EURONEXT

Document title

EURONEXT CASH MARKETS – OPTIQ® OEG THROTTLING MECHANISM

Document type or subject

Optiq OEG Throttling Mechanism - Functionality Description

Version number	Date
1.0.0	8 Feb 2018

Number of pages	Author
16	EURONEXT

PREFACE

PURPOSE

This document provides details of the throttling mechanism for the Optiq® Order Entry Gateways (OEG).

This document is a supporting document for other technical specifications made available for Optiq, some of which are listed in the section Associated documents.

ASSOCIATED DOCUMENTS

The following list identifies the associated documents, which either should be read in conjunction with this document, or which provide other relevant information to the clients:

- Euronext Cash Markets – Optiq OEG Client Specifications – SBE Interface
- Euronext Cash Markets – Optiq OEG Client Specifications – FIX 5.0 Interface
- Euronext Cash Markets – Optiq Kinematics Specifications
- Euronext Cash Markets – Optiq & TCS Error List
- Euronext Cash Markets – Optiq OEG Connectivity Configuration Specifications
- Euronext Cash and Derivatives Markets – Optiq File Specification
- Euronext Cash Markets - OEG Cancel on Disconnect (CoD) Functional Overview

Clients are advised to also refer to the Euronext Rules and Regulations documents for more details.

For the latest version of documentation please visit http://www.euronext.com/optiq

SUPPORT

Optiq Support Desk
Tel: +33 1 70 48 25 55
Email: optiq@euronext.com

DOCUMENT & REVISION HISTORY

Version	Change Description
1.0.0	First Release for Optiq Phase 2 (Cash markets)

CONTENTS

1. INTRODUCTION

The Optiq Order Entry Gateway (OEG) provides high-speed and real-time connection to the Exchange markets. Objectives of OEG throttling are similar to those found in any system that has to manage high amount of message exchange with multiple participants, to

- Regulate message / data traffic by evening out the concentration flow of messages, and distributing the use of available system processing ability and bandwidth across all users of the system
- Help minimize or limit message exchange and processing congestion, which assists in ensuring latency of the trading system remains stable and predicable
- Reduce the risk of disruptive events

Furthermore, for a trading venue OEG throttling mechanism allows to:

- Prevent disorderly trading conditions and detect potential threats to the orderly functioning of the market
- Ensure compliance with the articles in MIFID II supplementing Directive 2014/65/EU of the European Parliament and of the Council with regard to regulatory technical standards specifying organizational requirements of trading venues *(a.k.a. RTS7)*

This document provides information about the Throttling mechanisms for the OEG implemented in Optiq and details of:

- throttling mechanism supported by the OEG
- concepts, limits and formulas that support client's ability to manage their message flow in order to:
  - avoid being throttled,
  - be informed of the reason their messages were throttled
- recommended practices in using the OEG in most optimal way

The scope of this version of the document covers OEG throttling mechanism for the Cash markets hosted on Optiq.

1.1 GLOSSARY

This section provides a list of some terms & abbreviations commonly used in this document. Please note that some of these terms are described in more details in the dedicated sections within this document or in the associated Optiq specifications documents.

- **Order Entry Gateway (OEG):** is the software that manages the access for exchanges' clients, and acts as the private interface between the clients and the Optiq matching engine.
- **Market Data Gateway (MDG):** is the software that provides high-speed, real-time market data (public messages) for the Exchange's markets.
- **Matching Engine (ME):** is the software that manages the trading services for the Exchange's markets.
- **Optiq Segment:** defines a universe of instruments habitually sharing common trading properties. An Optiq Segment can contain one or several asset classes. An Optiq Segment access is setup through a Logical Access.
- **Partition:** is a technical subdivision of an Optiq Segment. An Optiq Segment may be comprised of at least one or several partitions, physically independent one from another, but connected to each other within the context of the Optiq Segment. Instruments may move from one partition to another within an Optiq segment.

- **Logical Access:** is an OEG (Order Entry Gateway) entry point, setup for clients to connect to a single Optiq Segment, containing the technical configuration for the client's connectivity. Multiple logical accesses can share the same SFTI line.
- **OE Session:** the individual physical connection, to a single Partition. A single Logical access may have as many OE sessions as there are partitions in the Optiq segment.

1.2   WORK IN PROGRESS

Some of the topics of OEG throttling mechanism are not yet finalized and will be provided or further clarified, in a later version of the document:

- Examples and Kinematics

                                          1.0.0

2.   OEG THROTTLING MECHANISM IN OPTIQ

Sections below provides high level description of the OEG throttling concepts and features in Optiq.

