
aiofstream Documentation

Release 0.2.1

Robert Marki

Jul 23, 2018

Contents

1	Features	3
2	Usage	5
3	Contents	7
3.1	User's guide	7
3.1.1	Installation	7
3.1.2	Quickstart	7
3.1.3	Advanced Usage	10
3.2	API Reference	14
3.2.1	Client	14
3.2.2	Authenticators	17
3.2.3	Replay	19
3.2.4	Exceptions	20
3.3	Changelog	21
3.3.1	0.2.1 (2018-05-25)	21
3.3.2	0.2.0 (2018-05-05)	21
3.3.3	0.1.0 (2018-04-26)	21
4	Indices and tables	23
	Python Module Index	25

aiofsstream is a [Salesforce Streaming API](#) client for [asyncio](#). It can be used to receive push notifications about changes on Salesforce objects or notifications of general events sent through the [Streaming API](#).

For detailed guidance on how to work with [PushTopics](#) or how to create [Generic Streaming Channels](#) please consult the [Streaming API](#) documentation.

CHAPTER 1

Features

- **Supported authentication types:**
 - using a username and password
 - using a refresh token
- **Subscribe to and receive messages on:**
 - [PushTopics](#)
 - [Generic Streaming Channels](#)
- Support for [durable messages and replay of events](#)
- Automatic recovery from replay errors

CHAPTER 2

Usage

```
import asyncio

from aiosfstream import SalesforceStreamingClient

async def stream_events():
    # connect to Streaming API
    async with SalesforceStreamingClient(
        consumer_key="<consumer key>",
        consumer_secret="<consumer secret>",
        username="<username>",
        password="<password>") as client:

        # subscribe to topics
        await client.subscribe("/topic/one")
        await client.subscribe("/topic/two")

        # listen for incoming messages
        async for message in client:
            topic = message["channel"]
            data = message["data"]
            print(f"{topic}: {data}")

if __name__ == "__main__":
    loop = asyncio.get_event_loop()
    loop.run_until_complete(stream_events())
```


3.1 User's guide

3.1.1 Installation

```
pip install aiosfstream
```

Install extras

aiosfstream defines several groups of optional requirements:

- `tests` for running unit tests
- `docs` for building the documentation
- `dev` for creating a complete development environment

Any combination of these options can be specified during installation.

```
pip install aiosfstream[tests,docs,dev]
```

3.1.2 Quickstart

Authentication

To connect to the [Salesforce Streaming API](#) all clients must authenticate themselves. The library supports the [username-password](#) based OAuth2 authentication flow as well as the [refresh token](#) based authentication.

Whichever technique you end up using, you must first create a [Connected App](#) on Salesforce to acquire a Consumer Key and Consumer Secret value. Which are actually the `client_id` and `client_secret` parameters in OAuth2 terminology.

Username-Password authentication

For username-password based authentication you can use the *SalesforceStreamingClient* class, with the Salesforce user's username and password:

```
client = SalesforceStreamingClient(  
    consumer_key="<consumer key>",  
    consumer_secret="<consumer secret>",  
    username="<username>",  
    password="<password>"  
)
```

SalesforceStreamingClient is actually just a convenience class, based on *Client*. It enables you to create a client object with the most common authentication technique, without having to create a separate *PasswordAuthenticator* object. You can actually use the *Client* class to create client that would be equivalent with the example above:

```
auth = PasswordAuthenticator(  
    consumer_key="<consumer key>",  
    consumer_secret="<consumer secret>",  
    username="<username>",  
    password="<password>"  
)  
client = Client(auth)
```

Refresh token authentication

The refresh token base authentication technique can be used by creating a *RefreshTokenAuthenticator* and passing it to the *Client* class:

```
auth = RefreshTokenAuthenticator(  
    consumer_key="<consumer key>",  
    consumer_secret="<consumer secret>",  
    refresh_token="<refresh_token>"  
)  
client = Client(auth)
```

You can get a refresh token using several different authentication techniques supported by Salesforce, the most commonly used one is probably the [web server authentication flow](#).

