**aws**

API Reference

# Amazon Data Firehose

**API Version 2015-08-04**

# Amazon Data Firehose: API Reference

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

# Table of Contents

# Welcome

> **ⓘ Note**
>
> Amazon Data Firehose was previously known as Amazon Kinesis Data Firehose.

Amazon Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon Simple Storage Service (Amazon S3), Amazon OpenSearch Service, Amazon Redshift, Splunk, and various other supported destinations.

This document was last published on July 2, 2024.

# Actions

The following actions are supported:

- CreateDeliveryStream
- DeleteDeliveryStream
- DescribeDeliveryStream
- ListDeliveryStreams
- ListTagsForDeliveryStream
- PutRecord
- PutRecordBatch
- StartDeliveryStreamEncryption
- StopDeliveryStreamEncryption
- TagDeliveryStream
- UntagDeliveryStream
- UpdateDestination

# CreateDeliveryStream

Creates a Firehose delivery stream.

By default, you can create up to 50 delivery streams per AWS 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. If the delivery stream creation fails, the status transitions to CREATING_FAILED. 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.

If the status of a delivery stream is CREATING_FAILED, this status doesn't change, and you can't invoke CreateDeliveryStream again on it. However, you can invoke the DeleteDeliveryStream operation to delete it.

A Firehose delivery stream can be configured to receive records directly from providers using PutRecord or PutRecordBatch, or it can be configured to use an existing Kinesis stream as its source. To specify a Kinesis data 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.

To create a delivery stream with server-side encryption (SSE) enabled, include DeliveryStreamEncryptionConfigurationInput in your request. This is optional. You can also invoke StartDeliveryStreamEncryption to turn on SSE for an existing delivery stream that doesn't have SSE enabled.

A delivery stream is configured with a single destination, such as Amazon Simple Storage Service (Amazon S3), Amazon Redshift, Amazon OpenSearch Service, Amazon OpenSearch Serverless, Splunk, and any custom HTTP endpoint or HTTP endpoints owned by or supported by third-party service providers, including Datadog, Dynatrace, LogicMonitor, MongoDB, New Relic, and Sumo Logic. 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, Firehose buffers data up to 5 MB or for 5 minutes, whichever condition is satisfied first. BufferingHints is a hint, so there are some cases where the service

cannot adhere to these conditions strictly. For example, record boundaries might be 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. Firehose first delivers data to Amazon 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 Firehose, 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 the destination. The role should allow the Firehose principal to assume the role, and the role should have permissions that allow the service to deliver the data. For more information, see [Grant Firehose Access to an Amazon S3 Destination](#) in the *Amazon Firehose Developer Guide*.

## Request Syntax

```
{
    "AmazonOpenSearchServerlessDestinationConfiguration": {
        "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
        },
        "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
        },
        "CollectionEndpoint": "string",
        "IndexName": "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"
        },
        "ErrorOutputPrefix": "string",
        "Prefix": "string",
        "RoleARN": "string"
    },
    "VpcConfiguration": {
        "RoleARN": "string",
        "SecurityGroupIds": [ "string" ],
        "SubnetIds": [ "string" ]
    }
```

```
      },
      "AmazonopensearchserviceDestinationConfiguration": {
         "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
         },
         "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
         },
         "ClusterEndpoint": "string",
         "DocumentIdOptions": {
            "DefaultDocumentIdFormat": "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"
            },
            "ErrorOutputPrefix": "string",
            "Prefix": "string",
            "RoleARN": "string"
        },
        "TypeName": "string",
        "VpcConfiguration": {
            "RoleARN": "string",
            "SecurityGroupIds": [ "string" ],
            "SubnetIds": [ "string" ]
        }
    },
    "DeliveryStreamEncryptionConfigurationInput": {
        "KeyARN": "string",
        "KeyType": "string"
    },
    "DeliveryStreamName": "string",
    "DeliveryStreamType": "string",
    "ElasticsearchDestinationConfiguration": {
        "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
        },
        "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
        },
        "ClusterEndpoint": "string",
        "DocumentIdOptions": {
            "DefaultDocumentIdFormat": "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"
         },
         "ErrorOutputPrefix": "string",
         "Prefix": "string",
         "RoleARN": "string"
      },
      "TypeName": "string",
      "VpcConfiguration": {
```

```
            "RoleARN": "string",
            "SecurityGroupIds": [ "string" ],
            "SubnetIds": [ "string" ]
        }
    },
    "ExtendedS3DestinationConfiguration": {
        "BucketARN": "string",
        "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
        },
        "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
        },
        "CompressionFormat": "string",
        "CustomTimeZone": "string",
        "DataFormatConversionConfiguration": {
            "Enabled": boolean,
            "InputFormatConfiguration": {
                "Deserializer": {
                    "HiveJsonSerDe": {
                        "TimestampFormats": [ "string" ]
                    },
                    "OpenXJsonSerDe": {
                        "CaseInsensitive": boolean,
                        "ColumnToJsonKeyMappings": {
                            "string" : "string"
                        },
                        "ConvertDotsInJsonKeysToUnderscores": boolean
                    }
                }
            },
            "OutputFormatConfiguration": {
                "Serializer": {
                    "OrcSerDe": {
                        "BlockSizeBytes": number,
                        "BloomFilterColumns": [ "string" ],
                        "BloomFilterFalsePositiveProbability": number,
                        "Compression": "string",
                        "DictionaryKeyThreshold": number,
                        "EnablePadding": boolean,
                        "FormatVersion": "string",
```

```
                    "PaddingTolerance": number,
                    "RowIndexStride": number,
                    "StripeSizeBytes": number
                },
                "ParquetSerDe": {
                    "BlockSizeBytes": number,
                    "Compression": "string",
                    "EnableDictionaryCompression": boolean,
                    "MaxPaddingBytes": number,
                    "PageSizeBytes": number,
                    "WriterVersion": "string"
                }
            }
        },
        "SchemaConfiguration": {
            "CatalogId": "string",
            "DatabaseName": "string",
            "Region": "string",
            "RoleARN": "string",
            "TableName": "string",
            "VersionId": "string"
        }
    },
    "DynamicPartitioningConfiguration": {
        "Enabled": boolean,
        "RetryOptions": {
            "DurationInSeconds": number
        }
    },
    "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
            "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
    },
    "ErrorOutputPrefix": "string",
    "FileExtension": "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"
    },
    "ErrorOutputPrefix": "string",
    "Prefix": "string",
    "RoleARN": "string"
},
"S3BackupMode": "string"
},
"HttpEndpointDestinationConfiguration": {
    "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
        "LogStreamName": "string"
    },
    "EndpointConfiguration": {
```

```
            "AccessKey": "string",
            "Name": "string",
            "Url": "string"
        },
        "ProcessingConfiguration": {
            "Enabled": boolean,
            "Processors": [
                {
                    "Parameters": [
                        {
                            "ParameterName": "string",
                            "ParameterValue": "string"
                        }
                    ],
                    "Type": "string"
                }
            ]
        },
        "RequestConfiguration": {
            "CommonAttributes": [
                {
                    "AttributeName": "string",
                    "AttributeValue": "string"
                }
            ],
            "ContentEncoding": "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"
            },
            "ErrorOutputPrefix": "string",
            "Prefix": "string",
            "RoleARN": "string"
        },
        "SecretsManagerConfiguration": {
            "Enabled": boolean,
            "RoleARN": "string",
            "SecretARN": "string"
        }
    },
    "KinesisStreamSourceConfiguration": {
        "KinesisStreamARN": "string",
        "RoleARN": "string"
    },
    "MSKSourceConfiguration": {
        "AuthenticationConfiguration": {
            "Connectivity": "string",
            "RoleARN": "string"
        },
        "MSKClusterARN": "string",
        "TopicName": "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"
        },
        "ErrorOutputPrefix": "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"
            },
            "ErrorOutputPrefix": "string",
            "Prefix": "string",
            "RoleARN": "string"
        },
        "SecretsManagerConfiguration": {
            "Enabled": boolean,
            "RoleARN": "string",
            "SecretARN": "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"
        },
        "ErrorOutputPrefix": "string",
        "Prefix": "string",
        "RoleARN": "string"
```

```
      },
   "SnowflakeDestinationConfiguration": {
      "AccountUrl": "string",
      "CloudWatchLoggingOptions": {
         "Enabled": boolean,
         "LogGroupName": "string",
         "LogStreamName": "string"
      },
      "ContentColumnName": "string",
      "Database": "string",
      "DataLoadingOption": "string",
      "KeyPassphrase": "string",
      "MetaDataColumnName": "string",
      "PrivateKey": "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"
            },
            "ErrorOutputPrefix": "string",
            "Prefix": "string",
            "RoleARN": "string"
        },
        "Schema": "string",
        "SecretsManagerConfiguration": {
            "Enabled": boolean,
            "RoleARN": "string",
            "SecretARN": "string"
        },
        "SnowflakeRoleConfiguration": {
            "Enabled": boolean,
            "SnowflakeRole": "string"
        },
        "SnowflakeVpcConfiguration": {
            "PrivateLinkVpceId": "string"
        },
        "Table": "string",
        "User": "string"
    },
    "SplunkDestinationConfiguration": {
        "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
        },
        "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"
            },
            "ErrorOutputPrefix": "string",
            "Prefix": "string",
            "RoleARN": "string"
        },
        "SecretsManagerConfiguration": {
            "Enabled": boolean,
            "RoleARN": "string",
            "SecretARN": "string"
        }
    },
    "Tags": [
```

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

## Request Parameters

The request accepts the following data in JSON format.

**AmazonOpenSearchServerlessDestinationConfiguration**

The destination in the Serverless offering for Amazon OpenSearch Service. You can specify only one destination.

Type: AmazonOpenSearchServerlessDestinationConfiguration object

Required: No

**AmazonopensearchserviceDestinationConfiguration**

The destination in Amazon OpenSearch Service. You can specify only one destination.

Type: AmazonopensearchserviceDestinationConfiguration object

Required: No

**DeliveryStreamEncryptionConfigurationInput**

Used to specify the type and Amazon Resource Name (ARN) of the KMS key needed for Server-Side Encryption (SSE).

Type: DeliveryStreamEncryptionConfigurationInput object

Required: No

**DeliveryStreamName**

The name of the delivery stream. This name must be unique per AWS account in the same AWS 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

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 data stream as a source.

Type: String

Valid Values: `DirectPut | KinesisStreamAsSource | MSKAsSource`

Required: No

## ElasticsearchDestinationConfiguration

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

Type: ElasticsearchDestinationConfiguration object

Required: No

## ExtendedS3DestinationConfiguration

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

Type: ExtendedS3DestinationConfiguration object

Required: No

## HttpEndpointDestinationConfiguration

Enables configuring Kinesis Firehose to deliver data to any HTTP endpoint destination. You can specify only one destination.

Type: HttpEndpointDestinationConfiguration object

Required: No

## KinesisStreamSourceConfiguration

When a Kinesis data stream is used as the source for the delivery stream, a KinesisStreamSourceConfiguration containing the Kinesis data stream Amazon Resource Name (ARN) and the role ARN for the source stream.

Type: [KinesisStreamSourceConfiguration](#) object

Required: No

## MSKSourceConfiguration

The configuration for the Amazon MSK cluster to be used as the source for a delivery stream.

Type: [MSKSourceConfiguration](#) object

Required: No

## RedshiftDestinationConfiguration

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

Type: [RedshiftDestinationConfiguration](#) object

Required: No

## S3DestinationConfiguration

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

Type: [S3DestinationConfiguration](#) object

Required: No

## SnowflakeDestinationConfiguration

Configure Snowflake destination

Type: [SnowflakeDestinationConfiguration](#) object

Required: No

## SplunkDestinationConfiguration

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

Type: [SplunkDestinationConfiguration](#) object

Required: No

## Tags

A set of tags to assign to the delivery stream. A tag is a key-value pair that you can define and assign to AWS resources. Tags are metadata. For example, you can add friendly names and descriptions or other types of information that can help you distinguish the delivery stream.

For more information about tags, see [Using Cost Allocation Tags](#) in the AWS Billing and Cost Management User Guide.

You can specify up to 50 tags when creating a delivery stream.

If you specify tags in the `CreateDeliveryStream` action, Amazon Data Firehose performs an additional authorization on the `firehose:TagDeliveryStream` action to verify if users have permissions to create tags. If you do not provide this permission, requests to create new Firehose delivery streams with IAM resource tags will fail with an `AccessDeniedException` such as following.

**AccessDeniedException**

User: arn:aws:sts::x:assumed-role/x/x is not authorized to perform: firehose:TagDeliveryStream on resource: arn:aws:firehose:us-east-1:x:deliverystream/x with an explicit deny in an identity-based policy.

For an example IAM policy, see [Tag example.](#)

Type: Array of [Tag](#) objects

Array Members: Minimum number of 1 item. Maximum number of 50 items.

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**

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](#).

**InvalidArgumentException**

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

HTTP Status Code: 400

**InvalidKMSResourceException**

Firehose throws this exception when an attempt to put records or to start or stop delivery stream encryption fails. This happens when the KMS service throws one of the following exception types: `AccessDeniedException`, `InvalidStateException`, `DisabledException`, or `NotFoundException`.

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

## Examples

### Example

The following JSON example creates a delivery stream named `exampleStreamName` with an Amazon S3 destination. To use this example, first replace the placeholders for the `RoleARN` and `BucketARN` keys with valid strings. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

## 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"
}
```

# Example IAM policy to create a delivery stream and apply tags

The following example demonstrates a policy that allows users to create a delivery stream and apply tags.

```
{
```

```
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "firehose:CreateDeliveryStream",
            "Resource": "*",
            }
        },
        {
            "Effect": "Allow",
            "Action": "firehose:TagDeliveryStream",
            "Resource": "*",
            }
        }
    ]
}
```

## 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 v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V3

# DeleteDeliveryStream

Deletes a delivery stream and its data.

You can delete a delivery stream only if it is in one of the following states: ACTIVE, DELETING, CREATING_FAILED, or DELETING_FAILED. You can't delete a delivery stream that is in the CREATING state. To check the state of a delivery stream, use [DescribeDeliveryStream](#).

DeleteDeliveryStream is an asynchronous API. When an API request to DeleteDeliveryStream succeeds, the delivery stream is marked for deletion, and it goes into the DELETING state.While the delivery stream is in the DELETING state, the service might continue to accept records, but it doesn't make any guarantees with respect to delivering the data. Therefore, as a best practice, first stop any applications that are sending records before you delete a delivery stream.

Removal of a delivery stream that is in the DELETING state is a low priority operation for the service. A stream may remain in the DELETING state for several minutes. Therefore, as a best practice, applications should not wait for streams in the DELETING state to be removed.

## Request Syntax

