



API Reference

Amazon Data Firehose



API Version 2015-08-04

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon Data Firehose: API Reference

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

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

Table of Contents

Welcome	1
Actions	2
CreateDeliveryStream	3
Request Syntax	4
Request Parameters	21
Response Syntax	25
Response Elements	26
Errors	26
Examples	27
See Also	28
DeleteDeliveryStream	30
Request Syntax	30
Request Parameters	30
Response Elements	31
Errors	31
Examples	31
See Also	32
DescribeDeliveryStream	33
Request Syntax	33
Request Parameters	33
Response Syntax	34
Response Elements	52
Errors	52
Examples	52
See Also	54
ListDeliveryStreams	56
Request Syntax	56
Request Parameters	56
Response Syntax	57
Response Elements	57
Errors	58
Examples	58
See Also	59
ListTagsForDeliveryStream	60

Request Syntax	60
Request Parameters	60
Response Syntax	61
Response Elements	61
Errors	62
Examples	62
See Also	63
PutRecord	65
Request Syntax	66
Request Parameters	66
Response Syntax	67
Response Elements	67
Errors	67
Examples	68
See Also	69
PutRecordBatch	71
Request Syntax	72
Request Parameters	73
Response Syntax	73
Response Elements	74
Errors	74
Examples	75
See Also	76
StartDeliveryStreamEncryption	78
Request Syntax	79
Request Parameters	79
Response Elements	79
Errors	80
Examples	80
See Also	81
StopDeliveryStreamEncryption	83
Request Syntax	83
Request Parameters	83
Response Elements	84
Errors	84
Examples	84

See Also	85
TagDeliveryStream	86
Request Syntax	86
Request Parameters	86
Response Elements	87
Errors	87
Examples	88
See Also	89
UntagDeliveryStream	90
Request Syntax	90
Request Parameters	90
Response Elements	91
Errors	91
Examples	91
See Also	92
UpdateDestination	93
Request Syntax	93
Request Parameters	109
Response Elements	111
Errors	112
Examples	112
See Also	113
Data Types	115
AmazonOpenSearchServerlessBufferingHints	119
Contents	119
See Also	119
AmazonOpenSearchServerlessDestinationConfiguration	121
Contents	121
See Also	123
AmazonOpenSearchServerlessDestinationDescription	124
Contents	124
See Also	126
AmazonOpenSearchServerlessDestinationUpdate	127
Contents	127
See Also	128
AmazonOpenSearchServerlessRetryOptions	130

Contents	130
See Also	130
AmazonopensearchserviceBufferingHints	131
Contents	131
See Also	131
AmazonopensearchserviceDestinationConfiguration	133
Contents	133
See Also	136
AmazonopensearchserviceDestinationDescription	137
Contents	137
See Also	140
AmazonopensearchserviceDestinationUpdate	141
Contents	141
See Also	144
AmazonopensearchserviceRetryOptions	145
Contents	145
See Also	145
AuthenticationConfiguration	146
Contents	146
See Also	146
BufferingHints	147
Contents	147
See Also	147
CatalogConfiguration	149
Contents	149
See Also	149
CloudWatchLoggingOptions	151
Contents	151
See Also	151
CopyCommand	153
Contents	153
See Also	154
DatabaseColumnList	155
Contents	155
See Also	155
DatabaseList	157

Contents	157
See Also	157
DatabaseSnapshotInfo	159
Contents	159
See Also	160
DatabaseSourceAuthenticationConfiguration	162
Contents	162
See Also	162
DatabaseSourceConfiguration	163
Contents	163
See Also	166
DatabaseSourceDescription	167
Contents	167
See Also	170
DatabaseSourceVPCCConfiguration	171
Contents	171
See Also	171
DatabaseTableList	172
Contents	172
See Also	172
DataFormatConversionConfiguration	174
Contents	174
See Also	175
DeliveryStreamDescription	176
Contents	176
See Also	178
DeliveryStreamEncryptionConfiguration	180
Contents	180
See Also	181
DeliveryStreamEncryptionConfigurationInput	182
Contents	182
See Also	183
Deserializer	184
Contents	184
See Also	184
DestinationDescription	185

Contents	185
See Also	187
DestinationTableConfiguration	188
Contents	188
See Also	189
DirectPutSourceConfiguration	190
Contents	190
See Also	190
DirectPutSourceDescription	191
Contents	191
See Also	191
DocumentIdOptions	192
Contents	192
See Also	192
DynamicPartitioningConfiguration	193
Contents	193
See Also	193
ElasticsearchBufferingHints	194
Contents	194
See Also	194
ElasticsearchDestinationConfiguration	196
Contents	196
See Also	199
ElasticsearchDestinationDescription	200
Contents	200
See Also	203
ElasticsearchDestinationUpdate	204
Contents	204
See Also	207
ElasticsearchRetryOptions	208
Contents	208
See Also	208
EncryptionConfiguration	209
Contents	209
See Also	209
ExtendedS3DestinationConfiguration	210

Contents	210
See Also	213
ExtendedS3DestinationDescription	214
Contents	214
See Also	217
ExtendedS3DestinationUpdate	218
Contents	218
See Also	221
FailureDescription	222
Contents	222
See Also	222
HiveJsonSerDe	224
Contents	224
See Also	224
HttpEndpointBufferingHints	225
Contents	225
See Also	225
HttpEndpointCommonAttribute	227
Contents	227
See Also	227
HttpEndpointConfiguration	228
Contents	228
See Also	229
HttpEndpointDescription	230
Contents	230
See Also	230
HttpEndpointDestinationConfiguration	231
Contents	231
See Also	233
HttpEndpointDestinationDescription	234
Contents	234
See Also	236
HttpEndpointDestinationUpdate	237
Contents	237
See Also	239
HttpEndpointRequestConfiguration	240

Contents	240
See Also	240
HttpEndpointRetryOptions	241
Contents	241
See Also	241
IcebergDestinationConfiguration	242
Contents	242
See Also	244
IcebergDestinationDescription	245
Contents	245
See Also	247
IcebergDestinationUpdate	248
Contents	248
See Also	250
InputFormatConfiguration	251
Contents	251
See Also	251
KinesisStreamSourceConfiguration	252
Contents	252
See Also	252
KinesisStreamSourceDescription	254
Contents	254
See Also	254
KMSEncryptionConfig	256
Contents	256
See Also	256
MSKSourceConfiguration	257
Contents	257
See Also	258
MSKSourceDescription	259
Contents	259
See Also	260
OpenXJsonSerDe	261
Contents	261
See Also	262
OrcSerDe	263

Contents	263
See Also	265
OutputFormatConfiguration	266
Contents	266
See Also	266
ParquetSerDe	267
Contents	267
See Also	268
PartitionField	269
Contents	269
See Also	269
PartitionSpec	270
Contents	270
See Also	270
ProcessingConfiguration	271
Contents	271
See Also	271
Processor	272
Contents	272
See Also	272
ProcessorParameter	274
Contents	274
See Also	274
PutRecordBatchResponseEntry	276
Contents	276
See Also	276
Record	278
Contents	278
See Also	278
RedshiftDestinationConfiguration	279
Contents	279
See Also	281
RedshiftDestinationDescription	283
Contents	283
See Also	285
RedshiftDestinationUpdate	286

