

Network Working Group
Internet-Draft

R. Friend
Hifn
September 7, 2003

Expires: March 7, 2004

Transport Layer Security Protocol Compression Using LZS
draft-friend-tls-lzs-compression-00.txt

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

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 March 7, 2004.

It is intended that a future version of this draft be submitted to the IESG for publication as an Informational RFC. 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. [[TLSCOMP](#)] defines another compression method. This document describes an additional compression method associated with the LZS data compression algorithm for use with TLS. This document defines the application of the LZS algorithm to the TLS Record Protocol [[RFC2246](#)].

Friend

Expires March 7, 2004

[Page 1]

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](#)
- [1.1](#) General. [3](#)
- [1.2](#) Licensing. [3](#)
- [1.3](#) Specification of Requirements. [4](#)
- [2.](#) Compression Methods [4](#)
- [2.1](#) LZS CompressionMethod [5](#)
- [2.2](#) Security Issues with Single History Compression. [5](#)
- [3.](#) LZS Compression [5](#)
- [3.1](#) Background of LZS Compression [5](#)
- [3.2](#) LZS Compression History and Packet Processing [6](#)
- [3.3](#) LZS Compressed Record Format [7](#)
- [3.4](#) TLSComp Header Format [7](#)
- [3.4.1](#) Flags [7](#)
- [3.5](#) LZS Compression Encoding Format [8](#)
- [3.6](#) Padding [9](#)
- [4.](#) Sending Compressed Records [9](#)
- [4.1](#) Transmitter Process [9](#)
- [4.2](#) Receiver Process [10](#)
- [4.3](#) Anti-Expansion Mechanism [10](#)
- [5.](#) Intellectual Property Considerations [12](#)
- [6.](#) Internationalization Considerations [13](#)
- [7.](#) IANA Considerations [14](#)
- [8.](#) Security Considerations [15](#)
- [9.](#) Acknowledgements [16](#)
- References [17](#)
- Author's Address [18](#)
- Intellectual Property and Copyright Statements [19](#)

Friend

Expires March 7, 2004

[Page 2]

1. Introduction

1.1 General

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. For example, SSL is now increasingly being used as an alternative VPN connection. Compression services have long been associated with IPsec and PPTP VPN connections, so extending compression services to SSL VPN connections preserves the user experience for any VPN connection. 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. This document specifies the application of LZS compression, a lossless compression algorithm, to TLS record payloads. This specification also assumes a thorough understanding of the TLS protocol [[TLS](#)].

1.3 Licensing

Hifn, Inc. holds patents on the LZS algorithm. Software licenses for a reference implementation are available for use in IPPCP, IPsec, PPP, in addition to TLS applications.

Software licenses are available on a non-discriminatory basis. Hardware implementations are also available. For more information, contact Hifn at the address listed with the author's address.

Friend

Expires March 7, 2004

[Page 3]

1.4 Specification of Requirements

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](#)].

2. Compression Methods

TLS [[TLS](#)] 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. [TLSComp] updates this definition 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 7](#).

