
Amazon Kinesis Firehose API Reference

API Reference

API Version 2015-08-04



Amazon Kinesis Firehose API Reference: API Reference

Copyright © 2018 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
CreateDeliveryStream	3
Request Syntax	3
Request Parameters	8
Response Syntax	9
Response Elements	9
Errors	9
Example	10
See Also	10
DeleteDeliveryStream	12
Request Syntax	12
Request Parameters	12
Response Elements	12
Errors	12
Example	13
See Also	13
DescribeDeliveryStream	14
Request Syntax	14
Request Parameters	14
Response Syntax	14
Response Elements	19
Errors	19
Example	19
See Also	21
ListDeliveryStreams	22
Request Syntax	22
Request Parameters	22
Response Syntax	23
Response Elements	23
Errors	23
Example	23
See Also	24
PutRecord	25
Request Syntax	25
Request Parameters	25
Response Syntax	26
Response Elements	26
Errors	26
Example	26
See Also	27
PutRecordBatch	28
Request Syntax	28
Request Parameters	29
Response Syntax	29
Response Elements	29
Errors	30
Example	30
See Also	31
UpdateDestination	32
Request Syntax	32
Request Parameters	36
Response Elements	37
Errors	37

Example	38
See Also	39
Data Types	40
BufferingHints	42
Contents	42
See Also	42
CloudWatchLoggingOptions	43
Contents	43
See Also	43
CopyCommand	44
Contents	44
See Also	44
DeliveryStreamDescription	45
Contents	45
See Also	46
DestinationDescription	47
Contents	47
See Also	47
ElasticsearchBufferingHints	49
Contents	49
See Also	49
ElasticsearchDestinationConfiguration	50
Contents	50
See Also	51
ElasticsearchDestinationDescription	53
Contents	53
See Also	54
ElasticsearchDestinationUpdate	55
Contents	55
See Also	56
ElasticsearchRetryOptions	57
Contents	57
See Also	57
EncryptionConfiguration	58
Contents	58
See Also	58
ExtendedS3DestinationConfiguration	59
Contents	59
See Also	60
ExtendedS3DestinationDescription	61
Contents	61
See Also	62
ExtendedS3DestinationUpdate	63
Contents	63
See Also	64
KinesisStreamSourceConfiguration	65
Contents	65
See Also	65
KinesisStreamSourceDescription	66
Contents	66
See Also	66
KMSEncryptionConfig	67
Contents	67
See Also	67
ProcessingConfiguration	68
Contents	68
See Also	68

Processor	69
Contents	69
See Also	69
ProcessorParameter	70
Contents	70
See Also	70
PutRecordBatchResponseEntry	71
Contents	71
See Also	71
Record	72
Contents	72
See Also	72
RedshiftDestinationConfiguration	73
Contents	73
See Also	74
RedshiftDestinationDescription	75
Contents	75
See Also	76
RedshiftDestinationUpdate	77
Contents	77
See Also	78
RedshiftRetryOptions	79
Contents	79
See Also	79
S3DestinationConfiguration	80
Contents	80
See Also	81
S3DestinationDescription	82
Contents	82
See Also	83
S3DestinationUpdate	84
Contents	84
See Also	85
SourceDescription	86
Contents	86
See Also	86
SplunkDestinationConfiguration	87
Contents	87
See Also	88
SplunkDestinationDescription	89
Contents	89
See Also	90
SplunkDestinationUpdate	91
Contents	91
See Also	92
SplunkRetryOptions	93
Contents	93
See Also	93
Common Errors	94

Welcome

Amazon Kinesis Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon Simple Storage Service (Amazon S3), Amazon Elasticsearch Service (Amazon ES), Amazon Redshift, and Splunk.

This document was last published on February 17, 2018.

Actions

The following actions are supported:

- [CreateDeliveryStream](#) (p. 3)
- [DeleteDeliveryStream](#) (p. 12)
- [DescribeDeliveryStream](#) (p. 14)
- [ListDeliveryStreams](#) (p. 22)
- [PutRecord](#) (p. 25)
- [PutRecordBatch](#) (p. 28)
- [UpdateDestination](#) (p. 32)

CreateDeliveryStream

Creates a Kinesis Data Firehose delivery stream.

By default, you can create up to 50 delivery streams per region.

This is an asynchronous operation that immediately returns. The initial status of the delivery stream is `CREATING`. After the delivery stream is created, its status is `ACTIVE` and it now accepts data. Attempts to send data to a delivery stream that is not in the `ACTIVE` state cause an exception. To check the state of a delivery stream, use [DescribeDeliveryStream](#) (p. 14).

A Kinesis Firehose delivery stream can be configured to receive records directly from providers using [PutRecord](#) (p. 25) or [PutRecordBatch](#) (p. 28), or it can be configured to use an existing Kinesis stream as its source. To specify a Kinesis stream as input, set the `DeliveryStreamType` parameter to `KinesisStreamAsSource`, and provide the Kinesis stream Amazon Resource Name (ARN) and role ARN in the `KinesisStreamSourceConfiguration` parameter.

A delivery stream is configured with a single destination: Amazon S3, Amazon ES, Amazon Redshift, or Splunk. You must specify only one of the following destination configuration parameters: **ExtendedS3DestinationConfiguration**, **S3DestinationConfiguration**, **ElasticsearchDestinationConfiguration**, **RedshiftDestinationConfiguration**, or **SplunkDestinationConfiguration**.

When you specify **S3DestinationConfiguration**, you can also provide the following optional values: **BufferingHints**, **EncryptionConfiguration**, and **CompressionFormat**. By default, if no **BufferingHints** value is provided, Kinesis Firehose buffers data up to 5 MB or for 5 minutes, whichever condition is satisfied first. Note that **BufferingHints** is a hint, so there are some cases where the service cannot adhere to these conditions strictly; for example, record boundaries are such that the size is a little over or under the configured buffering size. By default, no encryption is performed. We strongly recommend that you enable encryption to ensure secure data storage in Amazon S3.

A few notes about Amazon Redshift as a destination:

- An Amazon Redshift destination requires an S3 bucket as intermediate location, as Kinesis Firehose first delivers data to S3 and then uses `COPY` syntax to load data into an Amazon Redshift table. This is specified in the **RedshiftDestinationConfiguration.S3Configuration** parameter.
- The compression formats `SNAPPY` or `ZIP` cannot be specified in **RedshiftDestinationConfiguration.S3Configuration** because the Amazon Redshift `COPY` operation that reads from the S3 bucket doesn't support these compression formats.
- We strongly recommend that you use the user name and password you provide exclusively with Kinesis Firehose, and that the permissions for the account are restricted for Amazon Redshift `INSERT` permissions.

Kinesis Firehose assumes the IAM role that is configured as part of the destination. The role should allow the Kinesis Firehose principal to assume the role, and the role should have permissions that allow the service to deliver the data. For more information, see [Amazon S3 Bucket Access](#) in the *Amazon Kinesis Firehose Developer Guide*.

Request Syntax

```
{
  "DeliveryStreamName": "string",
  "DeliveryStreamType": "string",
  "ElasticsearchDestinationConfiguration": {
    "BufferingHints": {
      "IntervalInSeconds": number,
      "SizeInMBs": number
    }
  }
}
```



```

    },
    "CloudWatchLoggingOptions": {
      "Enabled": boolean,
      "LogGroupName": "string",
      "LogStreamName": "string"
    },
    "DomainARN": "string",
    "IndexName": "string",
    "IndexRotationPeriod": "string",
    "ProcessingConfiguration": {
      "Enabled": boolean,
      "Processors": [
        {
          "Parameters": [
            {
              "ParameterName": "string",
              "ParameterValue": "string"
            }
          ],
          "Type": "string"
        }
      ]
    },
    "RetryOptions": {
      "DurationInSeconds": number
    },
    "RoleARN": "string",
    "S3BackupMode": "string",
    "S3Configuration": {
      "BucketARN": "string",
      "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
      },
      "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
      },
      "CompressionFormat": "string",
      "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
          "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
      },
      "Prefix": "string",
      "RoleARN": "string"
    },
    "TypeName": "string"
  },
  "ExtendedS3DestinationConfiguration": {
    "BucketARN": "string",
    "BufferingHints": {
      "IntervalInSeconds": number,
      "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
      "Enabled": boolean,
      "LogGroupName": "string",
      "LogStreamName": "string"
    },
    "CompressionFormat": "string",
    "EncryptionConfiguration": {
      "KMSEncryptionConfig": {
        "AWSKMSKeyARN": "string"
      }
    }
  }
}

```

```

    },
    "NoEncryptionConfig": "string"
  },
  "Prefix": "string",
  "ProcessingConfiguration": {
    "Enabled": boolean,
    "Processors": [
      {
        "Parameters": [
          {
            "ParameterName": "string",
            "ParameterValue": "string"
          }
        ],
        "Type": "string"
      }
    ]
  },
  "RoleARN": "string",
  "S3BackupConfiguration": {
    "BucketARN": "string",
    "BufferingHints": {
      "IntervalInSeconds": number,
      "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
      "Enabled": boolean,
      "LogGroupName": "string",
      "LogStreamName": "string"
    },
    "CompressionFormat": "string",
    "EncryptionConfiguration": {
      "KMSEncryptionConfig": {
        "AWSKMSKeyARN": "string"
      },
      "NoEncryptionConfig": "string"
    },
    "Prefix": "string",
    "RoleARN": "string"
  },
  "S3BackupMode": "string"
},
"KinesisStreamSourceConfiguration": {
  "KinesisStreamARN": "string",
  "RoleARN": "string"
},
"RedshiftDestinationConfiguration": {
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "ClusterJDBCURL": "string",
  "CopyCommand": {
    "CopyOptions": "string",
    "DataTableColumns": "string",
    "DataTableName": "string"
  },
  "Password": "string",
  "ProcessingConfiguration": {
    "Enabled": boolean,
    "Processors": [
      {
        "Parameters": [
          {
            "ParameterName": "string",