Connecting

After creating a *Client* object the *open()* method should be called to establish a connection with the server. The connection is closed and the session is terminated by calling the *close()* method.

```
client = SalesforceStreamingClient(  
    consumer_key="<consumer key>",  
    consumer_secret="<consumer secret>",  
    username="<username>",  
    password="<password>"  
)  
await client.open()  
# subscribe and receive messages...  
await client.close()
```

Client objects can be also used as asynchronous context managers.

```
async with SalesforceStreamingClient(
    consumer_key="<consumer key>",
    consumer_secret="<consumer secret>",
    username="<username>",
    password="<password>") as client:
    # subscribe and receive messages...
```

Channels

A channel is a string that looks like a URL path such as `/topic/foo` or `/topic/bar`.

For detailed guidance on how to work with [PushTopics](#) or how to create [Generic Streaming Channels](#) please consult the [Streaming API](#) documentation.

Subscriptions

To receive notification messages the client must subscribe to the channels it's interested in.

```
await client.subscribe("/topic/foo")
```

If you no longer want to receive messages from one of the channels you're subscribed to then you must unsubscribe from the channel.

```
await client.unsubscribe("/topic/foo")
```

The current set of subscriptions can be obtained from the `Client.subscriptions` attribute.

Receiving messages

To receive messages broadcasted by Salesforce after *subscribing* to these *channels* the `receive()` method should be used.

```
message = await client.receive()
```

The `receive()` method will wait until a message is received or it will raise a `TransportTimeoutError` in case the connection is lost with the server and the client can't re-establish the connection or a `ServerError` if the connection gets closed by the server.

The client can also be used as an asynchronous iterator in a for loop to wait for incoming messages.

```
async for message in client:
    # process message
```

Replay of events

The great thing about streaming is that the client gets instantly notified about events as they occur. The downside is that if the client becomes temporarily disconnected, due to hardware, software or network failure, then it might miss some of the messages emitted by the server. This is where Salesforce's message durability comes in handy.

Salesforce stores events for 24 hours. Events outside the 24-hour retention period are discarded. Salesforce extends the event messages with `replayId` and `createdAt` fields (called as *ReplayMarker* by aiosfstream). These fields can be used by the client to request the missed event messages from the server when it reconnects.

The default behavior of the client is to receive only the new events sent after subscribing. To take advantage of message durability, all you have to do is to pass an object capable of storing the most recent *ReplayMarker* objects, so the next time the client reconnects, it can continue to process event messages from the point where it left off. The most convenient choice is a *Shelf* object, which can store *ReplayMarkers* on the disk, between application restarts.

```
with shelve.open("replay.db") as replay:

    async with SalesforceStreamingClient(
        consumer_key="<consumer key>",
        consumer_secret="<consumer secret>",
        username="<username>",
        password="<password>",
        replay=replay) as client:

        await client.subscribe("/topic/foo")

        async for message in client:
            # process message
```

Besides *Shelf* objects you can pass a lot of different kind of objects to the replay parameter, and you can configure different aspects of replay behavior as well. For a full description of replay configuration options check out the *Replay configuration* section.

3.1.3 Advanced Usage


Replay configuration

Salesforce stores events for 24 hours. Events outside the 24-hour retention period are discarded.

ReplayOption

A subscriber can choose which events to receive, such as all events within the retention window or starting after a particular event. In the *Client* object this can be specified with the *replay* parameter. The default is to receive only the new events sent after subscribing, the default *replay* parameter is *ReplayOption.NEW_EVENTS*

This high-level diagram shows how event consumers can read a stream of events by using various replay options.



If you want to receive all events within the retention window every time the *Client* connects, before receiving new events, then the *ReplayOption.ALL_EVENTS* value should be passed to the *Client*.

```
async with SalesforceStreamingClient(
    consumer_key="<consumer key>",
```

(continues on next page)

(continued from previous page)

```

        consumer_secret="<consumer secret>",
        username="<username>",
        password="<password>",
        replay=ReplayOption.ALL_EVENTS) as client:

    await client.subscribe("/topic/foo")

    async for message in client:
        # process message

```

ReplayMarkerStorage

Although using a fixed `ReplayOption` can be sometimes useful, the real advantage of using Salesforce's `replay` extension comes from being able to continue to process event messages from the point where the client left off. To take advantage of this feature, all you have to do is to pass an object capable of storing the most recent `ReplayMarker` for every channel.

Salesforce extends the event messages with `replayId` and `createdDate` fields (called as `ReplayMarker` by aiosfstream).