Contents	286
See Also	288
RedshiftRetryOptions	290
Contents	290
See Also	290
RetryOptions	291
Contents	291
See Also	291
S3DestinationConfiguration	292
Contents	292
See Also	294
S3DestinationDescription	295
Contents	295
See Also	297
S3DestinationUpdate	298
Contents	298
See Also	300
SchemaConfiguration	301
Contents	301
See Also	303
SchemaEvolutionConfiguration	304
Contents	304
See Also	304
SecretsManagerConfiguration	305
Contents	305
See Also	306
Serializer	307
Contents	307
See Also	307
SnowflakeBufferingHints	308
Contents	308
See Also	308
SnowflakeDestinationConfiguration	309
Contents	309
See Also	313
SnowflakeDestinationDescription	315

Contents	315
See Also	318
SnowflakeDestinationUpdate	320
Contents	320
See Also	324
SnowflakeRetryOptions	325
Contents	325
See Also	325
SnowflakeRoleConfiguration	326
Contents	326
See Also	326
SnowflakeVpcConfiguration	327
Contents	327
See Also	327
SourceDescription	328
Contents	328
See Also	328
SplunkBufferingHints	330
Contents	330
See Also	330
SplunkDestinationConfiguration	331
Contents	331
See Also	333
SplunkDestinationDescription	334
Contents	334
See Also	336
SplunkDestinationUpdate	337
Contents	337
See Also	339
SplunkRetryOptions	340
Contents	340
See Also	340
TableCreationConfiguration	341
Contents	341
See Also	341
Tag	342

Contents	342
See Also	342
VpcConfiguration	343
Contents	343
See Also	345
VpcConfigurationDescription	346
Contents	346
See Also	348
Common Errors	349

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 September 3, 2025.

Actions

The following actions are supported:

- [CreateDeliveryStream](#)
- [DeleteDeliveryStream](#)
- [DescribeDeliveryStream](#)
- [ListDeliveryStreams](#)
- [ListTagsForDeliveryStream](#)
- [PutRecord](#)
- [PutRecordBatch](#)
- [StartDeliveryStreamEncryption](#)
- [StopDeliveryStreamEncryption](#)
- [TagDeliveryStream](#)
- [UntagDeliveryStream](#)
- [UpdateDestination](#)

CreateDeliveryStream

Creates a Firehose stream.

By default, you can create up to 5,000 Firehose streams per AWS Region.

This is an asynchronous operation that immediately returns. The initial status of the Firehose stream is `CREATING`. After the Firehose stream is created, its status is `ACTIVE` and it now accepts data. If the Firehose 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 Firehose stream, use [DescribeDeliveryStream](#).

If the status of a Firehose 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 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 Firehose 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 Firehose stream that doesn't have SSE enabled.

A Firehose 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" ]
}
},
"DatabaseSourceConfiguration": {
  "Columns": {
    "Exclude": [ "string" ],
    "Include": [ "string" ]
  },
  "Databases": {
    "Exclude": [ "string" ],
    "Include": [ "string" ]
  },
  "DatabaseSourceAuthenticationConfiguration": {
    "SecretsManagerConfiguration": {
      "Enabled": boolean,
      "RoleARN": "string",
      "SecretARN": "string"
    }
  },
  "DatabaseSourceVPCConfiguration": {
    "VpcEndpointServiceName": "string"
  },
  "Endpoint": "string",
  "Port": number,
  "SnapshotWatermarkTable": "string",

```

```

    "SSLMode": "string",
    "SurrogateKeys": [ "string" ],
    "Tables": {
        "Exclude": [ "string" ],
        "Include": [ "string" ]
    },
    "Type": "string"
},
"DeliveryStreamEncryptionConfigurationInput": {
    "KeyARN": "string",
    "KeyType": "string"
},
"DeliveryStreamName": "string",
"DeliveryStreamType": "string",
"DirectPutSourceConfiguration": {
    "ThroughputHintInMBs": number
},
"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"
  }
},
"IcebergDestinationConfiguration": {
  "AppendOnly": boolean,
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
  "CatalogConfiguration": {
    "CatalogARN": "string",
    "WarehouseLocation": "string"
  },
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "DestinationTableConfigurationList": [
    {
      "DestinationDatabaseName": "string",
      "DestinationTableName": "string",
      "PartitionSpec": {
        "Identity": [
          {
            "SourceName": "string"
          }
        ]
      },
      "S3ErrorOutputPrefix": "string",
      "UniqueKeys": [ "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"
},
"SchemaEvolutionConfiguration": {
    "Enabled": boolean
},
"TableCreationConfiguration": {
    "Enabled": boolean
}
},
"KinesisStreamSourceConfiguration": {
    "KinesisStreamARN": "string",
    "RoleARN": "string"
},
"MSKSourceConfiguration": {

```

```
"AuthenticationConfiguration": {
  "Connectivity": "string",
  "RoleARN": "string"
},
"MSKClusterARN": "string",
"ReadFromTimestamp": number,
"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",
    "BufferingHints": {
      "IntervalInSeconds": number,
      "SizeInMBs": number
    },
    "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

[DatabaseSourceConfiguration](#)

The top level object for configuring streams with database as a source.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseSourceConfiguration](#) 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 Firehose stream. This name must be unique per AWS account in the same AWS Region. If the Firehose streams are in different accounts or different Regions, you can have multiple Firehose 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 Firehose stream type. This parameter can be one of the following values:

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

Type: String

Valid Values: `DirectPut` | `KinesisStreamAsSource` | `MSKAsSource` | `DatabaseAsSource`

Required: No

DirectPutSourceConfiguration

The structure that configures parameters such as `ThroughputHintInMBs` for a stream configured with Direct PUT as a source.

Type: [DirectPutSourceConfiguration](#) object

Required: No

ElasticsearchDestinationConfiguration

The destination in Amazon OpenSearch Service. 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

IcebergDestinationConfiguration

Configure Apache Iceberg Tables destination.

Type: [IcebergDestinationConfiguration](#) object

Required: No

[KinesisStreamSourceConfiguration](#)

When a Kinesis data stream is used as the source for the Firehose 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 Firehose 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 Firehose 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 Firehose 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 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 Firehose stream.

Type: String

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

Pattern: `arn:.*:firehose:[a-zA-Z0-9\-\+]:\d{12}:deliverystream/[a-zA-Z0-9._-]+`

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 Firehose 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 Firehose 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 Firehose stream and apply tags

The following example demonstrates a policy that allows users to create a Firehose 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 Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DeleteDeliveryStream

Deletes a Firehose stream and its data.

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

DeleteDeliveryStream is an asynchronous API. When an API request to DeleteDeliveryStream succeeds, the Firehose stream is marked for deletion, and it goes into the DELETING state. While the Firehose 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 Firehose stream.