```

```

        "ParameterValue": "string"
    }
  ],
  "Type": "string"
}
]
},
"RetryOptions": {
  "DurationInSeconds": number
},
"RoleARN": "string",
"S3BackupConfiguration": {
  "BucketARN": "string",
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "CompressionFormat": "string",
  "EncryptionConfiguration": {
    "KMSEncryptionConfig": {
      "AWSKMSKeyARN": "string"
    },
    "NoEncryptionConfig": "string"
  },
  "Prefix": "string",
  "RoleARN": "string"
},
"S3BackupMode": "string",
"S3Configuration": {
  "BucketARN": "string",
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "CompressionFormat": "string",
  "EncryptionConfiguration": {
    "KMSEncryptionConfig": {
      "AWSKMSKeyARN": "string"
    },
    "NoEncryptionConfig": "string"
  },
  "Prefix": "string",
  "RoleARN": "string"
},
"Username": "string"
},
"S3DestinationConfiguration": {
  "BucketARN": "string",
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  }
}

```

```

    },
    "CompressionFormat": "string",
    "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
            "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
    },
    "Prefix": "string",
    "RoleARN": "string"
},
"SplunkDestinationConfiguration": {
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "HECAcknowledgmentTimeoutInSeconds": number,
    "HECEndpoint": "string",
    "HECEndpointType": "string",
    "HECToken": "string",
    "ProcessingConfiguration": {
        "Enabled": boolean,
        "Processors": [
            {
                "Parameters": [
                    {
                        "ParameterName": "string",
                        "ParameterValue": "string"
                    }
                ],
                "Type": "string"
            }
        ]
    },
    "RetryOptions": {
        "DurationInSeconds": number
    },
    "S3BackupMode": "string",
    "S3Configuration": {
        "BucketARN": "string",
        "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
        },
        "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
        },
        "CompressionFormat": "string",
        "EncryptionConfiguration": {
            "KMSEncryptionConfig": {
                "AWSKMSKeyARN": "string"
            },
            "NoEncryptionConfig": "string"
        },
        "Prefix": "string",
        "RoleARN": "string"
    }
}
}

```

Request Parameters

The request accepts the following data in JSON format.

DeliveryStreamName (p. 3)

The name of the delivery stream. This name must be unique per AWS account in the same region. If the delivery streams are in different accounts or different regions, you can have multiple delivery streams with the same name.

Type: String

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

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

Required: Yes

DeliveryStreamType (p. 3)

The delivery stream type. This parameter can be one of the following values:

- `DirectPut`: Provider applications access the delivery stream directly.
- `KinesisStreamAsSource`: The delivery stream uses a Kinesis stream as a source.

Type: String

Valid Values: `DirectPut` | `KinesisStreamAsSource`

Required: No

ElasticsearchDestinationConfiguration (p. 3)

The destination in Amazon ES. You can specify only one destination.

Type: [ElasticsearchDestinationConfiguration \(p. 50\)](#) object

Required: No

ExtendedS3DestinationConfiguration (p. 3)

The destination in Amazon S3. You can specify only one destination.

Type: [ExtendedS3DestinationConfiguration \(p. 59\)](#) object

Required: No

KinesisStreamSourceConfiguration (p. 3)

When a Kinesis stream is used as the source for the delivery stream, a [KinesisStreamSourceConfiguration \(p. 65\)](#) containing the Kinesis stream Amazon Resource Name (ARN) and the role ARN for the source stream.

Type: [KinesisStreamSourceConfiguration \(p. 65\)](#) object

Required: No

RedshiftDestinationConfiguration (p. 3)

The destination in Amazon Redshift. You can specify only one destination.

Type: [RedshiftDestinationConfiguration \(p. 73\)](#) object

Required: No

[S3DestinationConfiguration \(p. 3\)](#)

[Deprecated] The destination in Amazon S3. You can specify only one destination.

Type: [S3DestinationConfiguration \(p. 80\)](#) object

Required: No

[SplunkDestinationConfiguration \(p. 3\)](#)

The destination in Splunk. You can specify only one destination.

Type: [SplunkDestinationConfiguration \(p. 87\)](#) object

Required: No

Response Syntax

```
{  
  "DeliveryStreamARN": "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.

[DeliveryStreamARN \(p. 9\)](#)

The ARN of the delivery stream.

Type: String

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

Pattern: arn:.*

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 94\)](#).

InvalidArgumentException

The specified input parameter has a value that is not valid.

HTTP Status Code: 400

LimitExceededException

You have already reached the limit for a requested resource.

HTTP Status Code: 400

ResourceInUseException

The resource is already in use and not available for this operation.

HTTP Status Code: 400

Example

The following JSON example creates a delivery stream named `exampleStreamName` with an S3 destination. To use this example, you must first replace the placeholders for the `RoleARN` and `BucketARN` keys with valid strings.

Sample Request

```
POST / HTTP/1.1
Host: firehose.<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: Firehose_20150804.CreateDeliveryStream
{
  "DeliveryStreamName": "exampleStreamName",
  "S3DestinationConfiguration": {
    "RoleARN": "insert-role-ARN",
    "BucketARN": "insert-bucket-ARN",
    "BufferingHints": {
      "SizeInMBs": 3,
      "IntervalInSeconds": 60
    },
    "CompressionFormat": "ZIP"
  }
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "DeliveryStreamARN": "arn:aws:firehose:us-east-1:814985986679:deliverystream/
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](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DeleteDeliveryStream

Deletes a delivery stream and its data.

You can delete a delivery stream only if it is in `ACTIVE` or `DELETING` state, and not in the `CREATING` state. While the deletion request is in process, the delivery stream is in the `DELETING` state.

To check the state of a delivery stream, use [DescribeDeliveryStream \(p. 14\)](#).

While the delivery stream is `DELETING` state, the service may continue to accept the records, but the service doesn't make any guarantees with respect to delivering the data. Therefore, as a best practice, you should first stop any applications that are sending records before deleting a delivery stream.

Request Syntax

```
{  
  "DeliveryStreamName": "string"  
}
```

Request Parameters

The request accepts the following data in JSON format.

DeliveryStreamName (p. 12)

The name of the delivery stream.

Type: String

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

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

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 Errors \(p. 94\)](#).

ResourceInUseException

The resource is already in use and not available for this operation.

HTTP Status Code: 400

ResourceNotFoundException

The specified resource could not be found.

HTTP Status Code: 400

Example

The following JSON example deletes a delivery stream named `exampleStreamName`.

Sample Request

```
POST / HTTP/1.1
Host: firehose.<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: Firehose_20150804.DeleteDeliveryStream
{
  "DeliveryStreamName": "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](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

DescribeDeliveryStream

Describes the specified delivery stream and gets the status. For example, after your delivery stream is created, call [DescribeDeliveryStream \(p. 14\)](#) to see if the delivery stream is `ACTIVE` and therefore ready for data to be sent to it.

Request Syntax

```
{
  "DeliveryStreamName": "string",
  "ExclusiveStartDestinationId": "string",
  "Limit": number
}
```

Request Parameters

The request accepts the following data in JSON format.

[DeliveryStreamName \(p. 14\)](#)

The name of the delivery stream.

Type: String

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

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

Required: Yes

[ExclusiveStartDestinationId \(p. 14\)](#)

The ID of the destination to start returning the destination information. Currently, Kinesis Firehose currently supports one destination per delivery stream.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Required: No

[Limit \(p. 14\)](#)

The limit on the number of destinations to return. Currently, you can have one destination per delivery stream.

Type: Integer

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

Required: No

Response Syntax

```
{
  "DeliveryStreamDescription": {
```

```

"CreateTimestamp": number,
"DeliveryStreamARN": "string",
"DeliveryStreamName": "string",
"DeliveryStreamStatus": "string",
"DeliveryStreamType": "string",
"Destinations": [
  {
    "DestinationId": "string",
    "ElasticsearchDestinationDescription": {
      "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
      },
      "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
      },
      "DomainARN": "string",
      "IndexName": "string",
      "IndexRotationPeriod": "string",
      "ProcessingConfiguration": {
        "Enabled": boolean,
        "Processors": [
          {
            "Parameters": [
              {
                "ParameterName": "string",
                "ParameterValue": "string"
              }
            ],
            "Type": "string"
          }
        ]
      }
    },
    "RetryOptions": {
      "DurationInSeconds": number
    },
    "RoleARN": "string",
    "S3BackupMode": "string",
    "S3DestinationDescription": {
      "BucketARN": "string",
      "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
      },
      "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
      },
      "CompressionFormat": "string",
      "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
          "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
      },
      "Prefix": "string",
      "RoleARN": "string"
    },
    "TypeName": "string"
  },
  "ExtendedS3DestinationDescription": {
    "BucketARN": "string",
    "BufferingHints": {

```

```

        "IntervalInSeconds": number,
        "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": string,
        "LogStreamName": string
    },
    "CompressionFormat": string,
    "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
            "AWSKMSKeyARN": string
        },
        "NoEncryptionConfig": string
    },
    "Prefix": string,
    "ProcessingConfiguration": {
        "Enabled": boolean,
        "Processors": [
            {
                "Parameters": [
                    {
                        "ParameterName": string,
                        "ParameterValue": string
                    }
                ],
                "Type": string
            }
        ]
    },
    "RoleARN": string,
    "S3BackupDescription": {
        "BucketARN": string,
        "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
        },
        "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": string,
            "LogStreamName": string
        },
        "CompressionFormat": string,
        "EncryptionConfiguration": {
            "KMSEncryptionConfig": {
                "AWSKMSKeyARN": string
            },
            "NoEncryptionConfig": string
        },
        "Prefix": string,
        "RoleARN": string
    },
    "S3BackupMode": string
},
"RedshiftDestinationDescription": {
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": string,
        "LogStreamName": string
    },
    "ClusterJDBCURL": string,
    "CopyCommand": {
        "CopyOptions": string,
        "DataTableColumns": string,
        "DataTableName": string
    }
},

```

```

"ProcessingConfiguration": {
  "Enabled": boolean,
  "Processors": [
    {
      "Parameters": [
        {
          "ParameterName": "string",
          "ParameterValue": "string"
        }
      ],
      "Type": "string"
    }
  ]
},
"RetryOptions": {
  "DurationInSeconds": number
},
"RoleARN": "string",
"S3BackupDescription": {
  "BucketARN": "string",
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "CompressionFormat": "string",
  "EncryptionConfiguration": {
    "KMSEncryptionConfig": {
      "AWSKMSKeyARN": "string"
    },
    "NoEncryptionConfig": "string"
  },
  "Prefix": "string",
  "RoleARN": "string"
},
"S3BackupMode": "string",
"S3DestinationDescription": {
  "BucketARN": "string",
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "CompressionFormat": "string",
  "EncryptionConfiguration": {
    "KMSEncryptionConfig": {
      "AWSKMSKeyARN": "string"
    },
    "NoEncryptionConfig": "string"
  },
  "Prefix": "string",
  "RoleARN": "string"
},
"Username": "string"
},
"S3DestinationDescription": {
  "BucketARN": "string",
  "BufferingHints": {

```

