Internet Engineering Task Force (IETF)
Request for Comments: 9659
Updates: 8878
Category: Informational
ISSN: 2070-1721

N. Jaju, Ed.
Google
W. F. Handte, Ed.
Meta Platforms, Inc.

September 2024

Window Sizing for Zstandard Content Encoding

Abstract

Deployments of Zstandard, or "zstd", can use different window sizes to limit memory usage during compression and decompression. Some browsers and user agents limit window sizes to mitigate memory usage concerns, thereby causing interoperability issues. This document updates the window size limit in RFC 8878 from a recommendation to a requirement in HTTP contexts.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc9659.

Table of Contents

1. Introduction
2. Conventions and Definitions
3. Window Size
4. Security Considerations
5. IANA Considerations
   5.1. Content Encoding
6. Normative References
Acknowledgments
Authors' Addresses

1. Introduction

Zstandard, or "zstd", specified in [RFC8878], is a lossless data compression mechanism similar to gzip. When used with HTTP, the "zstd" content coding token signals to the decoder that the content is Zstandard-compressed.

An important property of Zstandard-compressed content is its Window_Size ([RFC8878], Section 3.1.1.1.2), which describes the maximum distance for back-references and therefore how much of the content must be kept in memory during decompression.

The minimum Window_Size is 1 KB. The maximum Window_Size is (1<<41) + 7*(1<<38) bytes, where "<<" denotes a bitwise left shift, which is 3.75 TB. Larger Window_Size values tend to improve the compression ratio but at the cost of increased memory usage.

To protect against unreasonable memory usage, some browsers and user agents limit the maximum Window_Size they will handle. This causes failures to decode responses when the content is compressed with a larger Window_Size than the recipient allows, leading to decreased interoperability.

[RFC8878], Section 3.1.1.1.2 recommends that decoders support a Window_Size of up to 8 MB, and that encoders not generate frames using a Window_Size larger than 8 MB. However, it imposes no requirements.

This document updates [RFC8878] to enforce Window_Size limits on the encoder and decoder for the "zstd" HTTP content coding.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Window Size

To ensure interoperability, when using the "zstd" content coding, decoders MUST support a Window_Size of up to and including 8 MB, and encoders MUST NOT generate frames requiring a Window_Size larger than 8 MB (see Section 5.1).

4. Security Considerations

This document introduces no new security considerations beyond those discussed in [RFC8878].

Note that decoders still need to take into account that they can receive oversized frames that do not follow the window size limit specified in this document and fail decoding when such invalid frames are received.

5. IANA Considerations

5.1. Content Encoding

This document updates the following entry in the "HTTP Content Coding Registry" in the "Hypertext Transfer Protocol (HTTP) Parameters" registry group (https://www.iana.org/assignments/http-parameters):

Name: zstd

Description: A stream of bytes compressed using the Zstandard

protocol with a Window_Size of not more than 8 MB.

Reference: This document and [RFC8878]

6. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8878] Collet, Y. and M. Kucherawy, Ed., "Zstandard Compression and the 'application/zstd' Media Type", RFC 8878, DOI 10.17487/RFC8878, February 2021, <https://www.rfc-editor.org/info/rfc8878>.

Acknowledgments

Zstandard was developed by Yann Collet.

The authors would like to thank Yann Collet, Klaus Post, Adam Rice, and members of the Web Performance Working Group in the W3C for collaborating on the window size issue and helping to formulate a solution.

Authors' Addresses

Nidhi Jaju (editor)
Google
Shibuya Stream, 3 Chome-21-3 Shibuya, Shibuya City, Tokyo
150-0002
Japan
Email: nidhijaju@google.com

W. Felix P. Handte (editor)
Meta Platforms, Inc.
380 W 33rd St
New York, NY 10001
United States of America
Email: felixh@meta.com