

RObust Header Compression (ROHC):
Requirements on TCP/IP Header Compression

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.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document contains requirements on the TCP/IP header compression scheme (profile) to be developed by the RObust Header Compression (ROHC) Working Group. The document discusses the scope of TCP compression, performance considerations, assumptions about the surrounding environment, as well as Intellectual Property Rights concerns. The structure of this document is inherited from RFC 3096, which defines IP/UDP/RTP requirements for ROHC.

Table of Contents

1. Introduction	2
2. Header Compression Requirements	2
2.1. Impact on Internet Infrastructure	2
2.2. Supported Headers and Kinds of TCP Streams	3
2.3. Performance Issues	4
2.4. Requirements Related to Link Layer Characteristics	6
2.5. Intellectual Property Rights (IPR)	7
3. Security Consideration	7
4. IANA Considerations	7
5. Acknowledgements	7
6. Informative References	7

1. Introduction

The goal of the ROHC WG is to develop header compression schemes that perform well over links with high error rates and long link roundtrip times. The schemes must perform well for cellular links that use technologies such as Wideband Code Division Multiple Access (W-CDMA), Enhanced Data rates for GSM Evolution (EDGE), and CDMA2000. However, the schemes should also be applicable to other link technologies with high loss and long roundtrip times.

The main objective for ROHC has been robust compression of IP/UDP/RTP [5], but the WG is also chartered to develop new header compression solutions for IP/TCP [1], [2]. Because TCP traffic, in contrast to RTP, has usually been sent over reliable links, existing schemes for TCP, [3] and [4], have not experienced the same robustness problems as RTP compression. However, there are still many scenarios where TCP header compression will be implemented over less reliable links [11], [12], making robustness an important objective for the new TCP compression scheme. Other, equally important, objectives for ROHC TCP compression are: improved compression efficiency, enhanced capabilities for compression of header fields including TCP options, and finally incorporation of TCP compression into the ROHC framework [6].

2. Header Compression Requirements

The following requirements have, more or less arbitrarily, been divided into five groups. The first group deals with requirements concerning the impact of a header compression scheme on the rest of the Internet infrastructure. The second group defines what kind of headers must be compressed efficiently. The third and fourth groups concern performance requirements and capability requirements that stem from the properties of link technologies where ROHC TCP is expected to be used. Finally, the fifth section discusses Intellectual Property Rights related to ROHC TCP compression.

2.1. Impact on Internet Infrastructure

1. Transparency: When a header is compressed and then decompressed, the resulting header must be semantically identical to the original header. If this cannot be achieved, the packet containing the erroneous header must be discarded.

Justification: The header compression process must not produce headers that might cause problems for any current or future part of the Internet infrastructure.

Note: The ROHC WG has not found a case where "semantically identical" is not the same as "bitwise identical".

2. Ubiquity: Must not require modifications to existing IP (v4 or v6) or TCP implementations.

Justification: Ease of deployment.

Note: The ROHC WG may recommend changes that would increase the compression efficiency for the TCP streams emitted by implementations. However, ROHC cannot assume such recommendations will be followed.

Note: Several TCP variants are currently in use on the Internet. This requirement implies that the header compression scheme must work efficiently and correctly for all expected TCP variants.

2.2. Supported Headers and Kinds of TCP Streams

1. IPv4 and IPv6: Must support both IPv4 and IPv6. This means that all expected changes in the IP header fields must be handled by the compression scheme, and commonly changing fields should be compressed efficiently. Compression must still be possible when IPv6 Extensions are present in the header. When designing the compression scheme, the usage of Explicit Congestion Notification (ECN) [10] should be considered as a common behavior. Therefore, the scheme must also compress efficiently in the case when the ECN bits are used.

Justification: IPv4 and IPv6 will both be around for the foreseeable future, and Options/Extensions are expected to be more commonly used. ECN is expected to have a breakthrough and be widely deployed, especially in combination with TCP.

2. Mobile IP: The kinds of headers used by Mobile IP{v4,v6} must be supported and should be compressed efficiently. For IPv4 these include headers of tunneled packets. For IPv6 they include headers containing the Routing Header and the Home Address Option.