```

        "IntervalInSeconds": number,
        "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "CompressionFormat": "string",
    "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
            "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
    },
    "Prefix": "string",
    "RoleARN": "string"
},
"SplunkDestinationDescription": {
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "HECAcknowledgmentTimeoutInSeconds": number,
    "HECEndpoint": "string",
    "HECEndpointType": "string",
    "HECToken": "string",
    "ProcessingConfiguration": {
        "Enabled": boolean,
        "Processors": [
            {
                "Parameters": [
                    {
                        "ParameterName": "string",
                        "ParameterValue": "string"
                    }
                ],
                "Type": "string"
            }
        ]
    },
    "RetryOptions": {
        "DurationInSeconds": number
    },
    "S3BackupMode": "string",
    "S3DestinationDescription": {
        "BucketARN": "string",
        "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
        },
        "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
        },
        "CompressionFormat": "string",
        "EncryptionConfiguration": {
            "KMSEncryptionConfig": {
                "AWSKMSKeyARN": "string"
            },
            "NoEncryptionConfig": "string"
        },
        "Prefix": "string",
        "RoleARN": "string"
    }
}

```

```

    }
  }
],
"HasMoreDestinations": boolean,
"LastUpdateTimestamp": number,
"Source": {
  "KinesisStreamSourceDescription": {
    "DeliveryStartTimestamp": number,
    "KinesisStreamARN": "string",
    "RoleARN": "string"
  }
},
"VersionId": "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.

DeliveryStreamDescription (p. 14)

Information about the delivery stream.

Type: [DeliveryStreamDescription \(p. 45\)](#) object

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 94\)](#).

ResourceNotFoundException

The specified resource could not be found.

HTTP Status Code: 400

Example

The following JSON example describes a delivery stream.

Sample Request

```

POST / HTTP/1.1
Host: firehose.<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: Firehose_20150804.DescribeDeliveryStream
{
  "DeliveryStreamName": "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>
{
  "DeliveryStreamDescription": {
    "DeliveryStreamType": "DirectPut",
    "HasMoreDestinations": false,
    "VersionId": "1",
    "CreateTimestamp": 1517595920.596,
    "DeliveryStreamARN": "arn:aws:firehose:us-east-1:111222333444:deliverystream/
exampleStreamName",
    "DeliveryStreamStatus": "ACTIVE",
    "DeliveryStreamName": "exampleStreamName",
    "Destinations": [
      {
        "DestinationId": "destinationId-000000000001",
        "ExtendedS3DestinationDescription": {
          "RoleARN": "arn:aws:iam::111222333444:role/exampleStreamName",
          "Prefix": "",
          "BufferingHints": {
            "IntervalInSeconds": 60,
            "SizeInMBs": 1
          },
          "EncryptionConfiguration": {
            "NoEncryptionConfig": "NoEncryption"
          },
          "CompressionFormat": "UNCOMPRESSED",
          "S3BackupMode": "Disabled",
          "CloudWatchLoggingOptions": {
            "Enabled": true,
            "LogStreamName": "S3Delivery",
            "LogGroupName": "/aws/kinesisfirehose/exampleStreamName"
          },
          "BucketARN": "arn:aws:s3:::somebucket",
          "ProcessingConfiguration": {
            "Enabled": false,
            "Processors": []
          }
        },
        "S3DestinationDescription": {
          "RoleARN": "arn:aws:iam::111222333444:role/exampleStreamName",
          "Prefix": "",
          "BufferingHints": {
            "IntervalInSeconds": 60,
            "SizeInMBs": 1
          },
          "EncryptionConfiguration": {
            "NoEncryptionConfig": "NoEncryption"
          },
          "CompressionFormat": "UNCOMPRESSED",
          "CloudWatchLoggingOptions": {
            "Enabled": true,
            "LogStreamName": "S3Delivery",
            "LogGroupName": "/aws/kinesisfirehose/exampleStreamName"
          },
          "BucketARN": "arn:aws:s3:::somebucket"
        }
      }
    ]
  }
}
```

```
}  
  }  
}
```

See Also

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

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

ListDeliveryStreams

Lists your delivery streams.

The number of delivery streams might be too large to return using a single call to [ListDeliveryStreams](#) (p. 22). You can limit the number of delivery streams returned, using the **Limit** parameter. To determine whether there are more delivery streams to list, check the value of **HasMoreDeliveryStreams** in the output. If there are more delivery streams to list, you can request them by specifying the name of the last delivery stream returned in the call in the **ExclusiveStartDeliveryStreamName** parameter of a subsequent call.

Request Syntax

```
{
  "DeliveryStreamType": "string",
  "ExclusiveStartDeliveryStreamName": "string",
  "Limit": number
}
```

Request Parameters

The request accepts the following data in JSON format.

DeliveryStreamType (p. 22)

The delivery stream type. This can be one of the following values:

- **DirectPut**: Provider applications access the delivery stream directly.
- **KinesisStreamAsSource**: The delivery stream uses a Kinesis stream as a source.

This parameter is optional. If this parameter is omitted, delivery streams of all types are returned.

Type: String

Valid Values: `DirectPut` | `KinesisStreamAsSource`

Required: No

ExclusiveStartDeliveryStreamName (p. 22)

The name of the delivery stream to start the list with.

Type: String

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

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

Required: No

Limit (p. 22)

The maximum number of delivery streams to list. The default value is 10.

Type: Integer

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

Required: No

Response Syntax

```
{  
  "DeliveryStreamNames": [ "string" ],  
  "HasMoreDeliveryStreams": boolean  
}
```

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.

DeliveryStreamNames (p. 23)

The names of the delivery streams.

Type: Array of strings

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

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

HasMoreDeliveryStreams (p. 23)

Indicates whether there are more delivery streams available to list.

Type: Boolean

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 94\)](#).

Example

The following JSON example lists up to 3 delivery streams configured for `DirectPut`. The response indicates that there are more `DirectPut` delivery streams to be listed. To list the remaining `DirectPut` delivery streams, set `DeliveryStreamType` to `DirectPut` and `ExclusiveStartDeliveryStreamName` to `last_stream_in_first_listing`, and then run the operation again.

Sample Request

```
POST / HTTP/1.1  
Host: firehose.<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: Firehose_20150804.ListDeliveryStreams  
{  
  "DeliveryStreamType": "DirectPut",  
  "Limit": 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>
{
  "DeliveryStreamNames": [
    "some_delivery_stream",
    "another_example_delivery_stream",
    "last_stream_in_first_listing"
  ],
  "HasMoreDeliveryStreams": true
}
```

See Also

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

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

PutRecord

Writes a single data record into an Amazon Kinesis Firehose delivery stream. To write multiple data records into a delivery stream, use [PutRecordBatch \(p. 28\)](#). Applications using these operations are referred to as producers.

By default, each delivery stream can take in up to 2,000 transactions per second, 5,000 records per second, or 5 MB per second. Note that if you use [PutRecord \(p. 25\)](#) and [PutRecordBatch \(p. 28\)](#), the limits are an aggregate across these two operations for each delivery stream. For more information about limits and how to request an increase, see [Amazon Kinesis Firehose Limits](#).

You must specify the name of the delivery stream and the data record when using [PutRecord \(p. 25\)](#). The data record consists of a data blob that can be up to 1,000 KB in size, and any kind of data, for example, a segment from a log file, geographic location data, website clickstream data, and so on.

Kinesis Firehose buffers records before delivering them to the destination. To disambiguate the data blobs at the destination, a common solution is to use delimiters in the data, such as a newline (`\n`) or some other character unique within the data. This allows the consumer application to parse individual data items when reading the data from the destination.

The [PutRecord \(p. 25\)](#) operation returns a **RecordId**, which is a unique string assigned to each record. Producer applications can use this ID for purposes such as auditability and investigation.

If the [PutRecord \(p. 25\)](#) operation throws a **ServiceUnavailableException**, back off and retry. If the exception persists, it is possible that the throughput limits have been exceeded for the delivery stream.

Data records sent to Kinesis Firehose are stored for 24 hours from the time they are added to a delivery stream as it attempts to send the records to the destination. If the destination is unreachable for more than 24 hours, the data is no longer available.

Request Syntax

```
{
  "DeliveryStreamName": "string",
  "Record": {
    "Data": blob
  }
}
```

Request Parameters

The request accepts the following data in JSON format.

[DeliveryStreamName \(p. 25\)](#)

The name of the delivery stream.

Type: String

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

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

Required: Yes

[Record \(p. 25\)](#)

The record.