The simplest way is to pass an object for the `replay` parameter that inherits from `collections.abc.MutableMapping`. This can be a simple `dict`, `OrderedDict` or if you want to use persistent storage then a `Shelf` object, or maybe one of the key-value database drivers that inherit from `collections.abc.MutableMapping`.

```

with shelve.open("replay.db") as replay:

    async with SalesforceStreamingClient(
        consumer_key="<consumer key>",
        consumer_secret="<consumer secret>",
        username="<username>",
        password="<password>",
        replay=replay) as client:

        await client.subscribe("/topic/foo")

        async for message in client:
            # process message

```

By using a `collections.abc.MutableMapping` object, the client on the first connection will receive only new events, and on reconnection will continue from the last unretrieved message. If you want to receive all events from the retention window before continuing with new events, combined with the advantage of continuation on the next reconnect, then you can pass a `DefaultMappingStorage` object to the `replay` parameter.

```

with shelve.open("replay.db") as replay:

    default_mapping = DefaultMappingStorage(
        replay,
        ReplayOption.ALL_EVENTS
    )

    async with SalesforceStreamingClient(
        consumer_key="<consumer key>",
        consumer_secret="<consumer secret>",
        username="<username>",

```

(continues on next page)

(continued from previous page)

```
password="<password>",
replay=default_mapping) as client:

    await client.subscribe("/topic/foo")

    async for message in client:
        # process message
```

If you want complete control over how *ReplayMarkers* are stored and retrieved or you want to use your favorite database whose driver doesn't inherit from `collections.abc.MutableMapping` then you can provide your own *ReplayMarkerStorage* implementation.

```
class MyReplayMarkerStorage(ReplayMarkerStorage):
    async def set_replay_marker(self, subscription, replay_marker):
        # store *replay_marker* for the given *subscription*

    async def get_replay_marker(self, subscription):
        # retrieve the replay marker for the given *subscription*

replay = MyReplayMarkerStorage()

async with SalesforceStreamingClient(
    consumer_key="<consumer key>",
    consumer_secret="<consumer secret>",
    username="<username>",
    password="<password>",
    replay=replay) as client:

    await client.subscribe("/topic/foo")

    async for message in client:
        # process message
```

Subscription errors

Events outside the 24-hour retention period are discarded. If you're using some form of *ReplayMarkerStorage* or a *MutableMapping* object, and if your client doesn't connect to the Streaming API for more than 24 hours, then it's possible that the client will try to continue retrieving messages from a very old message outside the retention window. Since Salesforce no longer has the event message that the client would try to retrieve, it would raise *ServerError*.

```
try:
    await client.subscribe("/topic/foo")
except ServerError as error:
    print(error.error_message)
```

The above code would print the following message, if the client would request an event outside the retention window:

```
The replayId {1} you provided was invalid. Please provide a valid ID, -2
to replay all events, or -1 to replay only new events.
```

To recover from an error like the above, you would have to discard the *ReplayMarker* for the problematic channel, and try to subscribe again.


```

try:
    await client.subscribe("/topic/foo")
except ServerError as error:
    del replay["/topic/foo"]
    await client.subscribe("/topic/foo")

```

To spare you the hassle of recovering from errors like the one above, you can pass a *ReplayOption* for the `replay_fallback` parameter. If a subscription error occurs, then *Client* will try to resubscribe using the specified *ReplayOption*.

```

with shelve.open("replay.db") as replay:

    async with SalesforceStreamingClient(
        consumer_key="<consumer key>",
        consumer_secret="<consumer secret>",
        username="<username>",
        password="<password>",
        replay=replay,
        replay_fallback=ReplayOption.ALL_EVENTS) as client:

        await client.subscribe("/topic/foo")

        async for message in client:
            # process message

```

Network failures

When a *Client* object is opened, it will try to maintain a continuous connection in the background with the server. If any network failures happen while waiting to *receive()* messages, the client will reconnect to the server transparently, it will resubscribe to the subscribed channels, and continue to wait for incoming messages.

To avoid waiting for a server which went offline permanently, or in case of a permanent network failure, a `connection_timeout` can be passed to the *Client*, to limit how many seconds the client object should wait before raising a *TransportTimeoutError* if it can't reconnect to the server.

```

client = SalesforceStreamingClient(
    consumer_key="<consumer key>",
    consumer_secret="<consumer secret>",
    username="<username>",
    password="<password>",
    connection_timeout=60
)
await client.open()

try:
    message = await client.receive()
except TransportTimeoutError:
    print("Connection is lost with the server. "
          "Couldn't reconnect in 60 seconds.")

