

Workgroup: Network Working Group
Internet-Draft: draft-irtf-cfrg-aead-limits-03
Published: 12 July 2021
Intended Status: Informational
Expires: 13 January 2022
Authors: F. Günther M. Thomson C.A. Wood
 ETH Zurich Mozilla Cloudflare
Usage Limits on AEAD Algorithms

Abstract

An Authenticated Encryption with Associated Data (AEAD) algorithm provides confidentiality and integrity. Excessive use of the same key can give an attacker advantages in breaking these properties. This document provides simple guidance for users of common AEAD functions about how to limit the use of keys in order to bound the advantage given to an attacker. It considers limits in both single- and multi-key settings.

Discussion Venues

This note is to be removed before publishing as an RFC.

Source for this draft and an issue tracker can be found at <https://github.com/cfrg/draft-irtf-cfrg-aead-limits>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 13 January 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Requirements Notation](#)
- [3. Notation](#)
- [4. Calculating Limits](#)
- [5. Single-Key AEAD Limits](#)
 - [5.1. AEAD AES 128 GCM and AEAD AES 256 GCM](#)
 - [5.1.1. Confidentiality Limit](#)
 - [5.1.2. Integrity Limit](#)
 - [5.2. AEAD CHACHA20 POLY1305](#)
 - [5.3. AEAD AES 128 CCM](#)
 - [5.3.1. Confidentiality Limit](#)
 - [5.3.2. Integrity Limit](#)
 - [5.4. AEAD AES 128 CCM 8](#)
 - [5.5. Single-Key Examples](#)
- [6. Multi-Key AEAD Limits](#)
 - [6.1. AEAD AES 128 GCM and AEAD AES 256 GCM](#)
 - [6.1.1. Authenticated Encryption Security Limit](#)
 - [6.1.2. Confidentiality Limit](#)
 - [6.1.3. Integrity Limit](#)
 - [6.2. AEAD CHACHA20 POLY1305, AEAD AES 128 CCM, and AEAD AES 128 CCM 8](#)
 - [6.2.1. AEAD CHACHA20 POLY1305](#)
 - [6.2.2. AEAD AES 128 CCM and AEAD AES 128 CCM 8](#)
- [7. Security Considerations](#)
- [8. IANA Considerations](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

An Authenticated Encryption with Associated Data (AEAD) algorithm provides confidentiality and integrity. [RFC5116] specifies an AEAD as a function with four inputs -- secret key, nonce, plaintext, associated data (of which plaintext and associated data can optionally be zero-length) -- that produces ciphertext output and an

error code indicating success or failure. The ciphertext is typically composed of the encrypted plaintext bytes and an authentication tag.

The generic AEAD interface does not describe usage limits. Each AEAD algorithm does describe limits on its inputs, but these are formulated as strict functional limits, such as the maximum length of inputs, which are determined by the properties of the underlying AEAD composition. Degradation of the security of the AEAD as a single key is used multiple times is not given a thorough treatment.

These limits might also be influenced by the number of "users" of a given key. In the traditional setting, there is one key shared between two parties. Any limits on the maximum length of inputs or encryption operations apply to that single key. The attacker's goal is to break security (confidentiality or integrity) of that specific key. However, in practice, there are often many users with independent keys. This multi-key security setting, often referred to as the multi-user setting in the academic literature, hence considers an attacker's advantage in breaking security of any of these many keys, further assuming the attacker may have done some offline work to help break security. As a result, AEAD algorithm limits may depend on offline work and the number of keys. However, given that a multi-key attacker does not target any specific key, acceptable advantages may differ from that of the single-key setting.

The number of times a single pair of key and nonce can be used might also be relevant to security. For some algorithms, such as AEAD_AES_128_GCM or AEAD_AES_256_GCM, this limit is 1 and using the same pair of key and nonce has serious consequences for both confidentiality and integrity; see [[NonceDisrespecting](#)]. Nonce-reuse resistant algorithms like AEAD_AES_128_GCM_SIV can tolerate a limited amount of nonce reuse.