Type: [Record \(p. 72\)](#) object

Required: Yes

Response Syntax

```
{  
  "RecordId": "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.

RecordId (p. 26)

The ID of the record.

Type: String

Length Constraints: Minimum length of 1.

Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 94\)](#).

InvalidArgumentException

The specified input parameter has a value that is not valid.

HTTP Status Code: 400

ResourceNotFoundException

The specified resource could not be found.

HTTP Status Code: 400

ServiceUnavailableException

The service is unavailable. Back off and retry the operation. If you continue to see the exception, throughput limits for the delivery stream may have been exceeded. For more information about limits and how to request an increase, see [Amazon Kinesis Firehose Limits](#).

HTTP Status Code: 500

Example

The following JSON puts a record in the delivery stream named `some_delivery_stream`.

Sample Request

```
POST / HTTP/1.1
```

```
Host: firehose.<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: Firehose_20150804.PutRecord
{
  "DeliveryStreamName": "some_delivery_stream",
  "Record": {
    "Data": "..."
  }
}
```

Sample Response

```
HTTP/1.1 200 OK
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "RecordId": "CGojNMJq3msHbGoc+1mgSpifCmfFm71Fhuts//4Ft6sFVokyE6t+5ioEAjNm+sgQ6iVf/
YePEXBK6epIW4QeXqJp2xsbfZUNXsfOY1QrYXgRBCKznkjMMTP0BqJGOBm3fB//dHgEE0XDtc4wW065i/
tJyYI1Vy8qn8FMhpkZuh5bvG482XkkBxFmMGnhPTQwQ4A1IOP0sEOX99YnBK8RECdeQ2zxynvZ"
}
```

See Also

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

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

PutRecordBatch

Writes multiple data records into a delivery stream in a single call, which can achieve higher throughput per producer than when writing single records. To write single data records into a delivery stream, use [PutRecord \(p. 25\)](#). Applications using these operations are referred to as producers.

By default, each delivery stream can take in up to 2,000 transactions per second, 5,000 records per second, or 5 MB per second. If you use [PutRecord \(p. 25\)](#) and [PutRecordBatch \(p. 28\)](#), the limits are an aggregate across these two operations for each delivery stream. For more information about limits, see [Amazon Kinesis Firehose Limits](#).

Each [PutRecordBatch \(p. 28\)](#) request supports up to 500 records. Each record in the request can be as large as 1,000 KB (before 64-bit encoding), up to a limit of 4 MB for the entire request. These limits cannot be changed.

You must specify the name of the delivery stream and the data record when using [PutRecord \(p. 25\)](#). The data record consists of a data blob that can be up to 1,000 KB in size, and any kind of data. For example, it could be a segment from a log file, geographic location data, web site clickstream data, and so on.

Kinesis Firehose buffers records before delivering them to the destination. To disambiguate the data blobs at the destination, a common solution is to use delimiters in the data, such as a newline (`\n`) or some other character unique within the data. This allows the consumer application to parse individual data items when reading the data from the destination.

The [PutRecordBatch \(p. 28\)](#) response includes a count of failed records, **FailedPutCount**, and an array of responses, **RequestResponses**. Each entry in the **RequestResponses** array provides additional information about the processed record. It directly correlates with a record in the request array using the same ordering, from the top to the bottom. The response array always includes the same number of records as the request array. **RequestResponses** includes both successfully and unsuccessfully processed records. Kinesis Firehose attempts to process all records in each [PutRecordBatch \(p. 28\)](#) request. A single record failure does not stop the processing of subsequent records.

A successfully processed record includes a **RecordId** value, which is unique for the record. An unsuccessfully processed record includes **ErrorCode** and **ErrorMessage** values. **ErrorCode** reflects the type of error, and is one of the following values: `ServiceUnavailable` or `InternalFailure`. **ErrorMessage** provides more detailed information about the error.

If there is an internal server error or a timeout, the write might have completed or it might have failed. If **FailedPutCount** is greater than 0, retry the request, resending only those records that might have failed processing. This minimizes the possible duplicate records and also reduces the total bytes sent (and corresponding charges). We recommend that you handle any duplicates at the destination.

If [PutRecordBatch \(p. 28\)](#) throws **ServiceUnavailableException**, back off and retry. If the exception persists, it is possible that the throughput limits have been exceeded for the delivery stream.

Data records sent to Kinesis Firehose are stored for 24 hours from the time they are added to a delivery stream as it attempts to send the records to the destination. If the destination is unreachable for more than 24 hours, the data is no longer available.

Request Syntax

```
{
  "DeliveryStreamName": "string",
  "Records": [
    {
      "Data": blob
    }
  ]
}
```

```
}
```

Request Parameters

The request accepts the following data in JSON format.

[DeliveryStreamName \(p. 28\)](#)

The name of the delivery stream.

Type: String

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

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

Required: Yes

[Records \(p. 28\)](#)

One or more records.

Type: Array of [Record \(p. 72\)](#) objects

Array Members: Minimum number of 1 item. Maximum number of 500 items.

Required: Yes

Response Syntax

```
{  
  "FailedPutCount": number,  
  "RequestResponses": [  
    {  
      "ErrorCode": "string",  
      "ErrorMessage": "string",  
      "RecordId": "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.

[FailedPutCount \(p. 29\)](#)

The number of records that might have failed processing.

Type: Integer

Valid Range: Minimum value of 0.

[RequestResponses \(p. 29\)](#)

The results array. For each record, the index of the response element is the same as the index used in the request array.

Type: Array of [PutRecordBatchResponseEntry \(p. 71\)](#) 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 Errors \(p. 94\)](#).

InvalidArgumentException

The specified input parameter has a value that is not valid.

HTTP Status Code: 400

ResourceNotFoundException

The specified resource could not be found.

HTTP Status Code: 400

ServiceUnavailableException

The service is unavailable. Back off and retry the operation. If you continue to see the exception, throughput limits for the delivery stream may have been exceeded. For more information about limits and how to request an increase, see [Amazon Kinesis Firehose Limits](#).

HTTP Status Code: 500

Example

The following JSON puts two records in the delivery stream named `some_delivery_stream`.

Sample Request

```
POST / HTTP/1.1
Host: firehose.<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: Firehose_20150804.PutRecordBatch
{
  "DeliveryStreamName": "some_delivery_stream",
  "Records": [
    {
      "Data": "Some data blob."
    },
    {
      "Data": "Another blob of data."
    }
  ]
}
```

Sample Response

```
HTTP/1.1 200 OK
```

```
x-amzn-RequestId: <RequestId>
Content-Type: application/x-amz-json-1.1
Content-Length: <PayloadSizeBytes>
Date: <Date>
{
  "FailedPutCount": 0,
  "RequestResponses": [
    {
      "RecordId": "AJJBAL1fiFN9HyhPj6Dc+XqcR1TjgyIbr927TsEmWWpN39EK/
JbRTbXDZFNCCrWIs/4YUd3gjYUs6giZn76TuI4vv0ljDOMFJvLzqLuupu93RscZRmpW1CP8DeifzJjGvqIr1LRE/
MDozYenOz+v1ZiQeJECvfMwIz6silvDaGuYtIhVzwd8yWaDGFCX4ODes6W07W9W7Q5ViMPUKSt6F0nn7OhVkJ3/"
    },
    {
      "RecordId": "goGaFS919Mmv7lYET0oMaw+UL9iFpzi100o
+csoIc31SmvkqpWzQuT0RPqZ7QqfR1FJ
+HxJciW/8paFMWPByJ6qVDhiE7TtJxARKaP4YDccvWHXhD7x6Y4bc9AHZ0uzy
+BEuTTnH5Zsip1IrEPgu8lj8a7fd6sluTe/kZ54GVG+Sc105IrxIu08kPbriUTWL
+Wg0A3HzOQeePqokRKKqmoXLG0pzx/80"
    }
  ]
}
```

See Also

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

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

UpdateDestination

Updates the specified destination of the specified delivery stream.

You can use this operation to change the destination type (for example, to replace the Amazon S3 destination with Amazon Redshift) or change the parameters associated with a destination (for example, to change the bucket name of the Amazon S3 destination). The update might not occur immediately. The target delivery stream remains active while the configurations are updated, so data writes to the delivery stream can continue during this process. The updated configurations are usually effective within a few minutes.

Note that switching between Amazon ES and other services is not supported. For an Amazon ES destination, you can only update to another Amazon ES destination.

If the destination type is the same, Kinesis Firehose merges the configuration parameters specified with the destination configuration that already exists on the delivery stream. If any of the parameters are not specified in the call, the existing values are retained. For example, in the Amazon S3 destination, if [EncryptionConfiguration \(p. 58\)](#) is not specified, then the existing [EncryptionConfiguration \(p. 58\)](#) is maintained on the destination.

If the destination type is not the same, for example, changing the destination from Amazon S3 to Amazon Redshift, Kinesis Firehose does not merge any parameters. In this case, all parameters must be specified.

Kinesis Firehose uses **CurrentDeliveryStreamVersionId** to avoid race conditions and conflicting merges. This is a required field, and the service updates the configuration only if the existing configuration has a version ID that matches. After the update is applied successfully, the version ID is updated, and can be retrieved using [DescribeDeliveryStream \(p. 14\)](#). Use the new version ID to set **CurrentDeliveryStreamVersionId** in the next call.

Request Syntax

```
{
  "CurrentDeliveryStreamVersionId": "string",
  "DeliveryStreamName": "string",
  "DestinationId": "string",
  "ElasticsearchDestinationUpdate": {
    "BufferingHints": {
      "IntervalInSeconds": number,
      "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
      "Enabled": boolean,
      "LogGroupName": "string",
      "LogStreamName": "string"
    },
    "DomainARN": "string",
    "IndexName": "string",
    "IndexRotationPeriod": "string",
    "ProcessingConfiguration": {
      "Enabled": boolean,
      "Processors": [
        {
          "Parameters": [
            {
              "ParameterName": "string",
              "ParameterValue": "string"
            }
          ]
        }
      ],
      "Type": "string"
    }
  }
}
```

```

    }
  ]
},
"RetryOptions": {
  "DurationInSeconds": number
},
"RoleARN": "string",
"S3Update": {
  "BucketARN": "string",
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "CompressionFormat": "string",
  "EncryptionConfiguration": {
    "KMSEncryptionConfig": {
      "AWSKMSKeyARN": "string"
    },
    "NoEncryptionConfig": "string"
  },
  "Prefix": "string",
  "RoleARN": "string"
},
"TypeName": "string"
},
"ExtendedS3DestinationUpdate": {
  "BucketARN": "string",
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "CompressionFormat": "string",
  "EncryptionConfiguration": {
    "KMSEncryptionConfig": {
      "AWSKMSKeyARN": "string"
    },
    "NoEncryptionConfig": "string"
  },
  "Prefix": "string",
  "ProcessingConfiguration": {
    "Enabled": boolean,
    "Processors": [
      {
        "Parameters": [
          {
            "ParameterName": "string",
            "ParameterValue": "string"
          }
        ],
        "Type": "string"
      }
    ]
  },
  "RoleARN": "string",
  "S3BackupMode": "string",
  "S3BackupUpdate": {

```

```

    "BucketARN": "string",
    "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "CompressionFormat": "string",
    "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
            "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
    },
    "Prefix": "string",
    "RoleARN": "string"
}
},
"RedshiftDestinationUpdate": {
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "ClusterJDBCURL": "string",
    "CopyCommand": {
        "CopyOptions": "string",
        "DataTableColumns": "string",
        "DataTableName": "string"
    },
    "Password": "string",
    "ProcessingConfiguration": {
        "Enabled": boolean,
        "Processors": [
            {
                "Parameters": [
                    {
                        "ParameterName": "string",
                        "ParameterValue": "string"
                    }
                ]
            },
            {
                "Type": "string"
            }
        ]
    }
},
"RetryOptions": {
    "DurationInSeconds": number
},
"RoleARN": "string",
"S3BackupMode": "string",
"S3BackupUpdate": {
    "BucketARN": "string",
    "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "CompressionFormat": "string",
    "EncryptionConfiguration": {

```

```

        "KMSEncryptionConfig": {
            "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
    },
    "Prefix": "string",
    "RoleARN": "string"
},
"S3Update": {
    "BucketARN": "string",
    "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "CompressionFormat": "string",
    "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
            "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
    },
    "Prefix": "string",
    "RoleARN": "string"
},
"Username": "string"
},
"S3DestinationUpdate": {
    "BucketARN": "string",
    "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "CompressionFormat": "string",
    "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
            "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
    },
    "Prefix": "string",
    "RoleARN": "string"
},
"SplunkDestinationUpdate": {
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "HECAcknowledgmentTimeoutInSeconds": number,
    "HECEndpoint": "string",
    "HECEndpointType": "string",
    "HECToken": "string",
    "ProcessingConfiguration": {
        "Enabled": boolean,
        "Processors": [
            {

```



```

        "Parameters": [
            {
                "ParameterName": "string",
                "ParameterValue": "string"
            }
        ],
        "Type": "string"
    }
]
},
"RetryOptions": {
    "DurationInSeconds": number
},
"S3BackupMode": "string",
"S3Update": {
    "BucketARN": "string",
    "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "CompressionFormat": "string",
    "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
            "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
    },
    "Prefix": "string",
    "RoleARN": "string"
}
}
}

```

Request Parameters

The request accepts the following data in JSON format.