```

The default value is 10 seconds. If you pass `None` as the `connection_timeout` value, then the client will keep on trying indefinitely.

Prefetching

When a *Client* is opened it will start and maintain a connection in the background with the server. It will start to fetch messages from the server as soon as it's connected, even before *receive()* is called.

Prefetching messages has the advantage, that incoming messages will wait in a buffer for users to consume them when *receive()* is called, without any delay.

To avoid consuming all the available memory by the incoming messages, which are not consumed yet, the number of prefetched messages can be limited with the *max_pending_count* parameter of the *Client*. The default value is 100.

```
client = SalesforceStreamingClient(  
    consumer_key="<consumer key>",  
    consumer_secret="<consumer secret>",  
    username="<username>",  
    password="<password>",  
    max_pending_count=42  
)
```

The current number of messages waiting to be consumed can be obtained from the *Client.pending_count* attribute.

JSON encoder/decoder

Besides the standard *json* module, many third party libraries offer JSON serialization/deserialization functionality. To use a different library for handling JSON data types, you can specify the callable to use for serialization with the *json_dumps* and the callable for deserialization with the *json_loads* parameters of the *Client*.

```
import ujson  
  
client = SalesforceStreamingClient(  
    consumer_key="<consumer key>",  
    consumer_secret="<consumer secret>",  
    username="<username>",  
    password="<password>",  
    json_dumps=ujson.dumps,  
    json_loads=ujson.loads  
)
```

3.2 API Reference

3.2.1 Client

```
class aiosfstream.SalesforceStreamingClient(*,      consumer_key,      consumer_secret,  
                                             username,      password,      re-  
                                             play=<ReplayOption.NEW_EVENTS:  
-1>,      replay_fallback=None,  
                                             connection_timeout=10.0,  
                                             max_pending_count=100,  
                                             json_dumps=<function      dumps>,  
                                             json_loads=<function loads>, loop=None)  
    Salesforce Streaming API client with username/password authentication
```

This is a convenience class which is suitable for the most common use case. To use a different authentication method, use the general `Client` class with a different `Authenticator`

Parameters

- **consumer_key** (*str*) – Consumer key from the Salesforce connected app definition
- **consumer_secret** (*str*) – Consumer secret from the Salesforce connected app definition
- **username** (*str*) – Salesforce username
- **password** (*str*) – Salesforce password
- **replay** (*ReplayOption, ReplayMarkerStorage, collections.abc.MutableMapping or None*) – A `ReplayOption` or an object capable of storing replay ids if you want to take advantage of Salesforce’s replay extension. You can use one of the `ReplayOptions`, or an object that supports the `MutableMapping` protocol like `dict`, `defaultdict`, `Shelf` etc. or a custom `ReplayMarkerStorage` implementation.
- **replay_fallback** (*ReplayOption*) – Replay fallback policy, for when a subscribe operation fails because a replay id was specified for a message outside the retention window
- **connection_timeout** (*int, float or None*) – The maximum amount of time to wait for the transport to re-establish a connection with the server when the connection fails.
- **max_pending_count** (*int*) – The maximum number of messages to prefetch from the server. If the number of prefetched messages reach this size then the connection will be suspended, until messages are consumed. If it is less than or equal to zero, the count is infinite.
- **json_dumps** (*callable()*) – Function for JSON serialization, the default is `json.dumps()`
- **json_loads** (*callable()*) – Function for JSON deserialization, the default is `json.loads()`
- **loop** – Event loop used to schedule tasks. If `loop` is `None` then `asyncio.get_event_loop()` is used to get the default event loop.

```
class aiosfstream.Client (authenticator, *, replay=<ReplayOption.NEW_EVENTS:
                        -1>, replay_fallback=None, connection_timeout=10.0,
                        max_pending_count=100, json_dumps=<function dumps>,
                        json_loads=<function loads>, loop=None)
```

Salesforce Streaming API client

Parameters

- **authenticator** (*AuthenticatorBase*) – An authenticator object
- **replay** (*ReplayOption, ReplayMarkerStorage, collections.abc.MutableMapping or None*) – A `ReplayOption` or an object capable of storing replay ids if you want to take advantage of Salesforce’s replay extension. You can use one of the `ReplayOptions`, or an object that supports the `MutableMapping` protocol like `dict`, `defaultdict`, `Shelf` etc. or a custom `ReplayMarkerStorage` implementation.
- **replay_fallback** (*ReplayOption*) – Replay fallback policy, for when a subscribe operation fails because a replay id was specified for a message outside the retention window
- **connection_timeout** (*int, float or None*) – The maximum amount of time to wait for the transport to re-establish a connection with the server when the connection fails.
- **max_pending_count** (*int*) – The maximum number of messages to prefetch from the server. If the number of prefetched messages reach this size then the connection will be

suspended, until messages are consumed. If it is less than or equal to zero, the count is infinite.