⚠️ **Important note**: Clients are strongly urged to review the explanations provided in the various OEG and MDG specification, Connectivity specifications and kinematic documents in detail before continuing with this document

2.1   HOW OEG THROTTLING WORKS

Throttling mechanism is in place to:

■ Regulate message / data traffic by evening out the concentration flow of messages, and distributing the use of available system processing ability and bandwidth across all users of the system
■ Help minimize or limit message exchange and processing congestion,
■ Ensure latency of the trading system remains stable and predicable
■ Reduce the risk of disruptive events (incidents)
■ Prevent disorderly trading conditions and detect potential threats to the orderly functioning of the market
■ Ensure compliance with the articles in MIFID II requirements

In Optiq OEG throttling is accomplished by either rejecting or queueing messages that go over the limit set for the Logical Access.

Queue vs. Reject

With Optiq clients have the ability to choose whether their "over the limit" messages are Queued or Rejected. This selection is done in a **Logon** (100) / (FIX A) message, and is not a "pre-configured" setting. If clients disconnects and logs back in, they can change the selection intra-session.

Throttling behavior, when triggered, reflects the choice made in the **Logon** (100) / (FIX A) message, i.e. over the limit messages are either:

■ rejected, or
■ queued up to a pre-defined limit

Flexibility & Predictability

Optiq throttling solution is built to provide flexibility and predictability.

■ Flexibility: clients can choose how many messages they can send (vs. how often) to avoid being throttled
■ Predictability: in most cases clients can use the concepts and formulas provided to calculate the number of messages and time required, without having to wait for Optiq to send them specific messages

Using the solution clients could send a steady stream of few messages on a frequent basis or send larger bursts of messages but less frequently as required, based on the client's needs and market conditions

Limits for OEG Throttling

The limits and variables used for throttling are based on the per second message rate set for the individual Logical access.

To provide a steady and predictable way to calculate the OEG throttling limits they are based on the client's per second message rate ("throughput"). The approach to setting the throttling limits are the same for all Logical accesses, irrespective of the subscribed to rate.

While the solution uses a more granular approach, at worst its target is allow clients to be able to send $1/10^{th}$ of their per second rate, within every $1/10^{th}$ of a second.

Below this limit clients can choose to adjust continuously how many messages are sent at once, and to calculate how long they have before they can send more messages or a burst of messages.

Clients that do not wish to build the full solution for the formulas may choose to use a simplified approach for calculation, where a physical connection can send at max 1/10th of their rate in 1/10th of a second.

If the throttling limits or their calculations concepts are adjusted a communication will be issued to clients prior to such change being made.

Excessive Breaches of Per Second Rate (or Anti-Flooding)

In addition to the limits set out for throttling, OEG also enforces "anti-flooding" limits and controls, that are triggered in case excessive amount of messages or data is sent by the client. As with throttling, the excessive breaches of the connection rate use limits that are based on the per second message rate set for the individual Logical access.

For more details about the cases of disconnection initiated by exchange, including those for excessive breaches of connection rate, please review the *Euronext Cash Markets – Optiq OEG Connectivity Configuration Specifications* document.

2.2 OEG THROTTLING INTERACTION WITH OTHER MECHANISMS & COMPONENTS OF OPTIQ

High Availability: OEG throttling mechanism works in the same manner on the Mirror instance of the OEG, in any HA scenario, as it does on the Primary instance, with identical settings and limits.

In case of disruptive incident that triggers an HA event, any messages that were in a throttling queue are dropped, without further notification to the client.

Business Continuity: OEG throttling mechanism works in the same manner in the DR environment as it does in Production, with identical settings and limits.

In case of a Business Continuity event, any messages that were in a throttling queue in Production environment are dropped, without further notification to the client.

MDG: OEG throttling and associated messages are not reflected in any public messages.

Drop Copy: OEG throttling does not apply to Drop Copy gateways.

Sections below provides more details on concepts and limits that the mechanism uses, and could be useful to clients in setting up their configuration and systems in order to avoid being throttled.

　　　　　　　　1.0.0

2.3   DETAILS OF OEG THROTTLING

Sections below provide the details associated to the OEG Throttling mechanism, associated concepts and limits.

2.3.1   Queue vs. Reject

Client choose on every logon whether upon breaching their throttling limit their messages should be rejected or queued.