CurrentDeliveryStreamVersionId (p. 32)

Obtain this value from the **VersionId** result of [DeliveryStreamDescription \(p. 45\)](#). This value is required, and helps the service to perform conditional operations. For example, if there is an interleaving update and this value is null, then the update destination fails. After the update is successful, the **VersionId** value is updated. The service then performs a merge of the old configuration with the new configuration.

Type: String

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

Pattern: [0-9]+

Required: Yes

DeliveryStreamName (p. 32)

The name of the delivery stream.

Type: String

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

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

Required: Yes

DestinationId (p. 32)

The ID of the destination.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Required: Yes

ElasticsearchDestinationUpdate (p. 32)

Describes an update for a destination in Amazon ES.

Type: [ElasticsearchDestinationUpdate \(p. 55\)](#) object

Required: No

ExtendedS3DestinationUpdate (p. 32)

Describes an update for a destination in Amazon S3.

Type: [ExtendedS3DestinationUpdate \(p. 63\)](#) object

Required: No

RedshiftDestinationUpdate (p. 32)

Describes an update for a destination in Amazon Redshift.

Type: [RedshiftDestinationUpdate \(p. 77\)](#) object

Required: No

S3DestinationUpdate (p. 32)

[Deprecated] Describes an update for a destination in Amazon S3.

Type: [S3DestinationUpdate \(p. 84\)](#) object

Required: No

SplunkDestinationUpdate (p. 32)

Describes an update for a destination in Splunk.

Type: [SplunkDestinationUpdate \(p. 91\)](#) object

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 Errors \(p. 94\)](#).

ConcurrentModificationException

Another modification has already happened. Fetch **VersionId** again and use it to update the destination.

HTTP Status Code: 400

InvalidArgumentException

The specified input parameter has a value that is not valid.

HTTP Status Code: 400

ResourceInUseException

The resource is already in use and not available for this operation.

HTTP Status Code: 400

ResourceNotFoundException

The specified resource could not be found.

HTTP Status Code: 400

Example

The following JSON updates the buffering hints for the destination to 150 seconds and 2 MiB. In this example, you can obtain the other values you need to use in the following JSON from the response to a [DescribeDeliveryStream \(p. 14\)](#) invocation for the delivery stream in question here.

