



Document title

THROTTLING MECHANISM

Document type or subject

Flow Control Mechanism on order entry sessions

Revision number/ Version number

1.0

Date

12 Feb 2016

This document is for information purposes only. The information and materials contained in this document are provided 'as is' and Euronext does not warrant the accuracy, adequacy or completeness and expressly disclaims liability for any errors or omissions. This document is not intended to be, and shall not constitute in any way a binding or legal agreement, or impose any legal obligation on Euronext. This document and any contents thereof, as well as any prior or subsequent information exchanged with Euronext in relation to the subject matter of this presentation, are confidential and are for the sole attention of the intended recipient. Except as described below, all proprietary rights and interest in or connected with this publication shall vest in Euronext. No part of it may be redistributed or reproduced without the prior written permission of Euronext. Portions of this presentation may contain materials or information copyrighted, trademarked or otherwise owned by a third party. No permission to use these third party materials should be inferred from this presentation.

Euronext refers to Euronext N.V. and its affiliates. Information regarding trademarks and intellectual property rights of Euronext is located at <https://www.euronext.com/terms-use>.

© 2016 Euronext N.V. - All rights reserved.

PREFACE

PURPOSE

This white paper describes the flow control mechanism that is currently active on Euronext Cash Market order entry sessions - Cash order entry access (SLE) session - and the impact of this mechanism on member's trading flow.

This mechanism imposes throttling on a Cash order entry access (SLE) session if this session introduces more orders per second than permitted by its specific configuration.

SCOPE

The Flow Control Mechanism applies to the following environments:

- Euronext Cash Markets: Regulated Markets, Warrants, TCS, BondMatch, SmartPool

ASSOCIATED DOCUMENTS

The following lists the associated documents, which either should be read in conjunction with this document or which provide other relevant information for the user, at the following addresses:

- Cash CCG creation form [HERE](#)
This form allows the member to specify the default maximum throughput per CCG session
- CCG Binary Client and Appendix Specifications, [HERE](#)
- CCG FIX 4.2 Client and Appendix Specifications, [HERE](#)

CONTACT INFORMATION

For any questions about this document please contact Euronext Customer Technical Support Group (CTSG):

- Email: CTSG@euronext.com
- Tel.: +33 (0)1 8514 8588

CONTENTS

1.	THROTTLING MECHANISM OVERVIEW.....	4
2.	THROTTLING MECHANISM PARAMETERS.....	5
2.1	Theoretical Throughput	5
2.2	Counting Window	5
2.3	Throttling Counter	5
3.	THROTTLING MECHANISM EXPLAINED.....	6
4.	FLOW CONTROL REPORTS	8

1. THROTTLING MECHANISM OVERVIEW

Each Cash order entry access (SLE) session is configured on a specific Common Customer Gateway (CCG). The Cash order entry access (SLE) session is represented by a dedicated target IP address and a port, which are allocated per member. Each session is defined by several parameters such as the protocol supported, the platform(s) it can access, the authorized member code (s) or the permitted rate of injection. This last parameter, also referred to as the *theoretical throughput*, defines the number of messages that the customer is entitled to inject without any constraint. This rate is expressed in messages per second.

The *theoretical throughput* applies at the Cash order entry access (SLE) level even if the Cash order entry access (SLE) session is configured for multiple member codes.

As long as the customer injects orders at a rate below the *theoretical throughput* limit, the order flow remains untouched. The CCG processes the packets as they come in, translates them and redirects them to the correct internal trading unit where the traded instrument is located. However, if the customer injects a message quantity exceeding the *theoretical throughput* limit, the throttling mechanism kicks in.

Through the throttling mechanism, the CCG applies a delay to the incoming messages it receives above the *theoretical throughput* limit. **No** messages are dropped, but while the messages are being queued, they are not acknowledged, as the acknowledgement is performed by the trading unit. Therefore incoming messages remain stored for some time on the CCG before they are released to the matching engine.

To obtain the real throughput or the rate at which the customer is really entering messages onto the platform, only the **incoming** flow is taken into account. All acknowledgements of orders, trades or status messages sent to the customer are not part of the calculation.

2. THROTTLING MECHANISM PARAMETERS

2.1 THEORETICAL THROUGHPUT

The theoretical throughput is the maximum number of messages per second allowed for a given Cash order entry access (SLE) as defined in the [CCG form](#) which CCG customers must fill and send back to Euronext.

2.2 COUNTING WINDOW

The number of messages per second at any given time t is the number of messages received by the CCG in a one-second interval (or window), referred to here as the *counting window*.

The first interval (counting window) starts when the Cash order entry access (SLE) session initializes the connection to the CCG.

If a message is received during the first counting window, which starts when the Cash order entry access (SLE) is established, the timer is not reset. If a message is received outside the first counting window the timer is reset and a new counting window starts at the time when the message has been received.

2.3 THROTTLING COUNTER

The CCG use a *throttling counter* to log the number of occurrences that the throttling mechanism kicks in. A different *throttling counter* is associated with each Cash order entry access (SLE) session. The *counting window* starts on the CCG at the time a Cash order entry access (SLE) session is established with the CCG. The *throttling counter* is reset at the beginning of each new interval or *counting window*.

Since Cash order entry access (SLE) sessions are independent and connect at different times, each *throttling counter* will be linked to a different reference timestamp.

3. THROTTLING MECHANISM EXPLAINED

During a *counting window*, throttling occurs when the number of messages received by the CCG has exceeded the *theoretical throughput*. In other words, if the *theoretical throughput* is set to 100 messages per second, then the throttling mechanism is applied starting with the 101st message received during the one-second interval.

Each message received afterwards is delayed by a different length of time. This delay is the time until the next *counting window*. Therefore this delay varies depending on exactly when each message was received during the *counting window*.