```
{
    "AllowForceDelete": boolean,
    "DeliveryStreamName": "string"
}
```

## Request Parameters

The request accepts the following data in JSON format.

**AllowForceDelete**

Set this to true if you want to delete the delivery stream even if Firehose is unable to retire the grant for the CMK. Firehose might be unable to retire the grant due to a customer error, such as when the CMK or the grant are in an invalid state. If you force deletion, you can then use the [RevokeGrant](#) operation to revoke the grant you gave to Firehose. If a failure to retire the grant happens due to an AWS KMS issue, Firehose keeps retrying the delete operation.

The default value is false.

Type: Boolean

Required: No

## 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

# 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.

**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

# Examples

## 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 v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V3

# DescribeDeliveryStream

Describes the specified delivery stream and its status. For example, after your delivery stream is created, call `DescribeDeliveryStream` to see whether the delivery stream is `ACTIVE` and therefore ready for data to be sent to it.

If the status of a delivery stream is `CREATING_FAILED`, this status doesn't change, and you can't invoke [CreateDeliveryStream](#) again on it. However, you can invoke the [DeleteDeliveryStream](#) operation to delete it. If the status is `DELETING_FAILED`, you can force deletion by invoking [DeleteDeliveryStream](#) again but with [DeleteDeliveryStream:AllowForceDelete](#) set to true.

## Request Syntax

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

## Request Parameters

The request accepts the following data in JSON format.

**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

**ExclusiveStartDestinationId**

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

Type: String

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

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

Required: No

## Limit

The limit on the number of destinations to return. 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",
        "DeliveryStreamEncryptionConfiguration": {
            "FailureDescription": {
                "Details": "string",
                "Type": "string"
            },
            "KeyARN": "string",
            "KeyType": "string",
            "Status": "string"
        },
        "DeliveryStreamName": "string",
        "DeliveryStreamStatus": "string",
        "DeliveryStreamType": "string",
        "Destinations": [
            {
                "AmazonOpenSearchServerlessDestinationDescription": {
                    "BufferingHints": {
                        "IntervalInSeconds": number,
                        "SizeInMBs": number
                    },
                    "CloudWatchLoggingOptions": {
                        "Enabled": boolean,
                        "LogGroupName": "string",
                        "LogStreamName": "string"
```

```
            },
            "CollectionEndpoint": "string",
            "IndexName": "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"
                },
                "ErrorOutputPrefix": "string",
                "Prefix": "string",
                "RoleARN": "string"
            },
```

```
            "VpcConfigurationDescription": {
                "RoleARN": "string",
                "SecurityGroupIds": [ "string" ],
                "SubnetIds": [ "string" ],
                "VpcId": "string"
            }
        },
        "AmazonopensearchserviceDestinationDescription": {
            "BufferingHints": {
                "IntervalInSeconds": number,
                "SizeInMBs": number
            },
            "CloudWatchLoggingOptions": {
                "Enabled": boolean,
                "LogGroupName": "string",
                "LogStreamName": "string"
            },
            "ClusterEndpoint": "string",
            "DocumentIdOptions": {
                "DefaultDocumentIdFormat": "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"
            },
            "ErrorOutputPrefix": "string",
            "Prefix": "string",
            "RoleARN": "string"
         },
         "TypeName": "string",
         "VpcConfigurationDescription": {
            "RoleARN": "string",
            "SecurityGroupIds": [ "string" ],
            "SubnetIds": [ "string" ],
            "VpcId": "string"
         }
      },
      "DestinationId": "string",
      "ElasticsearchDestinationDescription": {
         "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
         },
         "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
         },
         "ClusterEndpoint": "string",
         "DocumentIdOptions": {
            "DefaultDocumentIdFormat": "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"
                },
                "ErrorOutputPrefix": "string",
                "Prefix": "string",
                "RoleARN": "string"
            },
```

```
            "TypeName": "string",
            "VpcConfigurationDescription": {
               "RoleARN": "string",
               "SecurityGroupIds": [ "string" ],
               "SubnetIds": [ "string" ],
               "VpcId": "string"
            }
         },
         "ExtendedS3DestinationDescription": {
            "BucketARN": "string",
            "BufferingHints": {
               "IntervalInSeconds": number,
               "SizeInMBs": number
            },
            "CloudWatchLoggingOptions": {
               "Enabled": boolean,
               "LogGroupName": "string",
               "LogStreamName": "string"
            },
            "CompressionFormat": "string",
            "CustomTimeZone": "string",
            "DataFormatConversionConfiguration": {
               "Enabled": boolean,
               "InputFormatConfiguration": {
                  "Deserializer": {
                     "HiveJsonSerDe": {
                        "TimestampFormats": [ "string" ]
                     },
                     "OpenXJsonSerDe": {
                        "CaseInsensitive": boolean,
                        "ColumnToJsonKeyMappings": {
                           "string" : "string"
                        },
                        "ConvertDotsInJsonKeysToUnderscores": boolean
                     }
                  }
               },
               "OutputFormatConfiguration": {
                  "Serializer": {
                     "OrcSerDe": {
                        "BlockSizeBytes": number,
                        "BloomFilterColumns": [ "string" ],
                        "BloomFilterFalsePositiveProbability": number,
                        "Compression": "string",
```

```
                    "DictionaryKeyThreshold": number,
                    "EnablePadding": boolean,
                    "FormatVersion": "string",
                    "PaddingTolerance": number,
                    "RowIndexStride": number,
                    "StripeSizeBytes": number
                },
                "ParquetSerDe": {
                    "BlockSizeBytes": number,
                    "Compression": "string",
                    "EnableDictionaryCompression": boolean,
                    "MaxPaddingBytes": number,
                    "PageSizeBytes": number,
                    "WriterVersion": "string"
                }
            }
        },
        "SchemaConfiguration": {
            "CatalogId": "string",
            "DatabaseName": "string",
            "Region": "string",
            "RoleARN": "string",
            "TableName": "string",
            "VersionId": "string"
        }
    },
    "DynamicPartitioningConfiguration": {
        "Enabled": boolean,
        "RetryOptions": {
            "DurationInSeconds": number
        }
    },
    "EncryptionConfiguration": {
        "KMSEncryptionConfig": {
            "AWSKMSKeyARN": "string"
        },
        "NoEncryptionConfig": "string"
    },
    "ErrorOutputPrefix": "string",
    "FileExtension": "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"
        },
        "ErrorOutputPrefix": "string",
        "Prefix": "string",
        "RoleARN": "string"
    },
    "S3BackupMode": "string"
},
"HttpEndpointDestinationDescription": {
    "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
        "LogGroupName": "string",
```

```
                    "LogStreamName": "string"
                },
                "EndpointConfiguration": {
                    "Name": "string",
                    "Url": "string"
                },
                "ProcessingConfiguration": {
                    "Enabled": boolean,
                    "Processors": [
                        {
                            "Parameters": [
                                {
                                    "ParameterName": "string",
                                    "ParameterValue": "string"
                                }
                            ],
                            "Type": "string"
                        }
                    ]
                },
                "RequestConfiguration": {
                    "CommonAttributes": [
                        {
                            "AttributeName": "string",
                            "AttributeValue": "string"
                        }
                    ],
                    "ContentEncoding": "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"
                },
                "ErrorOutputPrefix": "string",
                "Prefix": "string",
                "RoleARN": "string"
            },
            "SecretsManagerConfiguration": {
                "Enabled": boolean,
                "RoleARN": "string",
                "SecretARN": "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"
                },
                "ErrorOutputPrefix": "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"
```

```
            },
            "ErrorOutputPrefix": "string",
            "Prefix": "string",
            "RoleARN": "string"
        },
        "SecretsManagerConfiguration": {
            "Enabled": boolean,
            "RoleARN": "string",
            "SecretARN": "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"
        },
        "ErrorOutputPrefix": "string",
        "Prefix": "string",
        "RoleARN": "string"
    },
    "SnowflakeDestinationDescription": {
        "AccountUrl": "string",
        "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
        },
        "ContentColumnName": "string",
        "Database": "string",
        "DataLoadingOption": "string",
```

```
            "MetaDataColumnName": "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"
                },
                "ErrorOutputPrefix": "string",
                "Prefix": "string",
                "RoleARN": "string"
            },
            "Schema": "string",
            "SecretsManagerConfiguration": {
```

```
            "Enabled": boolean,
            "RoleARN": "string",
            "SecretARN": "string"
        },
        "SnowflakeRoleConfiguration": {
            "Enabled": boolean,
            "SnowflakeRole": "string"
        },
        "SnowflakeVpcConfiguration": {
            "PrivateLinkVpceId": "string"
        },
        "Table": "string",
        "User": "string"
    },
    "SplunkDestinationDescription": {
        "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
        },
        "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"
                },
                "ErrorOutputPrefix": "string",
                "Prefix": "string",
                "RoleARN": "string"
            },
            "SecretsManagerConfiguration": {
                "Enabled": boolean,
                "RoleARN": "string",
                "SecretARN": "string"
            }
        }
    }
],
"FailureDescription": {
    "Details": "string",
    "Type": "string"
},
"HasMoreDestinations": boolean,
"LastUpdateTimestamp": number,
"Source": {
    "KinesisStreamSourceDescription": {
        "DeliveryStartTimestamp": number,
        "KinesisStreamARN": "string",
        "RoleARN": "string"
    },
```

```
            "MSKSourceDescription": {
                "AuthenticationConfiguration": {
                    "Connectivity": "string",
                    "RoleARN": "string"
                },
                "DeliveryStartTimestamp": number,
                "MSKClusterARN": "string",
                "TopicName": "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**

Information about the delivery stream.

Type: DeliveryStreamDescription object

## Errors

For information about the errors that are common to all actions, see Common Errors.

**ResourceNotFoundException**

The specified resource could not be found.

HTTP Status Code: 400

## Examples

### 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",
    "DeliveryStreamEncryptionConfiguration": {
      "Status": "DISABLED"
    },
    "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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ListDeliveryStreams

Lists your delivery streams in alphabetical order of their names.

The number of delivery streams might be too large to return using a single call to `ListDeliveryStreams`. 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 calling this operation again and setting the `ExclusiveStartDeliveryStreamName` parameter to the name of the last delivery stream returned in the last call.

## Request Syntax

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

## Request Parameters

The request accepts the following data in JSON format.

### 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 data 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 | MSKAsSource`

Required: No

## ExclusiveStartDeliveryStreamName

The list of delivery streams returned by this call to `ListDeliveryStreams` will start with the delivery stream whose name comes alphabetically immediately after the name you specify in `ExclusiveStartDeliveryStreamName`.

Type: String

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

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

Required: No

## Limit

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

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**

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.

# Examples

## Example

The following JSON example lists up to three 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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# ListTagsForDeliveryStream

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

## Request Syntax

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

## Request Parameters

The request accepts the following data in JSON format.

**DeliveryStreamName**

The name of the delivery stream whose tags you want to list.

Type: String

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

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

Required: Yes

**ExclusiveStartTagKey**

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

Type: String

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

Pattern: `^(?!aws:)[\p{L}\p{Z}\p{N}_.:\/=+\-@%]*$`

Required: No

## Limit

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

Type: Integer

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

Required: No

# Response Syntax

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

# Response Elements

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

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

## HasMoreTags

If this is `true` in the response, more tags are available. To list the remaining tags, set ExclusiveStartTagKey to the key of the last tag returned and call ListTagsForDeliveryStream again.

Type: Boolean

## Tags

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

Type: Array of [Tag](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

# Errors

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

**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

**ResourceNotFoundException**

The specified resource could not be found.

HTTP Status Code: 400

# Examples

## To list the tags for a stream

The following JSON example lists the tags for the specified 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.ListTagsForDeliveryStream
{
   "DeliveryStreamName": "exampleDeliveryStreamName"
}
```

## Sample Response

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

# See Also

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

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

- [AWS SDK for Ruby V3](#)

# PutRecord

Writes a single data record into an Amazon Firehose delivery stream. To write multiple data records into a delivery stream, use PutRecordBatch. 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 and PutRecordBatch, 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 Firehose Limits.

Firehose accumulates and publishes a particular metric for a customer account in one minute intervals. It is possible that the bursts of incoming bytes/records ingested to a delivery stream last only for a few seconds. Due to this, the actual spikes in the traffic might not be fully visible in the customer's 1 minute CloudWatch metrics.

You must specify the name of the delivery stream and the data record when using PutRecord. The data record consists of a data blob that can be up to 1,000 KiB in size, and any kind of data. For example, it can be a segment from a log file, geographic location data, website 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 to parse individual data items when reading the data from the destination.

The `PutRecord` 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` operation throws a `ServiceUnavailableException`, the API is automatically reinvoked (retried) 3 times. If the exception persists, it is possible that the throughput limits have been exceeded for the delivery stream.

Re-invoking the Put API operations (for example, PutRecord and PutRecordBatch) can result in data duplicates. For larger data assets, allow for a longer time out before retrying Put API operations.

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

> **⚠ Important**
>
> Don't concatenate two or more base64 strings to form the data fields of your records.
> Instead, concatenate the raw data, then perform base64 encoding.

## Request Syntax

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

## Request Parameters

The request accepts the following data in JSON format.

### 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

### Record

The record.

Type: Record object

Required: Yes

## Response Syntax

```
{
```

```
      "Encrypted": boolean,
      "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.

### Encrypted

Indicates whether server-side encryption (SSE) was enabled during this operation.

Type: Boolean

### RecordId

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.

**InvalidArgumentException**

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

HTTP Status Code: 400

**InvalidKMSResourceException**

Firehose throws this exception when an attempt to put records or to start or stop delivery stream encryption fails. This happens when the KMS service throws one of the following exception types: `AccessDeniedException`, `InvalidStateException`, `DisabledException`, or `NotFoundException`.

HTTP Status Code: 400

**InvalidSourceException**

Only requests from CloudWatch Logs are supported when CloudWatch Logs decompression is enabled.

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 Firehose Limits.

HTTP Status Code: 500

# Examples

## 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+1mgSpifCmfFm7lFhuts//4Ft6sFVokyE6t+5ioEAjNm+sgQ6iVf/
YePEXBK6epIW4QeXqJp2xsbfZUNXsfOY1QrYXgRBCKznkjMMTP0BqJGObM3fB//dHgEE0XDTc4wW065i/
tJyYI1Vy8qn8FMhpkZuh5bvG482XkkBxFmMGnhPTQwQ4A1IOP0sE0X99YnBK8RECdeQ2zxynvZ"
}
```

## Example

The following example shows how you can use the [AWS CLI](#) to put a record in a delivery stream.

```
aws firehose put-record --delivery-stream-name mystream --record="{\"Data\":\"1\"}"
```

# 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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# 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. Applications using these operations are referred to as producers.

Firehose accumulates and publishes a particular metric for a customer account in one minute intervals. It is possible that the bursts of incoming bytes/records ingested to a delivery stream last only for a few seconds. Due to this, the actual spikes in the traffic might not be fully visible in the customer's 1 minute CloudWatch metrics.

For information about service quota, see Amazon Firehose Quota.

Each PutRecordBatch request supports up to 500 records. Each record in the request can be as large as 1,000 KB (before base64 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. 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, website 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 to parse individual data items when reading the data from the destination.

The PutRecordBatch response includes a count of failed records, FailedPutCount, and an array of responses, RequestResponses. Even if the PutRecordBatch call succeeds, the value of FailedPutCount may be greater than 0, indicating that there are records for which the operation didn't succeed. 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. Firehose tries to process all records in each PutRecordBatch 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: `ServiceUnavailableException` 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](#) throws `ServiceUnavailableException`, the API is automatically reinvoked (retried) 3 times. If the exception persists, it is possible that the throughput limits have been exceeded for the delivery stream.

Re-invoking the Put API operations (for example, PutRecord and PutRecordBatch) can result in data duplicates. For larger data assets, allow for a longer time out before retrying Put API operations.

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.

> ⚠️ **Important**
>
> Don't concatenate two or more base64 strings to form the data fields of your records. Instead, concatenate the raw data, then perform base64 encoding.

## Request Syntax

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

## Request Parameters

The request accepts the following data in JSON format.

**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

**Records**

One or more records.

Type: Array of Record objects

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

Required: Yes

## Response Syntax

```
{
   "Encrypted": boolean,
   "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.

## Encrypted

Indicates whether server-side encryption (SSE) was enabled during this operation.

Type: Boolean

## FailedPutCount

The number of records that might have failed processing. This number might be greater than 0 even if the [PutRecordBatch](#) call succeeds. Check `FailedPutCount` to determine whether there are records that you need to resend.

Type: Integer

Valid Range: Minimum value of 0.

## RequestResponses

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](#) 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](#).

**InvalidArgumentException**

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

HTTP Status Code: 400

**InvalidKMSResourceException**

Firehose throws this exception when an attempt to put records or to start or stop delivery stream encryption fails. This happens when the KMS service throws one of the following exception types: `AccessDeniedException`, `InvalidStateException`, `DisabledException`, or `NotFoundException`.

HTTP Status Code: 400

**InvalidSourceException**

Only requests from CloudWatch Logs are supported when CloudWatch Logs decompression is enabled.

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 Firehose Limits.

HTTP Status Code: 500

# Examples

## 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": "AJJBALlfiFN9HyhPj6Dc+XqcRlTjqyIbr927TsEmWWpN39EK/
JbRTbXDZFNCCrWIs/4YUd3gjYUs6giZn76TuI4vv0ljDOMFJvLzqLuupu93RscZRmpW1CP8DeiFzJJGvqIr1LRE/
MDozYenOz
+v1ZIqEjECvfMwIz6silvDaGuYtIhVzwd8yWaDGFCX4ODes6W07W9W7Q5ViMPUKSt6F0nn7OhVkJ3/"
        },
        {
            "RecordId": "goGaFS919Mmv7lYET0oMaw+UL9iFpzi1O0o
+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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# StartDeliveryStreamEncryption

Enables server-side encryption (SSE) for the delivery stream.

This operation is asynchronous. It returns immediately. When you invoke it, Firehose first sets the encryption status of the stream to ENABLING, and then to ENABLED. The encryption status of a delivery stream is the Status property in [DeliveryStreamEncryptionConfiguration](). If the operation fails, the encryption status changes to ENABLING_FAILED. You can continue to read and write data to your delivery stream while the encryption status is ENABLING, but the data is not encrypted. It can take up to 5 seconds after the encryption status changes to ENABLED before all records written to the delivery stream are encrypted. To find out whether a record or a batch of records was encrypted, check the response elements [PutRecord:Encrypted]() and [PutRecordBatch:Encrypted](), respectively.

To check the encryption status of a delivery stream, use [DescribeDeliveryStream]().

Even if encryption is currently enabled for a delivery stream, you can still invoke this operation on it to change the ARN of the CMK or both its type and ARN. If you invoke this method to change the CMK, and the old CMK is of type CUSTOMER_MANAGED_CMK, Firehose schedules the grant it had on the old CMK for retirement. If the new CMK is of type CUSTOMER_MANAGED_CMK, Firehose creates a grant that enables it to use the new CMK to encrypt and decrypt data and to manage the grant.

For the KMS grant creation to be successful, the Firehose API operations StartDeliveryStreamEncryption and CreateDeliveryStream should not be called with session credentials that are more than 6 hours old.

If a delivery stream already has encryption enabled and then you invoke this operation to change the ARN of the CMK or both its type and ARN and you get ENABLING_FAILED, this only means that the attempt to change the CMK failed. In this case, encryption remains enabled with the old CMK.

If the encryption status of your delivery stream is ENABLING_FAILED, you can invoke this operation again with a valid CMK. The CMK must be enabled and the key policy mustn't explicitly deny the permission for Firehose to invoke KMS encrypt and decrypt operations.

You can enable SSE for a delivery stream only if it's a delivery stream that uses DirectPut as its source.

The StartDeliveryStreamEncryption and StopDeliveryStreamEncryption operations have a combined limit of 25 calls per delivery stream per 24 hours. For

example, you reach the limit if you call `StartDeliveryStreamEncryption` 13 times and
`StopDeliveryStreamEncryption` 12 times for the same delivery stream in a 24-hour period.

## Request Syntax

```
{
    "DeliveryStreamEncryptionConfigurationInput": {
       "KeyARN": "string",
       "KeyType": "string"
    },
    "DeliveryStreamName": "string"
}
```

## Request Parameters

The request accepts the following data in JSON format.

**DeliveryStreamEncryptionConfigurationInput**

Used to specify the type and Amazon Resource Name (ARN) of the KMS key needed for Server-
Side Encryption (SSE).

Type: DeliveryStreamEncryptionConfigurationInput object

Required: No

**DeliveryStreamName**

The name of the delivery stream for which you want to enable server-side encryption (SSE).

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.

**InvalidArgumentException**

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

HTTP Status Code: 400

**InvalidKMSResourceException**

Firehose throws this exception when an attempt to put records or to start or stop delivery stream encryption fails. This happens when the KMS service throws one of the following exception types: `AccessDeniedException`, `InvalidStateException`, `DisabledException`, or `NotFoundException`.

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

**ResourceNotFoundException**

The specified resource could not be found.

HTTP Status Code: 400

# Examples

## To start server-side encryption for a stream

The following JSON example starts server-side encryption (SSE) for the specified 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.StartDeliveryStreamEncryption
{
   "DeliveryStreamName": "exampleDeliveryStreamName"
}
```

## 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 v2
- AWS SDK for Java V2
- AWS SDK for JavaScript V3
- AWS SDK for PHP V3
- AWS SDK for Python
- AWS SDK for Ruby V3

# StopDeliveryStreamEncryption

Disables server-side encryption (SSE) for the delivery stream.