Removal of a Firehose 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 Firehose 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 Firehose 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 Firehose 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 Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

DescribeDeliveryStream

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

If the status of a Firehose 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 Firehose 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 Firehose 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 Firehose 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"
  }
},
"IcebergDestinationDescription": {
  "AppendOnly": boolean,
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
  "CatalogConfiguration": {
    "CatalogARN": "string",
    "WarehouseLocation": "string"
  },
  "CloudWatchLoggingOptions": {
    "Enabled": boolean,
    "LogGroupName": "string",
    "LogStreamName": "string"
  },
  "DestinationTableConfigurationList": [
    {
      "DestinationDatabaseName": "string",
      "DestinationTableName": "string",
      "PartitionSpec": {
        "Identity": [
          {
            "SourceName": "string"
          }
        ]
      }
    }
  ]
},

```

```

        "S3ErrorOutputPrefix": "string",
        "UniqueKeys": [ "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"
}

```

```
    },
    "SchemaEvolutionConfiguration": {
      "Enabled": boolean
    },
    "TableCreationConfiguration": {
      "Enabled": boolean
    }
  },
  "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",
    "BufferingHints": {
      "IntervalInSeconds": number,
      "SizeInMBs": number
    },
    "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": {
  "DatabaseSourceDescription": {
    "Columns": {
      "Exclude": [ "string" ],
      "Include": [ "string" ]
    },
    "Databases": {
      "Exclude": [ "string" ],
      "Include": [ "string" ]
    },
    "DatabaseSourceAuthenticationConfiguration": {
      "SecretsManagerConfiguration": {

```



```

        "Enabled": boolean,
        "RoleARN": "string",
        "SecretARN": "string"
    },
},
"DatabaseSourceVPCConfiguration": {
    "VpcEndpointServiceName": "string"
},
"Endpoint": "string",
"Port": number,
"SnapshotInfo": [
    {
        "FailureDescription": {
            "Details": "string",
            "Type": "string"
        },
        "Id": "string",
        "RequestedBy": "string",
        "RequestTimestamp": number,
        "Status": "string",
        "Table": "string"
    }
],
"SnapshotWatermarkTable": "string",
"SSLMode": "string",
"SurrogateKeys": [ "string " ],
"Tables": {
    "Exclude": [ "string " ],
    "Include": [ "string " ]
},
"Type": "string"
},
"DirectPutSourceDescription": {
    "ThroughputHintInMBs": number
},
"KinesisStreamSourceDescription": {
    "DeliveryStartTimestamp": number,
    "KinesisStreamARN": "string",
    "RoleARN": "string"
},
"MSKSourceDescription": {
    "AuthenticationConfiguration": {
        "Connectivity": "string",
        "RoleARN": "string"
    }
}

```

```
    },
    "DeliveryStartTimestamp": number,
    "MSKClusterARN": "string",
    "ReadFromTimestamp": number,
    "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 Firehose 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 Firehose 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-0000000000001",
        "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 Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListDeliveryStreams

Lists your Firehose streams in alphabetical order of their names.

The number of Firehose streams might be too large to return using a single call to `ListDeliveryStreams`. You can limit the number of Firehose 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 Firehose streams to list, you can request them by calling this operation again and setting the `ExclusiveStartDeliveryStreamName` parameter to the name of the last Firehose 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 Firehose stream type. This can be one of the following values:

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

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

Type: String

Valid Values: `DirectPut` | `KinesisStreamAsSource` | `MSKAsSource` | `DatabaseAsSource`

Required: No

ExclusiveStartDeliveryStreamName

The list of Firehose streams returned by this call to `ListDeliveryStreams` will start with the Firehose 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 Firehose 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 Firehose 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 Firehose 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 Firehose streams configured for DirectPut. The response indicates that there are more DirectPut Firehose streams to be listed. To list the remaining DirectPut Firehose 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 Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

ListTagsForDeliveryStream

Lists the tags for the specified Firehose 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 Firehose 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 Firehose 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 Firehose 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 Kotlin](#)
- [AWS SDK for PHP V3](#)

- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

PutRecord

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

By default, each Firehose 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 Firehose 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 Firehose 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 Firehose 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.

For multi record de-aggregation, you can not put more than 500 records even if the data blob length is less than 1000 KiB. If you include more than 500 records, the request succeeds but the record de-aggregation doesn't work as expected and transformation lambda is invoked with the complete base64 encoded data blob instead of de-aggregated base64 decoded records.