- **json_dumps** (*callable()*) – Function for JSON serialization, the default is `json.dumps()`
- **json_loads** (*callable()*) – Function for JSON deserialization, the default is `json.loads()`
- **loop** – Event `loop` used to schedule tasks. If *loop* is `None` then `asyncio.get_event_loop()` is used to get the default event loop.

coroutine open()

Establish a connection with the Streaming API endpoint

Raises

- **ClientError** – If none of the connection types offered by the server are supported
- **ClientInvalidOperation** – If the client is already open, or in other words if it isn't *closed*
- **TransportError** – If a network or transport related error occurs
- **ServerError** – If the handshake or the first connect request gets rejected by the server.
- **AuthenticationError** – If the server rejects the authentication request or if a network failure occurs during the authentication

coroutine close()

Disconnect from the CometD server

coroutine publish(channel, data)

Publish *data* to the given *channel*

Warning: The Streaming API is implemented on top of CometD. The publish operation is a CometD operation. While it's still a legal operation, Salesforce chose not to implement the publishing of Generic Streaming and Platform events with CometD.

You should use the [REST API](#) to generate Generic Streaming events, or use the [REST](#) or [SOAP API](#) to publish Platform events.

Parameters

- **channel** (*str*) – Name of the channel
- **data** (*dict*) – Data to send to the server

Returns Publish response

Return type `dict`

Raises

- **ClientInvalidOperation** – If the client is *closed*
- **TransportError** – If a network or transport related error occurs
- **ServerError** – If the publish request gets rejected by the server

coroutine subscribe(channel)

Subscribe to *channel*

Parameters `channel` (*str*) – Name of the channel

Raises

- *ClientInvalidOperation* – If the client is *closed*
- *TransportError* – If a network or transport related error occurs
- *ServerError* – If the subscribe request gets rejected by the server

coroutine `unsubscribe` (*channel*)

Unsubscribe from *channel*

Parameters `channel` (*str*) – Name of the channel

Raises

- *ClientInvalidOperation* – If the client is *closed*
- *TransportError* – If a network or transport related error occurs
- *ServerError* – If the unsubscribe request gets rejected by the server

coroutine `receive` ()

Wait for incoming messages from the server

Returns Incoming message

Return type *dict*

Raises

- *ClientInvalidOperation* – If the client is closed, and has no more pending incoming messages
- *ServerError* – If the client receives a confirmation message which is not successful
- *TransportTimeoutError* – If the transport can't re-establish connection with the server in `connection_timeout` time.

closed

Marks whether the client is open or closed

subscriptions

Set of subscribed channels

connection_type

The current connection type in use if the client is open, otherwise *None*

pending_count

The number of pending incoming messages

Once *open* is called the client starts listening for messages from the server. The incoming messages are retrieved and stored in an internal queue until they get consumed by calling *receive*.

has_pending_messages

Marks whether the client has any pending incoming messages

3.2.2 Authenticators

```
class aiosfstream.auth.AuthenticatorBase (json_dumps=<function dumps>,
                                           json_loads=<function loads>)
```

Abstract base class to serve as a base for implementing concrete authenticators

Parameters

- **json_dumps** (*callable()*) – Function for JSON serialization, the default is `json.dumps()`
- **json_loads** (*callable()*) – Function for JSON deserialization, the default is `json.loads()`

```
class aiosfstream.PasswordAuthenticator(consumer_key, consumer_secret, username,  
                                         password, json_dumps=<function dumps>, json_loads=<function loads>)
```

Authenticator for using the OAuth 2.0 Username-Password Flow

Parameters

- **consumer_key** (*str*) – Consumer key from the Salesforce connected app definition
- **consumer_secret** (*str*) – Consumer secret from the Salesforce connected app definition
- **username** (*str*) – Salesforce username
- **password** (*str*) – Salesforce password
- **json_dumps** (*callable()*) – Function for JSON serialization, the default is `json.dumps()`
- **json_loads** (*callable()*) – Function for JSON deserialization, the default is `json.loads()`

client_id = None

OAuth2 client id

client_secret = None

OAuth2 client secret

password = None

Salesforce password

username = None

Salesforce username

```
class aiosfstream.RefreshTokenAuthenticator(consumer_key, consumer_secret, re-  
                                             fresh_token, json_dumps=<function  
                                             dumps>, json_loads=<function loads>)
```

