | UPDATING

Required: Yes

StreamCreationTimestamp

The timestamp at which the stream was created.

Type: Timestamp

Required: No

StreamModeDetails

Specifies the capacity mode to which you want to set your data stream. Currently, in Kinesis Data Streams, you can choose between an **on-demand** capacity mode and a **provisioned** capacity mode for your data streams.

Type: [StreamModeDetails](#) object

Required: No

See Also

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

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

SubscribeToShardEvent

After you call [SubscribeToShard](#), Kinesis Data Streams sends events of this type over an HTTP/2 connection to your consumer.

Contents

ContinuationSequenceNumber

Use this as `SequenceNumber` in the next call to [SubscribeToShard](#), with `StartingPosition` set to `AT_SEQUENCE_NUMBER` or `AFTER_SEQUENCE_NUMBER`. Use `ContinuationSequenceNumber` for checkpointing because it captures your shard progress even when no data is written to the shard.

Type: String

Pattern: `^(0|([1-9]\d{0,128}))$`

Required: Yes

MillisBehindLatest

The number of milliseconds the read records are from the tip of the stream, indicating how far behind current time the consumer is. A value of zero indicates that record processing is caught up, and there are no new records to process at this moment.

Type: Long

Valid Range: Minimum value of 0.

Required: Yes

Records

Type: Array of [Record](#) objects

Required: Yes

ChildShards

The list of the child shards of the current shard, returned only at the end of the current shard.

Type: Array of [ChildShard](#) objects

Required: No

See Also

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

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

SubscribeToShardEventStream

This is a tagged union for all of the types of events an enhanced fan-out consumer can receive over HTTP/2 after a call to [SubscribeToShard](#).

Contents

SubscribeToShardEvent

After you call [SubscribeToShard](#), Kinesis Data Streams sends events of this type to your consumer. For an example of how to handle these events, see [Enhanced Fan-Out Using the Kinesis Data Streams API](#).

Type: [SubscribeToShardEvent](#) object

Required: Yes

InternalFailureException

The processing of the request failed because of an unknown error, exception, or failure.

Type: Exception

HTTP Status Code:

Required: No

KMSAccessDeniedException

The ciphertext references a key that doesn't exist or that you don't have access to.

Type: Exception

HTTP Status Code:

Required: No

KMSDisabledException

The request was rejected because the specified customer master key (CMK) isn't enabled.

Type: Exception

HTTP Status Code:

Required: No

KMSInvalidStateException

The request was rejected because the state of the specified resource isn't valid for this request. For more information, see [How Key State Affects Use of a Customer Master Key](#) in the *AWS Key Management Service Developer Guide*.

Type: Exception

HTTP Status Code:

Required: No

KMSNotFoundException

The request was rejected because the specified entity or resource can't be found.

Type: Exception

HTTP Status Code:

Required: No

KMSOptInRequired

The AWS access key ID needs a subscription for the service.

Type: Exception

HTTP Status Code:

Required: No

KMSThrottlingException

The request was denied due to request throttling. For more information about throttling, see [Limits](#) in the *AWS Key Management Service Developer Guide*.

Type: Exception

HTTP Status Code:

Required: No

ResourceInUseException

The resource is not available for this operation. For successful operation, the resource must be in the ACTIVE state.

Type: Exception

HTTP Status Code:

Required: No

ResourceNotFoundException

The requested resource could not be found. The stream might not be specified correctly.

Type: Exception

HTTP Status Code:

Required: No

See Also

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

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

Tag

Metadata assigned to the stream or consumer, consisting of a key-value pair.

Contents

Key

A unique identifier for the tag. Maximum length: 128 characters. Valid characters: Unicode letters, digits, white space, `_` `.` `/` `=` `+` `-` `%` `@`

Type: String

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

Required: Yes

Value

An optional string, typically used to describe or define the tag. Maximum length: 256 characters. Valid characters: Unicode letters, digits, white space, `_` `.` `/` `=` `+` `-` `%` `@`

Type: String

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

Required: No

See Also

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

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

WarmThroughputObject

Represents the warm throughput configuration on the stream. This is only present for On-Demand Kinesis Data Streams in accounts that have `MinimumThroughputBillingCommitment` enabled.

Contents

CurrentMiBps

The current warm throughput value on the stream. This is the write throughput in MiBps that the stream is currently scaled to handle.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

TargetMiBps

The target warm throughput value on the stream. This indicates that the stream is currently scaling towards this target value.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

See Also

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

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

Common Error Types

This section lists common error types that this AWS service may return. Not all services return all error types listed here. For errors specific to an API action for this service, see the topic for that API action.

AccessDeniedException

You don't have permission to perform this action. Verify that your IAM policy includes the required permissions.

HTTP Status Code: 403

ExpiredTokenException

The security token included in the request has expired. Request a new security token and try again.

HTTP Status Code: 403

IncompleteSignature

The request signature doesn't conform to AWS standards. Verify that you're using valid AWS credentials and that your request is properly formatted. If you're using an SDK, ensure it's up to date.

HTTP Status Code: 403

InternalFailure

The request can't be processed right now because of an internal server issue. Try again later. If the problem persists, contact AWS Support.

HTTP Status Code: 500

MalformedHttpRequestException

The request body can't be processed. This typically happens when the request body can't be decompressed using the specified content encoding algorithm. Verify that the content encoding header matches the compression format used.

HTTP Status Code: 400

NotAuthorized

You don't have permissions to perform this action. Verify that your IAM policy includes the required permissions.

HTTP Status Code: 401

OptInRequired

Your AWS account needs a subscription for this service. Verify that you've enabled the service in your account.

HTTP Status Code: 403

RequestAbortedException

The request was aborted before a response could be returned. This typically happens when the client closes the connection.

HTTP Status Code: 400

RequestEntityTooLargeException

The request entity is too large. Reduce the size of the request body and try again.

HTTP Status Code: 413

RequestTimeoutException

The request timed out. The server didn't receive the complete request within the expected time frame. Try again.

HTTP Status Code: 408

ServiceUnavailable

The service is temporarily unavailable. Try again later.

HTTP Status Code: 503

ThrottlingException

Your request rate is too high. The AWS SDKs automatically retry requests that receive this exception. Reduce the frequency of requests.

HTTP Status Code: 400

UnknownOperationException

The action or operation isn't recognized. Verify that the action name is spelled correctly and that it's supported by the API version you're using.

HTTP Status Code: 404

UnrecognizedClientException

The X.509 certificate or AWS access key ID you provided doesn't exist in our records. Verify that you're using valid credentials and that they haven't expired.

HTTP Status Code: 403

ValidationError

The input doesn't meet the required format or constraints. Check that all required parameters are included and that values are valid.

HTTP Status Code: 400