This operation is asynchronous. It returns immediately. When you invoke it, Firehose first sets the encryption status of the stream to `DISABLING`, and then to `DISABLED`. You can continue to read and write data to your stream while its status is `DISABLING`. It can take up to 5 seconds after the encryption status changes to `DISABLED` before all records written to the delivery stream are no longer subject to encryption. To find out whether a record or a batch of records was encrypted, check the response elements [PutRecord:Encrypted](#) and [PutRecordBatch:Encrypted](#), respectively.

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

If SSE is enabled using a customer managed CMK and then you invoke `StopDeliveryStreamEncryption`, Firehose schedules the related KMS grant for retirement and then retires it after it ensures that it is finished delivering records to the destination.

The `StartDeliveryStreamEncryption` and `StopDeliveryStreamEncryption` operations have a combined limit of 25 calls per delivery stream per 24 hours. For example, you reach the limit if you call `StartDeliveryStreamEncryption` 13 times and `StopDeliveryStreamEncryption` 12 times for the same delivery stream in a 24-hour period.

## Request Syntax

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

## Request Parameters

The request accepts the following data in JSON format.

**DeliveryStreamName**

The name of the delivery stream for which you want to disable server-side encryption (SSE).

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](#).

**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

**ResourceNotFoundException**

The specified resource could not be found.

HTTP Status Code: 400

## Examples

### To stop server-side encryption for a stream

The following JSON example stops server-side encryption (SSE) for the specified 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.StopDeliveryStreamEncryption
{
   "DeliveryStreamName": "exampleDeliveryStreamName"
}
```

**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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# TagDeliveryStream

Adds or updates tags for the specified delivery stream. A tag is a key-value pair that you can define and assign to AWS resources. If you specify a tag that already exists, the tag value is replaced with the value that you specify in the request. Tags are metadata. For example, you can add friendly names and descriptions or other types of information that can help you distinguish the delivery stream. For more information about tags, see [Using Cost Allocation Tags](#) in the *AWS Billing and Cost Management User Guide*.

Each delivery stream can have up to 50 tags.

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

## Request Syntax

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

## Request Parameters

The request accepts the following data in JSON format.

**DeliveryStreamName**

The name of the delivery stream to which you want to add the tags.

Type: String

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

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

Required: Yes

**Tags**

A set of key-value pairs to use to create the tags.

Type: Array of Tag objects

Array Members: Minimum number of 1 item. Maximum number of 50 items.

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.

**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

**ResourceNotFoundException**

The specified resource could not be found.

HTTP Status Code: 400

# Examples

## To add tags to a stream

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

### Sample Request

```
POST / HTTP/1.1
Host: 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.TagDeliveryStream
{
  "DeliveryStreamName": "exampleDeliveryStreamName",
  "Tags": [
    {
      "Key": "Project",
      "Value": "myProject"
    },
    {
      "Key": "Environment",
      "Value": "Production"
    }
  ]
}
```

### Sample Response

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

# 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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

# UntagDeliveryStream

Removes tags from the specified delivery stream. Removed tags are deleted, and you can't recover them after this operation successfully completes.

If you specify a tag that doesn't exist, the operation ignores it.

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

## Request Syntax

```
{
    "DeliveryStreamName": "string",
    "TagKeys": [ "string" ]
}
```

## Request Parameters

The request accepts the following data in JSON format.

**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

**TagKeys**

   A list of tag keys. Each corresponding tag is removed from the delivery stream.

   Type: Array of strings

   Array Members: Minimum number of 1 item. Maximum number of 50 items.

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

   Pattern: `^(?!aws:)[\p{L}\p{Z}\p{N}_.:\/=+\-@%]*$`

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](#).

**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

**ResourceNotFoundException**

The specified resource could not be found.

HTTP Status Code: 400

## Examples

### To remove tags from a stream

The following JSON example removes the specified tag from the specified stream.

**Sample Request**

```
POST / HTTP/1.1
```

```
Host: 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.UntagDeliveryStream
{
   "DeliveryStreamName": "exampleDeliveryStreamName",
   "TagKeys": ["Project", "Environment"]
}
```

**Sample Response**

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

# See Also

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

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

# UpdateDestination

Updates the specified destination of the specified delivery stream.

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.

Switching between Amazon OpenSearch Service and other services is not supported. For an Amazon OpenSearch Service destination, you can only update to another Amazon OpenSearch Service destination.

If the destination type is the same, 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 is not specified, then the existing EncryptionConfiguration 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 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. Use the new version ID to set CurrentDeliveryStreamVersionId in the next call.

## Request Syntax

```
{
    "AmazonOpenSearchServerlessDestinationUpdate": {
        "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
        },
        "CloudWatchLoggingOptions": {
```

```
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
        },
        "CollectionEndpoint": "string",
        "IndexName": "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"
            },
            "ErrorOutputPrefix": "string",
            "Prefix": "string",
```

```
            "RoleARN": "string"
        }
    },
    "AmazonopensearchserviceDestinationUpdate": {
        "BufferingHints": {
            "IntervalInSeconds": number,
            "SizeInMBs": number
        },
        "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
        },
        "ClusterEndpoint": "string",
        "DocumentIdOptions": {
            "DefaultDocumentIdFormat": "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
            },
```

```json
                "CloudWatchLoggingOptions": {
                    "Enabled": boolean,
                    "LogGroupName": "string",
                    "LogStreamName": "string"
                },
                "CompressionFormat": "string",
                "EncryptionConfiguration": {
                    "KMSEncryptionConfig": {
                        "AWSKMSKeyARN": "string"
                    },
                    "NoEncryptionConfig": "string"
                },
                "ErrorOutputPrefix": "string",
                "Prefix": "string",
                "RoleARN": "string"
            },
            "TypeName": "string"
        },
        "CurrentDeliveryStreamVersionId": "string",
        "DeliveryStreamName": "string",
        "DestinationId": "string",
        "ElasticsearchDestinationUpdate": {
            "BufferingHints": {
                "IntervalInSeconds": number,
                "SizeInMBs": number
            },
            "CloudWatchLoggingOptions": {
                "Enabled": boolean,
                "LogGroupName": "string",
                "LogStreamName": "string"
            },
            "ClusterEndpoint": "string",
            "DocumentIdOptions": {
                "DefaultDocumentIdFormat": "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"
        },
        "ErrorOutputPrefix": "string",
        "Prefix": "string",
        "RoleARN": "string"
    },
    "TypeName": "string"
},
"ExtendedS3DestinationUpdate": {
    "BucketARN": "string",
    "BufferingHints": {
        "IntervalInSeconds": number,
        "SizeInMBs": number
    },
    "CloudWatchLoggingOptions": {
        "Enabled": boolean,
```

```
                "LogGroupName": "string",
                "LogStreamName": "string"
            },
            "CompressionFormat": "string",
            "CustomTimeZone": "string",
            "DataFormatConversionConfiguration": {
                "Enabled": boolean,
                "InputFormatConfiguration": {
                    "Deserializer": {
                        "HiveJsonSerDe": {
                            "TimestampFormats": [ "string" ]
                        },
                        "OpenXJsonSerDe": {
                            "CaseInsensitive": boolean,
                            "ColumnToJsonKeyMappings": {
                                "string" : "string"
                            },
                            "ConvertDotsInJsonKeysToUnderscores": boolean
                        }
                    }
                },
                "OutputFormatConfiguration": {
                    "Serializer": {
                        "OrcSerDe": {
                            "BlockSizeBytes": number,
                            "BloomFilterColumns": [ "string" ],
                            "BloomFilterFalsePositiveProbability": number,
                            "Compression": "string",
                            "DictionaryKeyThreshold": number,
                            "EnablePadding": boolean,
                            "FormatVersion": "string",
                            "PaddingTolerance": number,
                            "RowIndexStride": number,
                            "StripeSizeBytes": number
                        },
                        "ParquetSerDe": {
                            "BlockSizeBytes": number,
                            "Compression": "string",
                            "EnableDictionaryCompression": boolean,
                            "MaxPaddingBytes": number,
                            "PageSizeBytes": number,
                            "WriterVersion": "string"
                        }
                    }
```

```
      },
      "SchemaConfiguration": {
        "CatalogId": "string",
        "DatabaseName": "string",
        "Region": "string",
        "RoleARN": "string",
        "TableName": "string",
        "VersionId": "string"
      }
    },
    "DynamicPartitioningConfiguration": {
      "Enabled": boolean,
      "RetryOptions": {
        "DurationInSeconds": number
      }
    },
    "EncryptionConfiguration": {
      "KMSEncryptionConfig": {
        "AWSKMSKeyARN": "string"
      },
      "NoEncryptionConfig": "string"
    },
    "ErrorOutputPrefix": "string",
    "FileExtension": "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"
      },
      "ErrorOutputPrefix": "string",
      "Prefix": "string",
      "RoleARN": "string"
   }
},
"HttpEndpointDestinationUpdate": {
   "BufferingHints": {
      "IntervalInSeconds": number,
      "SizeInMBs": number
   },
   "CloudWatchLoggingOptions": {
      "Enabled": boolean,
      "LogGroupName": "string",
      "LogStreamName": "string"
   },
   "EndpointConfiguration": {
      "AccessKey": "string",
      "Name": "string",
      "Url": "string"
   },
   "ProcessingConfiguration": {
      "Enabled": boolean,
      "Processors": [
         {
            "Parameters": [
               {
                  "ParameterName": "string",
                  "ParameterValue": "string"
               }
```

```
            ],
            "Type": "string"
        }
    ]
},
"RequestConfiguration": {
    "CommonAttributes": [
        {
            "AttributeName": "string",
            "AttributeValue": "string"
        }
    ],
    "ContentEncoding": "string"
},
"RetryOptions": {
    "DurationInSeconds": number
},
"RoleARN": "string",
"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"
    },
    "ErrorOutputPrefix": "string",
    "Prefix": "string",
    "RoleARN": "string"
},
"SecretsManagerConfiguration": {
    "Enabled": boolean,
    "RoleARN": "string",
```

```
                "SecretARN": "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"
            },
            "ErrorOutputPrefix": "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"
            },
            "ErrorOutputPrefix": "string",
            "Prefix": "string",
            "RoleARN": "string"
        },
        "SecretsManagerConfiguration": {
            "Enabled": boolean,
            "RoleARN": "string",
            "SecretARN": "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"
         },
         "ErrorOutputPrefix": "string",
         "Prefix": "string",
         "RoleARN": "string"
      },
      "SnowflakeDestinationUpdate": {
         "AccountUrl": "string",
         "CloudWatchLoggingOptions": {
            "Enabled": boolean,
            "LogGroupName": "string",
            "LogStreamName": "string"
         },
         "ContentColumnName": "string",
         "Database": "string",
         "DataLoadingOption": "string",
         "KeyPassphrase": "string",
         "MetaDataColumnName": "string",
         "PrivateKey": "string",
         "ProcessingConfiguration": {
            "Enabled": boolean,
            "Processors": [
               {
                  "Parameters": [
                     {
                        "ParameterName": "string",
                        "ParameterValue": "string"
                     }
                  ],
                  "Type": "string"
```

```
            }
        ]
    },
    "RetryOptions": {
        "DurationInSeconds": number
    },
    "RoleARN": "string",
    "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"
        },
        "ErrorOutputPrefix": "string",
        "Prefix": "string",
        "RoleARN": "string"
    },
    "Schema": "string",
    "SecretsManagerConfiguration": {
        "Enabled": boolean,
        "RoleARN": "string",
        "SecretARN": "string"
    },
    "SnowflakeRoleConfiguration": {
        "Enabled": boolean,
        "SnowflakeRole": "string"
    },
    "Table": "string",
    "User": "string"
    },
    "SplunkDestinationUpdate": {
```