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 Firehose 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 Firehose 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 Firehose 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 Firehose 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 Firehose 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 Firehose 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/  
YePEXbK6epIW4QeXqJp2xsbFZUNXsf0Y1QrYXgRBCKznkjMMTP0BqJG0bM3fB//dHgEE0XDTc4wW065i/  
tJyYI1Vy8qn8FMhpkZuh5bvG482XkkBxFmMGnhPTQwQ4A1I0P0sE0X99YnBK8RECdeQ2zxynvZ"  
}
```

Example

The following example shows how you can use the [AWS CLI](#) to put a record in a Firehose 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 Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

PutRecordBatch

Writes multiple data records into a Firehose stream in a single call, which can achieve higher throughput per producer than when writing single records. To write single data records into a Firehose 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 Firehose 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 Firehose 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.

For multi record de-aggregation, you can not put more than 500 records even if the data blob length is less than 1000 KiB. If you include more than 500 records, the request succeeds but the record de-aggregation doesn't work as expected and transformation lambda is invoked with the complete base64 encoded data blob instead of de-aggregated base64 decoded records.

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 Firehose 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 Firehose 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 Firehose 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 Firehose 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 Firehose 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 Firehose 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/
JbRTbXDZFNCrWIs/4YUd3gjYUs6giZn76TuI4vv0ljD0MFJvLzqLuupu93RscZRmpW1CP8DeiFzJJGvqIr1LRE/
MDozYenOz
+v1ZlIqEjECvfMwIz6silvDaGuYtIhVzwd8yWaDGFCX40Des6W07W9W7Q5ViMPUKSt6F0nn70hVkJ3/"
    },
    {
      "RecordId": "goGaFS919Mmv71YET0oMaw+UL9iFpzi100o
+csoIc31SmvkqpWzQuT0RPqZ7QqfR1FJ
+HxJciW/8paFMWPByJ6qVDhiE7TtJxARKaP4YDccvWHXhD7x6Y4bc9AHZ0uzy
+BEuTTnH5Zsip1IrEPgu8lj8a7fd6sluTe/kZ54GVG+Sc105IrxIu08kPbRiUTWL
+Wg0A3Hz0QeePqokRKKqmoXLG0pzx/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 Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

StartDeliveryStreamEncryption

Enables server-side encryption (SSE) for the Firehose 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 Firehose 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 Firehose 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 Firehose 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 Firehose stream, use [DescribeDeliveryStream](#).

Even if encryption is currently enabled for a Firehose 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 Firehose 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 Firehose 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 Firehose stream only if it's a Firehose stream that uses `DirectPut` as its source.

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

example, you reach the limit if you call `StartDeliveryStreamEncryption` 13 times and `StopDeliveryStreamEncryption` 12 times for the same Firehose 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 Firehose 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 Firehose 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 Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

StopDeliveryStreamEncryption

Disables server-side encryption (SSE) for the Firehose 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 Firehose 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 Firehose 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 Firehose stream per 24 hours. For example, you reach the limit if you call `StartDeliveryStreamEncryption` 13 times and `StopDeliveryStreamEncryption` 12 times for the same Firehose 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 Firehose 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 Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

TagDeliveryStream

Adds or updates tags for the specified Firehose 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 Firehose stream. For more information about tags, see [Using Cost Allocation Tags](#) in the *AWS Billing and Cost Management User Guide*.

Each Firehose 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 Firehose 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 Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UntagDeliveryStream

Removes tags from the specified Firehose 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 Firehose 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 Kotlin](#)
- [AWS SDK for PHP V3](#)
- [AWS SDK for Python](#)
- [AWS SDK for Ruby V3](#)

UpdateDestination

Updates the specified destination of the specified Firehose 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 Firehose stream remains active while the configurations are updated, so data writes to the Firehose 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
    }
  },
},
```

```

    "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"
  },
},
"IcebergDestinationUpdate": {
  "AppendOnly": boolean,
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
},
"CatalogConfiguration": {
  "CatalogARN": "string",
  "WarehouseLocation": "string"
},
"CloudWatchLoggingOptions": {
  "Enabled": boolean,
  "LogGroupName": "string",
  "LogStreamName": "string"
},
"DestinationTableConfigurationList": [
  {
    "DestinationDatabaseName": "string",
    "DestinationTableName": "string",
    "PartitionSpec": {
      "Identity": [
        {
          "SourceName": "string"
        }
      ]
    },
    "S3ErrorOutputPrefix": "string",
    "UniqueKeys": [ "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"
},
"SchemaEvolutionConfiguration": {
  "Enabled": boolean
},
"TableCreationConfiguration": {
  "Enabled": boolean
}
},
"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",
  "BufferingHints": {
    "IntervalInSeconds": number,
    "SizeInMBs": number
  },
  "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 Firehose 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 OpenSearch Service.

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

[IcebergDestinationUpdate](#)

Describes an update for a destination in Apache Iceberg Tables.

Type: [IcebergDestinationUpdate](#) 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](#).

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](#) invocation for the Firehose 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-0000000000001",
  "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 Kotlin](#)
- [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](#)
- [CatalogConfiguration](#)
- [CloudWatchLoggingOptions](#)
- [CopyCommand](#)
- [DatabaseColumnList](#)
- [DatabaseList](#)
- [DatabaseSnapshotInfo](#)
- [DatabaseSourceAuthenticationConfiguration](#)
- [DatabaseSourceConfiguration](#)

- [DatabaseSourceDescription](#)
- [DatabaseSourceVPCConfiguration](#)
- [DatabaseTableList](#)
- [DataFormatConversionConfiguration](#)
- [DeliveryStreamDescription](#)
- [DeliveryStreamEncryptionConfiguration](#)
- [DeliveryStreamEncryptionConfigurationInput](#)
- [Deserializer](#)
- [DestinationDescription](#)
- [DestinationTableConfiguration](#)
- [DirectPutSourceConfiguration](#)
- [DirectPutSourceDescription](#)
- [DocumentIdOptions](#)
- [DynamicPartitioningConfiguration](#)
- [ElasticsearchBufferingHints](#)
- [ElasticsearchDestinationConfiguration](#)
- [ElasticsearchDestinationDescription](#)
- [ElasticsearchDestinationUpdate](#)
- [ElasticsearchRetryOptions](#)
- [EncryptionConfiguration](#)
- [ExtendedS3DestinationConfiguration](#)
- [ExtendedS3DestinationDescription](#)
- [ExtendedS3DestinationUpdate](#)
- [FailureDescription](#)
- [HiveJsonSerDe](#)
- [HttpEndpointBufferingHints](#)
- [HttpEndpointCommonAttribute](#)
- [HttpEndpointConfiguration](#)
- [HttpEndpointDescription](#)
- [HttpEndpointDestinationConfiguration](#)

- [HttpEndpointDestinationDescription](#)
- [HttpEndpointDestinationUpdate](#)
- [HttpEndpointRequestConfiguration](#)
- [HttpEndpointRetryOptions](#)
- [IcebergDestinationConfiguration](#)
- [IcebergDestinationDescription](#)
- [IcebergDestinationUpdate](#)
- [InputFormatConfiguration](#)
- [KinesisStreamSourceConfiguration](#)
- [KinesisStreamSourceDescription](#)
- [KMSEncryptionConfig](#)
- [MSKSourceConfiguration](#)
- [MSKSourceDescription](#)
- [OpenXJsonSerDe](#)
- [OrcSerDe](#)
- [OutputFormatConfiguration](#)
- [ParquetSerDe](#)
- [PartitionField](#)
- [PartitionSpec](#)
- [ProcessingConfiguration](#)
- [Processor](#)
- [ProcessorParameter](#)
- [PutRecordBatchResponseEntry](#)
- [Record](#)
- [RedshiftDestinationConfiguration](#)
- [RedshiftDestinationDescription](#)
- [RedshiftDestinationUpdate](#)
- [RedshiftRetryOptions](#)
- [RetryOptions](#)
- [S3DestinationConfiguration](#)

- [S3DestinationDescription](#)
- [S3DestinationUpdate](#)
- [SchemaConfiguration](#)
- [SchemaEvolutionConfiguration](#)
- [SecretsManagerConfiguration](#)
- [Serializer](#)
- [SnowflakeBufferingHints](#)
- [SnowflakeDestinationConfiguration](#)
- [SnowflakeDestinationDescription](#)
- [SnowflakeDestinationUpdate](#)
- [SnowflakeRetryOptions](#)
- [SnowflakeRoleConfiguration](#)
- [SnowflakeVpcConfiguration](#)
- [SourceDescription](#)
- [SplunkBufferingHints](#)
- [SplunkDestinationConfiguration](#)
- [SplunkDestinationDescription](#)
- [SplunkDestinationUpdate](#)
- [SplunkRetryOptions](#)
- [TableCreationConfiguration](#)
- [Tag](#)
- [VpcConfiguration](#)
- [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 Firehose 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:.*:iam:.*:\d{12}:role/[a-zA-Z_0-9+=,.\@-\_/\]^+*`

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

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 OpenSearch Service 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 Firehose 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:.*:iam:.*:\d{12}:role/[a-zA-Z_0-9+=,.\@\-_/\]+`

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 Firehose 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 Elasticsearch 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 Amazon OpenSearch Service Configuration API and for indexing documents.

Type: String

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

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

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 Firehose 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:.*:es:[a-zA-Z0-9\-\_]+\:\d{12}:domain/[a-z]([-0-9a-z]{2,27})`

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 Firehose 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:.*:es:[a-zA-Z0-9\-\_]+\:\d{12}:domain/[a-z][-0-9a-z]{2,27}`

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](#) 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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@-\_/\]+`

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 OpenSearch Service 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 Firehose 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:.*:es:[a-zA-Z0-9\-\_]+\:\d{12}:domain/[a-z][-0-9a-z]{2,27}`

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:.*:iam:.*:role/([a-zA-Z0-9+=,.\@\-_\/])+`

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 Firehose stream, Firehose still delivers data to Elasticsearch with the old index name and type name. If you want to update your Firehose 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:.*:iam:.*:\d{12}:role/[a-zA-Z_0-9+=,.\@\-_/\]+`

Required: Yes

See Also

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

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

CatalogConfiguration

Describes the containers where the destination Apache Iceberg Tables are persisted.

Contents

CatalogARN

Specifies the Glue catalog ARN identifier of the destination Apache Iceberg Tables. You must specify the ARN in the format `arn:aws:glue:region:account-id:catalog`.

Type: String

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

Pattern: `arn:.*:glue:.*:\d{12}:catalog(?:(/[a-z0-9_-]{1,2}))?`

Required: No

WarehouseLocation

The warehouse location for Apache Iceberg tables. You must configure this when schema evolution and table creation is enabled.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

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

Pattern: `s3:\V\.`*

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 Firehose 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,venue:25,venuecity:12,venuestate:2,venue: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 10240.

Pattern: . *

Required: No

DataTableColumns

A comma-separated list of column names.

Type: String

Length Constraints: Minimum length of 0. Maximum length of 10240.

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](#)

DatabaseColumnList

The structure used to configure the list of column patterns in source database endpoint for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Contents

Exclude

The list of column patterns in source database to be excluded for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 194.

Pattern: `[\u0001-\uFFFF]*`

Required: No

Include

The list of column patterns in source database to be included for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 194.

Pattern: `[\u0001-\uFFFF]*`

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](#)

DatabaseList

The structure used to configure the list of database patterns in source database endpoint for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Contents

Exclude

The list of database patterns in source database endpoint to be excluded for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: Array of strings

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

Pattern: `[\u0001-\uFFFF]*`

Required: No

Include

The list of database patterns in source database endpoint to be included for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: Array of strings

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

Pattern: `[\u0001-\uFFFF]*`

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](#)

DatabaseSnapshotInfo

The structure that describes the snapshot information of a table in source database endpoint that Firehose reads.

Amazon Data Firehose is in preview release and is subject to change.

Contents

Id