Sample Request

```
POST / HTTP/1.1
Host: firehose.<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: Firehose_20150804.UpdateDestination
{
  "CurrentDeliveryStreamVersionId": "1",
  "DeliveryStreamName": "exampleStreamName",
  "DestinationId": "destinationId-000000000001",
  "ExtendedS3DestinationUpdate": {
    "BucketARN": "arn:aws:s3:::somebucket",
    "BufferingHints": {
      "IntervalInSeconds": 150,
      "SizeInMBs": 2
    },
    "RoleARN": "arn:aws:iam::111222333444:role/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](#)
- [AWS SDK for .NET](#)
- [AWS SDK for C++](#)
- [AWS SDK for Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for JavaScript](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V2](#)

Data Types

The Amazon Kinesis Firehose 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:

- [BufferingHints](#) (p. 42)
- [CloudWatchLoggingOptions](#) (p. 43)
- [CopyCommand](#) (p. 44)
- [DeliveryStreamDescription](#) (p. 45)
- [DestinationDescription](#) (p. 47)
- [ElasticsearchBufferingHints](#) (p. 49)
- [ElasticsearchDestinationConfiguration](#) (p. 50)
- [ElasticsearchDestinationDescription](#) (p. 53)
- [ElasticsearchDestinationUpdate](#) (p. 55)
- [ElasticsearchRetryOptions](#) (p. 57)
- [EncryptionConfiguration](#) (p. 58)
- [ExtendedS3DestinationConfiguration](#) (p. 59)
- [ExtendedS3DestinationDescription](#) (p. 61)
- [ExtendedS3DestinationUpdate](#) (p. 63)
- [KinesisStreamSourceConfiguration](#) (p. 65)
- [KinesisStreamSourceDescription](#) (p. 66)
- [KMSEncryptionConfig](#) (p. 67)
- [ProcessingConfiguration](#) (p. 68)
- [Processor](#) (p. 69)
- [ProcessorParameter](#) (p. 70)
- [PutRecordBatchResponseEntry](#) (p. 71)
- [Record](#) (p. 72)
- [RedshiftDestinationConfiguration](#) (p. 73)
- [RedshiftDestinationDescription](#) (p. 75)
- [RedshiftDestinationUpdate](#) (p. 77)
- [RedshiftRetryOptions](#) (p. 79)
- [S3DestinationConfiguration](#) (p. 80)
- [S3DestinationDescription](#) (p. 82)
- [S3DestinationUpdate](#) (p. 84)
- [SourceDescription](#) (p. 86)
- [SplunkDestinationConfiguration](#) (p. 87)
- [SplunkDestinationDescription](#) (p. 89)
- [SplunkDestinationUpdate](#) (p. 91)
- [SplunkRetryOptions](#) (p. 93)

BufferingHints

Describes hints for the buffering to perform before delivering data to the destination. Please note that these options are treated as hints, and therefore Kinesis Firehose may choose to use different values when it is optimal.

Contents

IntervalInSeconds

Buffer incoming data for the specified period of time, in seconds, before delivering it to the destination. The default value is 300.

Type: Integer

Valid Range: Minimum value of 60. Maximum value of 900.

Required: No

SizeInMBs

Buffer incoming data to the specified size, in MBs, before delivering it to the destination. The default value is 5.

We recommend setting this parameter to a value greater than the amount of data you typically ingest into the delivery stream in 10 seconds. For example, if you typically ingest data at 1 MB/sec, the value should be 10 MB or higher.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 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 Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

CloudWatchLoggingOptions

Describes the Amazon CloudWatch logging options for your delivery stream.

Contents

Enabled

Enables or disables CloudWatch logging.

Type: Boolean

Required: No

LogGroupName

The CloudWatch group name for logging. This value is required if CloudWatch logging is enabled.

Type: String

Required: No

LogStreamName

The CloudWatch log stream name for logging. This value is required if CloudWatch logging is enabled.

Type: String

Required: No

See Also

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

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

CopyCommand

Describes a COPY command for Amazon Redshift.

Contents

CopyOptions

Optional parameters to use with the Amazon Redshift COPY command. For more information, see the "Optional Parameters" section of [Amazon Redshift COPY command](#). Some possible examples that would apply to Kinesis Firehose are as follows:

`delimiter '\t' lzop;` - fields are delimited with "\t" (TAB character) and compressed using lzop.

`delimiter '|' -` fields are delimited with "|" (this is the default delimiter).

`delimiter '|' escape -` the delimiter should be escaped.

`fixedwidth 'venueid:3,venueid:25,venueid:12,venueid:2,venueid:6'` - fields are fixed width in the source, with each width specified after every column in the table.

`JSON 's3://mybucket/jsonpaths.txt'` - data is in JSON format, and the path specified is the format of the data.

For more examples, see [Amazon Redshift COPY command examples](#).

Type: String

Required: No

DataTableColumns

A comma-separated list of column names.

Type: String

Required: No

DataTableName

The name of the target table. The table must already exist in the database.

Type: String

Length Constraints: Minimum length of 1.

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

DeliveryStreamDescription

Contains information about a delivery stream.

Contents

CreateTimestamp

The date and time that the delivery stream was created.

Type: Timestamp

Required: No

DeliveryStreamARN

The Amazon Resource Name (ARN) of the delivery stream.

Type: String

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

Pattern: `arn:.*`

Required: Yes

DeliveryStreamName

The name of the delivery stream.

Type: String

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

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

Required: Yes

DeliveryStreamStatus

The status of the delivery stream.

Type: String

Valid Values: `CREATING` | `DELETING` | `ACTIVE`

Required: Yes

DeliveryStreamType

The delivery stream type. This can be one of the following values:

- `DirectPut`: Provider applications access the delivery stream directly.
- `KinesisStreamAsSource`: The delivery stream uses a Kinesis stream as a source.

Type: String

Valid Values: `DirectPut` | `KinesisStreamAsSource`

Required: Yes

Destinations

The destinations.

Type: Array of [DestinationDescription \(p. 47\)](#) objects

Required: Yes

HasMoreDestinations

Indicates whether there are more destinations available to list.

Type: Boolean

Required: Yes

LastUpdateTimestamp

The date and time that the delivery stream was last updated.

Type: Timestamp

Required: No

Source

If the `DeliveryStreamType` parameter is `KinesisStreamAsSource`, a [SourceDescription \(p. 86\)](#) object describing the source Kinesis stream.

Type: [SourceDescription \(p. 86\)](#) object

Required: No

VersionId

Each time the destination is updated for a delivery stream, the version ID is changed, and the current version ID is required when updating the destination. This is so that the service knows it is applying the changes to the correct version of the delivery stream.

Type: String

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

Pattern: [0-9]+

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

DestinationDescription

Describes the destination for a delivery stream.

Contents

DestinationId

The ID of the destination.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

Required: Yes

ElasticsearchDestinationDescription

The destination in Amazon ES.

Type: [ElasticsearchDestinationDescription \(p. 53\)](#) object

Required: No

ExtendedS3DestinationDescription

The destination in Amazon S3.

Type: [ExtendedS3DestinationDescription \(p. 61\)](#) object

Required: No

RedshiftDestinationDescription

The destination in Amazon Redshift.

Type: [RedshiftDestinationDescription \(p. 75\)](#) object

Required: No

S3DestinationDescription

[Deprecated] The destination in Amazon S3.

Type: [S3DestinationDescription \(p. 82\)](#) object

Required: No

SplunkDestinationDescription

The destination in Splunk.

Type: [SplunkDestinationDescription \(p. 89\)](#) 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 Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

ElasticsearchBufferingHints

Describes the buffering to perform before delivering data to the Amazon ES destination.

Contents

IntervalInSeconds

Buffer incoming data for the specified period of time, in seconds, before delivering it to the destination. The default value is 300 (5 minutes).

Type: Integer

Valid Range: Minimum value of 60. Maximum value of 900.

Required: No

SizeInMBs

Buffer incoming data to the specified size, in MBs, before delivering it to the destination. The default value is 5.

We recommend setting this parameter to a value greater than the amount of data you typically ingest into the delivery stream in 10 seconds. For example, if you typically ingest data at 1 MB/sec, the value should be 10 MB or higher.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 100.

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

ElasticsearchDestinationConfiguration

Describes the configuration of a destination in Amazon ES.

Contents

BufferingHints

The buffering options. If no value is specified, the default values for **ElasticsearchBufferingHints** are used.

Type: [ElasticsearchBufferingHints \(p. 49\)](#) object

Required: No

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

DomainARN

The ARN of the Amazon ES domain. The IAM role must have permissions for `DescribeElasticsearchDomain`, `DescribeElasticsearchDomains`, and `DescribeElasticsearchDomainConfig` after assuming the role specified in **RoleARN**.

Type: String

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

Pattern: `arn:.*`

Required: Yes

IndexName

The Elasticsearch index name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 80.

Required: Yes

IndexRotationPeriod

The Elasticsearch index rotation period. Index rotation appends a time stamp to the `IndexName` to facilitate the expiration of old data. For more information, see [Index Rotation for Amazon Elasticsearch Service Destination](#). The default value is `OneDay`.

Type: String

Valid Values: `NoRotation` | `OneHour` | `OneDay` | `OneWeek` | `OneMonth`

Required: No

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RetryOptions

The retry behavior in case Kinesis Firehose is unable to deliver documents to Amazon ES. The default value is 300 (5 minutes).

Type: [ElasticsearchRetryOptions \(p. 57\)](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the IAM role to be assumed by Kinesis Firehose for calling the Amazon ES Configuration API and for indexing documents. For more information, see [Amazon S3 Bucket Access](#).

Type: String

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

Pattern: arn:.*

Required: Yes

S3BackupMode

Defines how documents should be delivered to Amazon S3. When set to `FailedDocumentsOnly`, Kinesis Firehose writes any documents that could not be indexed to the configured Amazon S3 destination, with `elasticsearch-failed/` appended to the key prefix. When set to `AllDocuments`, Kinesis Firehose delivers all incoming records to Amazon S3, and also writes failed documents with `elasticsearch-failed/` appended to the prefix. For more information, see [Amazon S3 Backup for Amazon Elasticsearch Service Destination](#). Default value is `FailedDocumentsOnly`.

Type: String

Valid Values: `FailedDocumentsOnly` | `AllDocuments`

Required: No

S3Configuration

The configuration for the backup Amazon S3 location.

Type: [S3DestinationConfiguration \(p. 80\)](#) object

Required: Yes

TypeName

The Elasticsearch type name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

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

ElasticsearchDestinationDescription

The destination description in Amazon ES.

Contents

BufferingHints

The buffering options.

Type: [ElasticsearchBufferingHints \(p. 49\)](#) object

Required: No

CloudWatchLoggingOptions

The CloudWatch logging options.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

DomainARN

The ARN of the Amazon ES domain.

Type: String

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

Pattern: `arn:.*`

Required: No

IndexName

The Elasticsearch index name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 80.

Required: No

IndexRotationPeriod

The Elasticsearch index rotation period

Type: String

Valid Values: `NoRotation` | `OneHour` | `OneDay` | `OneWeek` | `OneMonth`

Required: No

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RetryOptions

The Amazon ES retry options.

Type: [ElasticsearchRetryOptions \(p. 57\)](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the AWS credentials.

Type: String

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

Pattern: arn:.*

Required: No

S3BackupMode

The Amazon S3 backup mode.

Type: String

Valid Values: `FailedDocumentsOnly` | `AllDocuments`

Required: No

S3DestinationDescription

The Amazon S3 destination.

Type: [S3DestinationDescription \(p. 82\)](#) object

Required: No

TypeName

The Elasticsearch type name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

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

ElasticsearchDestinationUpdate

Describes an update for a destination in Amazon ES.

Contents

BufferingHints

The buffering options. If no value is specified, **ElasticsearchBufferingHints** object default values are used.

Type: [ElasticsearchBufferingHints \(p. 49\)](#) object

Required: No

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

DomainARN

The ARN of the Amazon ES domain. The IAM role must have permissions for `DescribeElasticsearchDomain`, `DescribeElasticsearchDomains`, and `DescribeElasticsearchDomainConfig` after assuming the IAM role specified in **RoleARN**.

Type: String

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

Pattern: `arn:.*`

Required: No

IndexName

The Elasticsearch index name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 80.

Required: No

IndexRotationPeriod

The Elasticsearch index rotation period. Index rotation appends a time stamp to `IndexName` to facilitate the expiration of old data. For more information, see [Index Rotation for Amazon Elasticsearch Service Destination](#). Default value is `OneDay`.

Type: String

Valid Values: `NoRotation` | `OneHour` | `OneDay` | `OneWeek` | `OneMonth`

Required: No

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RetryOptions

The retry behavior in case Kinesis Firehose is unable to deliver documents to Amazon ES. The default value is 300 (5 minutes).

