

---

# Amazon Kinesis Firehose API Reference

**API Reference**

**API Version 2015-08-04**



## **Amazon Kinesis Firehose API Reference: API Reference**

Copyright © 2016 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 .....	5
Response Syntax .....	6
Response Elements .....	6
Errors .....	6
DeleteDeliveryStream .....	7
Request Syntax .....	7
Request Parameters .....	7
Response Elements .....	7
Errors .....	7
DescribeDeliveryStream .....	8
Request Syntax .....	8
Request Parameters .....	8
Response Syntax .....	8
Response Elements .....	10
Errors .....	11
ListDeliveryStreams .....	12
Request Syntax .....	12
Request Parameters .....	12
Response Syntax .....	12
Response Elements .....	12
Errors .....	13
PutRecord .....	14
Request Syntax .....	14
Request Parameters .....	14
Response Syntax .....	15
Response Elements .....	15
Errors .....	15
PutRecordBatch .....	16
Request Syntax .....	16
Request Parameters .....	17
Response Syntax .....	17
Response Elements .....	17
Errors .....	17
UpdateDestination .....	19
Request Syntax .....	19
Request Parameters .....	21
Response Elements .....	22
Errors .....	22
Data Types .....	23
BufferingHints .....	24
Contents .....	24
CloudWatchLoggingOptions .....	25
Contents .....	25
CopyCommand .....	26
Contents .....	26
DeliveryStreamDescription .....	27
Contents .....	27
DestinationDescription .....	28
Contents .....	28
ElasticsearchBufferingHints .....	29
Contents .....	29

ElasticsearchDestinationConfiguration .....	30
Contents .....	30
ElasticsearchDestinationDescription .....	32
Contents .....	32
ElasticsearchDestinationUpdate .....	34
Contents .....	34
ElasticsearchRetryOptions .....	36
Contents .....	36
EncryptionConfiguration .....	37
Contents .....	37
KMSEncryptionConfig .....	38
Contents .....	38
PutRecordBatchResponseEntry .....	39
Contents .....	39
Record .....	40
Contents .....	40
RedshiftDestinationConfiguration .....	41
Contents .....	41
RedshiftDestinationDescription .....	43
Contents .....	43
RedshiftDestinationUpdate .....	44
Contents .....	44
RedshiftRetryOptions .....	46
Contents .....	46
S3DestinationConfiguration .....	47
Contents .....	47
S3DestinationDescription .....	48
Contents .....	48
S3DestinationUpdate .....	49
Contents .....	49
Common Errors .....	50

# 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), and Amazon Redshift.

This document was last published on December 9, 2016.

# Actions

---

The following actions are supported:

- [CreateDeliveryStream](#) (p. 3)
- [DeleteDeliveryStream](#) (p. 7)
- [DescribeDeliveryStream](#) (p. 8)
- [ListDeliveryStreams](#) (p. 12)
- [PutRecord](#) (p. 14)
- [PutRecordBatch](#) (p. 16)
- [UpdateDestination](#) (p. 19)

## CreateDeliveryStream

Creates a delivery stream.

[CreateDeliveryStream \(p. 3\)](#) 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 \(p. 8\)](#).

The name of a delivery stream identifies it. You can't have two delivery streams with the same name in the same region. Two delivery streams in different AWS accounts or different regions in the same AWS account can have the same name.

By default, you can create up to 20 delivery streams per region.

A delivery stream can only be configured with a single destination, Amazon S3, Amazon Elasticsearch Service, or Amazon Redshift. For correct [CreateDeliveryStream \(p. 3\)](#) request syntax, specify only one destination configuration parameter: either **S3DestinationConfiguration**, **ElasticsearchDestinationConfiguration**, or **RedshiftDestinationConfiguration**.

As part of **S3DestinationConfiguration**, optional values **BufferingHints**, **EncryptionConfiguration**, and **CompressionFormat** can be provided. By default, if no **BufferingHints** value is provided, 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 **RedshiftDestinationConfiguration**:

- An Amazon Redshift destination requires an S3 bucket as intermediate location, as 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 element.
- 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 the username and password provided is used exclusively for Firehose purposes, and that the permissions for the account are restricted for Amazon Redshift `INSERT` permissions.