The identifier of the current snapshot of the table in source database endpoint.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

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

Pattern: `^\S+$`

Required: Yes

RequestedBy

The principal that sent the request to take the current snapshot on the table.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

Valid Values: `USER` | `FIREHOSE`

Required: Yes

RequestTimestamp

The timestamp when the current snapshot is taken on the table.

Amazon Data Firehose is in preview release and is subject to change.

Type: Timestamp

Required: Yes

Status

The status of the current snapshot of the table.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

Valid Values: IN_PROGRESS | COMPLETE | SUSPENDED

Required: Yes

Table

The fully qualified name of the table in source database endpoint that Firehose reads.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 129.

Pattern: `[\u0001-\uFFFF]*`

Required: Yes

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

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](#)

DatabaseSourceAuthenticationConfiguration

The structure to configure the authentication methods for Firehose to connect to source database endpoint.

Amazon Data Firehose is in preview release and is subject to change.

Contents

SecretsManagerConfiguration

The structure that defines how Firehose accesses the secret.

Type: [SecretsManagerConfiguration](#) object

Required: Yes

See Also

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

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

DatabaseSourceConfiguration

The top level object for configuring streams with database as a source.

Amazon Data Firehose is in preview release and is subject to change.

Contents

Databases

The list of database patterns in source database endpoint for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseList](#) object

Required: Yes

DatabaseSourceAuthenticationConfiguration

The structure to configure the authentication methods for Firehose to connect to source database endpoint.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseSourceAuthenticationConfiguration](#) object

Required: Yes

DatabaseSourceVPCConfiguration

The details of the VPC Endpoint Service which Firehose uses to create a PrivateLink to the database.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseSourceVPCConfiguration](#) object

Required: Yes

Endpoint

The endpoint of the database server.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

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

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

Required: Yes

Port

The port of the database. This can be one of the following values.

- 3306 for MySQL database type
- 5432 for PostgreSQL database type

Amazon Data Firehose is in preview release and is subject to change.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 65535.

Required: Yes

SnapshotWatermarkTable

The fully qualified name of the table in source database endpoint that Firehose uses to track snapshot progress.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 129.

Pattern: `[\\u0001-\\uFFFF]*`

Required: Yes

Tables

The list of table patterns in source database endpoint for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseTableList](#) object

Required: Yes

Type

The type of database engine. This can be one of the following values.

- MySQL
- PostgreSQL

Amazon Data Firehose is in preview release and is subject to change.

Type: String

Valid Values: MySQL | PostgreSQL

Required: Yes

Columns

The list of column patterns in source database endpoint for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseColumnList](#) object

Required: No

SSLMode

The mode to enable or disable SSL when Firehose connects to the database endpoint.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

Valid Values: Disabled | Enabled

Required: No

SurrogateKeys

The optional list of table and column names used as unique key columns when taking snapshot if the tables don't have primary keys configured.

Amazon Data Firehose is in preview release and is subject to change.

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](#)

DatabaseSourceDescription

The top level object for database source description.

Amazon Data Firehose is in preview release and is subject to change.

Contents

Columns

The list of column patterns in source database endpoint for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseColumnList](#) object

Required: No

Databases

The list of database patterns in source database endpoint for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseList](#) object

Required: No

DatabaseSourceAuthenticationConfiguration

The structure to configure the authentication methods for Firehose to connect to source database endpoint.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseSourceAuthenticationConfiguration](#) object

Required: No

DatabaseSourceVPCConfiguration

The details of the VPC Endpoint Service which Firehose uses to create a PrivateLink to the database.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseSourceVPCConfiguration](#) object

Required: No

Endpoint

The endpoint of the database server.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

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

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

Required: No

Port

The port of the database. This can be one of the following values.

- 3306 for MySQL database type
- 5432 for PostgreSQL database type

Amazon Data Firehose is in preview release and is subject to change.

Type: Integer

Valid Range: Minimum value of 0. Maximum value of 65535.

Required: No

SnapshotInfo

The structure that describes the snapshot information of a table in source database endpoint that Firehose reads.

Amazon Data Firehose is in preview release and is subject to change.

Type: Array of [DatabaseSnapshotInfo](#) objects

Required: No

SnapshotWatermarkTable

The fully qualified name of the table in source database endpoint that Firehose uses to track snapshot progress.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

Length Constraints: Minimum length of 1. Maximum length of 129.

Pattern: `[\u0001-\uFFFF]*`

Required: No

SSLMode

The mode to enable or disable SSL when Firehose connects to the database endpoint.

Amazon Data Firehose is in preview release and is subject to change.

Type: String

Valid Values: Disabled | Enabled

Required: No

SurrogateKeys

The optional list of table and column names used as unique key columns when taking snapshot if the tables don't have primary keys configured.

Amazon Data Firehose is in preview release and is subject to change.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 194.

Pattern: `[\u0001-\uFFFF]*`

Required: No

Tables

The list of table patterns in source database endpoint for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseTableList](#) object

Required: No

Type

The type of database engine. This can be one of the following values.

- MySQL
- PostgreSQL

Amazon Data Firehose is in preview release and is subject to change.

Type: String

Valid Values: MySQL | PostgreSQL

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](#)

DatabaseSourceVPCConfiguration

The structure for details of the VPC Endpoint Service which Firehose uses to create a PrivateLink to the database.

Amazon Data Firehose is in preview release and is subject to change.

Contents

VpcEndpointServiceName

The VPC endpoint service name which Firehose uses to create a PrivateLink to the database. The endpoint service must have the Firehose service principle `firehose.amazonaws.com` as an allowed principal on the VPC endpoint service. The VPC endpoint service name is a string that looks like `com.amazonaws.vpce.<region>.<vpc-endpoint-service-id>`.

Amazon Data Firehose is in preview release and is subject to change.

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](#)

DatabaseTableList

The structure used to configure the list of table patterns in source database endpoint for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Contents

Exclude

The list of table patterns in source database endpoint to be excluded for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 129.

Pattern: `[\u0001-\uFFFF]*`

Required: No

Include

The list of table patterns in source database endpoint to be included for Firehose to read from.

Amazon Data Firehose is in preview release and is subject to change.

Type: Array of strings

Length Constraints: Minimum length of 1. Maximum length of 129.

Pattern: `[\u0001-\uFFFF]*`

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 Firehose stream.

Contents

DeliveryStreamARN

The Amazon Resource Name (ARN) of the Firehose 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:.*:firehose:[a-zA-Z0-9\-\_]+\:\d{12}:deliverystream/[a-zA-Z0-9._-\_]+`

Required: Yes

DeliveryStreamName

The name of the Firehose 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 Firehose stream. If the status of a Firehose 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 Firehose stream type. This can be one of the following values:

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

Type: String

Valid Values: `DirectPut` | `KinesisStreamAsSource` | `MSKAsSource` | `DatabaseAsSource`

Required: Yes

Destinations

The destinations.

Type: Array of [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 Firehose 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 Firehose stream was created.

Type: Timestamp

Required: No

DeliveryStreamEncryptionConfiguration

Indicates the server-side encryption (SSE) status for the Firehose 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 Firehose 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](#) operation.

Contents

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

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:.*:kms:[a-zA-Z0-9\-\+]:\d{12}:(key|alias)/[a-zA-Z0-9+=,.\@\_/_]+`

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 \(CMKs\)](#).

Type: String

Valid Values: `AWS_OWNED_CMK` | `CUSTOMER_MANAGED_CMK`

Required: No

Status

This is the server-side encryption (SSE) status for the Firehose 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 Firehose 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 Firehose streams. If a [CreateDeliveryStream](#) or [StartDeliveryStreamEncryption](#) operation exceeds this limit, Firehose throws a `LimitExceededException`.

Important

To encrypt your Firehose 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:.*:kms:[a-zA-Z0-9\-\+]:\d{12}:(key|alias)/[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](#)

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 Firehose 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 OpenSearch Service.

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

IcebergDestinationDescription

Describes a destination in Apache Iceberg Tables.

Type: [IcebergDestinationDescription](#) 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](#)

DestinationTableConfiguration

Describes the configuration of a destination in Apache Iceberg Tables. This section is only needed for tables where you want to update or delete data.

Contents

DestinationDatabaseName

The name of the Apache Iceberg database.

Type: String

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

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

Required: Yes

DestinationTableName

Specifies the name of the Apache Iceberg Table.

Type: String

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

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

Required: Yes

PartitionSpec

The partition spec configuration for a table that is used by automatic table creation.

Amazon Data Firehose is in preview release and is subject to change.

Type: [PartitionSpec](#) object

Required: No

S3ErrorOutputPrefix

The table specific S3 error output prefix. All the errors that occurred while delivering to this table will be prefixed with this value in S3 destination.

Type: String

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

Pattern: .*

Required: No

UniqueKeys

A list of unique keys for a given Apache Iceberg table. Firehose will use these for running Create, Update, or Delete operations on the given Iceberg table.

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](#)

DirectPutSourceConfiguration

The structure that configures parameters such as `ThroughputHintInMBs` for a stream configured with Direct PUT as a source.

Contents

ThroughputHintInMBs

The value that you configure for this parameter is for information purpose only and does not affect Firehose delivery throughput limit. You can use the [Firehose Limits form](#) to request a throughput limit increase.

Type: Integer

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

Required: Yes

See Also

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

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

DirectPutSourceDescription

The structure that configures parameters such as `ThroughputHintInMBs` for a stream configured with Direct PUT as a source.

Contents

ThroughputHintInMBs

The value that you configure for this parameter is for information purpose only and does not affect Firehose delivery throughput limit. You can use the [Firehose Limits form](#) to request a throughput limit increase.

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](#)

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 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 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 Firehose 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 OpenSearch Service.

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

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 Firehose 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 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:.*:es:[a-zA-Z0-9\-\_]+\:\d{12}:domain/[a-z]([-0-9a-z]{2,27})`

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 OpenSearch Service Destination](#). The default value is `OneDay`.

Type: String

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

Required: No

ProcessingConfiguration

The data processing configuration.

Type: [ProcessingConfiguration](#) 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: [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 OpenSearch Service Destination](#). Default value is `FailedDocumentsOnly`.

You can't change this backup mode after you create the Firehose 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 OpenSearch Service.

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 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. For more information, see [Amazon Resource Names \(ARNs\) and AWS Service Namespaces](#).