Type: [ElasticsearchRetryOptions \(p. 57\)](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the IAM role to be assumed by Kinesis Firehose for calling the Amazon ES Configuration API and for indexing documents. For more information, see [Amazon S3 Bucket Access](#).

Type: String

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

Pattern: arn:.*

Required: No

S3Update

The Amazon S3 destination.

Type: [S3DestinationUpdate \(p. 84\)](#) object

Required: No

TypeName

The Elasticsearch type name.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 100.

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

ElasticsearchRetryOptions

Configures retry behavior in case Kinesis Firehose is unable to deliver documents to Amazon ES.

Contents

DurationInSeconds

After an initial failure to deliver to Amazon ES, the total amount of time during which Kinesis Firehose re-attempts delivery (including the first attempt). After this time has elapsed, the failed documents are written to Amazon S3. Default value is 300 seconds (5 minutes). A value of 0 (zero) results in no retries.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 7200.

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

EncryptionConfiguration

Describes the encryption for a destination in Amazon S3.

Contents

KMSEncryptionConfig

The encryption key.

Type: [KMSEncryptionConfig \(p. 67\)](#) object

Required: No

NoEncryptionConfig

Specifically override existing encryption information to ensure that no encryption is used.

Type: String

Valid Values: `NoEncryption`

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

ExtendedS3DestinationConfiguration

Describes the configuration of a destination in Amazon S3.

Contents

BucketARN

The ARN of the S3 bucket.

Type: String

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

Pattern: `arn:.*`

Required: Yes

BufferingHints

The buffering option.

Type: [BufferingHints \(p. 42\)](#) object

Required: No

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

CompressionFormat

The compression format. If no value is specified, the default is UNCOMPRESSED.

Type: String

Valid Values: UNCOMPRESSED | GZIP | ZIP | Snappy

Required: No

EncryptionConfiguration

The encryption configuration. If no value is specified, the default is no encryption.

Type: [EncryptionConfiguration \(p. 58\)](#) object

Required: No

Prefix

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered S3 files. You can specify an extra prefix to be added in front of the time format prefix. If the prefix ends with a slash, it appears as a folder in the S3 bucket. For more information, see [Amazon S3 Object Name Format](#) in the *Amazon Kinesis Firehose Developer Guide*.

Type: String

Required: No

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the AWS credentials.

Type: String

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

Pattern: `arn:.*`

Required: Yes

S3BackupConfiguration

The configuration for backup in Amazon S3.

Type: [S3DestinationConfiguration \(p. 80\)](#) object

Required: No

S3BackupMode

The Amazon S3 backup mode.

Type: String

Valid Values: `Disabled` | `Enabled`

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

ExtendedS3DestinationDescription

Describes a destination in Amazon S3.

Contents

BucketARN

The ARN of the S3 bucket.

Type: String

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

Pattern: `arn:.*`

Required: Yes

BufferingHints

The buffering option.

Type: [BufferingHints \(p. 42\)](#) object

Required: Yes

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

CompressionFormat

The compression format. If no value is specified, the default is UNCOMPRESSED.

Type: String

Valid Values: UNCOMPRESSED | GZIP | ZIP | Snappy

Required: Yes

EncryptionConfiguration

The encryption configuration. If no value is specified, the default is no encryption.

Type: [EncryptionConfiguration \(p. 58\)](#) object

Required: Yes

Prefix

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered S3 files. You can specify an extra prefix to be added in front of the time format prefix. If the prefix ends with a slash, it appears as a folder in the S3 bucket. For more information, see [Amazon S3 Object Name Format](#) in the *Amazon Kinesis Firehose Developer Guide*.

Type: String

Required: No

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the AWS credentials.

Type: String

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

Pattern: `arn:.*`

Required: Yes

S3BackupDescription

The configuration for backup in Amazon S3.

Type: [S3DestinationDescription \(p. 82\)](#) object

Required: No

S3BackupMode

The Amazon S3 backup mode.

Type: String

Valid Values: `Disabled` | `Enabled`

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

ExtendedS3DestinationUpdate

Describes an update for a destination in Amazon S3.

Contents

BucketARN

The ARN of the S3 bucket.

Type: String

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

Pattern: `arn:.*`

Required: No

BufferingHints

The buffering option.

Type: [BufferingHints \(p. 42\)](#) object

Required: No

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

CompressionFormat

The compression format. If no value is specified, the default is UNCOMPRESSED.

Type: String

Valid Values: UNCOMPRESSED | GZIP | ZIP | Snappy

Required: No

EncryptionConfiguration

The encryption configuration. If no value is specified, the default is no encryption.

Type: [EncryptionConfiguration \(p. 58\)](#) object

Required: No

Prefix

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered S3 files. You can specify an extra prefix to be added in front of the time format prefix. If the prefix ends with a slash, it appears as a folder in the S3 bucket. For more information, see [Amazon S3 Object Name Format](#) in the *Amazon Kinesis Firehose Developer Guide*.

Type: String

Required: No

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the AWS credentials.

Type: String

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

Pattern: `arn:.*`

Required: No

S3BackupMode

Enables or disables Amazon S3 backup mode.

Type: String

Valid Values: `Disabled` | `Enabled`

Required: No

S3BackupUpdate

The Amazon S3 destination for backup.

Type: [S3DestinationUpdate \(p. 84\)](#) 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 Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

KinesisStreamSourceConfiguration

The stream and role Amazon Resource Names (ARNs) for a Kinesis stream used as the source for a delivery stream.

Contents

KinesisStreamARN

The ARN of the source Kinesis stream.

Type: String

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

Pattern: `arn:.*`

Required: Yes

RoleARN

The ARN of the role that provides access to the source Kinesis stream.

Type: String

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

Pattern: `arn:.*`

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

KinesisStreamSourceDescription

Details about a Kinesis stream used as the source for a Kinesis Firehose delivery stream.

Contents

DeliveryStartTimestamp

Kinesis Firehose starts retrieving records from the Kinesis stream starting with this time stamp.

Type: Timestamp

Required: No

KinesisStreamARN

The Amazon Resource Name (ARN) of the source Kinesis stream.

Type: String

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

Pattern: `arn:.*`

Required: No

RoleARN

The ARN of the role used by the source Kinesis stream.

Type: String

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

Pattern: `arn:.*`

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

KMSEncryptionConfig

Describes an encryption key for a destination in Amazon S3.

Contents

AWSKMSKeyARN

The Amazon Resource Name (ARN) of the encryption key. Must belong to the same region as the destination Amazon S3 bucket.

Type: String

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

Pattern: arn:.*

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

ProcessingConfiguration

Describes a data processing configuration.

Contents

Enabled

Enables or disables data processing.

Type: Boolean

Required: No

Processors

The data processors.

Type: Array of [Processor \(p. 69\)](#) 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 Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

Processor

Describes a data processor.

Contents

Parameters

The processor parameters.

Type: Array of [ProcessorParameter \(p. 70\)](#) objects

Required: No

Type

The type of processor.

Type: String

Valid Values: `Lambda`

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

ProcessorParameter

Describes the processor parameter.

Contents

ParameterName

The name of the parameter.

Type: String

Valid Values: `LambdaArn` | `NumberOfRetries` | `RoleArn` | `BufferSizeInMBs` | `BufferIntervalInSeconds`

Required: Yes

ParameterValue

The parameter value.

Type: String

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

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

PutRecordBatchResponseEntry

Contains the result for an individual record from a [PutRecordBatch \(p. 28\)](#) request. If the record is successfully added to your delivery stream, it receives a record ID. If the record fails to be added to your delivery stream, the result includes an error code and an error message.

Contents

ErrorCode

The error code for an individual record result.

Type: String

Required: No

ErrorMessage

The error message for an individual record result.

Type: String

Required: No

RecordId

The ID of the record.

Type: String

Length Constraints: Minimum length of 1.

Required: No

See Also

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

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

Record

The unit of data in a delivery stream.

Contents

Data

The data blob, which is base64-encoded when the blob is serialized. The maximum size of the data blob, before base64-encoding, is 1,000 KB.

Type: Base64-encoded binary data object

Length Constraints: Minimum length of 0. Maximum length of 1024000.

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

RedshiftDestinationConfiguration

Describes the configuration of a destination in Amazon Redshift.

Contents

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

ClusterJDBCURL

The database connection string.

Type: String

Length Constraints: Minimum length of 1.

Pattern: jdbc:(redshift|postgresql)://((?!-)[A-Za-z0-9-]{1,63}(?
<!--)\.)+redshift\.amazonaws\.com:\d{1,5}/[a-zA-Z0-9_]+

Required: Yes

CopyCommand

The COPY command.

Type: [CopyCommand \(p. 44\)](#) object

Required: Yes

Password

The user password.

Type: String

Length Constraints: Minimum length of 6.

Required: Yes

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RetryOptions

The retry behavior in case Kinesis Firehose is unable to deliver documents to Amazon Redshift. Default value is 3600 (60 minutes).

Type: [RedshiftRetryOptions \(p. 79\)](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the AWS credentials.

Type: String

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

Pattern: arn:.*

Required: Yes

S3BackupConfiguration

The configuration for backup in Amazon S3.

Type: [S3DestinationConfiguration \(p. 80\)](#) object

Required: No

S3BackupMode

The Amazon S3 backup mode.

Type: String

Valid Values: `Disabled` | `Enabled`

Required: No

S3Configuration

The configuration for the intermediate Amazon S3 location from which Amazon Redshift obtains data. Restrictions are described in the topic for [CreateDeliveryStream \(p. 3\)](#).

The compression formats `SNAPPY` or `ZIP` cannot be specified in **RedshiftDestinationConfiguration.S3Configuration** because the Amazon Redshift `COPY` operation that reads from the S3 bucket doesn't support these compression formats.

Type: [S3DestinationConfiguration \(p. 80\)](#) object

Required: Yes

Username

The name of the user.

Type: String

Length Constraints: Minimum length of 1.

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

RedshiftDestinationDescription

Describes a destination in Amazon Redshift.

Contents

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

ClusterJDBCURL

The database connection string.

Type: String

Length Constraints: Minimum length of 1.

Pattern: jdbc:(redshift|postgresql)://((?!-)[A-Za-z0-9-]{1,63}(?!-)\.)+redshift\.amazonaws\.com:\d{1,5}/[a-zA-Z0-9_#\$]+

Required: Yes

CopyCommand

The COPY command.

Type: [CopyCommand \(p. 44\)](#) object

Required: Yes

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RetryOptions

The retry behavior in case Kinesis Firehose is unable to deliver documents to Amazon Redshift. Default value is 3600 (60 minutes).

Type: [RedshiftRetryOptions \(p. 79\)](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the AWS credentials.

Type: String

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

Pattern: arn:.*

Required: Yes

S3BackupDescription

The configuration for backup in Amazon S3.

Type: [S3DestinationDescription \(p. 82\)](#) object

Required: No

S3BackupMode

The Amazon S3 backup mode.

Type: String

Valid Values: `Disabled` | `Enabled`

Required: No

S3DestinationDescription

The Amazon S3 destination.

Type: [S3DestinationDescription \(p. 82\)](#) object

Required: Yes

Username

The name of the user.

Type: String

Length Constraints: Minimum length of 1.

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

RedshiftDestinationUpdate

Describes an update for a destination in Amazon Redshift.

Contents

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

ClusterJDBCURL

The database connection string.

Type: String

Length Constraints: Minimum length of 1.

Pattern: jdbc:(redshift|postgresql)://((?!-)[A-Za-z0-9-]{1,63}(?
<!--)\.)+redshift\.amazonaws\.com:\d{1,5}/[a-zA-Z0-9_]+

Required: No

CopyCommand

The COPY command.

Type: [CopyCommand \(p. 44\)](#) object

Required: No

Password

The user password.

Type: String

Length Constraints: Minimum length of 6.

Required: No

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RetryOptions

The retry behavior in case Kinesis Firehose is unable to deliver documents to Amazon Redshift. Default value is 3600 (60 minutes).

Type: [RedshiftRetryOptions \(p. 79\)](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the AWS credentials.

Type: String

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

Pattern: `arn:.*`

Required: No

S3BackupMode

The Amazon S3 backup mode.

Type: String

Valid Values: `Disabled` | `Enabled`

Required: No

S3BackupUpdate

The Amazon S3 destination for backup.

Type: [S3DestinationUpdate \(p. 84\)](#) object

Required: No

S3Update