Firehose assumes the IAM role that is configured as part of destinations. The IAM role should allow the Firehose principal to assume the role, and the role should have permissions that allows 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",
  "ElasticsearchDestinationConfiguration": {
    "BufferingHints": {
      "IntervalInSeconds": number,
      "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
      "Enabled": boolean,
      "LogGroupName": "string",
      "LogStreamName": "string"
    },
  },
  "DomainARN": "string",
```

```

"IndexName": "string",
"IndexRotationPeriod": "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"
},
"RedshiftDestinationConfiguration": {
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "ClusterJDBCURL": "string",
  "CopyCommand": {
    "CopyOptions": "string",
    "DataTableColumns": "string",
    "DataTableName": "string"
  },
  "Password": "string",
  "RetryOptions": {
    "DurationInSeconds": number
  },
  "RoleARN": "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"
}
}
```

## Request Parameters

The request accepts the following data in JSON format.

### **DeliveryStreamName (p. 3)**

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

### **ElasticsearchDestinationConfiguration (p. 3)**

The destination in Amazon ES. This value cannot be specified if Amazon S3 or Amazon Redshift is the desired destination (see restrictions listed above).

Type: [ElasticsearchDestinationConfiguration \(p. 30\)](#) object

Required: No

### **RedshiftDestinationConfiguration (p. 3)**

The destination in Amazon Redshift. This value cannot be specified if Amazon S3 or Amazon Elasticsearch is the desired destination (see restrictions listed above).

Type: [RedshiftDestinationConfiguration \(p. 41\)](#) object

Required: No

### **S3DestinationConfiguration** (p. 3)

The destination in Amazon S3. This value must be specified if **ElasticsearchDestinationConfiguration** or **RedshiftDestinationConfiguration** is specified (see restrictions listed above).

Type: [S3DestinationConfiguration](#) (p. 47) 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. 6)

The ARN of the delivery stream.

Type: String

## Errors

For information about the errors that are common to all actions, see [Common Errors](#) (p. 50).

### **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

## 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 \(p. 8\)](#).

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. 7)**

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 \(p. 50\)](#).

#### **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

## DescribeDeliveryStream

Describes the specified delivery stream and gets the status. For example, after your delivery stream is created, call [DescribeDeliveryStream \(p. 8\)](#) 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. 8)

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. 8)

Specifies the destination ID to start returning the destination information. Currently Firehose supports one destination per delivery stream.

Type: String

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

Required: No

#### Limit (p. 8)

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",
    "Destinations": [
      {
        "DestinationId": "string",
        "ElasticsearchDestinationDescription": {
          "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
          }
        }
      }
    ]
  }
}
```

```

    },
    "CloudWatchLoggingOptions": {
      "Enabled": boolean,
      "LogGroupName": "string",
      "LogStreamName": "string"
    },
    },
    "DomainARN": "string",
    "IndexName": "string",
    "IndexRotationPeriod": "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"
},
"RedshiftDestinationDescription": {
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  },
  "ClusterJDBCURL": "string",
  "CopyCommand": {
    "CopyOptions": "string",
    "DataTableColumns": "string",
    "DataTableName": "string"
  },
  },
  "RetryOptions": {
    "DurationInSeconds": number
  },
  },
  "RoleARN": "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"
}
},
}
],
"HasMoreDestinations": boolean,
"LastUpdateTimestamp": number,
"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. 8)

Information about the delivery stream.

Type: [DeliveryStreamDescription \(p. 27\)](#) object

## Errors

For information about the errors that are common to all actions, see [Common Errors \(p. 50\)](#).

### **ResourceNotFoundException**

The specified resource could not be found.

HTTP Status Code: 400

## ListDeliveryStreams

Lists your delivery streams.

The number of delivery streams might be too large to return using a single call to [ListDeliveryStreams \(p. 12\)](#). 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

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

### Request Parameters

The request accepts the following data in JSON format.

#### ExclusiveStartDeliveryStreamName (p. 12)

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. 12)

The maximum number of delivery streams to list.

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. 12)

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 \(p. 12\)](#)**

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 \(p. 50\)](#).

## 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. 16). 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. 14) and [PutRecordBatch](#) (p. 16), 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. 14). 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, web site clickstream data, etc.

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(s) to parse individual data items when reading the data from the destination.

The [PutRecord](#) (p. 14) 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. 14) 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 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. 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

### [Record](#) (p. 14)

The record.

Type: [Record](#) (p. 40) 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. 15)**

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 \(p. 50\)](#).

### **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

## 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 \(p. 14\)](#). 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 \(p. 14\)](#) and [PutRecordBatch \(p. 16\)](#), 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 \(p. 16\)](#) 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 \(p. 14\)](#). 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, web site clickstream data, and so on.

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(s) to parse individual data items when reading the data from the destination.

The [PutRecordBatch \(p. 16\)](#) 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, and 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. Firehose attempts to process all records in each [PutRecordBatch \(p. 16\)](#) 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 \(p. 16\)](#) 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 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. 16)

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. 16)

One or more records.

Type: array of [Record](#) (p. 40) 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. 17)

The number of records that might have failed processing.

Type: Integer

Valid Range: Minimum value of 0.

### **RequestResponses** (p. 17)

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. 39) 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. 50).

### **InvalidArgumentException**

The specified input parameter has an 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

## UpdateDestination

Updates the specified destination of the specified delivery stream. Note: Switching between Elasticsearch and other services is not supported. For Elasticsearch destination, you can only update an existing Elasticsearch destination with this operation.

This operation can be used to change the destination type (for example, to replace the Amazon S3 destination with Amazon Redshift) or change the parameters associated with a given destination (for example, to change the bucket name of the Amazon S3 destination). The update may 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 normally effective within a few minutes.

If the destination type is the same, Firehose merges the configuration parameters specified in the [UpdateDestination \(p. 19\)](#) request with the destination configuration that already exists on the delivery stream. If any of the parameters are not specified in the update request, then the existing configuration parameters are retained. For example, in the Amazon S3 destination, if [EncryptionConfiguration \(p. 37\)](#) is not specified then the existing [EncryptionConfiguration \(p. 37\)](#) is maintained on the destination.

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

Firehose uses the **CurrentDeliveryStreamVersionId** to avoid race conditions and conflicting merges. This is a required field in every request and the service only updates the configuration if the existing configuration matches the **VersionId**. After the update is applied successfully, the **VersionId** is updated, which can be retrieved with the [DescribeDeliveryStream \(p. 8\)](#) operation. The new **VersionId** should be used to set **CurrentDeliveryStreamVersionId** in the next [UpdateDestination \(p. 19\)](#) operation.

## 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",
    "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"
},
"RedshiftDestinationUpdate": {
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "ClusterJDBCURL": "string",
  "CopyCommand": {
    "CopyOptions": "string",
    "DataTableColumns": "string",
    "DataTableName": "string"
  },
  "Password": "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"
  },
  "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"
  }
}