```
      "BufferingHints": {
         "IntervalInSeconds": number,
         "SizeInMBs": number
      },
      "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"
            },
            "ErrorOutputPrefix": "string",
            "Prefix": "string",
            "RoleARN": "string"
        },
        "SecretsManagerConfiguration": {
            "Enabled": boolean,
            "RoleARN": "string",
            "SecretARN": "string"
        }
    }
}
```

# Request Parameters

The request accepts the following data in JSON format.

**AmazonOpenSearchServerlessDestinationUpdate**

Describes an update for a destination in the Serverless offering for Amazon OpenSearch
Service.

Type: AmazonOpenSearchServerlessDestinationUpdate object

Required: No

**AmazonopensearchserviceDestinationUpdate**

Describes an update for a destination in Amazon OpenSearch Service.

Type: AmazonopensearchserviceDestinationUpdate object

Required: No

**CurrentDeliveryStreamVersionId**

Obtain this value from the `VersionId` result of DeliveryStreamDescription. This value is
required, and helps the service 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

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

The ID of the destination.

Type: String

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

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

Required: Yes

## ElasticsearchDestinationUpdate

Describes an update for a destination in Amazon ES.

Type: ElasticsearchDestinationUpdate object

Required: No

## ExtendedS3DestinationUpdate

Describes an update for a destination in Amazon S3.

Type: ExtendedS3DestinationUpdate object

Required: No

## HttpEndpointDestinationUpdate

Describes an update to the specified HTTP endpoint destination.

Type: HttpEndpointDestinationUpdate object

Required: No

## RedshiftDestinationUpdate

Describes an update for a destination in Amazon Redshift.

Type: RedshiftDestinationUpdate object

Required: No

## S3DestinationUpdate

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

Type: S3DestinationUpdate object

Required: No

## SnowflakeDestinationUpdate

Update to the Snowflake destination configuration settings.

Type: SnowflakeDestinationUpdate object

Required: No

## SplunkDestinationUpdate

Describes an update for a destination in Splunk.

Type: SplunkDestinationUpdate 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](Common Errors).

**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

# Examples

## Example

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

**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 v2](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for JavaScript V3](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)

- [AWS SDK for Ruby V3](#)

# 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:

- AmazonOpenSearchServerlessBufferingHints
- AmazonOpenSearchServerlessDestinationConfiguration
- AmazonOpenSearchServerlessDestinationDescription
- AmazonOpenSearchServerlessDestinationUpdate
- AmazonOpenSearchServerlessRetryOptions
- AmazonopensearchserviceBufferingHints
- AmazonopensearchserviceDestinationConfiguration
- AmazonopensearchserviceDestinationDescription
- AmazonopensearchserviceDestinationUpdate
- AmazonopensearchserviceRetryOptions
- AuthenticationConfiguration
- BufferingHints
- CloudWatchLoggingOptions
- CopyCommand
- DataFormatConversionConfiguration
- DeliveryStreamDescription
- DeliveryStreamEncryptionConfiguration
- DeliveryStreamEncryptionConfigurationInput
- Deserializer
- DestinationDescription

- [DocumentIdOptions](DocumentIdOptions)

- [DynamicPartitioningConfiguration](DynamicPartitioningConfiguration)

- [ElasticsearchBufferingHints](ElasticsearchBufferingHints)

- [ElasticsearchDestinationConfiguration](ElasticsearchDestinationConfiguration)

- [ElasticsearchDestinationDescription](ElasticsearchDestinationDescription)

- [ElasticsearchDestinationUpdate](ElasticsearchDestinationUpdate)

- [ElasticsearchRetryOptions](ElasticsearchRetryOptions)

- [EncryptionConfiguration](EncryptionConfiguration)

- [ExtendedS3DestinationConfiguration](ExtendedS3DestinationConfiguration)

- [ExtendedS3DestinationDescription](ExtendedS3DestinationDescription)

- [ExtendedS3DestinationUpdate](ExtendedS3DestinationUpdate)

- [FailureDescription](FailureDescription)

- [HiveJsonSerDe](HiveJsonSerDe)

- [HttpEndpointBufferingHints](HttpEndpointBufferingHints)

- [HttpEndpointCommonAttribute](HttpEndpointCommonAttribute)

- [HttpEndpointConfiguration](HttpEndpointConfiguration)

- [HttpEndpointDescription](HttpEndpointDescription)

- [HttpEndpointDestinationConfiguration](HttpEndpointDestinationConfiguration)

- [HttpEndpointDestinationDescription](HttpEndpointDestinationDescription)

- [HttpEndpointDestinationUpdate](HttpEndpointDestinationUpdate)

- [HttpEndpointRequestConfiguration](HttpEndpointRequestConfiguration)

- [HttpEndpointRetryOptions](HttpEndpointRetryOptions)

- [InputFormatConfiguration](InputFormatConfiguration)

- [KinesisStreamSourceConfiguration](KinesisStreamSourceConfiguration)

- [KinesisStreamSourceDescription](KinesisStreamSourceDescription)

- [KMSEncryptionConfig](KMSEncryptionConfig)

- [MSKSourceConfiguration](MSKSourceConfiguration)

- [MSKSourceDescription](MSKSourceDescription)

- [OpenXJsonSerDe](OpenXJsonSerDe)

- [OrcSerDe](OrcSerDe)

- OutputFormatConfiguration

- ParquetSerDe

- ProcessingConfiguration

- Processor

- ProcessorParameter

- PutRecordBatchResponseEntry

- Record

- RedshiftDestinationConfiguration

- RedshiftDestinationDescription

- RedshiftDestinationUpdate

- RedshiftRetryOptions

- RetryOptions

- S3DestinationConfiguration

- S3DestinationDescription

- S3DestinationUpdate

- SchemaConfiguration

- SecretsManagerConfiguration

- Serializer

- SnowflakeDestinationConfiguration

- SnowflakeDestinationDescription

- SnowflakeDestinationUpdate

- SnowflakeRetryOptions

- SnowflakeRoleConfiguration

- SnowflakeVpcConfiguration

- SourceDescription

- SplunkBufferingHints

- SplunkDestinationConfiguration

- SplunkDestinationDescription

- SplunkDestinationUpdate

- SplunkRetryOptions

- [Tag](Tag)
- [VpcConfiguration](VpcConfiguration)
- [VpcConfigurationDescription](VpcConfigurationDescription)

# AmazonOpenSearchServerlessBufferingHints

Describes the buffering to perform before delivering data to the Serverless offering for Amazon OpenSearch Service 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 0. 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 Java V2
- AWS SDK for Ruby V3

# AmazonOpenSearchServerlessDestinationConfiguration

Describes the configuration of a destination in the Serverless offering for Amazon OpenSearch Service.

## Contents

**IndexName**

The Serverless offering for Amazon OpenSearch Service index name.

Type: String

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

Pattern: .*

Required: Yes

**RoleARN**

The Amazon Resource Name (ARN) of the IAM role to be assumed by Firehose for calling the Serverless offering for Amazon OpenSearch Service Configuration API and for indexing documents.

Type: String

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

Pattern: arn:.*

Required: Yes

**S3Configuration**

Describes the configuration of a destination in Amazon S3.

Type: [S3DestinationConfiguration](#) object

Required: Yes

**BufferingHints**

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

Type: [AmazonOpenSearchServerlessBufferingHints](#) object

Required: No

## CloudWatchLoggingOptions

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

## CollectionEndpoint

The endpoint to use when communicating with the collection in the Serverless offering for Amazon OpenSearch Service.

Type: String

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

Pattern: `https:.*`

Required: No

## ProcessingConfiguration

Describes a data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

## RetryOptions

The retry behavior in case Firehose is unable to deliver documents to the Serverless offering for Amazon OpenSearch Service. The default value is 300 (5 minutes).

Type: [AmazonOpenSearchServerlessRetryOptions](#) object

Required: No

## S3BackupMode

Defines how documents should be delivered to Amazon S3. When it is set to FailedDocumentsOnly, Firehose writes any documents that could not be indexed to the

configured Amazon S3 destination, with AmazonOpenSearchService-failed/ appended to the key prefix. When set to AllDocuments, Firehose delivers all incoming records to Amazon S3, and also writes failed documents with AmazonOpenSearchService-failed/ appended to the prefix.

Type: String

Valid Values: `FailedDocumentsOnly | AllDocuments`

Required: No

**VpcConfiguration**

The details of the VPC of the Amazon OpenSearch or Amazon OpenSearch Serverless destination.

Type: [VpcConfiguration](#) object

Required: No

# See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# AmazonOpenSearchServerlessDestinationDescription

The destination description in the Serverless offering for Amazon OpenSearch Service.

## Contents

**BufferingHints**

The buffering options.

Type: [AmazonOpenSearchServerlessBufferingHints](#) object

Required: No

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

**CollectionEndpoint**

The endpoint to use when communicating with the collection in the Serverless offering for Amazon OpenSearch Service.

Type: String

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

Pattern: `https:.*`

Required: No

**IndexName**

The Serverless offering for Amazon OpenSearch Service index name.

Type: String

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

Pattern: `.*`

Required: No

**ProcessingConfiguration**

Describes a data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

**RetryOptions**

The Serverless offering for Amazon OpenSearch Service retry options.

Type: [AmazonOpenSearchServerlessRetryOptions](#) 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**

Describes a destination in Amazon S3.

Type: [S3DestinationDescription](#) object

Required: No

**VpcConfigurationDescription**

The details of the VPC of the Amazon ES destination.

Type: [VpcConfigurationDescription](#) object

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# AmazonOpenSearchServerlessDestinationUpdate

Describes an update for a destination in the Serverless offering for Amazon OpenSearch Service.

## Contents

**BufferingHints**

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

Type: AmazonOpenSearchServerlessBufferingHints object

Required: No

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: CloudWatchLoggingOptions object

Required: No

**CollectionEndpoint**

The endpoint to use when communicating with the collection in the Serverless offering for Amazon OpenSearch Service.

Type: String

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

Pattern: `https:.*`

Required: No

**IndexName**

The Serverless offering for Amazon OpenSearch Service index name.

Type: String

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

Pattern: `.*`

Required: No

**ProcessingConfiguration**

Describes a data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

**RetryOptions**

The retry behavior in case Firehose is unable to deliver documents to the Serverless offering for Amazon OpenSearch Service. The default value is 300 (5 minutes).

Type: [AmazonOpenSearchServerlessRetryOptions](#) object

Required: No

**RoleARN**

The Amazon Resource Name (ARN) of the IAM role to be assumed by Firehose for calling the Serverless offering for Amazon OpenSearch Service Configuration API and for indexing documents.

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](#) object

Required: No

# See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# AmazonOpenSearchServerlessRetryOptions

Configures retry behavior in case Firehose is unable to deliver documents to the Serverless offering for Amazon OpenSearch Service.

## Contents

### DurationInSeconds

After an initial failure to deliver to the Serverless offering for Amazon OpenSearch Service, the total amount of time during which Firehose retries 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# AmazonopensearchserviceBufferingHints

Describes the buffering to perform before delivering data to the Amazon OpenSearch Service 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 0. 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 Java V2
- AWS SDK for Ruby V3

# AmazonopensearchserviceDestinationConfiguration

Describes the configuration of a destination in Amazon OpenSearch Service

## Contents

**IndexName**

The ElasticsearAmazon OpenSearch Service index name.

Type: String

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

Pattern: .*

Required: Yes

**RoleARN**

The Amazon Resource Name (ARN) of the IAM role to be assumed by Firehose for calling the Amazon OpenSearch Service Configuration API and for indexing documents.

Type: String

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

Pattern: arn:.*

Required: Yes

**S3Configuration**

Describes the configuration of a destination in Amazon S3.

Type: S3DestinationConfiguration object

Required: Yes

**BufferingHints**

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

Type: [AmazonopensearchserviceBufferingHints](#) object

Required: No

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

**ClusterEndpoint**

The endpoint to use when communicating with the cluster. Specify either this ClusterEndpoint or the DomainARN field.

Type: String

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

Pattern: `https:.*`

Required: No

**DocumentIdOptions**

Indicates the method for setting up document ID. The supported methods are Firehose generated document ID and OpenSearch Service generated document ID.

Type: [DocumentIdOptions](#) object

Required: No

**DomainARN**

The ARN of the Amazon OpenSearch Service 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: No

## IndexRotationPeriod

The Amazon OpenSearch Service index rotation period. Index rotation appends a timestamp to the IndexName to facilitate the expiration of old data.

Type: String

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

Required: No

## ProcessingConfiguration

Describes a data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

## RetryOptions

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

Type: [AmazonopensearchserviceRetryOptions](#) object

Required: No

## S3BackupMode

Defines how documents should be delivered to Amazon S3. When it is set to FailedDocumentsOnly, Firehose writes any documents that could not be indexed to the configured Amazon S3 destination, with AmazonOpenSearchService-failed/ appended to the key prefix. When set to AllDocuments, Firehose delivers all incoming records to Amazon S3, and also writes failed documents with AmazonOpenSearchService-failed/ appended to the prefix.

Type: String

Valid Values: `FailedDocumentsOnly | AllDocuments`

Required: No

**TypeName**

The Amazon OpenSearch Service type name. For Elasticsearch 6.x, there can be only one type per index. If you try to specify a new type for an existing index that already has another type, Firehose returns an error during run time.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 100.

Pattern:  . *

Required: No

**VpcConfiguration**

The details of the VPC of the Amazon OpenSearch or Amazon OpenSearch Serverless destination.

Type: VpcConfiguration object

Required: No

# See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# AmazonopensearchserviceDestinationDescription

The destination description in Amazon OpenSearch Service.

## Contents

**BufferingHints**

The buffering options.

Type: AmazonopensearchserviceBufferingHints object

Required: No

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: CloudWatchLoggingOptions object

Required: No

**ClusterEndpoint**

The endpoint to use when communicating with the cluster. Firehose uses either this ClusterEndpoint or the DomainARN field to send data to Amazon OpenSearch Service.

Type: String

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

Pattern: `https:.*`

Required: No

**DocumentIdOptions**

Indicates the method for setting up document ID. The supported methods are Firehose generated document ID and OpenSearch Service generated document ID.

Type: DocumentIdOptions object

Required: No

## DomainARN

The ARN of the Amazon OpenSearch Service domain.

Type: String

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

Pattern: `arn:.*`

Required: No

## IndexName

The Amazon OpenSearch Service index name.

Type: String

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

Pattern: `.*`

Required: No

## IndexRotationPeriod

The Amazon OpenSearch Service index rotation period

Type: String

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

Required: No

## ProcessingConfiguration

Describes a data processing configuration.

Type: [ProcessingConfiguration](ProcessingConfiguration) object

Required: No

## RetryOptions

The Amazon OpenSearch Service retry options.

Type: [AmazonopensearchserviceRetryOptions](#) 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**

Describes a destination in Amazon S3.

Type: [S3DestinationDescription](#) object

Required: No

**TypeName**

The Amazon OpenSearch Service type name. This applies to Elasticsearch 6.x and lower versions. For Elasticsearch 7.x and OpenSearch Service 1.x, there's no value for TypeName.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 100.

Pattern: `.*`

Required: No

**VpcConfigurationDescription**

The details of the VPC of the Amazon ES destination.

Type: [VpcConfigurationDescription](#) object

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# AmazonopensearchserviceDestinationUpdate

Describes an update for a destination in Amazon OpenSearch Service.

## Contents

**BufferingHints**

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

Type: [AmazonopensearchserviceBufferingHints](#) object

Required: No

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

**ClusterEndpoint**

The endpoint to use when communicating with the cluster. Specify either this ClusterEndpoint or the DomainARN field.

Type: String

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

Pattern: `https:.*`

Required: No

**DocumentIdOptions**

Indicates the method for setting up document ID. The supported methods are Firehose generated document ID and OpenSearch Service generated document ID.

Type: [DocumentIdOptions](#) object

Required: No

**DomainARN**

The ARN of the Amazon OpenSearch Service domain. The IAM role must have permissions for DescribeDomain, DescribeDomains, and DescribeDomainConfig 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 Amazon OpenSearch Service index name.

Type: String

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

Pattern: `.*`

Required: No

**IndexRotationPeriod**

The Amazon OpenSearch Service index rotation period. Index rotation appends a timestamp to IndexName to facilitate the expiration of old data.

Type: String

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

Required: No

**ProcessingConfiguration**

Describes a data processing configuration.

Type: ProcessingConfiguration object

Required: No

**RetryOptions**

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

Type: [AmazonopensearchserviceRetryOptions](#) object

Required: No

**RoleARN**

The Amazon Resource Name (ARN) of the IAM role to be assumed by Firehose for calling the Amazon OpenSearch Service Configuration API and for indexing documents.

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](#) object

Required: No

**TypeName**

The Amazon OpenSearch Service type name. For Elasticsearch 6.x, there can be only one type per index. If you try to specify a new type for an existing index that already has another type, Firehose returns an error during runtime.

If you upgrade Elasticsearch from 6.x to 7.x and don't update your delivery stream, Firehose still delivers data to Elasticsearch with the old index name and type name. If you want to update your delivery stream with a new index name, provide an empty string for TypeName.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 100.

Pattern: `.*`

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# AmazonopensearchserviceRetryOptions

Configures retry behavior in case Firehose is unable to deliver documents to Amazon OpenSearch Service.

## Contents

### DurationInSeconds

After an initial failure to deliver to Amazon OpenSearch Service, the total amount of time during which Firehose retries 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# AuthenticationConfiguration

The authentication configuration of the Amazon MSK cluster.

## Contents

**Connectivity**

The type of connectivity used to access the Amazon MSK cluster.

Type: String

Valid Values: `PUBLIC | PRIVATE`

Required: Yes

**RoleARN**

The ARN of the role used to access the Amazon MSK cluster.

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

# BufferingHints

Describes hints for the buffering to perform before delivering data to the destination. These options are treated as hints, and therefore Firehose might choose to use different values when it is optimal. The `SizeInMBs` and `IntervalInSeconds` parameters are optional. However, if specify a value for one of them, you must also provide a value for the other.

## Contents

**IntervalInSeconds**

Buffer incoming data for the specified period of time, in seconds, before delivering it to the destination. The default value is 300. This parameter is optional but if you specify a value for it, you must also specify a value for `SizeInMBs`, and vice versa.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 900.

Required: No

**SizeInMBs**

Buffer incoming data to the specified size, in MiBs, before delivering it to the destination. The default value is 5. This parameter is optional but if you specify a value for it, you must also specify a value for `IntervalInSeconds`, and vice versa.

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 MiB/sec, the value should be 10 MiB 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# 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

Length Constraints: Minimum length of 0. Maximum length of 512.

Pattern: [\.\-_/#A-Za-z0-9]*

Required: No

**LogStreamName**

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

Type: String

Length Constraints: Minimum length of 0. Maximum length of 512.

Pattern: [^:*]*

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# CopyCommand

Describes a COPY command for Amazon Redshift.

## Contents

**DataTableName**

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

Type: String

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

Pattern: .*

Required: Yes

**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,venuename:25,venuecity:12,venuestate:2,venueseats: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

Length Constraints: Minimum length of 0. Maximum length of 204800.

Pattern: . *

Required: No

**DataTableColumns**

A comma-separated list of column names.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 204800.

Pattern: . *

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# DataFormatConversionConfiguration

Specifies that you want Firehose to convert data from the JSON format to the Parquet or ORC format before writing it to Amazon S3. Firehose uses the serializer and deserializer that you specify, in addition to the column information from the AWS Glue table, to deserialize your input data from JSON and then serialize it to the Parquet or ORC format. For more information, see [Firehose Record Format Conversion](#).