Justification: It is very likely that Mobile IP will be used by cellular devices.

3. Generality: Must handle all headers from arbitrary TCP streams.

Justification: There must be a generic scheme that can compress reasonably well for any TCP traffic pattern. This does not preclude optimizations for certain traffic patterns.

4. IPSEC: The scheme should be able to compress headers containing IPSEC subheaders where the NULL encryption algorithm is used.

Justification: IPSEC is expected to be used to provide necessary end-to-end security.

Note: It is not possible to compress the encrypted part of an ESP header, nor the cryptographic data in an AH header.

5. TCP: All fields supported by [4] should be handled with efficient compression, as should be the cases when the SYN, FIN or TCP ECN [10] bits are set.

Justification: These bits are expected to be commonly used.

6. TCP options: The scheme must support compression of packets with any TCP option present, even if the option itself is not compressed. Further, for some commonly used options the scheme should also provide compression mechanisms for the options.

Justification: Because various TCP options are commonly used, applicability of the compression scheme would be significantly reduced if packets with options could not be compressed.

Note: Options that should be compressed are:

- Selective Acknowledgement (SACK), [8], [9]
- Timestamp, [7]

2.3. Performance Issues

1. Performance/Spectral Efficiency: The scheme must provide low relative overhead under expected operating conditions; compression efficiency should be better than for RFC 2507 [4] under equivalent operating conditions.

Justification: Spectrum efficiency is a primary goal.

Note: The relative overhead is the average header overhead relative to the payload. Any auxiliary (e.g., control or feedback) channels used by the scheme should be taken into account when calculating the header overhead.

2. Losses between compressor and decompressor: The scheme should make sure losses between compressor and decompressor do not result in losses of subsequent packets, or cause damage to the context that results in incorrect decompression of subsequent packet headers.

Justification: Even though link layer retransmission in most cases is expected to almost eliminate losses between compressor and decompressor, there are still many scenarios where TCP header compression will be implemented over less reliable links [11], [12]. In such cases, loss propagation due to header compression could affect certain TCP mechanisms that are capable of handling some losses; loss propagation could also have a negative impact on the performance of TCP loss recovery.

3. Residual errors in compressed headers: Residual errors in compressed headers may result in delivery of incorrectly decompressed headers not only for the damaged packet itself, but also for subsequent packets, because errors may be saved in the context state. For TCP, the compression scheme is not required to implement explicit mechanisms for residual error detection, but the compression scheme must not affect TCP's end-to-end mechanisms for error detection.

Justification: For links carrying TCP traffic, the residual error rate is expected to be insignificant. However, residual errors may still occur, especially in the end-to-end path. Therefore, it is crucial that TCP is not prevented from handling these.

Note: This requirement implies that the TCP checksum must be carried unmodified in all compressed headers.

Note: The error detection mechanism in TCP may be able to detect residual bit errors, but the mechanism is not designed for this purpose, and might actually provide rather weak protection. Therefore, although it is not a requirement of the compression scheme, it should be possible for the decompressor to detect residual errors and discard such packets.

4. Short-lived TCP transfers: The scheme should provide mechanisms for efficient compression of short-lived TCP transfers, minimizing the size of context initiation headers.

Justification: Many TCP transfers are short-lived. This may lead to a low gain for header compression schemes that, for each new packet stream, requires full headers to be sent initially and allows small compressed headers only after the initialization phase.

Note: This requirement implies that mechanisms for building new contexts that are based on information from previous contexts or for concurrent packet streams to share context information should be considered.

- 5a. Moderate Packet Misordering: The scheme should efficiently handle moderate misordering (2-3 packets) in the packet stream reaching the compressor.

Justification: This kind of misordering is common.

- 5b. Packet Misordering: The scheme must be able to correctly handle packet misordering and preferably compress when misordered packets are in the TCP stream reaching the compressor.

Justification: Misordering happens regularly in the Internet. However, because the Internet is engineered to run TCP reasonably well, excessive misordering will not be common and need not be handled with optimum efficiency.

6. Processing delay: The scheme should not contribute significantly to the system delay budget.

2.4. Requirements Related to Link Layer Characteristics

1. Unidirectional links: Must be possible to implement (possibly with less efficiency) without explicit feedback messages from decompressor to compressor.