```

## Request Parameters

The request accepts the following data in JSON format.

### CurrentDeliveryStreamVersionId (p. 19)

Obtain this value from the **VersionId** result of the [DeliveryStreamDescription](#) (p. 27) operation. This value is required, and helps the service to perform conditional operations. For example, if there is a 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. 19)

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. 19)

The ID of the destination.

Type: String

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

Required: Yes

### ElasticsearchDestinationUpdate (p. 19)

Describes an update for a destination in Amazon ES.

Type: [ElasticsearchDestinationUpdate](#) (p. 34) object

Required: No

### RedshiftDestinationUpdate (p. 19)

Describes an update for a destination in Amazon Redshift.

Type: [RedshiftDestinationUpdate](#) (p. 44) object

Required: No

### **S3DestinationUpdate** (p. 19)

Describes an update for a destination in Amazon S3.

Type: [S3DestinationUpdate](#) (p. 49) 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. 50).

### **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

# 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. 24)
- [CloudWatchLoggingOptions](#) (p. 25)
- [CopyCommand](#) (p. 26)
- [DeliveryStreamDescription](#) (p. 27)
- [DestinationDescription](#) (p. 28)
- [ElasticsearchBufferingHints](#) (p. 29)
- [ElasticsearchDestinationConfiguration](#) (p. 30)
- [ElasticsearchDestinationDescription](#) (p. 32)
- [ElasticsearchDestinationUpdate](#) (p. 34)
- [ElasticsearchRetryOptions](#) (p. 36)
- [EncryptionConfiguration](#) (p. 37)
- [KMSEncryptionConfig](#) (p. 38)
- [PutRecordBatchResponseEntry](#) (p. 39)
- [Record](#) (p. 40)
- [RedshiftDestinationConfiguration](#) (p. 41)
- [RedshiftDestinationDescription](#) (p. 43)
- [RedshiftDestinationUpdate](#) (p. 44)
- [RedshiftRetryOptions](#) (p. 46)
- [S3DestinationConfiguration](#) (p. 47)
- [S3DestinationDescription](#) (p. 48)
- [S3DestinationUpdate](#) (p. 49)

## 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 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 SizeInMBs 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 set SizeInMBs to be 10 MB or higher.