In the **Logon** (100) / (FIX A) message, the field used is

■ SBE: *Queueing Indicator*
■ FIX: *QueueingIndicator* (tag: 21020)

Where values used are:

■ **0** – False, which indicates clients wish to Reject over the limit messages
■ **1** – True, , which indicates clients wish to Queue over the limit messages

If client chooses to Reject – any limit breached will result in a rejection message.

If client chooses to Queue – over the limit messages are stored in a limited size Throttling queue.

2.3.2   Communication of Throttling Events

In case of any rejection due to OEG throttling Optiq sends the following rejection messages & codes:

In SBE, clients receive **TechnicalReject** (108) message, where the field *Error Code* is set with the value indicating the reason throttling event has occurred.

In FIX, clients receive **Reject** (3) messages, where the field *SessionRejectReason* (tag: 373) is set with the value indicating the reason throttling event has occurred.

Three types of rejection are possible in OEG throttling. Table below provides the correspondence of values used for these three types of rejection in SBE and FIX:

SBE Error Code	FIX SessionRejectReason (373)	Description of the OEG Throttling Rejection
2085 = Rate exceeded	26 = Throttling Rate exceeded	Individual message sent is over the limit allocated to the Logical Access. This rejection occur when client chose to Reject over the limit messages
2087 = Throttling queue full	25 = Throttling queue full	This rejection occur when client chose to Queue over the limit messages, and after the initial throttling queue limit based on the rate is breached, and the throttled messages are stored in the throttling queue, the limit of the queue has also been breached.
2086 = System busy	27 = System busy	This rejection may occur when client chose to Reject over the limit messages, and the system is overloaded by processing of previously sent messages and can't accept more messages until the processing has finished.

2.3.3 Behavior if Client Chose Rejection

As identified above, a rejection message will be sent to identify each message over the limit.

As throttling limit are rate and time based, messages reject over the rate limit would occur intra-session, and may occur with immediately preceding and following messages being accepted.

To provide indication on which message is rejected due to OEG throttling, the following fields are provided in the rejection messages:

■ <u>SBE</u>**: TechnicalReject** (108) message contains the following fields for this purpose:

Field	Short Description
Rejected Client Message Sequence Number	Indicates the Client Message Sequence Number of the rejected message
Rejected Message	Provides the Type of message rejected, by indicating its Template ID
Error Code	Indicates the type of OEG throttling rejection that occurred

■ <u>FIX</u>: **Reject (3)** message contains the following fields for this purpose:

Tag	Field	Short Description
45	RefSeqNum	Indicates the reference sequence number of the rejected message
372	RefMsgType	Provides the Type of message rejected, by indicating its MsgType (35)
373	SessionRejectReason	Indicates the type of OEG throttling rejection that occurred

The rejection messages do not contain any of the following identified fields: Client Order ID, Order ID, Quote Request ID nor Mass Status Request ID

2.3.4 Behavior if Client Chose Queueing

After client's OEG throttling limit is breached, and client chose to queue their messages, over the limit messages are stored (or queued) in a buffer called Throttling queue.

From the throttling queue messages are

■ processed one at a time at a pre-defined period of time (replenish time)
■ messages being processed from the queue / buffer are considered as activity, and during 1 second do not strictly require client sending a heartbeat

While messages in the throttling queue are processed:

■ until the queue limits has been reached, client may submit additional messages, which will also be stored in the throttling queue
■ as throttling queue is limited in size to the maximum of the connection rate,
    – any messages over the queue size are rejected, even if client chooses to Queue their messages
    – after soon as a single message is processed from the throttling queue, and a replenish time has passed, client's next submitted messages will be added to the throttling queue
    – client's do not receive a specific notification message for messages that were queued, however the acknowledgement messages have a field that indicates that the message in question was queued due to throttling
        ♦ SBE: **Ack** (03) message contains a field *Ack Qualifiers*, one of the positions of which is used for the Queue Indicator. For messages that were queued due to OEG throttling this position is set to one (1)

◆ FIX: **ExecutionReport** (8) message contains a field *AckQualifiers* (tag: 21014), one of the positions of which is used for the Queue Indicator. For the acknowledgement of new order messages that were queued due to OEG throttling this position is set to one (1)

■ Even if clients choose the Queue their over the limit messages, if their connections breach their rate either with a large number of messages, or data within the allocated time such are not processed

– Clients are strongly urged to send at maximum their rate per second to guarantee no rejections

2.3.5 Behavior for Excessive Breaches of Rate

Clients connections are assigned a maximum message rate, which are enforced by the OEG, in part, by the Exchange's throttling mechanism for inbound messages.