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

```
enum { null(0), LZS(TBD), (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 LZS compression method described

in this document is stateful.

Friend

Expires March 25, 2004

[Page 4]

Compression algorithms can occasionally expand, rather than compress, input data. The worst case expansion factor of the LZS compression method is only 12.5%. Thus, TLS records of 15K bytes can never exceed the expansion limits described in [section 6.2.2 of RFC 2246](#) [2]. If TLS records of 16K bytes expand to greater than 17K bytes, then the uncompressed version of the TLS record must be transmitted, as described below.

2.1 LZS CompressionMethod

The LZS CompressionMethod is a 16-bit index, and is negotiated as described in [\[TLS\]](#) and [\[TLSComp\]](#). The LZS CompressionMethod is stored in the TLS Record Layer connection state as described in [\[TLS\]](#).

The value for LZS CompressionMethod is TBD.

2.2 Security Issues with Compression Histories

Sharing compression histories between more than one TLS session may potentially cause information leakage between the TLS sessions, as pathological compressed data can potentially reference data prior to the beginning of the current record. LZS implementations guard against this situation. However, to avoid this potential situation from occurring, implementations that support TLS compression MUST use separate compression histories for each TLS session. This is not a limitation of LZS compression, but is an artifact for any compression algorithm.

Furthermore, the LZS compression history (as well as any compression history) contains plaintext. Specifically, the LZS history contains the last 2K bytes of plaintext of the TLS session. Thus, when the TLS session terminates, the implementation SHOULD treat the history as it does the session key material (e.g. free memory, zeroize contents, etc.).

3. LZS Compression

3.1 Background of LZS Compression

Starting with a sliding window compression history, similar to [\[LZ1\]](#), Hifn developed a new, enhanced compression algorithm identified as LZS. The LZS algorithm is a general purpose lossless compression algorithm for use with a wide variety of data types. Its encoding method is very efficient, providing compression for strings as short as two octets in length.

The LZS algorithm uses a sliding window of 2,048 bytes. During compression, redundant sequences of data are replaced with tokens that represent those sequences. During decompression, the original

sequences are substituted for the tokens in such a way that the

Friend

Expires March 25, 2004

[Page 5]

original data is exactly recovered. LZS differs from lossy compression algorithms, such as those often used for video compression, that do not exactly reproduce the original data. The details of LZS compression can be found in [[ANSI94](#)].

3.2 LZS Compression History and Packet Processing

This standard specifies "stateful" compression, that is, maintaining the compression history between records within a particular TLS compression session. Within each separate compression history, the LZS CompressionMethod has the ability to maintain compression history information when compressing and decompressing record payloads. Stateful compression provides a higher compression ratio to be achieved on the data stream as compared to stateless compression (resetting the compression history between every record), particularly for small records.

Stateful compression requires both a reliable link and sequenced packet delivery, to ensure all records can be decompressed in the same order they were compressed. Since TLS and lower-layer protocols provide reliable, sequenced packet delivery, compression history information MAY be maintained and exploited when using the LZS CompressionMethod.

Furthermore, there MUST be a separate LZS compression history associated with each open TLS session. This not only provides enhanced security (no potential information leakage between sessions via a shared compression history), but also enables superior compression ratio (bit bandwidth on the connection) across all open TLS sessions with compression. A shared history would require resetting the compression (and decompression) history when switching between TLS sessions, and a single history implementation would require resetting the compression (and decompression) history between each record.