## Contents

**Enabled**

Defaults to `true`. Set it to `false` if you want to disable format conversion while preserving the configuration details.

Type: Boolean

Required: No

**InputFormatConfiguration**

Specifies the deserializer that you want Firehose to use to convert the format of your data from JSON. This parameter is required if `Enabled` is set to true.

Type: [InputFormatConfiguration](#) object

Required: No

**OutputFormatConfiguration**

Specifies the serializer that you want Firehose to use to convert the format of your data to the Parquet or ORC format. This parameter is required if `Enabled` is set to true.

Type: [OutputFormatConfiguration](#) object

Required: No

**SchemaConfiguration**

Specifies the AWS Glue Data Catalog table that contains the column information. This parameter is required if `Enabled` is set to true.

Type: [SchemaConfiguration](#) object

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# DeliveryStreamDescription

Contains information about a delivery stream.

## Contents

**DeliveryStreamARN**

The Amazon Resource Name (ARN) of the delivery stream. For more information, see Amazon Resource Names (ARNs) and AWS Service Namespaces.

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. If the status of a delivery stream is `CREATING_FAILED`, this status doesn't change, and you can't invoke `CreateDeliveryStream` again on it. However, you can invoke the DeleteDeliveryStream operation to delete it.

Type: String

Valid Values: `CREATING | CREATING_FAILED | DELETING | DELETING_FAILED | 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 data stream as a source.

Type: String

Valid Values: `DirectPut | KinesisStreamAsSource | MSKAsSource`

Required: Yes

**Destinations**

The destinations.

Type: Array of [DestinationDescription](DestinationDescription) objects

Required: Yes

**HasMoreDestinations**

Indicates whether there are more destinations available to list.

Type: Boolean

Required: Yes

**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

**CreateTimestamp**

The date and time that the delivery stream was created.

Type: Timestamp

Required: No

**DeliveryStreamEncryptionConfiguration**

Indicates the server-side encryption (SSE) status for the delivery stream.

Type: DeliveryStreamEncryptionConfiguration object

Required: No

**FailureDescription**

Provides details in case one of the following operations fails due to an error related to KMS: CreateDeliveryStream, DeleteDeliveryStream, StartDeliveryStreamEncryption, StopDeliveryStreamEncryption.

Type: FailureDescription object

Required: No

**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 object describing the source Kinesis data stream.

Type: SourceDescription object

Required: No

# See Also

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

- AWS SDK for C++

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

# DeliveryStreamEncryptionConfiguration

Contains information about the server-side encryption (SSE) status for the delivery stream, the type customer master key (CMK) in use, if any, and the ARN of the CMK. You can get `DeliveryStreamEncryptionConfiguration` by invoking the [DescribeDeliveryStream](DescribeDeliveryStream) operation.

## Contents

**FailureDescription**

Provides details in case one of the following operations fails due to an error related to KMS: [CreateDeliveryStream](CreateDeliveryStream), [DeleteDeliveryStream](DeleteDeliveryStream), [StartDeliveryStreamEncryption](StartDeliveryStreamEncryption), [StopDeliveryStreamEncryption](StopDeliveryStreamEncryption).

Type: [FailureDescription](FailureDescription) object

Required: No

**KeyARN**

If `KeyType` is `CUSTOMER_MANAGED_CMK`, this field contains the ARN of the customer managed CMK. If `KeyType` is `AWS_OWNED_CMK`, `DeliveryStreamEncryptionConfiguration` doesn't contain a value for `KeyARN`.

Type: String

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

Pattern: `arn:.*`

Required: No

**KeyType**

Indicates the type of customer master key (CMK) that is used for encryption. The default setting is `AWS_OWNED_CMK`. For more information about CMKs, see [Customer Master Keys (CMKs)](Customer Master Keys (CMKs)).

Type: String

Valid Values: `AWS_OWNED_CMK | CUSTOMER_MANAGED_CMK`

Required: No

**Status**

This is the server-side encryption (SSE) status for the delivery stream. For a full description of the different values of this status, see StartDeliveryStreamEncryption and StopDeliveryStreamEncryption. If this status is ENABLING_FAILED or DISABLING_FAILED, it is the status of the most recent attempt to enable or disable SSE, respectively.

Type: String

Valid Values: ENABLED | ENABLING | ENABLING_FAILED | DISABLED | DISABLING | DISABLING_FAILED

Required: No

# See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# DeliveryStreamEncryptionConfigurationInput

Specifies the type and Amazon Resource Name (ARN) of the CMK to use for Server-Side Encryption (SSE).

## Contents

**KeyType**

Indicates the type of customer master key (CMK) to use for encryption. The default setting is `AWS_OWNED_CMK`. For more information about CMKs, see [Customer Master Keys (CMKs)](#). When you invoke [CreateDeliveryStream](#) or [StartDeliveryStreamEncryption](#) with KeyType set to CUSTOMER_MANAGED_CMK, Firehose invokes the Amazon KMS operation [CreateGrant](#) to create a grant that allows the Firehose service to use the customer managed CMK to perform encryption and decryption. Firehose manages that grant.

When you invoke [StartDeliveryStreamEncryption](#) to change the CMK for a delivery stream that is encrypted with a customer managed CMK, Firehose schedules the grant it had on the old CMK for retirement.

You can use a CMK of type CUSTOMER_MANAGED_CMK to encrypt up to 500 delivery streams. If a [CreateDeliveryStream](#) or [StartDeliveryStreamEncryption](#) operation exceeds this limit, Firehose throws a `LimitExceededException`.

> ⚠️ **Important**
>
> To encrypt your delivery stream, use symmetric CMKs. Firehose doesn't support asymmetric CMKs. For information about symmetric and asymmetric CMKs, see [About Symmetric and Asymmetric CMKs](#) in the AWS Key Management Service developer guide.

Type: String

Valid Values: `AWS_OWNED_CMK` | `CUSTOMER_MANAGED_CMK`

Required: Yes

**KeyARN**

If you set KeyType to CUSTOMER_MANAGED_CMK, you must specify the Amazon Resource Name (ARN) of the CMK. If you set KeyType to   AWS_OWNED_CMK, Firehose uses a service-account CMK.

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 Java V2
- AWS SDK for Ruby V3

# Deserializer

The deserializer you want Firehose to use for converting the input data from JSON. Firehose then serializes the data to its final format using the [Serializer](). Firehose supports two types of deserializers: the [Apache Hive JSON SerDe]() and the [OpenX JSON SerDe]().

## Contents

**HiveJsonSerDe**

The native Hive / HCatalog JsonSerDe. Used by Firehose for deserializing data, which means converting it from the JSON format in preparation for serializing it to the Parquet or ORC format. This is one of two deserializers you can choose, depending on which one offers the functionality you need. The other option is the OpenX SerDe.

Type: [HiveJsonSerDe]() object

Required: No

**OpenXJsonSerDe**

The OpenX SerDe. Used by Firehose for deserializing data, which means converting it from the JSON format in preparation for serializing it to the Parquet or ORC format. This is one of two deserializers you can choose, depending on which one offers the functionality you need. The other option is the native Hive / HCatalog JsonSerDe.

Type: [OpenXJsonSerDe]() object

Required: No

## See Also

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

- [AWS SDK for C++]()
- [AWS SDK for Java V2]()
- [AWS SDK for Ruby V3]()

# 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.

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

Required: Yes

**AmazonOpenSearchServerlessDestinationDescription**

The destination in the Serverless offering for Amazon OpenSearch Service.

Type: [AmazonOpenSearchServerlessDestinationDescription](#) object

Required: No

**AmazonopensearchserviceDestinationDescription**

The destination in Amazon OpenSearch Service.

Type: [AmazonopensearchserviceDestinationDescription](#) object

Required: No

**ElasticsearchDestinationDescription**

The destination in Amazon ES.

Type: [ElasticsearchDestinationDescription](#) object

Required: No

**ExtendedS3DestinationDescription**

The destination in Amazon S3.

Type: [ExtendedS3DestinationDescription](#) object

Required: No

**HttpEndpointDestinationDescription**

Describes the specified HTTP endpoint destination.

Type: HttpEndpointDestinationDescription object

Required: No

**RedshiftDestinationDescription**

The destination in Amazon Redshift.

Type: RedshiftDestinationDescription object

Required: No

**S3DestinationDescription**

[Deprecated] The destination in Amazon S3.

Type: S3DestinationDescription object

Required: No

**SnowflakeDestinationDescription**

Optional description for the destination

Type: SnowflakeDestinationDescription object

Required: No

**SplunkDestinationDescription**

The destination in Splunk.

Type: SplunkDestinationDescription object

Required: No

# See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# DocumentIdOptions

Indicates the method for setting up document ID. The supported methods are Firehose generated document ID and OpenSearch Service generated document ID.

## Contents

**DefaultDocumentIdFormat**

When the `FIREHOSE_DEFAULT` option is chosen, Firehose generates a unique document ID for each record based on a unique internal identifier. The generated document ID is stable across multiple delivery attempts, which helps prevent the same record from being indexed multiple times with different document IDs.

When the `NO_DOCUMENT_ID` option is chosen, Firehose does not include any document IDs in the requests it sends to the Amazon OpenSearch Service. This causes the Amazon OpenSearch Service domain to generate document IDs. In case of multiple delivery attempts, this may cause the same record to be indexed more than once with different document IDs. This option enables write-heavy operations, such as the ingestion of logs and observability data, to consume less resources in the Amazon OpenSearch Service domain, resulting in improved performance.

Type: String

Valid Values: `FIREHOSE_DEFAULT | NO_DOCUMENT_ID`

Required: Yes

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# DynamicPartitioningConfiguration

The configuration of the dynamic partitioning mechanism that creates smaller data sets from the streaming data by partitioning it based on partition keys. Currently, dynamic partitioning is only supported for Amazon S3 destinations.

## Contents

### Enabled

Specifies that the dynamic partitioning is enabled for this Firehose delivery stream.

Type: Boolean

Required: No

### RetryOptions

The retry behavior in case Firehose is unable to deliver data to an Amazon S3 prefix.

Type: RetryOptions object

Required: No

## See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# 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 0. 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# ElasticsearchDestinationConfiguration

Describes the configuration of a destination in Amazon ES.

## Contents

**IndexName**

The Elasticsearch index name.

Type: String

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

Pattern: `.*`

Required: Yes

**RoleARN**

The Amazon Resource Name (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 [Grant Firehose Access to an Amazon S3 Destination](#) and [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: Yes

**S3Configuration**

The configuration for the backup Amazon S3 location.

Type: [S3DestinationConfiguration](#) object

Required: Yes

**BufferingHints**

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

Type: [ElasticsearchBufferingHints](#) object

Required: No

## CloudWatchLoggingOptions

The Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

## ClusterEndpoint

The endpoint to use when communicating with the cluster. Specify either this `ClusterEndpoint` or the `DomainARN` field.

Type: String

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

Pattern: `https:.*`

Required: No

## DocumentIdOptions

Indicates the method for setting up document ID. The supported methods are Firehose generated document ID and OpenSearch Service generated document ID.

Type: [DocumentIdOptions](#) object

Required: No

## DomainARN

The ARN of the Amazon ES domain. The IAM role must have permissions for `DescribeDomain`, `DescribeDomains`, and `DescribeDomainConfig` after assuming the role specified in **RoleARN**. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Specify either `ClusterEndpoint` or `DomainARN`.

Type: String

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

Pattern: `arn:.*`

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 the Amazon ES Destination](#). The default value is `OneDay`.

Type: String

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

Required: No

**ProcessingConfiguration**

The data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

**RetryOptions**

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

Type: [ElasticsearchRetryOptions](#) object

Required: No

**S3BackupMode**

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

You can't change this backup mode after you create the delivery stream.

Type: String

Valid Values: `FailedDocumentsOnly | AllDocuments`

Required: No

**TypeName**

The Elasticsearch type name. For Elasticsearch 6.x, there can be only one type per index. If you try to specify a new type for an existing index that already has another type, Firehose returns an error during run time.

For Elasticsearch 7.x, don't specify a `TypeName`.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 100.

Pattern: `. *`

Required: No

**VpcConfiguration**

The details of the VPC of the Amazon destination.

Type: [VpcConfiguration](#) object

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# ElasticsearchDestinationDescription

The destination description in Amazon ES.

## Contents

**BufferingHints**

The buffering options.

Type: [ElasticsearchBufferingHints](#) object

Required: No

**CloudWatchLoggingOptions**

The Amazon CloudWatch logging options.

Type: [CloudWatchLoggingOptions](#) object

Required: No

**ClusterEndpoint**

The endpoint to use when communicating with the cluster. Firehose uses either this `ClusterEndpoint` or the `DomainARN` field to send data to Amazon ES.

Type: String

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

Pattern: `https:.*`

Required: No

**DocumentIdOptions**

Indicates the method for setting up document ID. The supported methods are Firehose generated document ID and OpenSearch Service generated document ID.

Type: [DocumentIdOptions](#) object

Required: No

**DomainARN**

The ARN of the Amazon ES domain. For more information, see Amazon Resource Names (ARNs) and AWS Service Namespaces.

Firehose uses either `ClusterEndpoint` or `DomainARN` to send data to Amazon ES.

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.

Pattern: `.*`

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 object

Required: No

**RetryOptions**

The Amazon ES retry options.

Type: [ElasticsearchRetryOptions](#) object

Required: No

**RoleARN**

The Amazon Resource Name (ARN) of the AWS credentials. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

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](#) object

Required: No

**TypeName**

The Elasticsearch type name. This applies to Elasticsearch 6.x and lower versions. For Elasticsearch 7.x and OpenSearch Service 1.x, there's no value for TypeName.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 100.

Pattern: . *

Required: No

**VpcConfigurationDescription**

The details of the VPC of the Amazon OpenSearch or the Amazon OpenSearch Serverless destination.

Type: VpcConfigurationDescription object

Required: No

# See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# 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 object

Required: No

**CloudWatchLoggingOptions**

The CloudWatch logging options for your delivery stream.

Type: CloudWatchLoggingOptions object

Required: No

**ClusterEndpoint**

The endpoint to use when communicating with the cluster. Specify either this `ClusterEndpoint` or the `DomainARN` field.

Type: String

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

Pattern: `https:.*`

Required: No

**DocumentIdOptions**

Indicates the method for setting up document ID. The supported methods are Firehose generated document ID and OpenSearch Service generated document ID.

Type: DocumentIdOptions object

Required: No

**DomainARN**

The ARN of the Amazon ES domain. The IAM role must have permissions for `DescribeDomain`, `DescribeDomains`, and `DescribeDomainConfig` after assuming the IAM role specified in RoleARN. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Specify either `ClusterEndpoint` or `DomainARN`.

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.

Pattern: `.*`

Required: No

**IndexRotationPeriod**

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

Type: String

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

Required: No

**ProcessingConfiguration**

The data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

**RetryOptions**

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

Type: [ElasticsearchRetryOptions](#) object

Required: No

**RoleARN**

The Amazon Resource Name (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 [Grant Firehose Access to an Amazon S3 Destination](#) and [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: No

**S3Update**

The Amazon S3 destination.

Type: [S3DestinationUpdate](#) object

Required: No

**TypeName**

The Elasticsearch type name. For Elasticsearch 6.x, there can be only one type per index. If you try to specify a new type for an existing index that already has another type, Firehose returns an error during runtime.

If you upgrade Elasticsearch from 6.x to 7.x and don't update your delivery stream, Firehose still delivers data to Elasticsearch with the old index name and type name. If you want to update your delivery stream with a new index name, provide an empty string for TypeName.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 100.

Pattern: . *

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# ElasticsearchRetryOptions

Configures retry behavior in case 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 retries 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# EncryptionConfiguration

Describes the encryption for a destination in Amazon S3.

## Contents

**KMSEncryptionConfig**

The encryption key.

Type: [KMSEncryptionConfig](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# ExtendedS3DestinationConfiguration

Describes the configuration of a destination in Amazon S3.

## Contents

**BucketARN**

The ARN of the S3 bucket. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: Yes

**RoleARN**

The Amazon Resource Name (ARN) of the AWS credentials. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: Yes

**BufferingHints**

The buffering option.

Type: [BufferingHints](#) object

Required: No

**CloudWatchLoggingOptions**

The Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

**CompressionFormat**

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

Type: String

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

Required: No

**CustomTimeZone**

The time zone you prefer. UTC is the default.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 50.

Required: No

**DataFormatConversionConfiguration**

The serializer, deserializer, and schema for converting data from the JSON format to the Parquet or ORC format before writing it to Amazon S3.

Type: [DataFormatConversionConfiguration](#) object

Required: No

**DynamicPartitioningConfiguration**

The configuration of the dynamic partitioning mechanism that creates smaller data sets from the streaming data by partitioning it based on partition keys. Currently, dynamic partitioning is only supported for Amazon S3 destinations.

Type: [DynamicPartitioningConfiguration](#) object

Required: No

**EncryptionConfiguration**

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

Type: [EncryptionConfiguration](#) object

Required: No

**ErrorOutputPrefix**

A prefix that Firehose evaluates and adds to failed records before writing them to S3. This prefix appears immediately following the bucket name. For information about how to specify this prefix, see [Custom Prefixes for Amazon S3 Objects](#).

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: . *

Required: No

**FileExtension**

Specify a file extension. It will override the default file extension

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: ^$|\.[0-9a-z!\-_.*'()]+

Required: No

**Prefix**

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered Amazon S3 files. You can also specify a custom prefix, as described in [Custom Prefixes for Amazon S3 Objects](#).

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: . *

Required: No

**ProcessingConfiguration**

The data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

**S3BackupConfiguration**

The configuration for backup in Amazon S3.

Type: S3DestinationConfiguration object

Required: No

**S3BackupMode**

The Amazon S3 backup mode. After you create a delivery stream, you can update it to enable Amazon S3 backup if it is disabled. If backup is enabled, you can't update the delivery stream to disable it.

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 Java V2
- AWS SDK for Ruby V3

# ExtendedS3DestinationDescription

Describes a destination in Amazon S3.

## Contents

**BucketARN**

The ARN of the S3 bucket. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: Yes

**BufferingHints**

The buffering option.

Type: [BufferingHints](#) object

Required: Yes

**CompressionFormat**

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

Type: String

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

Required: Yes

**EncryptionConfiguration**

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

Type: [EncryptionConfiguration](#) object

Required: Yes

**RoleARN**

The Amazon Resource Name (ARN) of the AWS credentials. For more information, see Amazon Resource Names (ARNs) and AWS Service Namespaces.

Type: String

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

Pattern: `arn:.*`

Required: Yes

**CloudWatchLoggingOptions**

The Amazon CloudWatch logging options for your delivery stream.

Type: CloudWatchLoggingOptions object

Required: No

**CustomTimeZone**

The time zone you prefer. UTC is the default.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 50.

Required: No

**DataFormatConversionConfiguration**

The serializer, deserializer, and schema for converting data from the JSON format to the Parquet or ORC format before writing it to Amazon S3.

Type: DataFormatConversionConfiguration object

Required: No