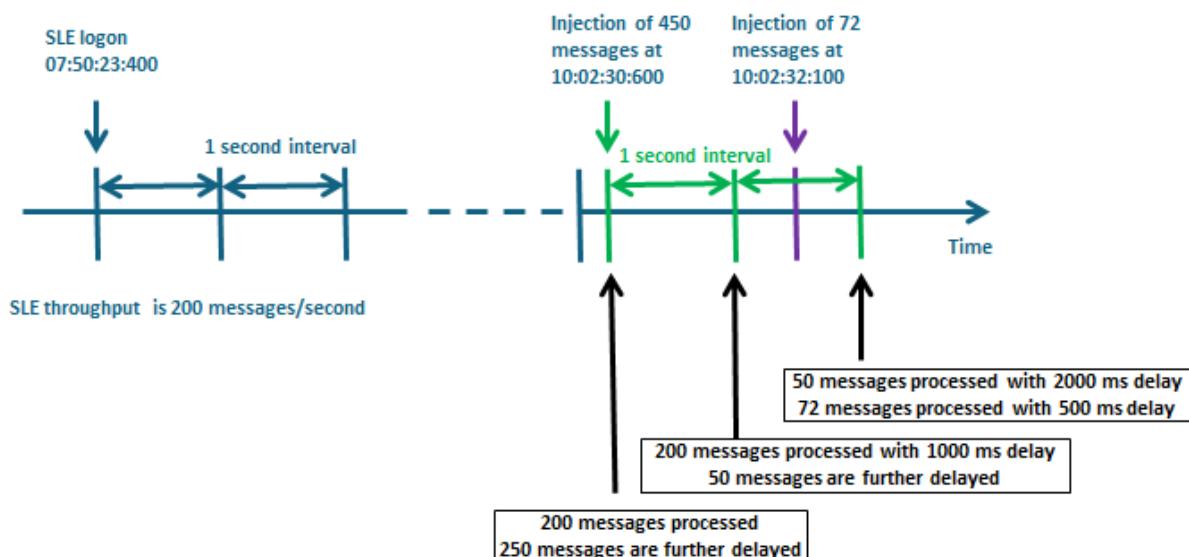
Example:

The Cash order entry access (SLE) session is established and thus the *counting window* start at 07:50:23:400.

If this Cash order entry access (SLE) has a capacity of 200 messages per second and 450 messages are injected at 10:02:30:600, then 200 messages are processed by the CCG transparently, and the following 250 messages are throttled. The next *counting window* starts at 10:02:30:600 when the 450 messages were received. At 10:02:31:600, the CCG absorbs another 200 messages. As the 50 messages remain and exceed the allowed throughput, these messages will remain unprocessed until the next *counting window*, one second later, at 10:02:32:600.

Meanwhile, at 10:02:32:100, some 72 new messages are sent by the customer. These messages remain in the TCP (Transmission Control Protocol) buffer until the next *counting window* at 10:02:32:600.

At the next *counting window* (10:02:32:600) the 50 messages waiting in the system are processed first and only after the other 72 messages are retrieved from the TCP buffer and processed.



Throttling Mechanism

Inside a counting window, the CCG processes messages for a given Cash order entry access (SLE) session as follows:

- CCG waits until a message is available from the TCP connection.
- When a message is available, the CCG retrieves it and processes the message.
- If the number of messages received is greater than the *theoretical throughput*, the CCG places the thread in *sleep mode* until the next *counting window* starts.
- When the next *counting window* starts, the sleeping thread is being awakened. The first action it takes is to record in a log file that throttling has occurred. Then, and only then, does the CCG start to retrieve/process/count messages again.

- New messages received while the throttling mechanism is running will not be retrieved from the TCP buffer. They will be stored in the TCP buffer of the CCG. If more messages are sent by the customer then this buffer may become full, and if this is the case the customer will receive TCP error messages.
- If the customer does not send any new messages while throttling is effective, and if the first message sent after the throttling mechanism has kicked in arrives in the second following the throttling event, then no messages are delayed.
- If a customer injects orders above the throughput limit at the start of the *counting window*, then the sleeping time is close to 999 ms. This means that there is a higher probability that some subsequent messages will be delayed.

4. FLOW CONTROL REPORTS

Flow control reports display for each direct customer (Member, ISP, ASP) how often throttling occurs for their Cash order entry access (SLE)'s. Three types of reports can be generated :

- Client reports (global, detailed hour by hour or fully detailed)
- Summaries per CCGs or customers
- Detailed reports (history, top100)

The following information can be found within these three reports: Throughput, Timestamp, Throttling Count and Throttling delay per each Cash order entry access (SLE).

The *timestamp* displayed in the detailed Flow Control Reports is actually a *log timestamp*.

This *timestamp* is calculated from the beginning of the *counting window* following the throttling event (when the throttling warning should be logged), plus if necessary any extra time needed by the logging process to handle other log messages already in its queue. This extra time depends on the load of the server and especially on the load of the logging process. Hence, this *timestamp* cannot be fully reliable.

If the timestamp would be accurate, the exact throttling time could be calculated using the following formula:

$$\text{Throttling timestamp} = \text{log timestamp} - \text{sleeping time}$$

In reality we get the following:

$$\text{Throttling timestamp} = \text{log timestamp} - \varepsilon - \text{sleeping time}$$

where ε is the time spent by the logging process in logging the throttling information.

However, the *throttling delay* displayed in the reports is an accurate value: it details the sleeping time of the message-reading process mentioned above.

Flow control reports are being created on daily basis and can be provided to the members via e-mail. In order to receive a report (or all), the member needs to subscribe via an e-mail to CTSG@euronext.com.

Before subscribing, the member can request an example of each of these reports at the same e-mail address provided above.