The Amazon S3 destination.

The compression formats `SNAPPY` or `ZIP` cannot be specified in **RedshiftDestinationUpdate.S3Update** because the Amazon Redshift `COPY` operation that reads from the S3 bucket doesn't support these compression formats.

Type: [S3DestinationUpdate \(p. 84\)](#) object

Required: No

Username

The name of the user.

Type: String

Length Constraints: Minimum length of 1.

Required: No

See Also

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

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

RedshiftRetryOptions

Configures retry behavior in case Kinesis Firehose is unable to deliver documents to Amazon Redshift.

Contents

DurationInSeconds

The length of time during which Kinesis Firehose retries delivery after a failure, starting from the initial request and including the first attempt. The default value is 3600 seconds (60 minutes). Kinesis Firehose does not retry if the value of `DurationInSeconds` is 0 (zero) or if the first delivery attempt takes longer than the current value.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 7200.

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

S3DestinationConfiguration

Describes the configuration of a destination in Amazon S3.

Contents

BucketARN

The ARN of the S3 bucket.

Type: String

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

Pattern: `arn:.*`

Required: Yes

BufferingHints

The buffering option. If no value is specified, **BufferingHints** object default values are used.

Type: [BufferingHints \(p. 42\)](#) object

Required: No

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

CompressionFormat

The compression format. If no value is specified, the default is `UNCOMPRESSED`.

The compression formats `SNAPPY` or `ZIP` cannot be specified for Amazon Redshift destinations because they are not supported by the Amazon Redshift `COPY` operation that reads from the S3 bucket.

Type: String

Valid Values: `UNCOMPRESSED` | `GZIP` | `ZIP` | `Snappy`

Required: No

EncryptionConfiguration

The encryption configuration. If no value is specified, the default is no encryption.

Type: [EncryptionConfiguration \(p. 58\)](#) object

Required: No

Prefix

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered S3 files. You can specify an extra prefix to be added in front of the time format prefix. If the prefix ends with a slash, it appears as a folder in the S3 bucket. For more information, see [Amazon S3 Object Name Format](#) in the *Amazon Kinesis Firehose Developer Guide*.

Type: String

Required: No

RoleARN

The Amazon Resource Name (ARN) of the AWS credentials.

Type: String

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

Pattern: arn:.**

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

S3DestinationDescription

Describes a destination in Amazon S3.

Contents

BucketARN

The ARN of the S3 bucket.

Type: String

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

Pattern: arn:.*

Required: Yes

BufferingHints

The buffering option. If no value is specified, **BufferingHints** object default values are used.

Type: [BufferingHints \(p. 42\)](#) object

Required: Yes

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

CompressionFormat

The compression format. If no value is specified, the default is UNCOMPRESSED.

Type: String

Valid Values: UNCOMPRESSED | GZIP | ZIP | Snappy

Required: Yes

EncryptionConfiguration

The encryption configuration. If no value is specified, the default is no encryption.

Type: [EncryptionConfiguration \(p. 58\)](#) object

Required: Yes

Prefix

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered S3 files. You can specify an extra prefix to be added in front of the time format prefix. If the prefix ends with a slash, it appears as a folder in the S3 bucket. For more information, see [Amazon S3 Object Name Format](#) in the *Amazon Kinesis Firehose Developer Guide*.

Type: String

Required: No

RoleARN

The Amazon Resource Name (ARN) of the AWS credentials.

Type: String

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

Pattern: arn:.*

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

S3DestinationUpdate

Describes an update for a destination in Amazon S3.

Contents

BucketARN

The ARN of the S3 bucket.

Type: String

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

Pattern: `arn:.*`

Required: No

BufferingHints

The buffering option. If no value is specified, **BufferingHints** object default values are used.

Type: [BufferingHints \(p. 42\)](#) object

Required: No

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

CompressionFormat

The compression format. If no value is specified, the default is `UNCOMPRESSED`.

The compression formats `SNAPPY` or `ZIP` cannot be specified for Amazon Redshift destinations because they are not supported by the Amazon Redshift `COPY` operation that reads from the S3 bucket.

Type: String

Valid Values: `UNCOMPRESSED` | `GZIP` | `ZIP` | `Snappy`

Required: No

EncryptionConfiguration

The encryption configuration. If no value is specified, the default is no encryption.

Type: [EncryptionConfiguration \(p. 58\)](#) object

Required: No

Prefix

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered S3 files. You can specify an extra prefix to be added in front of the time format prefix. If the prefix ends with a slash, it appears as a folder in the S3 bucket. For more information, see [Amazon S3 Object Name Format](#) in the *Amazon Kinesis Firehose Developer Guide*.

Type: String

Required: No

RoleARN

The Amazon Resource Name (ARN) of the AWS credentials.

Type: String

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

Pattern: arn:.**

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

SourceDescription

Details about a Kinesis stream used as the source for a Kinesis Firehose delivery stream.

Contents

KinesisStreamSourceDescription

The [KinesisStreamSourceDescription \(p. 66\)](#) value for the source Kinesis stream.

Type: [KinesisStreamSourceDescription \(p. 66\)](#) 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 Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

SplunkDestinationConfiguration

Describes the configuration of a destination in Splunk.

Contents

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

HECAcknowledgmentTimeoutInSeconds

The amount of time that Kinesis Firehose waits to receive an acknowledgment from Splunk after it sends it data. At the end of the timeout period, Kinesis Firehose either tries to send the data again or considers it an error, based on your retry settings.

Type: Integer

Valid Range: Minimum value of 180. Maximum value of 600.

Required: No

HECEndpoint

The HTTP Event Collector (HEC) endpoint to which Kinesis Firehose sends your data.

Type: String

Required: Yes

HECEndpointType

This type can be either "Raw" or "Event".

Type: String

Valid Values: `Raw` | `Event`

Required: Yes

HECToken

This is a GUID that you obtain from your Splunk cluster when you create a new HEC endpoint.

Type: String

Required: Yes

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RetryOptions

The retry behavior in case Kinesis Firehose is unable to deliver data to Splunk, or if it doesn't receive an acknowledgment of receipt from Splunk.

Type: [SplunkRetryOptions \(p. 93\)](#) object

Required: No

S3BackupMode

Defines how documents should be delivered to Amazon S3. When set to `FailedDocumentsOnly`, Kinesis Firehose writes any data that could not be indexed to the configured Amazon S3 destination. When set to `AllDocuments`, Kinesis Firehose delivers all incoming records to Amazon S3, and also writes failed documents to Amazon S3. Default value is `FailedDocumentsOnly`.

Type: String

Valid Values: `FailedEventsOnly` | `AllEvents`

Required: No

S3Configuration

The configuration for the backup Amazon S3 location.

Type: [S3DestinationConfiguration \(p. 80\)](#) 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 Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

SplunkDestinationDescription

Describes a destination in Splunk.

Contents

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

HECAcknowledgmentTimeoutInSeconds

The amount of time that Kinesis Firehose waits to receive an acknowledgment from Splunk after it sends it data. At the end of the timeout period, Kinesis Firehose either tries to send the data again or considers it an error, based on your retry settings.

Type: Integer

Valid Range: Minimum value of 180. Maximum value of 600.

Required: No

HECEndpoint

The HTTP Event Collector (HEC) endpoint to which Kinesis Firehose sends your data.

Type: String

Required: No

HECEndpointType

This type can be either "Raw" or "Event".

Type: String

Valid Values: `Raw` | `Event`

Required: No

HECToken

This is a GUID you obtain from your Splunk cluster when you create a new HEC endpoint.

Type: String

Required: No

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RetryOptions

The retry behavior in case Kinesis Firehose is unable to deliver data to Splunk or if it doesn't receive an acknowledgment of receipt from Splunk.

Type: [SplunkRetryOptions \(p. 93\)](#) object

Required: No

S3BackupMode

Defines how documents should be delivered to Amazon S3. When set to `FailedDocumentsOnly`, Kinesis Firehose writes any data that could not be indexed to the configured Amazon S3 destination. When set to `AllDocuments`, Kinesis Firehose delivers all incoming records to Amazon S3, and also writes failed documents to Amazon S3. Default value is `FailedDocumentsOnly`.

Type: String

Valid Values: `FailedEventsOnly` | `AllEvents`

Required: No

S3DestinationDescription

The Amazon S3 destination.>

Type: [S3DestinationDescription \(p. 82\)](#) 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 Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

SplunkDestinationUpdate

Describes an update for a destination in Splunk.

Contents

CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 43\)](#) object

Required: No

HECAcknowledgmentTimeoutInSeconds

The amount of time that Kinesis Firehose waits to receive an acknowledgment from Splunk after it sends data. At the end of the timeout period, Kinesis Firehose either tries to send the data again or considers it an error, based on your retry settings.

Type: Integer

Valid Range: Minimum value of 180. Maximum value of 600.

Required: No

HECEndpoint

The HTTP Event Collector (HEC) endpoint to which Kinesis Firehose sends your data.

Type: String

Required: No

HECEndpointType

This type can be either "Raw" or "Event".

Type: String

Valid Values: `Raw` | `Event`

Required: No

HECToken

This is a GUID that you obtain from your Splunk cluster when you create a new HEC endpoint.

Type: String

Required: No

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration \(p. 68\)](#) object

Required: No

RetryOptions

The retry behavior in case Kinesis Firehose is unable to deliver data to Splunk or if it doesn't receive an acknowledgment of receipt from Splunk.

Type: [SplunkRetryOptions \(p. 93\)](#) object

Required: No

S3BackupMode

Defines how documents should be delivered to Amazon S3. When set to `FailedDocumentsOnly`, Kinesis Firehose writes any data that could not be indexed to the configured Amazon S3 destination. When set to `AllDocuments`, Kinesis Firehose delivers all incoming records to Amazon S3, and also writes failed documents to Amazon S3. Default value is `FailedDocumentsOnly`.

Type: String

Valid Values: `FailedEventsOnly` | `AllEvents`

Required: No

S3Update

Your update to the configuration of the backup Amazon S3 location.

Type: [S3DestinationUpdate \(p. 84\)](#) 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 Go](#)
- [AWS SDK for Java](#)
- [AWS SDK for Ruby V2](#)

SplunkRetryOptions

Configures retry behavior in case Kinesis Firehose is unable to deliver documents to Splunk or if it doesn't receive an acknowledgment from Splunk.

Contents

DurationInSeconds

The total amount of time that Kinesis Firehose spends on retries. This duration starts after the initial attempt to send data to Splunk fails and doesn't include the periods during which Kinesis Firehose waits for acknowledgment from Splunk after each attempt.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 7200.

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

Common Errors

This section lists the errors common to the API actions of all AWS services. For errors specific to an API action for this service, see the topic for that API action.

AccessDeniedException

You do not have sufficient access to perform this action.

HTTP Status Code: 400

IncompleteSignature

The request signature does not conform to AWS standards.

HTTP Status Code: 400

InternalFailure

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

HTTP Status Code: 500

InvalidAction

The action or operation requested is invalid. Verify that the action is typed correctly.

HTTP Status Code: 400

InvalidClientTokenId

The X.509 certificate or AWS access key ID provided does not exist in our records.

HTTP Status Code: 403

InvalidParameterCombination

Parameters that must not be used together were used together.

HTTP Status Code: 400

InvalidParameterValue

An invalid or out-of-range value was supplied for the input parameter.

HTTP Status Code: 400

InvalidQueryParameter

The AWS query string is malformed or does not adhere to AWS standards.

HTTP Status Code: 400

MalformedQueryString

The query string contains a syntax error.

HTTP Status Code: 404

MissingAction

The request is missing an action or a required parameter.

HTTP Status Code: 400

MissingAuthenticationToken

The request must contain either a valid (registered) AWS access key ID or X.509 certificate.

HTTP Status Code: 403

MissingParameter

A required parameter for the specified action is not supplied.

HTTP Status Code: 400

OptInRequired

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

HTTP Status Code: 403

RequestExpired

The request reached the service more than 15 minutes after the date stamp on the request or more than 15 minutes after the request expiration date (such as for pre-signed URLs), or the date stamp on the request is more than 15 minutes in the future.

HTTP Status Code: 400

ServiceUnavailable

The request has failed due to a temporary failure of the server.

HTTP Status Code: 503

ThrottlingException

The request was denied due to request throttling.

HTTP Status Code: 400

ValidationError

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

HTTP Status Code: 400