

Transport Layer Security Protocol Compression Methods

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

The Transport Layer Security (TLS) protocol (RFC 2246) includes features to negotiate selection of a lossless data compression method as part of the TLS Handshake Protocol and to then apply the algorithm associated with the selected method as part of the TLS Record Protocol. TLS defines one standard compression method which specifies that data exchanged via the record protocol will not be compressed. This document describes an additional compression method associated with a lossless data compression algorithm for use with TLS, and it describes a method for the specification of additional TLS compression methods.

Table of Contents

1.	Introduction	2
2.	Compression Methods	2
2.1.	DEFLATE Compression.	3
3.	Compression History and Packet Processing	4
4.	Internationalization Considerations	4
5.	IANA Considerations	4
6.	Security Considerations	5
7.	Acknowledgements	6
8.	References	6
8.1.	Normative References	6
8.2.	Informative References	6
	Author's Address	7
	Full Copyright Statement	8

1. Introduction

The Transport Layer Security (TLS) protocol (RFC 2246, [2]) includes features to negotiate selection of a lossless data compression method as part of the TLS Handshake Protocol and to then apply the algorithm associated with the selected method as part of the TLS Record Protocol. TLS defines one standard compression method, `CompressionMethod.null`, which specifies that data exchanged via the record protocol will not be compressed. While this single compression method helps ensure that TLS implementations are interoperable, the lack of additional standard compression methods has limited the ability of implementers to develop interoperable implementations that include data compression.

TLS is used extensively to secure client-server connections on the World Wide Web. While these connections can often be characterized as short-lived and exchanging relatively small amounts of data, TLS is also being used in environments where connections can be long-lived and the amount of data exchanged can extend into thousands or millions of octets. XML [4], for example, is increasingly being used as a data representation method on the Internet, and XML tends to be verbose. Compression within TLS is one way to help reduce the bandwidth and latency requirements associated with exchanging large amounts of data while preserving the security services provided by TLS.

This document describes an additional compression method associated with a lossless data compression algorithm for use with TLS. Standardization of the compressed data formats and compression algorithms associated with this compression method is beyond the scope of this document.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [1].

2. Compression Methods

TLS [2] includes the following compression method structure in sections 6.1 and 7.4.1.2 and Appendix sections A.4.1 and A.6:

```
enum { null(0), (255) } CompressionMethod;
```

which allows for later specification of up to 256 different compression methods. This definition is updated to segregate the range of allowable values into three zones:

1. Values from 0 (zero) through 63 decimal (0x3F) inclusive are reserved for IETF Standards Track protocols.
2. Values from 64 decimal (0x40) through 223 decimal (0xDF) inclusive are reserved for assignment for non-Standards Track methods.
3. Values from 224 decimal (0xE0) through 255 decimal (0xFF) inclusive are reserved for private use.

Additional information describing the role of the IANA in the allocation of compression method identifiers is described in Section 5.

In addition, this definition is updated to include assignment of an identifier for the DEFLATE compression method:

```
enum { null(0), DEFLATE(1), (255) } CompressionMethod;
```

As described in section 6 of RFC 2246 [2], TLS is a stateful protocol. Compression methods used with TLS can be either stateful (the compressor maintains its state through all compressed records) or stateless (the compressor compresses each record independently), but there seems to be little known benefit in using a stateless compression method within TLS.

The DEFLATE compression method described in this document is stateful. It is RECOMMENDED that other compression methods that might be standardized in the future be stateful as well.

Compression algorithms can occasionally expand, rather than compress, input data. A compression method that exceeds the expansion limits described in section 6.2.2 of RFC 2246 [2] MUST NOT be used with TLS.

2.1. DEFLATE Compression

The DEFLATE compression method and encoding format is described in RFC 1951 [5]. Examples of DEFLATE use in IETF protocols can be found in RFC 1979 [6], RFC 2394 [7], and RFC 3274 [8].

DEFLATE allows the sending compressor to select from among several options to provide varying compression ratios, processing speeds, and memory requirements. The receiving decompressor MUST automatically adjust to the parameters selected by the sender. All data that was submitted for compression MUST be included in the compressed output,

with no data retained to be included in a later output payload. Flushing ensures that each compressed packet payload can be decompressed completely.