Authenticator for using the OAuth 2.0 Refresh Token Flow

Parameters

- **consumer_key** (*str*) – Consumer key from the Salesforce connected app definition
- **consumer_secret** (*str*) – Consumer secret from the Salesforce connected app definition
- **refresh_token** (*str*) – A refresh token obtained from Salesforce by using one of its authentication methods (for example with the OAuth 2.0 Web Server Authentication Flow)
- **json_dumps** (*callable()*) – Function for JSON serialization, the default is `json.dumps()`
- **json_loads** (*callable()*) – Function for JSON deserialization, the default is `json.loads()`

client_id = None

OAuth2 client id

client_secret = None

OAuth2 client secret

refresh_token = None

Salesforce refresh token

3.2.3 Replay

class aiosfstream.ReplayOption

Replay options supported by Salesforce

ALL_EVENTS = -2

Receive all events, including past events that are within the 24-hour retention window and new events sent after subscription

NEW_EVENTS = -1

Receive new events that are broadcast after the client subscribes

class aiosfstream.ReplayMarker(*date*, *replay_id*)

Bases: `tuple`

Class for storing a message replay id and its creation date

Parameters

- **date** (*str*) – Creation date of a message, as a ISO 8601 formatted datetime string
- **replay_id** (*int*) – Replay id of a message

date

Alias for field number 0

replay_id

Alias for field number 1

class aiosfstream.ReplayMarkerStorage

Abstract base class for replay marker storage implementations

coroutine **get_replay_marker**(*subscription*)

Retrieve a stored replay marker for the given *subscription*

Parameters **subscription** (*str*) – Name of the subscribed channel

Returns A replay marker or None if there is nothing stored for the given *subscription*

Return type *ReplayMarker* or None

coroutine **set_replay_marker**(*subscription*, *replay_marker*)

Store the *replay_marker* for the given *subscription*

Parameters

- **subscription** (*str*) – Name of the subscribed channel
- **replay_marker** (*ReplayMarker*) – A replay marker

class aiosfstream.MappingStorage(*mapping*)

Mapping based replay marker storage

Parameters **mapping** (*collections.abc.MutableMapping*) – A MutableMapping object for storing replay markers

class aiosfstream.**DefaultMappingStorage** (*mapping, default_id*)

Mapping based replay marker storage which will return a default replay id if there is not replay marker for the given subscription

Parameters

- **mapping** (*collections.abc.MutableMapping*) – A MutableMapping object for storing replay markers
- **default_id** (*int*) – A replay id

class aiosfstream.**ConstantReplayId** (*default_id, **kwargs*)

A replay marker storage which will return a constant replay id for every subscription

Note: This implementations doesn't actually stores anything for later retrieval.

Parameters **default_id** (*int*) – A replay id

3.2.4 Exceptions

Exception types

Exception hierarchy:

```
AiosfstreamException
  AuthenticationError
  ClientError
    ClientInvalidOperation
  TransportError
    TransportInvalidOperation
    TransportTimeoutError
    TransportConnectionClosed
  ServerError
```

exception aiosfstream.exceptions.**AiosfstreamException**

Base exception type.

All exceptions of the package inherit from this class.

exception aiosfstream.exceptions.**AuthenticationError**

Authentication failure

exception aiosfstream.exceptions.**ClientError**

Client side error

exception aiosfstream.exceptions.**ClientInvalidOperation**

The requested operation can't be executed on the current state of the client

exception aiosfstream.exceptions.**TransportError**

Error during the transportation of messages

exception aiosfstream.exceptions.**TransportInvalidOperation**

The requested operation can't be executed on the current state of the transport

exception aiosfstream.exceptions.**TransportTimeoutError**

Transport timeout

exception aiosfstream.exceptions.**TransportConnectionClosed**

The connection unexpectedly closed

exception `aiosfstream.exceptions.ServerError`

Streaming API server side error

If the *response* contains an error field it gets parsed according to the *specs*