The sender MUST reset the compression history prior to compressing the first TLS record of a TLS session after TLS handshake completes. It is advantageous for the sender to maintain the compression history for all subsequent records processed during the TLS session. This results in the greatest compression ratio for a given data set. In either case, this compression history MUST NOT be used for any other open TLS session, in order to ensure privacy between TLS sessions.

The sender MUST "flush" the compressor each time it transmits a compressed record. Flushing means that all data going into the compressor is included in the output, i.e., no data is retained in the hope of achieving better compression. Flushing ensures that each compressed packet payload can be decompressed completely.

Flushing is necessary to prevent a record's data from spilling over into a later record. This is important to synchronize compressed

Friend

Expires March 25, 2004

[Page 6]

data with the authenticated and encrypted data in a TLS record. Flushing is handled automatically in most LZS implementations.

When the TLS session terminates, the implementation SHOULD dispose of the memory resources associated with the related TLS compression history. That is, the compression history SHOULD be handled as the TLS key material is handled.


The LZS CompressionMethod also features "decompressing" uncompressed data in order to maintain the history in the case that the "compressed" data actually expanded. The LZS CompressionMethod record format facilitates identifying whether records contain compressed or uncompressed data. The LZS decoding process accommodates decompressing either compressed or uncompressed data.

3.3 LZS Compressed Record Format

Prior to compression, the uncompressed data (TLSPlaintext.fragment) comprises a plaintext TLS record. After compression, the compressed data (TLSCompressed.fragment) comprises a 8-bit TLSComp header followed by the compressed (or uncompressed) data.


3.4. TLSComp Header Format

The one-octet header has the following structure:



3.4.1 Flags

The format of the 8-bit Flags TLSComp field is as follows:



Res - Reserved

Reserved for future use. MUST be set to zero. MUST be ignored by the receiving node.

RST - Reset Compression History

Friend

Expires March 25, 2004

[Page 7]

The RST bit is used to inform the decompressing peer that the compression history in this TLS session was reset prior to the data contained in this TLS record being compressed. When the RST bit is set to "1" a compression history reset is performed, when RST is set to "0", a compression history reset is not performed.

This bit MUST be set to a value of "1" for the first compressed TLS transmitted record of a TLS session. This bit may also be used by the transmitter for other exception cases when the compression history must be reset.

C/U - Compressed/Uncompressed Bit

The C/U indicates whether the data field contains compressed or uncompressed data. A value of 1 indicates compressed data (often referred to as a compressed packet), and a value of 0 indicates uncompressed data (or an uncompressed packet).

3.5 LZS Compression Encoding Format

The LZS compression method, encoding format, and application examples are described in [RFC 1967](#) [7], [RFC 1974](#) [6], [RFC 2395](#)[5], and [[ANSI94](#)].

Some implementations of LZS allow the sending compressor to select from among several options to provide varying compression ratios, processing speeds, and memory requirements. Other implementations of LZS provide optimal compression ratio at byte per clock speeds.

The receiving LZS decompressor automatically adjusts to the settings selected by the sender. Also, receiving LZS decompressors will update the decompression history with uncompressed data. This facilitates never obtaining less than 1:1 compression ratio (never transmit with expanded data) in the session.

The input to the payload compression algorithm is TLSPlaintext data destined to an active TLS session with compression negotiated. The output of the algorithm is a new (and hopefully smaller) TLSCompressed record. The output payload contains the input payload's data in either compressed or uncompressed format. The input and output payloads are each an integral number of bytes in length.

The output payload is always prepended with the TLSComp header. If the uncompressed form is used, the output payload is identical to the input payload, and the TLSComp header reflects uncompressed data.

If the compressed form is used, encoded as defined in [[ANSI94](#)], and the TLSComp header reflects compressed data. The LZS encoded format is repeated here for informational purposes ONLY.

Friend

Expires March 25, 2004

[Page 8]