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 128.

Required: No

# CloudWatchLoggingOptions

Describes 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 Enabled is true.

Type: String

Required: No

### **LogStreamName**

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

Type: String

Required: No

# 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 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

# 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

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

### Destinations

The destinations.

Type: array of [DestinationDescription \(p. 28\)](#) 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

### VersionId

Used when calling the [UpdateDestination \(p. 19\)](#) operation. Each time the destination is updated for the delivery stream, the VersionId is changed, and the current VersionId 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

# 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 \(p. 32\)](#) object

Required: No

### **RedshiftDestinationDescription**

The destination in Amazon Redshift.

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

Required: No

### **S3DestinationDescription**

The Amazon S3 destination.

Type: [S3DestinationDescription \(p. 48\)](#) object

Required: No



## 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 **SizeInMBs** 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, set **SizeInMBs** to be 10 MB or higher.

Type: Integer

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

Required: No

# ElasticsearchDestinationConfiguration

Describes the configuration of a destination in Amazon ES.

## Contents

### BufferingHints

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

Type: [ElasticsearchBufferingHints \(p. 29\)](#) object

Required: No

### CloudWatchLoggingOptions

Describes CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 25\)](#) object

Required: No

### DomainARN

The ARN of the Amazon ES domain. The IAM role must have permission for `DescribeElasticsearchDomain`, `DescribeElasticsearchDomains`, and `DescribeElasticsearchDomainConfig` after assuming **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 timestamp to the `IndexName` to facilitate 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

### RetryOptions

Configures retry behavior in the event that Firehose is unable to deliver documents to Amazon ES. Default value is 300 (5 minutes).

Type: [ElasticsearchRetryOptions \(p. 36\)](#) object

Required: No

### RoleARN

The ARN of the IAM role to be assumed by 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`, 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`, 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**

Describes the configuration of a destination in Amazon S3.

Type: [S3DestinationConfiguration \(p. 47\)](#) object

Required: Yes

### **TypeName**

The Elasticsearch type name.

Type: String

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

Required: Yes

# ElasticsearchDestinationDescription

The destination description in Amazon ES.

## Contents

### BufferingHints

Buffering options.

Type: [ElasticsearchBufferingHints \(p. 29\)](#) object

Required: No

### CloudWatchLoggingOptions

CloudWatch logging options.

Type: [CloudWatchLoggingOptions \(p. 25\)](#) 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

### RetryOptions

Elasticsearch retry options.

Type: [ElasticsearchRetryOptions \(p. 36\)](#) object

Required: No

### RoleARN

The ARN of the AWS credentials.

Type: String

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

Pattern: `arn:.*`

Required: No

### S3BackupMode

Amazon S3 backup mode.

Type: String

Valid Values: `FailedDocumentsOnly` | `AllDocuments`

Required: No

### S3DestinationDescription

Describes a destination in Amazon S3.

Type: [S3DestinationDescription \(p. 48\)](#) object

Required: No

**TypeName**

The Elasticsearch type name.

Type: String

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

Required: No

# ElasticsearchDestinationUpdate

Describes an update for a destination in Amazon ES.

## Contents

### BufferingHints

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

Type: [ElasticsearchBufferingHints \(p. 29\)](#) object

Required: No

### CloudWatchLoggingOptions

Describes CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 25\)](#) object

Required: No

### DomainARN

The ARN of the Amazon ES domain. The IAM role must have permission for `DescribeElasticsearchDomain`, `DescribeElasticsearchDomains`, and `DescribeElasticsearchDomainConfig` after assuming **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 timestamp to the `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

### RetryOptions

Configures retry behavior in the event that Firehose is unable to deliver documents to Amazon ES. Default value is 300 (5 minutes).

Type: [ElasticsearchRetryOptions \(p. 36\)](#) object

Required: No

### RoleARN

The ARN of the IAM role to be assumed by 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

Describes an update for a destination in Amazon S3.

Type: [S3DestinationUpdate \(p. 49\)](#) object

Required: No

**TypeName**

The Elasticsearch type name.

Type: String

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

Required: No

# ElasticsearchRetryOptions

Configures retry behavior in the event that 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 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



# EncryptionConfiguration

Describes the encryption for a destination in Amazon S3.