Parameters

- **message** (*str*) – Error description
- **response** (*dict*) – Server response message

message

Error description

response

Server response message

error

Error field in the *response*

error_code

Error code part of the error code part of the *error*, message field

error_args

Arguments part of the *error*, message field

error_message

Description part of the *error*, message field

3.3 Changelog

3.3.1 0.2.1 (2018-05-25)

- Fix replay issues on mass record delete operations
- Improve the documentation of the `Client.publish` method

3.3.2 0.2.0 (2018-05-05)

- Enable the usage of third party JSON libraries
- Expose authentication results as public attributes in Authenticator classes

3.3.3 0.1.0 (2018-04-26)

- **Supported authentication types:**
 - using a username and password
 - using a refresh token
- **Subscribe to and receive messages on:**
 - `PushTopics`
 - `Generic Streaming Channels`
- Support for *durable messages and replay of events*
- Automatic recovery from replay errors

CHAPTER 4

Indices and tables

- `genindex`
- `modindex`
- `search`

a

`aiosfstream.exceptions`, [20](#)

A

`aiosfstream.exceptions` (module), 20
`AiosfstreamException`, 20
`ALL_EVENTS` (`aiosfstream.ReplayOption` attribute), 19
`AuthenticationError`, 20
`AuthenticatorBase` (class in `aiosfstream.auth`), 17

C

`Client` (class in `aiosfstream`), 15
`client_id` (`aiosfstream.PasswordAuthenticator` attribute), 18
`client_id` (`aiosfstream.RefreshTokenAuthenticator` attribute), 18
`client_secret` (`aiosfstream.PasswordAuthenticator` attribute), 18
`client_secret` (`aiosfstream.RefreshTokenAuthenticator` attribute), 18
`ClientError`, 20
`ClientInvalidOperation`, 20
`close()` (`aiosfstream.Client` method), 16
`closed` (`aiosfstream.Client` attribute), 17
`connection_type` (`aiosfstream.Client` attribute), 17
`ConstantReplayId` (class in `aiosfstream`), 20

D

`date` (`aiosfstream.ReplayMarker` attribute), 19
`DefaultMappingStorage` (class in `aiosfstream`), 19

E

`error` (`aiosfstream.exceptions.ServerError` attribute), 21
`error_args` (`aiosfstream.exceptions.ServerError` attribute), 21
`error_code` (`aiosfstream.exceptions.ServerError` attribute), 21
`error_message` (`aiosfstream.exceptions.ServerError` attribute), 21

G

`get_replay_marker()` (`aiosfstream.ReplayMarkerStorage` method), 19

H

`has_pending_messages` (`aiosfstream.Client` attribute), 17

M

`MappingStorage` (class in `aiosfstream`), 19
`message` (`aiosfstream.exceptions.ServerError` attribute), 21

N

`NEW_EVENTS` (`aiosfstream.ReplayOption` attribute), 19

O

`open()` (`aiosfstream.Client` method), 16

P

`password` (`aiosfstream.PasswordAuthenticator` attribute), 18
`PasswordAuthenticator` (class in `aiosfstream`), 18
`pending_count` (`aiosfstream.Client` attribute), 17
`publish()` (`aiosfstream.Client` method), 16

R

`receive()` (`aiosfstream.Client` method), 17
`refresh_token` (`aiosfstream.RefreshTokenAuthenticator` attribute), 19
`RefreshTokenAuthenticator` (class in `aiosfstream`), 18
`replay_id` (`aiosfstream.ReplayMarker` attribute), 19
`ReplayMarker` (class in `aiosfstream`), 19
`ReplayMarkerStorage` (class in `aiosfstream`), 19
`ReplayOption` (class in `aiosfstream`), 19
`response` (`aiosfstream.exceptions.ServerError` attribute), 21

S

`SalesforceStreamingClient` (class in `aiosfstream`), 14
`ServerError`, 20
`set_replay_marker()` (`aiosfstream.ReplayMarkerStorage` method), 19
`subscribe()` (`aiosfstream.Client` method), 16

[subscriptions \(aiosfstream.Client attribute\)](#), [17](#)

T

[TransportConnectionClosed](#), [20](#)

[TransportError](#), [20](#)

[TransportInvalidOperation](#), [20](#)

[TransportTimeoutError](#), [20](#)

U

[unsubscribe\(\)](#) (aiosfstream.Client method), [17](#)

[username](#) (aiosfstream.PasswordAuthenticator attribute),
[18](#)