**DynamicPartitioningConfiguration**

The configuration of the dynamic partitioning mechanism that creates smaller data sets from the streaming data by partitioning it based on partition keys. Currently, dynamic partitioning is only supported for Amazon S3 destinations.

Type: [DynamicPartitioningConfiguration](#) object

Required: No

**ErrorOutputPrefix**

A prefix that Firehose evaluates and adds to failed records before writing them to S3. This prefix appears immediately following the bucket name. For information about how to specify this prefix, see [Custom Prefixes for Amazon S3 Objects](#).

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: .*

Required: No

**FileExtension**

Specify a file extension. It will override the default file extension

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: ^$|\.[0-9a-z!\-_.*'()]+

Required: No

**Prefix**

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered Amazon S3 files. You can also specify a custom prefix, as described in [Custom Prefixes for Amazon S3 Objects](#).

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: .*

Required: No

**ProcessingConfiguration**

The data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

**S3BackupDescription**

The configuration for backup in Amazon S3.

Type: [S3DestinationDescription](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# ExtendedS3DestinationUpdate

Describes an update for a destination in Amazon S3.

## Contents

**BucketARN**

The ARN of the S3 bucket. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: No

**BufferingHints**

The buffering option.

Type: [BufferingHints](#) object

Required: No

**CloudWatchLoggingOptions**

The Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

**CompressionFormat**

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

Type: String

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

Required: No

**CustomTimeZone**

The time zone you prefer. UTC is the default.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 50.

Required: No

**DataFormatConversionConfiguration**

The serializer, deserializer, and schema for converting data from the JSON format to the Parquet or ORC format before writing it to Amazon S3.

Type: DataFormatConversionConfiguration object

Required: No

**DynamicPartitioningConfiguration**

The configuration of the dynamic partitioning mechanism that creates smaller data sets from the streaming data by partitioning it based on partition keys. Currently, dynamic partitioning is only supported for Amazon S3 destinations.

Type: DynamicPartitioningConfiguration object

Required: No

**EncryptionConfiguration**

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

Type: EncryptionConfiguration object

Required: No

**ErrorOutputPrefix**

A prefix that Firehose evaluates and adds to failed records before writing them to S3. This prefix appears immediately following the bucket name. For information about how to specify this prefix, see Custom Prefixes for Amazon S3 Objects.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: . *

Required: No

## FileExtension

Specify a file extension. It will override the default file extension

Type: String

Length Constraints: Minimum length of 0. Maximum length of 128.

Pattern: ^$|\.[0-9a-z!\-_.*'()]+

Required: No

## Prefix

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered Amazon S3 files. You can also specify a custom prefix, as described in Custom Prefixes for Amazon S3 Objects.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: . *

Required: No

## ProcessingConfiguration

The data processing configuration.

Type: ProcessingConfiguration object

Required: No

## RoleARN

The Amazon Resource Name (ARN) of the AWS credentials. For more information, see Amazon Resource Names (ARNs) and AWS Service Namespaces.

Type: String

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

Pattern: `arn:.*`

Required: No

**S3BackupMode**

You can update a delivery stream to enable Amazon S3 backup if it is disabled. If backup is enabled, you can't update the delivery stream to disable it.

Type: String

Valid Values: `Disabled | Enabled`

Required: No

**S3BackupUpdate**

The Amazon S3 destination for backup.

Type: [S3DestinationUpdate](#) object

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# FailureDescription

Provides details in case one of the following operations fails due to an error related to KMS: [CreateDeliveryStream](), [DeleteDeliveryStream](), [StartDeliveryStreamEncryption](), [StopDeliveryStreamEncryption]().

## Contents

### Details

A message providing details about the error that caused the failure.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^(?!\s*$).+`

Required: Yes

### Type

The type of error that caused the failure.

Type: String

Valid Values: `RETIRE_KMS_GRANT_FAILED | CREATE_KMS_GRANT_FAILED | KMS_ACCESS_DENIED | DISABLED_KMS_KEY | INVALID_KMS_KEY | KMS_KEY_NOT_FOUND | KMS_OPT_IN_REQUIRED | CREATE_ENI_FAILED | DELETE_ENI_FAILED | SUBNET_NOT_FOUND | SECURITY_GROUP_NOT_FOUND | ENI_ACCESS_DENIED | SUBNET_ACCESS_DENIED | SECURITY_GROUP_ACCESS_DENIED | UNKNOWN_ERROR`

Required: Yes

## See Also

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

- [AWS SDK for C++]()

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

# HiveJsonSerDe

The native Hive / HCatalog JsonSerDe. Used by Firehose for deserializing data, which means converting it from the JSON format in preparation for serializing it to the Parquet or ORC format. This is one of two deserializers you can choose, depending on which one offers the functionality you need. The other option is the OpenX SerDe.

## Contents

**TimestampFormats**

Indicates how you want Firehose to parse the date and timestamps that may be present in your input data JSON. To specify these format strings, follow the pattern syntax of JodaTime's DateTimeFormat format strings. For more information, see [Class DateTimeFormat](). You can also use the special value `millis` to parse timestamps in epoch milliseconds. If you don't specify a format, Firehose uses `java.sql.Timestamp::valueOf` by default.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: `^(?!\s*$).+`

Required: No

## See Also

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

- [AWS SDK for C++]()
- [AWS SDK for Java V2]()
- [AWS SDK for Ruby V3]()

# HttpEndpointBufferingHints

Describes the buffering options that can be applied before data is delivered to the HTTP endpoint destination. Firehose treats these options as hints, and it might choose to use more optimal values. The `SizeInMBs` and `IntervalInSeconds` parameters are optional. However, if specify a value for one of them, you must also provide a value for the other.

## 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 0. 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 64.

Required: No

## See Also

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

- [AWS SDK for C++](#)

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

# HttpEndpointCommonAttribute

Describes the metadata that's delivered to the specified HTTP endpoint destination.

## Contents

**AttributeName**

The name of the HTTP endpoint common attribute.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^(?!\s*$).+`

Required: Yes

**AttributeValue**

The value of the HTTP endpoint common attribute.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: `.*`

Required: Yes

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# HttpEndpointConfiguration

Describes the configuration of the HTTP endpoint to which Kinesis Firehose delivers data.

## Contents

**Url**

The URL of the HTTP endpoint selected as the destination.

> ⚠️ **Important**
>
> If you choose an HTTP endpoint as your destination, review and follow the instructions in the [Appendix - HTTP Endpoint Delivery Request and Response Specifications](#).

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1000.

Pattern: `https://.*`

Required: Yes

**AccessKey**

The access key required for Kinesis Firehose to authenticate with the HTTP endpoint selected as the destination.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 4096.

Pattern: `.*`

Required: No

**Name**

The name of the HTTP endpoint selected as the destination.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^(?!\s*$).+`

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# HttpEndpointDescription

Describes the HTTP endpoint selected as the destination.

## Contents

**Name**

The name of the HTTP endpoint selected as the destination.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 256.

Pattern: `^(?!\s*$).+`

Required: No

**Url**

The URL of the HTTP endpoint selected as the destination.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1000.

Pattern: `https://.*`

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# HttpEndpointDestinationConfiguration

Describes the configuration of the HTTP endpoint destination.

## Contents

**EndpointConfiguration**

The configuration of the HTTP endpoint selected as the destination.

Type: HttpEndpointConfiguration object

Required: Yes

**S3Configuration**

Describes the configuration of a destination in Amazon S3.

Type: S3DestinationConfiguration object

Required: Yes

**BufferingHints**

The buffering options that can be used before data is delivered to the specified destination. Firehose treats these options as hints, and it might choose to use more optimal values. The `SizeInMBs` and `IntervalInSeconds` parameters are optional. However, if you specify a value for one of them, you must also provide a value for the other.

Type: HttpEndpointBufferingHints object

Required: No

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: CloudWatchLoggingOptions object

Required: No

**ProcessingConfiguration**

Describes a data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

**RequestConfiguration**

The configuration of the request sent to the HTTP endpoint that is specified as the destination.

Type: [HttpEndpointRequestConfiguration](#) object

Required: No

**RetryOptions**

Describes the retry behavior in case Firehose is unable to deliver data to the specified HTTP endpoint destination, or if it doesn't receive a valid acknowledgment of receipt from the specified HTTP endpoint destination.

Type: [HttpEndpointRetryOptions](#) object

Required: No

**RoleARN**

Firehose uses this IAM role for all the permissions that the delivery stream needs.

Type: String

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

Pattern: `arn:.*`

Required: No

**S3BackupMode**

Describes the S3 bucket backup options for the data that Firehose delivers to the HTTP endpoint destination. You can back up all documents (`AllData`) or only the documents that Firehose could not deliver to the specified HTTP endpoint destination (`FailedDataOnly`).

Type: String

Valid Values: `FailedDataOnly | AllData`

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for HTTP Endpoint destination.

Type: SecretsManagerConfiguration object

Required: No

## See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# HttpEndpointDestinationDescription

Describes the HTTP endpoint destination.

## Contents

**BufferingHints**

Describes buffering options that can be applied to the data before it is delivered to the HTTPS endpoint destination. Firehose teats these options as hints, and it might choose to use more optimal values. The `SizeInMBs` and `IntervalInSeconds` parameters are optional. However, if specify a value for one of them, you must also provide a value for the other.

Type: HttpEndpointBufferingHints object

Required: No

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: CloudWatchLoggingOptions object

Required: No

**EndpointConfiguration**

The configuration of the specified HTTP endpoint destination.

Type: HttpEndpointDescription object

Required: No

**ProcessingConfiguration**

Describes a data processing configuration.

Type: ProcessingConfiguration object

Required: No

**RequestConfiguration**

The configuration of request sent to the HTTP endpoint specified as the destination.

Type: [HttpEndpointRequestConfiguration](HttpEndpointRequestConfiguration) object

Required: No

**RetryOptions**

Describes the retry behavior in case Firehose is unable to deliver data to the specified HTTP endpoint destination, or if it doesn't receive a valid acknowledgment of receipt from the specified HTTP endpoint destination.

Type: [HttpEndpointRetryOptions](HttpEndpointRetryOptions) object

Required: No

**RoleARN**

Firehose uses this IAM role for all the permissions that the delivery stream needs.

Type: String

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

Pattern: `arn:.*`

Required: No

**S3BackupMode**

Describes the S3 bucket backup options for the data that Kinesis Firehose delivers to the HTTP endpoint destination. You can back up all documents (`AllData`) or only the documents that Firehose could not deliver to the specified HTTP endpoint destination (`FailedDataOnly`).

Type: String

Valid Values: `FailedDataOnly | AllData`

Required: No

**S3DestinationDescription**

Describes a destination in Amazon S3.

Type: [S3DestinationDescription](S3DestinationDescription) object

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for HTTP Endpoint destination.

Type: SecretsManagerConfiguration object

Required: No

## See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# HttpEndpointDestinationUpdate

Updates the specified HTTP endpoint destination.

## Contents

**BufferingHints**

Describes buffering options that can be applied to the data before it is delivered to the HTTPS endpoint destination. Firehose teats these options as hints, and it might choose to use more optimal values. The `SizeInMBs` and `IntervalInSeconds` parameters are optional. However, if specify a value for one of them, you must also provide a value for the other.

Type: HttpEndpointBufferingHints object

Required: No

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: CloudWatchLoggingOptions object

Required: No

**EndpointConfiguration**

Describes the configuration of the HTTP endpoint destination.

Type: HttpEndpointConfiguration object

Required: No

**ProcessingConfiguration**

Describes a data processing configuration.

Type: ProcessingConfiguration object

Required: No

**RequestConfiguration**

The configuration of the request sent to the HTTP endpoint specified as the destination.

Type: [HttpEndpointRequestConfiguration](#) object

Required: No

**RetryOptions**

Describes the retry behavior in case Firehose is unable to deliver data to the specified HTTP endpoint destination, or if it doesn't receive a valid acknowledgment of receipt from the specified HTTP endpoint destination.

Type: [HttpEndpointRetryOptions](#) object

Required: No

**RoleARN**

Firehose uses this IAM role for all the permissions that the delivery stream needs.

Type: String

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

Pattern: `arn:.*`

Required: No

**S3BackupMode**

Describes the S3 bucket backup options for the data that Kinesis Firehose delivers to the HTTP endpoint destination. You can back up all documents (`AllData`) or only the documents that Firehose could not deliver to the specified HTTP endpoint destination (`FailedDataOnly`).

Type: String

Valid Values: `FailedDataOnly | AllData`

Required: No

**S3Update**

Describes an update for a destination in Amazon S3.

Type: [S3DestinationUpdate](#) object

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for HTTP Endpoint destination.

Type: SecretsManagerConfiguration object

Required: No

# See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# HttpEndpointRequestConfiguration

The configuration of the HTTP endpoint request.

## Contents

**CommonAttributes**

Describes the metadata sent to the HTTP endpoint destination.

Type: Array of [HttpEndpointCommonAttribute](#) objects

Array Members: Minimum number of 0 items. Maximum number of 50 items.

Required: No

**ContentEncoding**

Firehose uses the content encoding to compress the body of a request before sending the request to the destination. For more information, see [Content-Encoding](#) in MDN Web Docs, the official Mozilla documentation.

Type: String

Valid Values: `NONE | GZIP`

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# HttpEndpointRetryOptions

Describes the retry behavior in case Firehose is unable to deliver data to the specified HTTP endpoint destination, or if it doesn't receive a valid acknowledgment of receipt from the specified HTTP endpoint destination.

## Contents

### DurationInSeconds

The total amount of time that Firehose spends on retries. This duration starts after the initial attempt to send data to the custom destination via HTTPS endpoint fails. It doesn't include the periods during which Firehose waits for acknowledgment from the specified destination 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# InputFormatConfiguration

Specifies the deserializer you want to use to convert the format of the input data. This parameter is required if `Enabled` is set to true.

## Contents

**Deserializer**

Specifies which deserializer to use. You can choose either the Apache Hive JSON SerDe or the OpenX JSON SerDe. If both are non-null, the server rejects the request.

Type: [Deserializer](#) object

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# KinesisStreamSourceConfiguration

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

## Contents

**KinesisStreamARN**

The ARN of the source Kinesis data stream. For more information, see Amazon Kinesis Data Streams ARN Format.

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 data stream. For more information, see AWS Identity and Access Management (IAM) ARN Format.

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 Java V2
- AWS SDK for Ruby V3

# KinesisStreamSourceDescription

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

## Contents

**DeliveryStartTimestamp**

Firehose starts retrieving records from the Kinesis data stream starting with this timestamp.

Type: Timestamp

Required: No

**KinesisStreamARN**

The Amazon Resource Name (ARN) of the source Kinesis data stream. For more information, see [Amazon Kinesis Data Streams ARN Format](#).

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 data stream. For more information, see [AWS Identity and Access Management (IAM) ARN Format](#).

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

# 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 AWS Region as the destination Amazon S3 bucket. For more information, see Amazon Resource Names (ARNs) and AWS Service Namespaces.

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 Java V2
- AWS SDK for Ruby V3

# MSKSourceConfiguration

The configuration for the Amazon MSK cluster to be used as the source for a delivery stream.

## Contents

**AuthenticationConfiguration**

The authentication configuration of the Amazon MSK cluster.

Type: AuthenticationConfiguration object

Required: Yes

**MSKClusterARN**

The ARN of the Amazon MSK cluster.

Type: String

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

Pattern: `arn:.*`

Required: Yes

**TopicName**

The topic name within the Amazon MSK cluster.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: Yes

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# MSKSourceDescription

Details about the Amazon MSK cluster used as the source for a Firehose delivery stream.

## Contents

**AuthenticationConfiguration**

The authentication configuration of the Amazon MSK cluster.

Type: AuthenticationConfiguration object

Required: No

**DeliveryStartTimestamp**

Firehose starts retrieving records from the topic within the Amazon MSK cluster starting with this timestamp.

Type: Timestamp

Required: No

**MSKClusterARN**

The ARN of the Amazon MSK cluster.

Type: String

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

Pattern: `arn:.*`

Required: No

**TopicName**

The topic name within the Amazon MSK cluster.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

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

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# OpenXJsonSerDe

The OpenX SerDe. Used by Firehose for deserializing data, which means converting it from the JSON format in preparation for serializing it to the Parquet or ORC format. This is one of two deserializers you can choose, depending on which one offers the functionality you need. The other option is the native Hive / HCatalog JsonSerDe.

## Contents

**CaseInsensitive**

When set to `true`, which is the default, Firehose converts JSON keys to lowercase before deserializing them.

Type: Boolean

Required: No

**ColumnToJsonKeyMappings**

Maps column names to JSON keys that aren't identical to the column names. This is useful when the JSON contains keys that are Hive keywords. For example, `timestamp` is a Hive keyword. If you have a JSON key named `timestamp`, set this parameter to `{"ts":` `"timestamp"}` to map this key to a column named `ts`.

Type: String to string map

Key Length Constraints: Minimum length of 1. Maximum length of 1024.

Key Pattern: ^\S+$

Value Length Constraints: Minimum length of 1. Maximum length of 1024.

Value Pattern: ^(?!\s*$).+

Required: No

**ConvertDotsInJsonKeysToUnderscores**

When set to `true`, specifies that the names of the keys include dots and that you want Firehose to replace them with underscores. This is useful because Apache Hive does not allow dots in column names. For example, if the JSON contains a key whose name is "a.b", you can define the column name to be "a_b" when using this option.

The default is `false`.

Type: Boolean

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# OrcSerDe

A serializer to use for converting data to the ORC format before storing it in Amazon S3. For more information, see [Apache ORC](#).

## Contents

**BlockSizeBytes**

The Hadoop Distributed File System (HDFS) block size. This is useful if you intend to copy the data from Amazon S3 to HDFS before querying. The default is 256 MiB and the minimum is 64 MiB. Firehose uses this value for padding calculations.

Type: Integer

Valid Range: Minimum value of 67108864.

Required: No

**BloomFilterColumns**

The column names for which you want Firehose to create bloom filters. The default is `null`.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: No

**BloomFilterFalsePositiveProbability**

The Bloom filter false positive probability (FPP). The lower the FPP, the bigger the Bloom filter. The default value is 0.05, the minimum is 0, and the maximum is 1.

Type: Double

Valid Range: Minimum value of 0. Maximum value of 1.

Required: No

**Compression**

The compression code to use over data blocks. The default is SNAPPY.

Type: String

Valid Values: NONE | ZLIB | SNAPPY

Required: No

**DictionaryKeyThreshold**

Represents the fraction of the total number of non-null rows. To turn off dictionary encoding, set this fraction to a number that is less than the number of distinct keys in a dictionary. To always use dictionary encoding, set this threshold to 1.

Type: Double

Valid Range: Minimum value of 0. Maximum value of 1.

Required: No

**EnablePadding**

Set this to `true` to indicate that you want stripes to be padded to the HDFS block boundaries. This is useful if you intend to copy the data from Amazon S3 to HDFS before querying. The default is `false`.

Type: Boolean

Required: No

**FormatVersion**

The version of the file to write. The possible values are `V0_11` and `V0_12`. The default is `V0_12`.

Type: String