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

Type: String

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

Pattern: `arn:.*:es:[a-zA-Z0-9\-\+]:\d{12}:domain/[a-z][-0-9a-z]{2,27}`

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 OpenSearch Service 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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@-\_/\]+`

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 OpenSearch Service.

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 Firehose 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`. 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:.*:es:[a-zA-Z0-9\-\_]+\:\d{12}:domain/[a-z] [-0-9a-z]{2,27}`

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 OpenSearch Service 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 OpenSearch Service. 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 OpenSearch Service 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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@\-_/\]+`

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 Firehose stream, Firehose still delivers data to Elasticsearch with the old index name and type name. If you want to update your Firehose 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 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](#)

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:.*:s3::[\w\.-]{1,255}`

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

Required: Yes

BufferingHints

The buffering option.

Type: [BufferingHints](#) object

Required: No

CloudWatchLoggingOptions

The Amazon CloudWatch logging options for your Firehose 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.

Pattern: ^\$| [a-zA-Z/_]+

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 Firehose stream, you can update it to enable Amazon S3 backup if it is disabled. If backup is enabled, you can't update the Firehose 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:.*:s3::[\w\.\-]{1,255}`

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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@-_/]+`

Required: Yes

CloudWatchLoggingOptions

The Amazon CloudWatch logging options for your Firehose 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.

Pattern: `^$| [a-zA-Z/_]+`

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:.*:s3::[\w\.\-]{1,255}`

Required: No

BufferingHints

The buffering option.

Type: [BufferingHints](#) object

Required: No

CloudWatchLoggingOptions

The Amazon CloudWatch logging options for your Firehose 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.

Pattern: `^$| [a-zA-Z/_]+`

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

Required: No

S3BackupMode

You can update a Firehose stream to enable Amazon S3 backup if it is disabled. If backup is enabled, you can't update the Firehose 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: VPC_ENDPOINT_SERVICE_NAME_NOT_FOUND | VPC_INTERFACE_ENDPOINT_SERVICE_ACCESS_DENIED | 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 Firehose 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 Firehose 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:.*:iam:.*:\d{12}:role/[a-zA-Z_0-9+=,.\@\-_/\]+`

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 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.

Type: [HttpEndpointBufferingHints](#) object

Required: No

CloudWatchLoggingOptions

Describes the Amazon CloudWatch logging options for your Firehose 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](#) 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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@\-_/\]+`

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](#) 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 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.

Type: [HttpEndpointBufferingHints](#) object

Required: No

CloudWatchLoggingOptions

Describes the Amazon CloudWatch logging options for your Firehose 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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@-\_/\]+`

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](#)

IcebergDestinationConfiguration

Specifies the destination configure settings for Apache Iceberg Table.

Contents

CatalogConfiguration

Configuration describing where the destination Apache Iceberg Tables are persisted.

Type: [CatalogConfiguration](#) object

Required: Yes

RoleARN

The Amazon Resource Name (ARN) of the IAM role to be assumed by Firehose for calling Apache Iceberg Tables.

Type: String

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

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

Required: Yes

S3Configuration

Describes the configuration of a destination in Amazon S3.

Type: [S3DestinationConfiguration](#) object

Required: Yes

AppendOnly

Describes whether all incoming data for this delivery stream will be append only (inserts only and not for updates and deletes) for Iceberg delivery. This feature is only applicable for Apache Iceberg Tables.

The default value is false. If you set this value to true, Firehose automatically increases the throughput limit of a stream based on the throttling levels of the stream. If you set this parameter to true for a stream with updates and deletes, you will see out of order delivery.

Type: Boolean

Required: No

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.

Type: [BufferingHints](#) object

Required: No

CloudWatchLoggingOptions

Describes the Amazon CloudWatch logging options for your Firehose stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

DestinationTableConfigurationList

Provides a list of `DestinationTableConfigurations` which Firehose uses to deliver data to Apache Iceberg Tables. Firehose will write data with insert if table specific configuration is not provided here.

Type: Array of [DestinationTableConfiguration](#) objects

Required: No

ProcessingConfiguration

Describes a data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

RetryOptions

The retry behavior in case Firehose is unable to deliver data to a destination.

Type: [RetryOptions](#) object

Required: No

S3BackupMode

Describes how Firehose will backup records. Currently, S3 backup only supports FailedDataOnly.

Type: String

Valid Values: FailedDataOnly | AllData

Required: No

SchemaEvolutionConfiguration

The configuration to enable automatic schema evolution.

Amazon Data Firehose is in preview release and is subject to change.

Type: [SchemaEvolutionConfiguration](#) object

Required: No

TableCreationConfiguration

The configuration to enable automatic table creation.

Amazon Data Firehose is in preview release and is subject to change.

Type: [TableCreationConfiguration](#) 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](#)

IcebergDestinationDescription

Describes a destination in Apache Iceberg Tables.

Contents

AppendOnly

Describes whether all incoming data for this delivery stream will be append only (inserts only and not for updates and deletes) for Iceberg delivery. This feature is only applicable for Apache Iceberg Tables.

The default value is false. If you set this value to true, Firehose automatically increases the throughput limit of a stream based on the throttling levels of the stream. If you set this parameter to true for a stream with updates and deletes, you will see out of order delivery.

Type: Boolean

Required: No

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.

Type: [BufferingHints](#) object

Required: No

CatalogConfiguration

Configuration describing where the destination Iceberg tables are persisted.

Type: [CatalogConfiguration](#) object

Required: No

CloudWatchLoggingOptions

Describes the Amazon CloudWatch logging options for your Firehose stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

DestinationTableConfigurationList

Provides a list of `DestinationTableConfigurations` which Firehose uses to deliver data to Apache Iceberg Tables. Firehose will write data with insert if table specific configuration is not provided here.

Type: Array of [DestinationTableConfiguration](#) objects

Required: No

ProcessingConfiguration

Describes a data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

RetryOptions

The retry behavior in case Firehose is unable to deliver data to a destination.

Type: [RetryOptions](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the IAM role to be assumed by Firehose for calling Apache Iceberg Tables.

Type: String

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

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

Required: No

S3BackupMode

Describes how Firehose will backup records. Currently, Firehose only supports `FailedDataOnly`.

Type: String

Valid Values: FailedDataOnly | AllData

Required: No

S3DestinationDescription

Describes a destination in Amazon S3.

Type: [S3DestinationDescription](#) object

Required: No

SchemaEvolutionConfiguration

The description of automatic schema evolution configuration.

Amazon Data Firehose is in preview release and is subject to change.

Type: [SchemaEvolutionConfiguration](#) object

Required: No

TableCreationConfiguration

The description of table creation configuration.

Amazon Data Firehose is in preview release and is subject to change.

Type: [TableCreationConfiguration](#) 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](#)

IcebergDestinationUpdate

Describes an update for a destination in Apache Iceberg Tables.

Contents

AppendOnly

Describes whether all incoming data for this delivery stream will be append only (inserts only and not for updates and deletes) for Iceberg delivery. This feature is only applicable for Apache Iceberg Tables.

The default value is false. If you set this value to true, Firehose automatically increases the throughput limit of a stream based on the throttling levels of the stream. If you set this parameter to true for a stream with updates and deletes, you will see out of order delivery.

Type: Boolean

Required: No

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.

Type: [BufferingHints](#) object

Required: No

CatalogConfiguration

Configuration describing where the destination Iceberg tables are persisted.

Type: [CatalogConfiguration](#) object

Required: No

CloudWatchLoggingOptions

Describes the Amazon CloudWatch logging options for your Firehose stream.

Type: [CloudWatchLoggingOptions](#) object

Required: No

DestinationTableConfigurationList

Provides a list of `DestinationTableConfigurations` which Firehose uses to deliver data to Apache Iceberg Tables. Firehose will write data with insert if table specific configuration is not provided here.

Type: Array of [DestinationTableConfiguration](#) objects

Required: No

ProcessingConfiguration

Describes a data processing configuration.

Type: [ProcessingConfiguration](#) object

Required: No

RetryOptions

The retry behavior in case Firehose is unable to deliver data to a destination.

Type: [RetryOptions](#) object

Required: No

RoleARN

The Amazon Resource Name (ARN) of the IAM role to be assumed by Firehose for calling Apache Iceberg Tables.

Type: String

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

Pattern: `arn:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@\-_/\]+`

Required: No

S3BackupMode

Describes how Firehose will backup records. Currently, Firehose only supports `FailedDataOnly`.

Type: String

Valid Values: FailedDataOnly | AllData

Required: No

S3Configuration

Describes the configuration of a destination in Amazon S3.

Type: [S3DestinationConfiguration](#) object

Required: No

SchemaEvolutionConfiguration

The configuration to enable automatic schema evolution.

Amazon Data Firehose is in preview release and is subject to change.

Type: [SchemaEvolutionConfiguration](#) object

Required: No

TableCreationConfiguration

The configuration to enable automatic table creation.

Amazon Data Firehose is in preview release and is subject to change.

Type: [TableCreationConfiguration](#) 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](#)

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 Firehose 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:.*:kinesis:[a-zA-Z0-9\-]+\:\d{12}:stream/[a-zA-Z0-9_\-\.]+\`

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

Required: Yes

See Also

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

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

KinesisStreamSourceDescription

Details about a Kinesis data stream used as the source for a Firehose 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:.*:kinesis:[a-zA-Z0-9\-\+]:\d{12}:stream/[a-zA-Z0-9_\. -]+`

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

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](#)

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:.*:kms:[a-zA-Z0-9\-]+\:\d{12}:(key|alias)/[a-zA-Z_0-9+=,.\@\_/-\_/\]+`

Required: Yes

See Also

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

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

ReadFromTimestamp

The start date and time in UTC for the offset position within your MSK topic from where Firehose begins to read. By default, this is set to timestamp when Firehose becomes Active.

If you want to create a Firehose stream with Earliest start position from SDK or CLI, you need to set the `ReadFromTimestamp` parameter to Epoch (1970-01-01T00:00:00Z).

Type: Timestamp

Required: No

See Also

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

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

MSKSourceDescription

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

ReadFromTimestamp

The start date and time in UTC for the offset position within your MSK topic from where Firehose begins to read. By default, this is set to timestamp when Firehose becomes Active.

If you want to create a Firehose stream with Earliest start position from SDK or CLI, you need to set the `ReadFromTimestampUTC` parameter to Epoch (1970-01-01T00:00:00Z).

Type: Timestamp

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](#)

PartitionField

Represents a single field in a `PartitionSpec`.

Amazon Data Firehose is in preview release and is subject to change.

Contents

SourceName

The column name to be configured in partition spec.

Amazon Data Firehose is in preview release and is subject to change.

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](#)

PartitionSpec

Represents how to produce partition data for a table. Partition data is produced by transforming columns in a table. Each column transform is represented by a named `PartitionField`.

Here is an example of the schema in JSON.

