EURONEXT OPTIQ TECHNICAL NOTES

INTRODUCTION TO LZ4 COMPRESSION

Scope and audience: This technical note intends to provide general technical information about LZ4 compression on Optiq®. For additional information regarding the LZ4 implementation on the Euronext Optiq platform please consult the specifications **here.**

This document aims to familiarize developers who will work with LZ4 for the first time, specifically for the MDG compressed and snapshot channels.

Optiq MDG will use LZ4 compression, and will be available for real-time market data used on low bandwidth connections (100Mbps) and for all snapshots. Only the body of the Market Data packets will be compressed, excluding the packet header. It should be noted that a compressed market data packet can contain several different messages, which are all compressed into a single compressed packet. LZ4 was specifically chosen for the speed with which it is able to compress and decompress, and tests were performed to validate the implementation of LZ4 on Optiq MDG.

HOW IT WORKS

Data is represented as a series of sequences in the LZ4 algorithm, with each sequence beginning with a one byte token broken down into two 4-bit fields. The number of literal bytes that are to be copied to the output are in the first field, while the second field contains the number of bytes to copy from the decoded output buffer. The value of '0' represents the minimum match length of 4 bytes. Essentially, a sequence is a series of literals (i.e., non-compressed bytes), followed by a match copy.

The string of literals comes after the token, and any extra bytes needed to indicate string length. An offset indicating at which point in the output buffer to begin copying follows, and the extra bytes (if any) of the match-length would be placed at the end of the sequence.

When a value of 15 is represented in either of the bitfields, this indicates that the length is larger and that an extra byte of data should be added to the length. Similarly, a value of 255 in the extra bytes suggests that another byte should be added. As such, arbitrary lengths are represented by series of extra bytes with the value 255.

Each sequence starts with a token. The token is a one byte value, separated into two 4-bits fields. Therefore each field ranges from 0 to 15.

With LZ4, compression can either be done in a stream or in blocks. Higher compression ratios can be achieved by focusing on finding best matches, which then results in both a smaller output and faster decompression.

PARSING RESTRICTIONS

There are specific parsing rules to respect in order to remain compatible with assumptions made by the decoder:

1. The last 5 bytes are always literals
2. The last match must start at least 12 bytes before end of block. Consequently, a block with less than 13 bytes cannot be compressed.

These rules are in place to ensure that the decoder will never read beyond the input buffer, nor write beyond the output buffer.

HOW LZ4 IS USED FOR OPTIQ MDG

MDG COMPRESSION LZ4

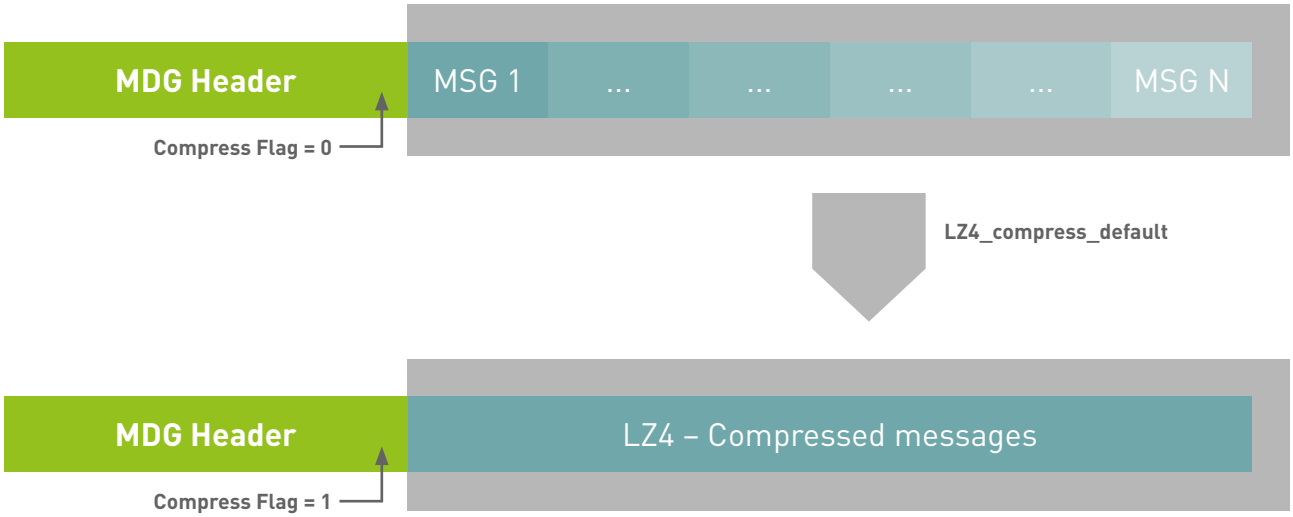

Q&A ON THE IMPLEMENTATION OF LZ4 IN OPTIQ MDG

Q: In the context of MDG, only the body of the messages will be compressed, and not the packet headers. What happens if the length of the packet is too small (i.e.less than 12 bytes): will it still be compressed?
The packet will not be compressed, and the Packet Flags field will indicate this with the flag COMPRESSED (bit 0) cleared with value 0.

Q: Is it possible to receive non-compressed messages or packets in the compressed channels?
Yes, this is possible and can happen when compression is not possible, or the compressed data would be larger than the initial data.

Q: Which language should be used to develop the decompression mechanism?
There are many possible implementations, and one of the references is C, although the API is simple to use from C++ as well.

Q: Is LZ4 is natively available in Linux distributions?
Yes

Q: What is the licensing for LZ4 library?
License differs for each implementation. We encourage developers to have their legal department review the license of the LZ4 library they plan to use.

Q: Is Euronext using the standard LZ4 Frame format, and which encoding is being used?
Euronext does not use LZ4 framing, only bare LZ4 (namely LZ4_compress_default). It uses little endian for encoding, the maximum value of an offset does not apply to the Euronext implementation, the packets MDG is sending are small (and not queued for latency purpose).

Q: LZ4 allows for very fast compression and is based on well-known LZ77 (Lempel-Ziv) algorithm. It features an extremely fast decoder. What is the recommended fastest decompression tool available?
Euronext uses API LZ4_decompress_safe, though clients are encouraged to explore all of the available options for decoding.

Disclaimers

This document contains information which is confidential and of value to Euronext. 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 or changes enabled to be made for any reason included correction, update and upgrade purpose. This document contains links to certain Internet Websites developed, sponsored or maintained by third parties unaffiliated with Euronext. The content you view therein is not provided or controlled by Euronext. Euronext is not responsible for that content, nor has it developed, checked for accuracy or otherwise reviewed the content or privacy policy of any such third party Website. This document is not intended to 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 document, 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 publication 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 publication. Implementation of the project may be subject to regulatory approval. 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.

LZ4 Library's license terms are the following: LZ4 Library Copyright (c) 2011-2014, Yann Collet All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER AND CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. " 2. ALSO ANY USE OF THE LZ4 LIBRARY SHALL BE MADE UNDER THE SOLE RESPONSIBILITY OF CUSTOMER AND THE EXCHANGE NV AND ITS AFFILIATES HEREBY DISCLAIM ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO ANY USE OF THE LZ4 LIBRARY BY THE CUSTOMER AND/OR ANY OF ITS AFFILIATES; IN NO EVENT SHALL THE EXCHANGE NV AND/OR ANY OF ITS AFFILIATES BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THE LZ4 LIBRARY EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. Additional information on LZ4 compression provided on this page was sourced from Wikipedia, https://en.wikipedia.org/wiki/LZ4_(compression_algorithm)