Upon breaching the throttling limit messages above the limit are either rejected or queued. In addition Exchange sets a limit for excessive breaching (either in number of messages or amount of data sent), of the assigned rate / size of connection.

Excessive breaching means that client attempted to submit a number of messages, or amount of data in bytes, that is multiple times over their allowed rate.

In case client's connection continuously breaches their rate limit, either in number of messages, or in the amount of data, such connections will be disconnected, and won't be allowed to reconnect for 15 seconds. If client's connections are continuously breaching limits in this manner, and are continuously disconnected Market Operations will contact the client and may choose to suspend client's access.

This case could occur due to various reasons, including a technical issue in the client's system. To assist clients in identifying the issue and correcting it as quickly as possible, the **Logout** (103) / (FIX 5) message sent on disconnection in such a case provides specific values identified below.

Case	How to Identify the Case in Logout message	
	SBE [Log Out Reason Code]	FIX [SessionStatus (tag: 1409)]
Excessive number of message	3 = Excessive Number of Messages	106 = Excessive Number of Messages
Excessive amount of data in bytes	4 = Excessive Amount of Data in Bytes	107 = Excessive Amount of Data in Bytes
Excessive number of messages and amount of data in bytes	5 = Excessive Number of Messages & Amount of Data in Bytes	108 = Excessive Number of Messages & Amount of Data in Bytes

2.3.6 Behavior on Intra-session Disconnection:

The OEG throttling behavior in case of intra-session disconnection is the same, independently of whether the disconnection occurred in the Exchange or Client systems, if it occurred on the same instance of the OEG continuously available, or if a disruptive incident triggered a HA or a Business Continuity event.

■ Any messages at the moment of disconnection present in the throttling queue are dropped, as if never received by the Exchange. Such messages do not receive acknowledgement or rejection from the OEG

■ On Reconnection, during usual sequence number processes and resynchronization mechanism, clients could receive throttling rejection messages, that serve as an indication of messages that were throttled, and as such were never processed prior to the disconnection

2.3.7   Main Concepts & Limits

Section below provides details on the concepts and limits used for OEG throttling. Please note that some of the concepts listed below have further explanation in dedicated sections within this document.

- Overall throughput limit / Per second Rate: the max number of messages setup as the rate for the Logical Access, and used by its physical connections
  Example: 100 messages per second
  - Client connections can't go over this limit defined as their rate within a second. This is accomplished mechanically, as the more granular size / time management of throttling limit will keep the overall throughput within the rate limit
- Max burst (or bucket) size: the max number of messages that a client may send at once, in a single "burst" of messages (measured in number of messages).
  - represents a constant figure, calculated as a fraction of the overall throughput limit
- Time to replenish [a single message]: time needed for a single token to be added into the bucket
  - This time equals to: 1 divided by the client's rate (1/rate)

  Example: if client's rate is 100 messages per second, then Time to replenish will be 1/100 = 0.01 second

- Token: Represents systems ability to process 1 message. 1 token gives the right to send 1 message
- Bucket: is a measuring mechanism used to calculate how many messages can still be sent and when, in order not to be throttled
  - Bucket size is the maximum number of messages that can be sent at once (or a burst), counted in tokens
  - Bucket doesn't store any messages. All messages are stored in the queues
  - Size of the bucket: As identified above, bucket is currently 1/10th of the client's per second rate, and is measured in tokens
- Scope of Messages for Throttling:
  - The main level of OEG throttling mechanism doesn't take into consideration administration messages (only application messages) and messages that are technically invalid
  - All application messages are in scope of throttling (no exclusion list)
  - "Excessive breaches" level of OEG throttling takes into consideration all messages received from the client (administrative, application, as well as technical or functionally valid or invalid)

2.3.7.1   Summary of Formulas

Concept / Limit	Formula
Throttling queue – Ts (messages)	For Queuing - Rate (size) of connection For Rejection - Zero
Size of the Bucket – Sb (messages) This is the allowed burst size, i.e. max number of messages a session can send in a row before the throttling kicks in.	1/10th of the rate or Minimum Limit*
Replenish Time (to replenish single token [Tr] (seconds))	1/rate size
Maximum # of "burst" messages to send at once (messages)	Equal to Sb
Wait time to send 1 (more) message (seconds)	1/rate
Wait time to send "X" messages	(1/rate) * X (messages)
Max messages before rejection (messages)	Ts + Sb **

Notes on Calculation

Calculated figures on number of messages and size of the bucket are rounded up to the closest whole number

1.0.0

* In case the rate setup for the Logical Access results in a calculation of the bucket size that is below 10, the bucket size will be equal to the rate of the Logical access.