Justification: There are links that do not provide a feedback channel or where feedback is not desirable for other reasons.

2. Link delay: Must operate under all expected link delay conditions.
3. Header compression coexistence: The scheme must fit into the ROHC framework together with other ROHC profiles (e.g., [6]).
4. Note on misordering between compressor and decompressor:

When compression is applied over tunnels, misordering often cannot be completely avoided. The header compression scheme should not prohibit misordering between compressor and decompressor, as it would therefore not be applicable in many tunneling scenarios. However, in the case of tunneling, it is usually possible to get misordering indications. Therefore, the compression scheme does not have to support detection of misordering, but can assume that such information is available from lower layers when misordering occurs.

2.5. Intellectual Property Rights (IPR)

The ROHC WG must spend effort to achieve a high degree of confidence that there are no known IPR claims that cover the final compression solution for TCP.

Justification: Currently there is no TCP header compression scheme available that can efficiently compress the packet headers of modern TCP, e.g., with SACK, ECN, etc. ROHC is expected to fill this gap by providing a ROHC TCP scheme that is applicable in the wide area Internet, not only over error-prone radio links. It must thus attempt to be as future-proof as possible, and only unencumbered solutions, or solutions where the terms of any IPR are such that there is no hindrance on implementation and deployment, will be acceptable to the Internet at large.

3. Security Consideration

A protocol specified to meet these requirements must be able to compress packets containing IPSEC headers according to the IPSEC requirement, 2.2.4. There may be other security aspects to consider in such protocols. This document by itself, however, does not add any security risks.

4. IANA Considerations

A protocol that meets these requirements will require the IANA to assign various numbers. This document by itself, however, does not require any IANA involvement.

5. Acknowledgements

This document has evolved through fruitful discussions with and input from Gorrry Fairhurst, Carsten Bormann, Mikael Degermark, Mark West, Jan Kullander, Qian Zhang, Richard Price, and Aaron Falk. The document quality was significantly improved thanks to Peter Eriksson, who made a thorough linguistic review.

Last, but not least, Ghyslain Pelletier and Kristofer Sandlund served as committed working group document reviewers.

6. Informative References

- [1] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [2] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.

- [3] Jacobson, V., "Compressing TCP/IP headers for low-speed serial links", RFC 1144, February 1990.
- [4] Degermark, M., Nordgren, B., and S. Pink, "IP Header Compression", RFC 2507, February 1999.
- [5] Degermark, M., "Requirements for robust IP/UDP/RTP header compression", RFC 3096, July 2001.
- [6] Bormann, C., Burmeister, C., Degermark, M., Fukushima, H., Hannu, H., Jonsson, L-E., Hakenberg, R., Koren, T., Le, K., Liu, Z., Martensson, A., Miyazaki, A., Svanbro, K., Wiebke, T., Yoshimura, T., and H. Zheng, "RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed", RFC 3095, July 2001.
- [7] Jacobson, V., Braden, R., and D. Borman, "TCP Extensions for High Performance", RFC 1323, May 1992.
- [8] Mathis, M., Mahdavi, J., Floyd, S., and A. Romanow, "TCP Selective Acknowledgement Options", RFC 2018, October 1996.
- [9] Floyd, S., Mahdavi, J., Mathis, M., and M. Podolsky, "An Extension to the Selective Acknowledgement (SACK) Option for TCP", RFC 2883, July 2000.
- [10] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, September 2001.
- [11] Dawkins, S., Montenegro, G., Kojo, M., and V. Magret, "End-to-end Performance Implications of Slow Links", BCP 48, RFC 3150, July 2001.
- [12] Fairhurst, G. and L. Wood, "Advice to link designers on link Automatic Repeat reQuest (ARQ)", BCP 62, RFC 3366, August 2002.

Author's Address

Lars-Erik Jonsson
Ericsson AB
Box 920
SE-971 28 Lulea
Sweden

Phone: +46 8 404 29 61
Fax: +46 920 996 21
EMail: lars-erik.jonsson@ericsson.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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.