```
"partitionSpec": { "identity": [ {"sourceName": "column1"}, {"sourceName": "column2"}, {"sourceName": "column3"} ] }
```

Amazon Data Firehose is in preview release and is subject to change.

Contents

Identity

List of identity [transforms](#) that performs an identity transformation. The transform takes the source value, and does not modify it. Result type is the source type.

Amazon Data Firehose is in preview release and is subject to change.

Type: Array of [PartitionField](#) 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](#)

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 Firehose stream, it receives a record ID. If the record fails to be added to your Firehose 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 Firehose 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:.*:iam::\d{12}:role/[a-zA-Z_0-9+=,.\@_-/]+

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 Firehose 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 Firehose stream, you can update it to enable Amazon S3 backup if it is disabled. If backup is enabled, you can't update the Firehose 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:.*:iam::\d{12}:role/[a-zA-Z_0-9+=, .@-_/]+

Required: Yes

S3DestinationDescription

The Amazon S3 destination.

Type: [S3DestinationDescription](#) object

Required: Yes

CloudWatchLoggingOptions

The Amazon CloudWatch logging options for your Firehose 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 Firehose 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:.*:iam:.*:\d{12}:role/[a-zA-Z_0-9+=,.\@-\_/\]+`

Required: No

S3BackupMode

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

Contents

DurationInSeconds

The period of time during which Firehose retries to deliver data to the specified destination.

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:.*:s3::[\w\.\-]{1,255}`

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:.*:iam::\d{12}:role/[a-zA-Z_0-9+=,.\@\-_/\r]{1,512}`

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 Firehose 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:.*:s3::[\w\.-]{1,255}`

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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@\-_/\]+`

Required: Yes

CloudWatchLoggingOptions

The Amazon CloudWatch logging options for your Firehose 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:.*:s3:::[\w\.-]{1,255}`

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

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](#)

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](#)

SchemaEvolutionConfiguration

The configuration to enable schema evolution.

Amazon Data Firehose is in preview release and is subject to change.

Contents

Enabled

Specify whether you want to enable schema evolution.

Amazon Data Firehose is in preview release and is subject to change.

Type: Boolean

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](#)

SecretsManagerConfiguration

The structure that defines how Firehose accesses the secret.

Contents

Enabled

Specifies whether you want to use 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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@\-_/\]^+*`

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 Firehose 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:.*:secretsmanager:[a-zA-Z0-9\-\+]:\d{12}:secret:[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](#)

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](#)

SnowflakeBufferingHints

Describes the buffering to perform before delivering data to the Snowflake destination. If you do not specify any value, Firehose uses the default values.

Contents

IntervalInSeconds

Buffer incoming data for the specified period of time, in seconds, before delivering it to the destination. The default value is 0.

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 128.

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](#)

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:.*:iam:.*:role/[a-zA-Z_0-9+=, .@-_/]+

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

BufferingHints

Describes the buffering to perform before delivering data to the Snowflake destination. If you do not specify any value, Firehose uses the default values.

Type: [SnowflakeBufferingHints](#) object

Required: No

CloudWatchLoggingOptions

Describes the Amazon CloudWatch logging options for your Firehose 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

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

Specify a column name in the table, where the metadata information has to be loaded. When you enable this field, you will see the following column in the snowflake table, which differs based on the source type.

For Direct PUT as source