** Clients are strongly urged to send at maximum their rate per second to guarantee no rejections

2.3.8   Private Messages Used by the OEG Throttling Mechanism

Optiq supports the following sub-set of private OEG messages that participate in the:

- process of throttling,
- settings of throttling configuration, or
- communicating the reasons for throttling event or disconnection.

Message Name	SBE Message Code	FIX Message Code	Details
Logon	100	A	Used to select queue vs. reject behavior for throttling
Technical Reject	108	3	Used to indicate rejection of messages due to throttling, and the reason throttling event occurred This message is not sent to Drop Copy
Logout	103	5	Used to indicate reason for disconnection in case of anti-flooding

For more details about these messages clients are advised to review the OEG client specifications for SBE and FIX protocols.

2.3.9   Use of the Bucket Concept:

- At the start-up bucket is full of tokens, to the maximum size of the bucket
- Each message sent to the OEG uses up a token
- Maximum size of the bucket is always constant. Size of bucket never goes above the maximum size, even if client doesn't send any messages for a long period of time
    - Any tokens above the size of the bucket disappear (fall out of the bucket and are not counted)
- Size of the bursts can vary depending on client's choice, but can't exceed the size of the bucket, and the time required to replenish the tokens in the bucket
    - *Example: for a connection size of 100 messages per second / bucket size of 10, time to replenish one token is 0.01 of a second. Client can send (i) 10 bursts of up to 10 messages within one second, every 0.1 of a second OR (2) 1 messages every 0.01 second OR (3) 2 messages every 0.02 seconds OR (4) bursts of 5 messages every 0.05 seconds OR (5) a combination of any of the bursts of messages that do not exceed the size of the bucket within the time allocated to replenish the tokens*
- If the bucket is empty, client's messages are either rejected or stored in a throttling queue
- 1 token is put back into the bucket (replenished) after a single period of time equal to "Time to Replenish" has passed
- The number of message in the bucket is replenished continuously, up to the maximum size of the bucket, independently of whether client submitted new messages or has been inactive for a long period of time
- When the time to replenish has passed, system automatically allocates one more message
    - If the Max limit of the bucket has been reached no other messages are added to the bucket
    - If the Max limit of the bucket is not yet reached, one more message is added to the bucket

2.4 ASSOCIATED GUIDELINES FOR CLIENTS

2.4.1 Measures to be Adopted Following an OEG Throttling Event

Various cases of OEG throttling may occur, and are identified by the fields provided in rejection and logout messages. Section below provides the guidelines on measures clients should adopt following an OEG throttling event .

2.4.1.1 Over the OEG Throttling Limit (Rejection)

For clients that choose to reject messages over the throttling limit. If a client receives a rejection indicating that their messages were throttled because they have exceeded their rate (flagged as follows):

SBE			FIX		
Message	**Feld**	**Value**	**Message**	**Feld**	**Value**
Technical Reject (108)	Error Code	2085 = Rate exceeded	Reject (3)	SessionRejectReason (373)	26 = Throttling Rate exceeded

the message rejected is not processed by the Exchange and, clients should:

■ take this rejection into consideration in their system, and if required resend an instruction to the exchange,
■ wait for at minimum 1 replenish time cycle before sending any further messages,
■ assess the speed and/or number of messages being sent by their system and either
  – reduce the frequency of sending to be in line with their replenish time, or
  – reduce the number of messages sent to be in line with the rate and associated throttling limits set for their logical access

2.4.1.2 System Busy (Rejection)

If a client receives a rejection indicating that their messages were throttled because System is Busy (flagged as follows):

SBE			FIX		
Message	**Feld**	**Value**	**Message**	**Feld**	**Value**
Technical Reject (108)	Error Code	2086 = System busy	Reject (3)	SessionRejectReason (373)	27 = System busy

This rejection may occur when client chose to reject messages over the throttling limit, but the system is overloaded by processing of previously sent messages and can't accept more messages until the processing has finished.

