

Network Working Group  
Internet-Draft  
Updates: [2246](#) (if approved)  
Expires: May 20, 2004

S. Hollenbeck  
VeriSign, Inc.  
November 20, 2003

Transport Layer Security Protocol Compression Methods  
draft-ietf-tls-compression-06.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on May 20, 2004.

Copyright Notice

Copyright (C) The Internet Society (2003). 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.

Internet-Draft

TLS Compression Methods

November 2003

Conventions Used In 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].

Table of Contents

- [1.](#) Introduction . . . . . [3](#)
- [2.](#) Compression Methods . . . . . [4](#)
- [2.1](#) Compression History and Packet Processing . . . . . [5](#)
- [2.2](#) ZLIB Compression . . . . . [5](#)
- [3.](#) Intellectual Property Considerations . . . . . [6](#)
- [4.](#) Internationalization Considerations . . . . . [7](#)
- [5.](#) IANA Considerations . . . . . [8](#)
- [6.](#) Security Considerations . . . . . [9](#)
- [7.](#) Acknowledgements . . . . . [10](#)
- Normative References . . . . . [11](#)
- Informative References . . . . . [12](#)
- Author's Address . . . . . [12](#)
- Intellectual Property and Copyright Statements . . . . . [13](#)

## 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.

## [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 future standardization efforts of the IETF TLS working group.
2. Values from 64 decimal (0x40) through 192 decimal (0xC0) are reserved for assignment by the IANA for specifications developed outside the TLS working group. Assignments from this range of values MUST be made by the IANA and MUST be associated with a formal reference that describes the compression method.
3. Values from 193 decimal (0xC1) through 255 decimal (0xFF) 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 ZLIB compression method:

```
enum { null(0), ZLIB(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 ZLIB 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](#) Compression History and Packet Processing

Some compression methods have the ability to maintain 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.

## [2.2](#) ZLIB Compression

The ZLIB compression method and encoding format is described in [RFC 1950 \[5\]](#) and [RFC 1951 \[6\]](#). Examples of ZLIB use in IETF protocols can

be found in [RFC 1979](#) [7], [RFC 2394](#) [8], and [RFC 3274](#) [9].

ZLIB 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](#). Intellectual Property Considerations

Many compression algorithms are subject to patent or other intellectual property rights claims. Implementers are encouraged to seek legal guidance to better understand the implications of developing implementations of the compression method described in this document or other documents that describe compression methods for use with TLS.

#### [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 ZLIB compression method. Identifier values from the range reserved for future standardization efforts of the TLS working group MUST be assigned according to the "Standards Action" policy described in [RFC 2434](#) [3]. Values from the range reserved for private use MUST be used according to the "Private Use" policy described in [RFC 2434](#). Values from the general IANA pool MUST be assigned according to the "IETF Consensus" policy described in [RFC 2434](#).

## 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 don't 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.

## [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, and Win Treese.

## Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Dierks, T., Allen, C., Treese, W., Karlton, P., Freier, A. and P. Kocher, "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.

## 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, L. and J-L. Gailly, "ZLIB Compressed Data Format Specification version 3.3", [RFC 1950](#), May 1996.
- [6] Deutsch, P., "DEFLATE Compressed Data Format Specification version 1.3", [RFC 1951](#), May 1996.
- [7] Woods, J., "PPP Deflate Protocol", [RFC 1979](#), August 1996.
- [8] Pereira, R., "IP Payload Compression Using DEFLATE", [RFC 2394](#), December 1998.
- [9] Gutmann, P., "Compressed Data Content Type for Cryptographic Message Syntax (CMS)", [RFC 3274](#), June 2002.

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

EMail: shollenbeck@verisign.com

Hollenbeck

Expires May 20, 2004

[Page 12]

---

Internet-Draft

TLS Compression Methods

November 2003

#### Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

#### Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

#### Acknowledgment

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