Valid Values: V0_11 | V0_12

Required: No

**PaddingTolerance**

A number between 0 and 1 that defines the tolerance for block padding as a decimal fraction of stripe size. The default value is 0.05, which means 5 percent of stripe size.

For the default values of 64 MiB ORC stripes and 256 MiB HDFS blocks, the default block padding tolerance of 5 percent reserves a maximum of 3.2 MiB for padding within the 256 MiB

block. In such a case, if the available size within the block is more than 3.2 MiB, a new, smaller stripe is inserted to fit within that space. This ensures that no stripe crosses block boundaries and causes remote reads within a node-local task.

Firehose ignores this parameter when [OrcSerDe:EnablePadding](#) is `false`.

Type: Double

Valid Range: Minimum value of 0. Maximum value of 1.

Required: No

**RowIndexStride**

The number of rows between index entries. The default is 10,000 and the minimum is 1,000.

Type: Integer

Valid Range: Minimum value of 1000.

Required: No

**StripeSizeBytes**

The number of bytes in each stripe. The default is 64 MiB and the minimum is 8 MiB.

Type: Integer

Valid Range: Minimum value of 8388608.

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# OutputFormatConfiguration

Specifies the serializer that you want Firehose to use to convert the format of your data before it writes it to Amazon S3. This parameter is required if `Enabled` is set to true.

## Contents

**Serializer**

Specifies which serializer to use. You can choose either the ORC SerDe or the Parquet SerDe. If both are non-null, the server rejects the request.

Type: [Serializer](#) object

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# ParquetSerDe

A serializer to use for converting data to the Parquet format before storing it in Amazon S3. For more information, see [Apache Parquet](#).

## Contents

**BlockSizeBytes**

The Hadoop Distributed File System (HDFS) block size. This is useful if you intend to copy the data from Amazon S3 to HDFS before querying. The default is 256 MiB and the minimum is 64 MiB. Firehose uses this value for padding calculations.

Type: Integer

Valid Range: Minimum value of 67108864.

Required: No

**Compression**

The compression code to use over data blocks. The possible values are UNCOMPRESSED, SNAPPY, and GZIP, with the default being SNAPPY. Use SNAPPY for higher decompression speed. Use GZIP if the compression ratio is more important than speed.

Type: String

Valid Values: UNCOMPRESSED | GZIP | SNAPPY

Required: No

**EnableDictionaryCompression**

Indicates whether to enable dictionary compression.

Type: Boolean

Required: No

**MaxPaddingBytes**

The maximum amount of padding to apply. This is useful if you intend to copy the data from Amazon S3 to HDFS before querying. The default is 0.

Type: Integer

Valid Range: Minimum value of 0.

Required: No

**PageSizeBytes**

The Parquet page size. Column chunks are divided into pages. A page is conceptually an indivisible unit (in terms of compression and encoding). The minimum value is 64 KiB and the default is 1 MiB.

Type: Integer

Valid Range: Minimum value of 65536.

Required: No

**WriterVersion**

Indicates the version of row format to output. The possible values are V1 and V2. The default is V1.

Type: String

Valid Values: V1 | V2

Required: No

# See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# 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](#) objects

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# Processor

Describes a data processor.

> **ⓘ Note**
>
> If you want to add a new line delimiter between records in objects that are delivered to Amazon S3, choose `AppendDelimiterToRecord` as a processor type. You don't have to put a processor parameter when you select `AppendDelimiterToRecord`.

## Contents

**Type**

The type of processor.

Type: String

Valid Values: `RecordDeAggregation` | `Decompression` | `CloudWatchLogProcessing` | `Lambda` | `MetadataExtraction` | `AppendDelimiterToRecord`

Required: Yes

**Parameters**

The processor parameters.

Type: Array of [ProcessorParameter](#) objects

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# ProcessorParameter

Describes the processor parameter.

## Contents

**ParameterName**

The name of the parameter. Currently the following default values are supported: 3 for `NumberOfRetries` and 60 for the `BufferIntervalInSeconds`. The `BufferSizeInMBs` ranges between 0.2 MB and up to 3MB. The default buffering hint is 1MB for all destinations, except Splunk. For Splunk, the default buffering hint is 256 KB.

Type: String

Valid Values: `LambdaArn | NumberOfRetries | MetadataExtractionQuery | JsonParsingEngine | RoleArn | BufferSizeInMBs | BufferIntervalInSeconds | SubRecordType | Delimiter | CompressionFormat | DataMessageExtraction`

Required: Yes

**ParameterValue**

The parameter value.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 5120.

Pattern: `^(?!\s*$).+`

Required: Yes

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# PutRecordBatchResponseEntry

Contains the result for an individual record from a [PutRecordBatch](#) 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# 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 KiB.

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

# RedshiftDestinationConfiguration

Describes the configuration of a destination in Amazon Redshift.

## Contents

**ClusterJDBCURL**

The database connection string.

Type: String

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

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

Required: Yes

**CopyCommand**

The COPY command.

Type: [CopyCommand](#) object

Required: Yes

**RoleARN**

The Amazon Resource Name (ARN) of the AWS credentials. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: Yes

**S3Configuration**

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

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](#) object

Required: Yes

## CloudWatchLoggingOptions

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

## Password

The user password.

Type: String

Length Constraints: Minimum length of 6. Maximum length of 512.

Pattern: `.*`

Required: No

## ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

## RetryOptions

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

Type: [RedshiftRetryOptions](#) object

Required: No

**S3BackupConfiguration**

The configuration for backup in Amazon S3.

Type: S3DestinationConfiguration object

Required: No

**S3BackupMode**

The Amazon S3 backup mode. After you create a delivery stream, you can update it to enable Amazon S3 backup if it is disabled. If backup is enabled, you can't update the delivery stream to disable it.

Type: String

Valid Values: `Disabled | Enabled`

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for Amazon Redshift.

Type: SecretsManagerConfiguration object

Required: No

**Username**

The name of the user.

Type: String

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

Pattern: `.*`

Required: No

# See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# RedshiftDestinationDescription

Describes a destination in Amazon Redshift.

## Contents

**ClusterJDBCURL**

The database connection string.

Type: String

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

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

Required: Yes

**CopyCommand**

The COPY command.

Type: CopyCommand object

Required: Yes

**RoleARN**

The Amazon Resource Name (ARN) of the AWS credentials. For more information, see Amazon Resource Names (ARNs) and AWS Service Namespaces.

Type: String

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

Pattern: `arn:.*`

Required: Yes

**S3DestinationDescription**

The Amazon S3 destination.

Type: [S3DestinationDescription](#) object

Required: Yes

## CloudWatchLoggingOptions

The Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

## ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

## RetryOptions

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

Type: [RedshiftRetryOptions](#) object

Required: No

## S3BackupDescription

The configuration for backup in Amazon S3.

Type: [S3DestinationDescription](#) object

Required: No

## S3BackupMode

The Amazon S3 backup mode.

Type: String

Valid Values: `Disabled | Enabled`

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for Amazon Redshift.

Type: SecretsManagerConfiguration object

Required: No

**Username**

The name of the user.

Type: String

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

Pattern: . *

Required: No

# See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# RedshiftDestinationUpdate

Describes an update for a destination in Amazon Redshift.

## Contents

**CloudWatchLoggingOptions**

The Amazon CloudWatch logging options for your delivery stream.

Type: CloudWatchLoggingOptions object

Required: No

**ClusterJDBCURL**

The database connection string.

Type: String

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

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

Required: No

**CopyCommand**

The COPY command.

Type: CopyCommand object

Required: No

**Password**

The user password.

Type: String

Length Constraints: Minimum length of 6. Maximum length of 512.

Pattern: `.*`

Required: No

**ProcessingConfiguration**

The data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

**RetryOptions**

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

Type: [RedshiftRetryOptions](#) object

Required: No

**RoleARN**

The Amazon Resource Name (ARN) of the AWS credentials. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: No

**S3BackupMode**

You can update a delivery stream to enable Amazon S3 backup if it is disabled. If backup is enabled, you can't update the delivery stream to disable it.

Type: String

Valid Values: `Disabled | Enabled`

Required: No

**S3BackupUpdate**

The Amazon S3 destination for backup.

Type: [S3DestinationUpdate](#) 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](#) object

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for Amazon Redshift.

Type: [SecretsManagerConfiguration](#) object

Required: No

**Username**

The name of the user.

Type: String

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

Pattern: `.*`

Required: No

# See Also

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

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

- [AWS SDK for Ruby V3](#)

# RedshiftRetryOptions

Configures retry behavior in case 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

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# RetryOptions

The retry behavior in case Firehose is unable to deliver data to an Amazon S3 prefix.

## Contents

**DurationInSeconds**

The period of time during which Firehose retries to deliver data to the specified Amazon S3 prefix.

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

# S3DestinationConfiguration

Describes the configuration of a destination in Amazon S3.

## Contents

**BucketARN**

The ARN of the S3 bucket. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: Yes

**RoleARN**

The Amazon Resource Name (ARN) of the AWS credentials. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: Yes

**BufferingHints**

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

Type: [BufferingHints](#) object

Required: No

**CloudWatchLoggingOptions**

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) 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 | HADOOP_SNAPPY

Required: No

**EncryptionConfiguration**

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

Type: EncryptionConfiguration object

Required: No

**ErrorOutputPrefix**

A prefix that Firehose evaluates and adds to failed records before writing them to S3. This prefix appears immediately following the bucket name. For information about how to specify this prefix, see Custom Prefixes for Amazon S3 Objects.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: .*

Required: No

**Prefix**

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered Amazon S3 files. You can also specify a custom prefix, as described in Custom Prefixes for Amazon S3 Objects.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: . *

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# S3DestinationDescription

Describes a destination in Amazon S3.

## Contents

**BucketARN**

The ARN of the S3 bucket. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: Yes

**BufferingHints**

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

Type: [BufferingHints](#) object

Required: Yes

**CompressionFormat**

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

Type: String

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

Required: Yes

**EncryptionConfiguration**

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

Type: [EncryptionConfiguration](#) object

Required: Yes

**RoleARN**

The Amazon Resource Name (ARN) of the AWS credentials. For more information, see Amazon Resource Names (ARNs) and AWS Service Namespaces.

Type: String

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

Pattern: `arn:.*`

Required: Yes

**CloudWatchLoggingOptions**

The Amazon CloudWatch logging options for your delivery stream.

Type: CloudWatchLoggingOptions object

Required: No

**ErrorOutputPrefix**

A prefix that Firehose evaluates and adds to failed records before writing them to S3. This prefix appears immediately following the bucket name. For information about how to specify this prefix, see Custom Prefixes for Amazon S3 Objects.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: `.*`

Required: No

**Prefix**

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered Amazon S3 files. You can also specify a custom prefix, as described in Custom Prefixes for Amazon S3 Objects.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: `.*`

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# S3DestinationUpdate

Describes an update for a destination in Amazon S3.

## Contents

**BucketARN**

The ARN of the S3 bucket. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

Type: String

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

Pattern: `arn:.*`

Required: No

**BufferingHints**

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

Type: [BufferingHints](#) object

Required: No

**CloudWatchLoggingOptions**

The CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) 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 | HADOOP_SNAPPY

Required: No

**EncryptionConfiguration**

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

Type: [EncryptionConfiguration](#) object

Required: No

**ErrorOutputPrefix**

A prefix that Firehose evaluates and adds to failed records before writing them to S3. This prefix appears immediately following the bucket name. For information about how to specify this prefix, see [Custom Prefixes for Amazon S3 Objects](#).

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: .*

Required: No

**Prefix**

The "YYYY/MM/DD/HH" time format prefix is automatically used for delivered Amazon S3 files. You can also specify a custom prefix, as described in [Custom Prefixes for Amazon S3 Objects](#).

Type: String

Length Constraints: Minimum length of 0. Maximum length of 1024.

Pattern: .*

Required: No

**RoleARN**

The Amazon Resource Name (ARN) of the AWS credentials. For more information, see [Amazon Resource Names (ARNs) and AWS Service Namespaces](#).

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

# SchemaConfiguration

Specifies the schema to which you want Firehose to configure your data before it writes it to Amazon S3. This parameter is required if `Enabled` is set to true.

## Contents

**CatalogId**

The ID of the AWS Glue Data Catalog. If you don't supply this, the AWS account ID is used by default.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: No

**DatabaseName**

Specifies the name of the AWS Glue database that contains the schema for the output data.

> ⚠️ **Important**
>
> If the `SchemaConfiguration` request parameter is used as part of invoking the `CreateDeliveryStream` API, then the `DatabaseName` property is required and its value must be specified.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: No

**Region**

If you don't specify an AWS Region, the default is the current Region.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: No

**RoleARN**

The role that Firehose can use to access AWS Glue. This role must be in the same account you use for Firehose. Cross-account roles aren't allowed.

> ⚠️ **Important**
>
> If the `SchemaConfiguration` request parameter is used as part of invoking the `CreateDeliveryStream` API, then the `RoleARN` property is required and its value must be specified.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: No

**TableName**

Specifies the AWS Glue table that contains the column information that constitutes your data schema.

> ⚠️ **Important**
>
> If the `SchemaConfiguration` request parameter is used as part of invoking the `CreateDeliveryStream` API, then the `TableName` property is required and its value must be specified.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: No

**VersionId**

Specifies the table version for the output data schema. If you don't specify this version ID, or if you set it to LATEST, Firehose uses the most recent version. This means that any updates to the table are automatically picked up.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# SecretsManagerConfiguration

The structure that defines how Firehose accesses the secret.

## Contents

**Enabled**

Specifies whether you want to use the the secrets manager feature. When set as `True` the secrets manager configuration overwrites the existing secrets in the destination configuration. When it's set to `False` Firehose falls back to the credentials in the destination configuration.

Type: Boolean

Required: Yes

**RoleARN**

Specifies the role that Firehose assumes when calling the Secrets Manager API operation. When you provide the role, it overrides any destination specific role defined in the destination configuration. If you do not provide the then we use the destination specific role. This parameter is required for Splunk.

Type: String

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

Pattern: `arn:.*`

Required: No

**SecretARN**

The ARN of the secret that stores your credentials. It must be in the same region as the Firehose stream and the role. The secret ARN can reside in a different account than the delivery stream and role as Firehose supports cross-account secret access. This parameter is required when **Enabled** is set to `True`.

Type: String

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

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

# Serializer

The serializer that you want Firehose to use to convert data to the target format before writing it to Amazon S3. Firehose supports two types of serializers: the ORC SerDe and the Parquet SerDe.

## Contents

**OrcSerDe**

A serializer to use for converting data to the ORC format before storing it in Amazon S3. For more information, see Apache ORC.

Type: OrcSerDe object

Required: No

**ParquetSerDe**

A serializer to use for converting data to the Parquet format before storing it in Amazon S3. For more information, see Apache Parquet.

Type: ParquetSerDe object

Required: No

## See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# SnowflakeDestinationConfiguration

Configure Snowflake destination

## Contents

**AccountUrl**

URL for accessing your Snowflake account. This URL must include your [account identifier](). Note that the protocol (https://) and port number are optional.

Type: String

Length Constraints: Minimum length of 24. Maximum length of 2048.

Pattern: `.+?\.snowflakecomputing\.com`

Required: Yes

**Database**

All data in Snowflake is maintained in databases.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

**RoleARN**

The Amazon Resource Name (ARN) of the Snowflake role

Type: String

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

Pattern: `arn:.*`

Required: Yes

**S3Configuration**

Describes the configuration of a destination in Amazon S3.

Type: [S3DestinationConfiguration]() object

Required: Yes

**Schema**

Each database consists of one or more schemas, which are logical groupings of database objects, such as tables and views

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

**Table**

All data in Snowflake is stored in database tables, logically structured as collections of columns and rows.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: Yes

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](CloudWatchLoggingOptions) object

Required: No

**ContentColumnName**

The name of the record content column

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

**DataLoadingOption**

Choose to load JSON keys mapped to table column names or choose to split the JSON payload where content is mapped to a record content column and source metadata is mapped to a record metadata column.

Type: String

Valid Values: `JSON_MAPPING` | `VARIANT_CONTENT_MAPPING` |
`VARIANT_CONTENT_AND_METADATA_MAPPING`

Required: No

**KeyPassphrase**

Passphrase to decrypt the private key when the key is encrypted. For information, see [Using Key
Pair Authentication & Key Rotation](#).

Type: String

Length Constraints: Minimum length of 7. Maximum length of 255.

Required: No

**MetaDataColumnName**

The name of the record metadata column

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

**PrivateKey**

The private key used to encrypt your Snowflake client. For information, see [Using Key Pair
Authentication & Key Rotation](#).

Type: String

Length Constraints: Minimum length of 256. Maximum length of 4096.

Pattern: `^(?:[A-Za-z0-9+\/]{4})*(?:[A-Za-z0-9+\/]{2}==|[A-Za-z0-9+\/]{3}=)?`
`$`

Required: No

**ProcessingConfiguration**

Describes a data processing configuration.

Type: [ProcessingConfiguration](ProcessingConfiguration) object

Required: No

**RetryOptions**

The time period where Firehose will retry sending data to the chosen HTTP endpoint.

Type: [SnowflakeRetryOptions](SnowflakeRetryOptions) object

Required: No

**S3BackupMode**

Choose an S3 backup mode

Type: String

Valid Values: `FailedDataOnly | AllData`

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for Snowflake.

Type: [SecretsManagerConfiguration](SecretsManagerConfiguration) object

Required: No

**SnowflakeRoleConfiguration**

Optionally configure a Snowflake role. Otherwise the default user role will be used.

Type: [SnowflakeRoleConfiguration](SnowflakeRoleConfiguration) object

Required: No

**SnowflakeVpcConfiguration**

The VPCE ID for Firehose to privately connect with Snowflake. The ID format is com.amazonaws.vpce.[region].vpce-svc-<[id]>. For more information, see [Amazon PrivateLink & Snowflake](Amazon PrivateLink & Snowflake)

Type: [SnowflakeVpcConfiguration](SnowflakeVpcConfiguration) object

Required: No

**User**

User login name for the Snowflake account.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# SnowflakeDestinationDescription

Optional Snowflake destination description

## Contents

**AccountUrl**

URL for accessing your Snowflake account. This URL must include your [account identifier](). Note that the protocol (https://) and port number are optional.

Type: String

Length Constraints: Minimum length of 24. Maximum length of 2048.