The message rejected is not processed by the Exchange and clients should:

■ take this rejection into consideration in their system, and if required resend an instruction to the exchange
■ wait for at minimum 1 second before sending any further messages,
■ review the Market Status page for information on a possible disruptive incident
■ may wish to assess the speed and/or number of messages being sent by their system and either
  – reduce the frequency of sending to be in line with their replenish time, or

 1.0.0

– reduce the number of messages sent to be in line with the rate and associated throttling limits set for their logical access

2.4.1.3   Over the OEG Throttling Limit (Queueing)

For clients that choose to queue throttled messages, over the OEG throttling limit, queued messages are indicated by a flag in the Ack message, as identified above. Such messages are processed by the Exchange with a delay associated to queuing.

As the maximum size of the OEG throttling queue is equal to the rate of client's connection for duration of 1 second, the processing of the queued messages, in normal conditions, is expected to last between 0 and 1 seconds, depending on the number of messages being queued.

To avoid being queued clients should:

■ assess the speed and/or number of messages being sent by their system and either
   – reduce the frequency of sending to be in line with their replenish time, or
   – reduce the number of messages sent to be in line with the rate and associated throttling limits set for their logical access

2.4.1.4   Over the Throttling Queue Size (Queueing)

For clients that choose to queue messages over the throttling limit. If a client receives a rejection indicating that their messages were throttled because their throttling queue is full (flagged as follows):

SBE			FIX		
Message	**Field**	**Value**	**Message**	**Field**	**Value**
Technical Reject (108)	Error Code	2087 = Throttling queue full	Reject (3)	SessionRejectReason (373)	25 = Throttling queue full

the message rejected is not processed by the Exchange and, clients should:

■ take this rejection into consideration in their system, and if required resend an instruction to the exchange
■ wait for at minimum 1 replenish time period before sending any further messages,
■ assess the speed and/or number of messages being sent by their system and either
   – reduce the frequency of sending to be in line with their replenish time, or
   – reduce the number of messages sent to be in line with the rate and associated throttling limits set for their logical access

2.4.1.5   Excessive Breaches of Rate

In case client is disconnected with one of the excessive breaches of message or data reasons, the reason for such disconnection is identified in the **Logout** (103) / (FIX 5) message in the fields identified in the table below for each protocol.

Case	How to Identify the Case in Logout message	
	SBE **[Log Out Reason Code]**	**FIX** **[SessionStatus (tag 1409)]**
Excessive number of message	3 = Excessive Number of Messages	106 = Excessive Number of Messages
Excessive amount of data in bytes	4 = Excessive Amount of Data in Bytes	107 = Excessive Amount of Data in Bytes
Excessive number of messages and amount of data in bytes	5 = Excessive Number of Messages & Amount of Data in Bytes	108 = Excessive Number of Messages & Amount of Data in Bytes

In case of such a disconnection clients are urged to review the rate at which they are sending messages to the OEG vs the allowed rate for the logical access, or assess if there is a technical issue in the client's system. For further assistance clients should contact one of the Exchange's support teams.

For more details about the cases of disconnection initiated by exchange, including those for excessive breaches of connection rate, please review the *Euronext Cash Markets – Optiq OEG Connectivity Configuration Specifications* document.

2.4.2   How to Avoid being Throttled

Granularity of sending one message at a time, would allow for all messages of a rate to be sent at least once every period which is equal to "time to replenish" (because it is 1/rate).

Granularity of sending more than one message at once depends on the max size of the "bucket" and number of messages (tokens) still remaining in the bucket

To avoid being throttled clients should:

- send at maximum the amount of messages they calculate as having in their bucket at the current moment at time, whether this is done with few messages or in bursts,
- send at once at maximum the number of messages equal to their bucket size, and wait for the period of time required to fully replenish it, before sending their full bucket size again, OR
- split the sending of bursts into fewer messages, which allows for shorter amount of time before the next sending is possible
- in case client chooses queueing, send at maximum their rate per second to guarantee no rejections

If sending fewer messages than the bucket size, and until the bucket is empty, clients do not to need to wait for a specific period of time between messages being sent.

APPENDIX A: REVISION HISTORY

SUMMARY OF CHANGES

Version	Change Description
1.0.0	First Release for Optiq Phase 2 (Cash markets)

DOCUMENT HISTORY

REVISION NO.	DATE	AUTHOR	CHANGE DESCRIPTION
1.0.0	8 February 2018	Euronext	First Release for Optiq Phase 2 (Cash markets)

 1.0.0