```

<Compressed Stream> := [<Compressed String>*] <End Marker>
<Compressed String> := 0 <Raw Byte> | 1 <Compressed Bytes>

<Raw Byte> := <b><b><b><b><b><b><b><b>          (8-bit byte)
<Compressed Bytes> := <Offset> <Length>

<Offset> := 1 <b><b><b><b><b><b><b> |          (7-bit offset)
            0 <b><b><b><b><b><b><b><b><b><b><b><b><b> (11-bit offset)
<End Marker> := 110000000
<b> := 1 | 0

<Length> :=
00      = 2      1111 0110      = 14
01      = 3      1111 0111      = 15
10      = 4      1111 1000      = 16
1100    = 5      1111 1001      = 17
1101    = 6      1111 1010      = 18
1110    = 7      1111 1011      = 19
1111 0000 = 8      1111 1100      = 20
1111 0001 = 9      1111 1101      = 21
1111 0010 = 10     1111 1110      = 22
1111 0011 = 11     1111 1111 0000 = 23
1111 0100 = 12     1111 1111 0001 = 24
1111 0101 = 13     ...

```

3.6 Padding

A datagram payload compressed using LZS always ends with the last compressed data byte (also known as the <end marker>), which is used to disambiguate padding. This allows trailing bits as well as bytes to be considered padding.

The size of a compressed payload MUST be in whole octet units.

4. Sending Compressed Datagrams

All TLS records processed with a TLS session state that includes LZS compression are processed as follows. The reliable and efficient transport of LZS compressed records in the TLS session depends on the following processes.

4.1. Transmitter Process

The compression operation results in either compressed or uncompressed data. When a TLS record is received, it is assigned to a particular TLS context that includes the LZS compression history buffer. It is processed according to ANSI X3.241-1994 to form compressed data or used as is to form uncompressed data. For the first record of the session, or for

exception conditions, the compression history MUST be cleared. In

Friend

Expires March 25, 2004

[Page 9]

performing the compression operation the compression history MUST be updated when either a compressed record or uncompressed record is produced. Uncompressed TLS records MAY be sent at any time. Uncompressed TLS records MUST be sent if compression causes enough expansion to cause the data compression TLS record size to exceed the MTU defined in [section 6.2.2 in RFC 2246](#). The output of the compression operation is placed in the fragment field of the TLSCompressed structure (TLSCompressed.fragment).

The TLSComp header byte is located just prior to the first byte of the compressed TLS record in TLSCompressed.fragment. The C/U bit in the TLSComp header is set according to whether the data field contains compressed or uncompressed data. The RST bit in the TLSComp header is set to "1" if the compression history was reset prior to compressing the TLSplaintext.fragment that comprises TLSCompressed.fragment. Uncompressed data MUST be transmitted (and the C/U bit set to 0) if the "compressed" (expanded) data exceeded 17K bytes.

[4.2.](#) Receiver Process

Prior to decompressing the first compressed TLS record in the TLS session, the receiver MUST reset the decompression history. Subsequent records are decompressed in the order received. The receiver decompresses the Payload Data field according to the encoding specified in section 3.2 of [\[ANSI94\]](#).

If the received datagram is not compressed, the receiver needs to perform no decompression processing and the Payload Data field of the datagram is ready for processing by the next protocol layer.

After a TLS record is received from the peer and decrypted, the RST and C/U bits MUST be checked.

If the C/U bit is set to "1", the resulting compressed data block MUST be decompressed according to [\[ANSI94\]](#).

If the C/U bit is set to "0", the specified decompression history MUST be updated with the received uncompressed data.

If the RST bit is set to "1", the receiving decompression history MAY be reset to an initial state prior to decompressing the TLS record. (However, due to the characteristics of the Hifn LZS algorithm, a decompression history reset is not required). After reset, any compressed or uncompressed data contained in the packet is processed.

[4.3.](#) Anti-Expansion Mechanism

During compression, there are two options on how to handle

Friend

Expires March 25, 2004

[Page 10]

packets that expand:

- 1) Send the expanded data (as long as TLSCompressed.length is 17K or less) and maintain the history, thus allowing loss of current bandwidth but preserving future bandwidth on the link.
- 2) Send the uncompressed data and do not clear the compression history; the decompressor will update its history, thus conserving the current bandwidth and future bandwidth on the link.

The second option is the preferred option, and SHOULD be implemented.

A third option:

- 3) Send the uncompressed data and clear the history, thus conserving current bandwidth, but allowing possible loss of future bandwidth on the link.

SHOULD NOT be implemented.

Friend

Expires March 25, 2004

[Page 11]

5. Intellectual Property Considerations

The LZS compression algorithm is subject to patent and 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.

Friend

Expires March 25, 2004

[Page 12]

6. 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.

7. 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 LZS 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](#).

Friend

Expires March 25, 2004

[Page 14]

8. Security Considerations

This document introduces several topics that may alter the threat model addressed by TLS. One security topic is the sharing of a single compression history across more than one TLS session, as addressed in [section 2.2](#) above.

Another security topic specific to compression is that the LZS compression history contains plaintext, so when the TLS session terminates, the implementation SHOULD treat the history as it does the session key material. This is described both in sections [2.2](#) and [3.2](#) above.

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. 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.

The security considerations described throughout [RFC 2246](#) [2] apply here as well.

Friend

Expires March 25, 2004

[Page 15]

9. Acknowledgements

The concepts described in this document were derived from [RFC 1967](#) [7], [RFC 1974](#) [6], [RFC 2395](#)[5], and [TLSComp] (work in progress). The author acknowledges the contributions of Scott Hollenbeck, Douglas Whiting, and Russell Dietz.

- . Later suggestions that have been incorporated into this document were provided by ...

Friend

Expires March 25, 2004

[Page 16]

References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [TLS] 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.
- [TLSCOMP] Hollenbeck, S. " Transport Layer Security Protocol Compression Methods", [draft-ietf-tls-compression-05.txt](#), September 2003, work in progress.
- [5] Friend, R., "IP Payload Compression Using LZS", [RFC 2395](#), December 1998.
- [6] Friend, R., "PPP Stac LZS Compression Protocol", [RFC 1974](#), August 1996.
- [7] Friend, R., "PPP LZS-DCP Compression Protocol (LZS-DCP)", [RFC 1967](#), August 1996.
- [ANSI94] American National Standards Institute, Inc., "Data Compression Method for Information Systems," ANSI X3.241-1994, August 1994.
- [LZ1] Lempel, A., and Ziv, J., "A Universal Algorithm for Sequential Data Compression", IEEE Transactions On Information Theory, Vol. IT-23, No. 3, September 1977.

Friend

Expires March 25, 2004

[Page 17]

Author's Address

Robert Friend
Hifn
5973 Avenida Encinas
Carlsbad, CA 92008
US

Email: rfriend@hifn.com

Friend

Expires March 25, 2004

[Page 18]

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.

Friend

Expires March 25, 2004

[Page 19]

Acknowledgement

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

Friend

Expires March 7, 2004

[Page 20]