Pattern: `.+?\.snowflakecomputing\.com`

Required: No

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions]() object

Required: No

**ContentColumnName**

The name of the record content column

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

**Database**

All data in Snowflake is maintained in databases.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

## DataLoadingOption

Choose to load JSON keys mapped to table column names or choose to split the JSON payload where content is mapped to a record content column and source metadata is mapped to a record metadata column.

Type: String

Valid Values: `JSON_MAPPING` | `VARIANT_CONTENT_MAPPING` | `VARIANT_CONTENT_AND_METADATA_MAPPING`

Required: No

## MetaDataColumnName

The name of the record metadata column

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

## ProcessingConfiguration

Describes a data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

## RetryOptions

The time period where Firehose will retry sending data to the chosen HTTP endpoint.

Type: [SnowflakeRetryOptions](#) object

Required: No

## RoleARN

The Amazon Resource Name (ARN) of the Snowflake role

Type: String

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

Pattern: `arn:.*`

Required: No

**S3BackupMode**

Choose an S3 backup mode

Type: String

Valid Values: `FailedDataOnly | AllData`

Required: No

**S3DestinationDescription**

Describes a destination in Amazon S3.

Type: [S3DestinationDescription](#) object

Required: No

**Schema**

Each database consists of one or more schemas, which are logical groupings of database objects, such as tables and views

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for Snowflake.

Type: [SecretsManagerConfiguration](#) object

Required: No

**SnowflakeRoleConfiguration**

Optionally configure a Snowflake role. Otherwise the default user role will be used.

Type: [SnowflakeRoleConfiguration](#) object

Required: No

**SnowflakeVpcConfiguration**

The VPCE ID for Firehose to privately connect with Snowflake. The ID format is com.amazonaws.vpce.[region].vpce-svc-<[id]>. For more information, see [Amazon PrivateLink & Snowflake](#)

Type: [SnowflakeVpcConfiguration](#) object

Required: No

**Table**

All data in Snowflake is stored in database tables, logically structured as collections of columns and rows.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

**User**

User login name for the Snowflake account.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

# See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# SnowflakeDestinationUpdate

Update to configuration settings

## Contents

**AccountUrl**

URL for accessing your Snowflake account. This URL must include your [account identifier](#). Note that the protocol (https://) and port number are optional.

Type: String

Length Constraints: Minimum length of 24. Maximum length of 2048.

Pattern: `.+?\.snowflakecomputing\.com`

Required: No

**CloudWatchLoggingOptions**

Describes the Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

**ContentColumnName**

The name of the content metadata column

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

**Database**

All data in Snowflake is maintained in databases.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

## DataLoadingOption

JSON keys mapped to table column names or choose to split the JSON payload where content is mapped to a record content column and source metadata is mapped to a record metadata column.

Type: String

Valid Values: `JSON_MAPPING` | `VARIANT_CONTENT_MAPPING` | `VARIANT_CONTENT_AND_METADATA_MAPPING`

Required: No

## KeyPassphrase

Passphrase to decrypt the private key when the key is encrypted. For information, see Using Key Pair Authentication & Key Rotation.

Type: String

Length Constraints: Minimum length of 7. Maximum length of 255.

Required: No

## MetaDataColumnName

The name of the record metadata column

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

## PrivateKey

The private key used to encrypt your Snowflake client. For information, see Using Key Pair Authentication & Key Rotation.

Type: String

Length Constraints: Minimum length of 256. Maximum length of 4096.

Pattern: `^(?:[A-Za-z0-9+\/]{4})*(?:[A-Za-z0-9+\/]{2}==|[A-Za-z0-9+\/]{3}=)?$`

Required: No

## ProcessingConfiguration

Describes a data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

## RetryOptions

Specify how long Firehose retries sending data to the New Relic HTTP endpoint. After sending data, Firehose first waits for an acknowledgment from the HTTP endpoint. If an error occurs or the acknowledgment doesn't arrive within the acknowledgment timeout period, Firehose starts the retry duration counter. It keeps retrying until the retry duration expires. After that, Firehose considers it a data delivery failure and backs up the data to your Amazon S3 bucket. Every time that Firehose sends data to the HTTP endpoint (either the initial attempt or a retry), it restarts the acknowledgement timeout counter and waits for an acknowledgement from the HTTP endpoint. Even if the retry duration expires, Firehose still waits for the acknowledgment until it receives it or the acknowledgement timeout period is reached. If the acknowledgment times out, Firehose determines whether there's time left in the retry counter. If there is time left, it retries again and repeats the logic until it receives an acknowledgment or determines that the retry time has expired. If you don't want Firehose to retry sending data, set this value to 0.

Type: [SnowflakeRetryOptions](#) object

Required: No

## RoleARN

The Amazon Resource Name (ARN) of the Snowflake role

Type: String

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

Pattern: `arn:.*`

Required: No

**S3BackupMode**

Choose an S3 backup mode

Type: String

Valid Values: `FailedDataOnly | AllData`

Required: No

**S3Update**

Describes an update for a destination in Amazon S3.

Type: [S3DestinationUpdate](#) object

Required: No

**Schema**

Each database consists of one or more schemas, which are logical groupings of database objects, such as tables and views

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

**SecretsManagerConfiguration**

Describes the Secrets Manager configuration in Snowflake.

Type: [SecretsManagerConfiguration](#) object

Required: No

**SnowflakeRoleConfiguration**

Optionally configure a Snowflake role. Otherwise the default user role will be used.

Type: [SnowflakeRoleConfiguration](#) object

Required: No

**Table**

All data in Snowflake is stored in database tables, logically structured as collections of columns and rows.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

**User**

User login name for the Snowflake account.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

# See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# SnowflakeRetryOptions

Specify how long Firehose retries sending data to the New Relic HTTP endpoint. After sending data, Firehose first waits for an acknowledgment from the HTTP endpoint. If an error occurs or the acknowledgment doesn't arrive within the acknowledgment timeout period, Firehose starts the retry duration counter. It keeps retrying until the retry duration expires. After that, Firehose considers it a data delivery failure and backs up the data to your Amazon S3 bucket. Every time that Firehose sends data to the HTTP endpoint (either the initial attempt or a retry), it restarts the acknowledgement timeout counter and waits for an acknowledgement from the HTTP endpoint. Even if the retry duration expires, Firehose still waits for the acknowledgment until it receives it or the acknowledgement timeout period is reached. If the acknowledgment times out, Firehose determines whether there's time left in the retry counter. If there is time left, it retries again and repeats the logic until it receives an acknowledgment or determines that the retry time has expired. If you don't want Firehose to retry sending data, set this value to 0.

## Contents

**DurationInSeconds**

the time period where Firehose will retry sending data to the chosen HTTP endpoint.

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 C++)
- [AWS SDK for Java V2](AWS SDK for Java V2)
- [AWS SDK for Ruby V3](AWS SDK for Ruby V3)

# SnowflakeRoleConfiguration

Optionally configure a Snowflake role. Otherwise the default user role will be used.

## Contents

**Enabled**

Enable Snowflake role

Type: Boolean

Required: No

**SnowflakeRole**

The Snowflake role you wish to configure

Type: String

Length Constraints: Minimum length of 1. Maximum length of 255.

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# SnowflakeVpcConfiguration

Configure a Snowflake VPC

## Contents

**PrivateLinkVpceId**

The VPCE ID for Firehose to privately connect with Snowflake. The ID format is com.amazonaws.vpce.[region].vpce-svc-<[id]>. For more information, see Amazon PrivateLink & Snowflake

Type: String

Length Constraints: Minimum length of 47. Maximum length of 255.

Pattern: `([a-zA-Z0-9\-\_]+\.){2,3}vpce\.[a-zA-Z0-9\-]*\.vpce-svc\-[a-zA-Z0-9\-]{17}$`

Required: Yes

## See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# SourceDescription

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

## Contents

**KinesisStreamSourceDescription**

The [KinesisStreamSourceDescription](#) value for the source Kinesis data stream.

Type: [KinesisStreamSourceDescription](#) object

Required: No

**MSKSourceDescription**

The configuration description for the Amazon MSK cluster to be used as the source for a delivery stream.

Type: [MSKSourceDescription](#) object

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# SplunkBufferingHints

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

## Contents

**IntervalInSeconds**

Buffer incoming data for the specified period of time, in seconds, before delivering it to the destination. The default value is 60 (1 minute).

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 60.

Required: No

**SizeInMBs**

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

Type: Integer

Valid Range: Minimum value of 1. Maximum value of 5.

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# SplunkDestinationConfiguration

Describes the configuration of a destination in Splunk.

## Contents

**HECEndpoint**

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

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: . *

Required: Yes

**HECEndpointType**

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

Type: String

Valid Values: Raw | Event

Required: Yes

**S3Configuration**

The configuration for the backup Amazon S3 location.

Type: S3DestinationConfiguration object

Required: Yes

**BufferingHints**

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

Type: SplunkBufferingHints object

Required: No

## CloudWatchLoggingOptions

The Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](CloudWatchLoggingOptions) object

Required: No

## HECAcknowledgmentTimeoutInSeconds

The amount of time that Firehose waits to receive an acknowledgment from Splunk after it sends it data. At the end of the timeout period, 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

## HECToken

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

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: . *

Required: No

## ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration](ProcessingConfiguration) object

Required: No

## RetryOptions

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

Type: [SplunkRetryOptions](SplunkRetryOptions) object

Required: No

**S3BackupMode**

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

You can update this backup mode from `FailedEventsOnly` to `AllEvents`. You can't update it from `AllEvents` to `FailedEventsOnly`.

Type: String

Valid Values: `FailedEventsOnly | AllEvents`

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for Splunk.

Type: SecretsManagerConfiguration object

Required: No

# See Also

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

- AWS SDK for C++
- AWS SDK for Java V2
- AWS SDK for Ruby V3

# SplunkDestinationDescription

Describes a destination in Splunk.

## Contents

**BufferingHints**

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

Type: [SplunkBufferingHints](#) object

Required: No

**CloudWatchLoggingOptions**

The Amazon CloudWatch logging options for your delivery stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

**HECAcknowledgmentTimeoutInSeconds**

The amount of time that Firehose waits to receive an acknowledgment from Splunk after it sends it data. At the end of the timeout period, 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 Firehose sends your data.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: .*

Required: No

**HECEndpointType**

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

Type: String

Valid Values: `Raw` | `Event`

Required: No

**HECToken**

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

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: `.*`

Required: No

**ProcessingConfiguration**

The data processing configuration.

Type: ProcessingConfiguration object

Required: No

**RetryOptions**

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

Type: SplunkRetryOptions object

Required: No

**S3BackupMode**

Defines how documents should be delivered to Amazon S3. When set to `FailedDocumentsOnly`, Firehose writes any data that could not be indexed to the configured Amazon S3 destination. When set to `AllDocuments`, 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](#) object

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for Splunk.

Type: [SecretsManagerConfiguration](#) object

Required: No

# See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# SplunkDestinationUpdate

Describes an update for a destination in Splunk.

## Contents

**BufferingHints**

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

Type: SplunkBufferingHints object

Required: No

**CloudWatchLoggingOptions**

The Amazon CloudWatch logging options for your delivery stream.

Type: CloudWatchLoggingOptions object

Required: No

**HECAcknowledgmentTimeoutInSeconds**

The amount of time that Firehose waits to receive an acknowledgment from Splunk after it sends data. At the end of the timeout period, 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 Firehose sends your data.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: .*

Required: No

**HECEndpointType**

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

Type: String

Valid Values: `Raw` | `Event`

Required: No

**HECToken**

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

Type: String

Length Constraints: Minimum length of 0. Maximum length of 2048.

Pattern: `.*`

Required: No

**ProcessingConfiguration**

The data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

**RetryOptions**

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

Type: [SplunkRetryOptions](#) object

Required: No

**S3BackupMode**

Specifies how you want Firehose to back up documents to Amazon S3. When set to `FailedDocumentsOnly`, Firehose writes any data that could not be indexed to the configured Amazon S3 destination. When set to `AllEvents`, Firehose delivers all incoming records to Amazon S3, and also writes failed documents to Amazon S3. The default value is `FailedEventsOnly`.

You can update this backup mode from `FailedEventsOnly` to `AllEvents`. You can't update it from `AllEvents` to `FailedEventsOnly`.

Type: String

Valid Values: `FailedEventsOnly | AllEvents`

Required: No

**S3Update**

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

Type: [S3DestinationUpdate](#) object

Required: No

**SecretsManagerConfiguration**

The configuration that defines how you access secrets for Splunk.

Type: [SecretsManagerConfiguration](#) object

Required: No

# See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# SplunkRetryOptions

Configures retry behavior in case 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 Firehose spends on retries. This duration starts after the initial attempt to send data to Splunk fails. It doesn't include the periods during which 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 Java V2](#)
- [AWS SDK for Ruby V3](#)

# Tag

Metadata that you can assign to a delivery stream, consisting of a key-value pair.

## Contents

**Key**

A unique identifier for the tag. Maximum length: 128 characters. Valid characters: Unicode letters, digits, white space, _ . / = + - % @

Type: String

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

Pattern: `^(?!aws:)[\p{L}\p{Z}\p{N}_.:\/=+\-@%]*$`

Required: Yes

**Value**

An optional string, which you can use to describe or define the tag. Maximum length: 256 characters. Valid characters: Unicode letters, digits, white space, _ . / = + - % @

Type: String

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

Pattern: `^[\p{L}\p{Z}\p{N}_.:\/=+\-@%]*$`

Required: No

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# VpcConfiguration

The details of the VPC of the Amazon OpenSearch or Amazon OpenSearch Serverless destination.

## Contents

**RoleARN**

The ARN of the IAM role that you want the delivery stream to use to create endpoints in the destination VPC. You can use your existing Firehose delivery role or you can specify a new role. In either case, make sure that the role trusts the Firehose service principal and that it grants the following permissions:

- `ec2:DescribeVpcs`
- `ec2:DescribeVpcAttribute`
- `ec2:DescribeSubnets`
- `ec2:DescribeSecurityGroups`
- `ec2:DescribeNetworkInterfaces`
- `ec2:CreateNetworkInterface`
- `ec2:CreateNetworkInterfacePermission`
- `ec2:DeleteNetworkInterface`

> ⚠️ **Important**
>
> When you specify subnets for delivering data to the destination in a private VPC, make sure you have enough number of free IP addresses in chosen subnets. If there is no available free IP address in a specified subnet, Firehose cannot create or add ENIs for the data delivery in the private VPC, and the delivery will be degraded or fail.

Type: String

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

Pattern: `arn:.*`

Required: Yes

**SecurityGroupIds**

The IDs of the security groups that you want Firehose to use when it creates ENIs in the VPC of the Amazon ES destination. You can use the same security group that the Amazon ES domain uses or different ones. If you specify different security groups here, ensure that they allow outbound HTTPS traffic to the Amazon ES domain's security group. Also ensure that the Amazon ES domain's security group allows HTTPS traffic from the security groups specified here. If you use the same security group for both your delivery stream and the Amazon ES domain, make sure the security group inbound rule allows HTTPS traffic. For more information about security group rules, see [Security group rules](#) in the Amazon VPC documentation.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 5 items.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: Yes

**SubnetIds**

The IDs of the subnets that you want Firehose to use to create ENIs in the VPC of the Amazon ES destination. Make sure that the routing tables and inbound and outbound rules allow traffic to flow from the subnets whose IDs are specified here to the subnets that have the destination Amazon ES endpoints. Firehose creates at least one ENI in each of the subnets that are specified here. Do not delete or modify these ENIs.

The number of ENIs that Firehose creates in the subnets specified here scales up and down automatically based on throughput. To enable Firehose to scale up the number of ENIs to match throughput, ensure that you have sufficient quota. To help you calculate the quota you need, assume that Firehose can create up to three ENIs for this delivery stream for each of the subnets specified here. For more information about ENI quota, see [Network Interfaces](#) in the Amazon VPC Quotas topic.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 16 items.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: Yes

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# VpcConfigurationDescription

The details of the VPC of the Amazon ES destination.

## Contents

**RoleARN**

The ARN of the IAM role that the delivery stream uses to create endpoints in the destination VPC. You can use your existing Firehose delivery role or you can specify a new role. In either case, make sure that the role trusts the Firehose service principal and that it grants the following permissions:

- `ec2:DescribeVpcs`
- `ec2:DescribeVpcAttribute`
- `ec2:DescribeSubnets`
- `ec2:DescribeSecurityGroups`
- `ec2:DescribeNetworkInterfaces`
- `ec2:CreateNetworkInterface`
- `ec2:CreateNetworkInterfacePermission`
- `ec2:DeleteNetworkInterface`

If you revoke these permissions after you create the delivery stream, Firehose can't scale out by creating more ENIs when necessary. You might therefore see a degradation in performance.

Type: String

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

Pattern: `arn:.*`

Required: Yes

**SecurityGroupIds**

The IDs of the security groups that Firehose uses when it creates ENIs in the VPC of the Amazon ES destination. You can use the same security group that the Amazon ES domain uses or different ones. If you specify different security groups, ensure that they allow outbound HTTPS traffic to the Amazon ES domain's security group. Also ensure that the Amazon ES domain's

security group allows HTTPS traffic from the security groups specified here. If you use the same security group for both your delivery stream and the Amazon ES domain, make sure the security group inbound rule allows HTTPS traffic. For more information about security group rules, see Security group rules in the Amazon VPC documentation.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 5 items.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: Yes

**SubnetIds**

The IDs of the subnets that Firehose uses to create ENIs in the VPC of the Amazon ES destination. Make sure that the routing tables and inbound and outbound rules allow traffic to flow from the subnets whose IDs are specified here to the subnets that have the destination Amazon ES endpoints. Firehose creates at least one ENI in each of the subnets that are specified here. Do not delete or modify these ENIs.

The number of ENIs that Firehose creates in the subnets specified here scales up and down automatically based on throughput. To enable Firehose to scale up the number of ENIs to match throughput, ensure that you have sufficient quota. To help you calculate the quota you need, assume that Firehose can create up to three ENIs for this delivery stream for each of the subnets specified here. For more information about ENI quota, see Network Interfaces in the Amazon VPC Quotas topic.

Type: Array of strings

Array Members: Minimum number of 1 item. Maximum number of 16 items.

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: Yes

**VpcId**

The ID of the Amazon ES destination's VPC.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 1024.

Pattern: ^\S+$

Required: Yes

## See Also

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

- [AWS SDK for C++](#)
- [AWS SDK for Java V2](#)
- [AWS SDK for Ruby V3](#)

# 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

**NotAuthorized**

You do not have permission to perform this action.

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