3. Compression History and Packet Processing

Some compression methods have the ability to maintain state/history information when compressing and decompressing packet payloads. The compression history allows a higher compression ratio to be achieved on a stream as compared to per-packet compression, but maintaining a history across packets implies that a packet might contain data needed to completely decompress data contained in a different packet. History maintenance thus requires both a reliable link and sequenced packet delivery. Since TLS and lower-layer protocols provide reliable, sequenced packet delivery, compression history information MAY be maintained and exploited if supported by the compression method.

As described in section 7 of RFC 2246 [2], TLS allows multiple connections to be instantiated using the same session through the resumption feature of the TLS Handshake Protocol. Session resumption has operational implications when multiple compression methods are available for use within the session. For example, load balancers will need to maintain additional state information if the compression state is not cleared when a session is resumed. As a result, the following restrictions MUST be observed when resuming a session:

1. The compression algorithm MUST be retained when resuming a session.
2. The compression state/history MUST be cleared when resuming a session.

4. Internationalization Considerations

The compression method identifiers specified in this document are machine-readable numbers. As such, issues of human internationalization and localization are not introduced.

5. IANA Considerations

Section 2 of this document describes a registry of compression method identifiers to be maintained by the IANA, including assignment of an identifier for the DEFLATE compression method. Identifier values from the range 0-63 (decimal) inclusive are assigned via RFC 2434 Standards Action [3]. Values from the range 64-223 (decimal)

inclusive are assigned via RFC 2434 Specification Required [3]. Identifier values from 224-255 (decimal) inclusive are reserved for RFC 2434 Private Use [3].

6. Security Considerations

This document does not introduce any topics that alter the threat model addressed by TLS. The security considerations described throughout RFC 2246 [2] apply here as well.

However, combining compression with encryption can sometimes reveal information that would not have been revealed without compression: data that is the same length before compression might be a different length after compression, so adversaries that observe the length of the compressed data might be able to derive information about the corresponding uncompressed data. Some symmetric encryption ciphersuites do not hide the length of symmetrically encrypted data at all. Others hide it to some extent, but still do not hide it fully. For example, ciphersuites that use stream cipher encryption without padding do not hide length at all; ciphersuites that use Cipher Block Chaining (CBC) encryption with padding provide some length hiding, depending on how the amount of padding is chosen. Use of TLS compression SHOULD take into account that the length of compressed data may leak more information than the length of the original uncompressed data.

Compression algorithms tend to be mathematically complex and prone to implementation errors. An implementation error that can produce a buffer overrun introduces a potential security risk for programming languages and operating systems that do not provide buffer overrun protections. Careful consideration should thus be given to protections against implementation errors that introduce security risks.

As described in Section 2, compression algorithms can occasionally expand, rather than compress, input data. This feature introduces the ability to construct rogue data that expands to some enormous size when compressed or decompressed. RFC 2246 describes several methods to ameliorate this kind of attack. First, compression has to be lossless. Second, a limit (1,024 bytes) is placed on the amount of allowable compression content length increase. Finally, a limit (2^{14} bytes) is placed on the total content length. See section 6.2.2 of RFC 2246 [2] for complete details.

7. Acknowledgements

The concepts described in this document were originally discussed on the IETF TLS working group mailing list in December, 2000. The author acknowledges the contributions to that discussion provided by Jeffrey Altman, Eric Rescorla, and Marc Van Heyningen. Later suggestions that have been incorporated into this document were provided by Tim Dierks, Pasi Eronen, Peter Gutmann, Elgin Lee, Nikos Mavroyanopoulos, Alexey Melnikov, Bodo Moeller, Win Treese, and the IESG.

8. References

8.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Dierks, T. and C. Allen, "The TLS Protocol Version 1.0", RFC 2246, January 1999.
- [3] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

8.2. Informative References

- [4] Bray, T., Paoli, J., Sperberg-McQueen, C. and E. Maler, "Extensible Markup Language (XML) 1.0 (2nd ed)", W3C REC-xml, October 2000, <<http://www.w3.org/TR/REC-xml>>.
- [5] Deutsch, P., "DEFLATE Compressed Data Format Specification version 1.3", RFC 1951, May 1996.
- [6] Woods, J., "PPP Deflate Protocol", RFC 1979, August 1996.
- [7] Pereira, R., "IP Payload Compression Using DEFLATE", RFC 2394, December 1998.
- [8] Gutmann, P., "Compressed Data Content Type for Cryptographic Message Syntax (CMS)", RFC 3274, June 2002.

Author's Address

Scott Hollenbeck
VeriSign, Inc.
21345 Ridgetop Circle
Dulles, VA 20166-6503
US

Email: shollenbeck@verisign.com

Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