## Contents

### **KMSEncryptionConfig**

The encryption key.

Type: [KMSEncryptionConfig \(p. 38\)](#) object

Required: No

### **NoEncryptionConfig**

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

Type: String

Valid Values: `NoEncryption`

Required: No

# KMSEncryptionConfig

Describes an encryption key for a destination in Amazon S3.

## Contents

### **AWSKMSKeyARN**

The 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

## PutRecordBatchResponseEntry

Contains the result for an individual record from a [PutRecordBatch](#) (p. 16) 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

# 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

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

Required: Yes

# RedshiftDestinationConfiguration

Describes the configuration of a destination in Amazon Redshift.

## Contents

### CloudWatchLoggingOptions

Describes CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 25\)](#) 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 \(p. 26\)](#) object

Required: Yes

### Password

The user password.

Type: String

Length Constraints: Minimum length of 6.

Required: Yes

### RetryOptions

Configures retry behavior in the event that Firehose is unable to deliver documents to Amazon Redshift. Default value is 3600 (60 minutes).

Type: [RedshiftRetryOptions \(p. 46\)](#) object

Required: No

### RoleARN

The ARN of the AWS credentials.

Type: String

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

Pattern: arn:.\*

Required: Yes

### S3Configuration

The S3 configuration for the intermediate 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 \(p. 47\)](#) object

Required: Yes

### Username

The name of the user.

Type: String

Length Constraints: Minimum length of 1.

Required: Yes



# RedshiftDestinationDescription

Describes a destination in Amazon Redshift.

## Contents

### CloudWatchLoggingOptions

Describes CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 25\)](#) 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 \(p. 26\)](#) object

Required: Yes

### RetryOptions

Configures retry behavior in the event that Firehose is unable to deliver documents to Amazon Redshift. Default value is 3600 (60 minutes).

Type: [RedshiftRetryOptions \(p. 46\)](#) object

Required: No

### RoleARN

The ARN of the AWS credentials.

Type: String

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

Pattern: arn:.\*

Required: Yes

### S3DestinationDescription

The Amazon S3 destination.

Type: [S3DestinationDescription \(p. 48\)](#) object

Required: Yes

### Username

The name of the user.

Type: String

Length Constraints: Minimum length of 1.

Required: Yes

# RedshiftDestinationUpdate

Describes an update for a destination in Amazon Redshift.

## Contents

### CloudWatchLoggingOptions

Describes CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) (p. 25) 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. 26) object

Required: No

### Password

The user password.

Type: String

Length Constraints: Minimum length of 6.

Required: No

### RetryOptions

Configures retry behavior in the event that Firehose is unable to deliver documents to Amazon Redshift. Default value is 3600 (60 minutes).

Type: [RedshiftRetryOptions](#) (p. 46) object

Required: No

### RoleARN

The ARN of the AWS credentials.

Type: String

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

Pattern: arn:.\*

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. 49) object

Required: No

### Username

The name of the user.

Type: String

Length Constraints: Minimum length of 1.

Required: No





# RedshiftRetryOptions

Configures retry behavior in the event that Firehose is unable to deliver documents to Amazon Redshift.

## Contents

### **DurationInSeconds**

The length of time during which Firehose retries delivery after a failure, starting from the initial request and including the first attempt. The default value is 3600 seconds (60 minutes). 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

# 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 \(p. 24\)](#) object

Required: No

### CloudWatchLoggingOptions

Describes CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 25\)](#) 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 \(p. 37\)](#) 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. Note that 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 ARN of the AWS credentials.

Type: String

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

Pattern: `arn:.*`

Required: Yes

# 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 \(p. 24\)](#) object

Required: Yes

### CloudWatchLoggingOptions

Describes CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 25\)](#) 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 \(p. 37\)](#) 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. Note that 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 ARN of the AWS credentials.

Type: String

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

Pattern: `arn:.*`

Required: Yes

# 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 \(p. 24\)](#) object

Required: No

### CloudWatchLoggingOptions

Describes CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions \(p. 25\)](#) 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 \(p. 37\)](#) 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. Note that 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 ARN of the AWS credentials.

Type: String

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

Pattern: `arn:.*`

Required: No

# Common Errors

---

This section lists the common errors that all actions return. Any action-specific errors are listed in the topic for the action.

**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

**Throttling**

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