```
{ "firehoseDeliveryStreamName" : "streamname", "IngestionTime" :  
"timestamp" }
```

For Kinesis Data Stream as source

```
"kinesisStreamName" : "streamname", "kinesisShardId" : "Id",  
"kinesisPartitionKey" : "key", "kinesisSequenceNumber" : "1234",  
"subsequenceNumber" : "2334", "IngestionTime" : "timestamp" }
```

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

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

Type: [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](#) 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

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

BufferingHints

Describes the buffering to perform before delivering data to the Snowflake destination. If you do not specify any value, Firehose uses the default values.

Type: [SnowflakeBufferingHints](#) object

Required: No

CloudWatchLoggingOptions

Describes the Amazon CloudWatch logging options for your Firehose 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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@\-_/\]+`

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

BufferingHints

Describes the buffering to perform before delivering data to the Snowflake destination.

Type: [SnowflakeBufferingHints](#) object

Required: No

CloudWatchLoggingOptions

Describes the Amazon CloudWatch logging options for your Firehose 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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@-\_/\]+`

Required: No

S3BackupMode

Choose an S3 backup mode. Once you set the mode as `AllData`, you can not change it to `FailedDataOnly`.

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 Java V2](#)
- [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

PrivateLinkVpcId

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\-\_\.]{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 stream.

Contents

DatabaseSourceDescription

Details about a database used as the source for a Firehose stream.

Amazon Data Firehose is in preview release and is subject to change.

Type: [DatabaseSourceDescription](#) object

Required: No

DirectPutSourceDescription

Details about Direct PUT used as the source for a Firehose stream.

Type: [DirectPutSourceDescription](#) object

Required: No

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 Firehose 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

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](#) 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 `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 Firehose 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 Firehose 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](#)

TableCreationConfiguration

The configuration to enable automatic table creation.

Amazon Data Firehose is in preview release and is subject to change.

Contents

Enabled

Specify whether you want to enable automatic table creation.

Amazon Data Firehose is in preview release and is subject to change.

Type: Boolean

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](#)

Tag

Metadata that you can assign to a Firehose 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 Firehose 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:.*:iam:.*:\d{12}:role/[a-zA-Z_0-9+=,.\@\-_/\]+`

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 OpenSearch Service destination. You can use the same security group that the Amazon OpenSearch Service domain uses or different ones. If you specify different security groups here, ensure that they allow outbound HTTPS traffic to the Amazon OpenSearch Service domain's security group. Also ensure that the Amazon OpenSearch Service 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 OpenSearch Service 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 OpenSearch Service 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 OpenSearch Service 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 Firehose 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 OpenSearch Service destination.

Contents

RoleARN

The ARN of the IAM role that the Firehose 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 Firehose 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:.*:iam:.*:role/[a-zA-Z_0-9+=,.\@-_/]+`

Required: Yes

SecurityGroupIds

The IDs of the security groups that Firehose uses when it creates ENIs in the VPC of the Amazon OpenSearch Service 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 OpenSearch Service domain's security group. Also ensure that the Amazon OpenSearch Service domain's security group allows HTTPS traffic from

the security groups specified here. If you use the same security group for both your Firehose stream and the Amazon OpenSearch Service 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 OpenSearch Service 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 OpenSearch Service 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 Firehose 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 OpenSearch Service 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