It is good practice to have limits on how many times the same key (or pair of key and nonce) are used. Setting a limit based on some measurable property of the usage, such as number of protected messages or amount of data transferred, ensures that it is easy to apply limits. This might require the application of simplifying assumptions. For example, TLS 1.3 specifies limits on the number of records that can be protected, using the simplifying assumption that records are the same size; see [Section 5.5](#) of [[TLS](#)].

Currently, AEAD limits and usage requirements are scattered among peer-reviewed papers, standards documents, and other RFCs. Determining the correct limits for a given setting is challenging as papers do not use consistent labels or conventions, and rarely apply any simplifications that might aid in reaching a simple limit.

The intent of this document is to collate all relevant information about the proper usage and limits of AEAD algorithms in one place. This may serve as a standard reference when considering which AEAD algorithm to use, and how to use it.

2. Requirements Notation

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. Notation

This document defines limitations in part using the quantities below.

Symbol	Description
n	AEAD block length (in bits)
k	AEAD key length (in bits)
r	AEAD nonce length (in bits)
t	Size of the authentication tag (in bits)
l	Maximum length of each message (in blocks)
s	Total plaintext length in all messages (in blocks)
q	Number of protected messages (AEAD encryption invocations)
v	Number of attacker forgery attempts (failed AEAD decryption invocations)
p	Upper bound on adversary attack probability
o	Offline adversary work (in number of encryption and decryption queries; multi-key setting only)
u	Number of keys (multi-key setting only)
B	Maximum number of blocks encrypted by any key (multi-key setting only)

Table 1

For each AEAD algorithm, we define the (passive) confidentiality and (active) integrity advantage roughly as the advantage an attacker has in breaking the corresponding classical security property for the algorithm. A passive attacker can query ciphertexts for arbitrary plaintexts. An active attacker can additionally query plaintexts for arbitrary ciphertexts. Moreover, we define the combined authenticated encryption advantage guaranteeing both confidentiality and integrity against an active attacker. Specifically:

*Confidentiality advantage (CA): The probability of a passive attacker succeeding in breaking the confidentiality properties (IND-CPA) of the AEAD scheme. In this document, the definition of

confidentiality advantage roughly is the probability that an attacker successfully distinguishes the ciphertext outputs of the AEAD scheme from the outputs of a random function.

*Integrity advantage (IA): The probability of a active attacker succeeding in breaking the integrity properties (INT-CTXT) of the AEAD scheme. In this document, the definition of integrity advantage roughly is the probability that an attacker is able to forge a ciphertext that will be accepted as valid.

*Authenticated Encryption advantage (AEA): The probability of a active attacker succeeding in breaking the authenticated-encryption properties of the AEAD scheme. In this document, the definition of authenticated encryption advantage roughly is the probability that an attacker successfully distinguishes the ciphertext outputs of the AEAD scheme from the outputs of a random function or is able to forge a ciphertext that will be accepted as valid.

See [[AECComposition](#)], [[AEAD](#)] for the formal definitions of and relations between passive confidentiality (IND-CPA), ciphertext integrity (INT-CTXT), and authenticated encryption security (AE). The authenticated encryption advantage subsumes, and can be derived as the combination of, both CA and IA:

$CA \leq AEA$

$IA \leq AEA$

$AEA \leq CA + IA$

Each application requires an individual determination of limits in order to keep CA and IA sufficiently small. For instance, TLS aims to keep CA below 2^{-60} and IA below 2^{-57} in the single-key setting; see [Section 5.5](#) of [[TLS](#)].

4. Calculating Limits

Once upper bounds on CA, IA, or AEA are determined, this document defines a process for determining three overall operational limits:

*Confidentiality limit (CL): The number of messages an application can encrypt before giving the adversary a confidentiality advantage higher than CA.

*Integrity limit (IL): The number ciphertexts an application can decrypt, either successfully or not, before giving the adversary an integrity advantage higher than IA.

*Authenticated encryption limit (AEL): The combined number of messages and number of ciphertexts an application can encrypt or

decrypt before giving the adversary an authenticated encryption advantage higher than AEA.

When limits are expressed as a number of messages an application can encrypt or decrypt, this requires assumptions about the size of messages and any authenticated additional data (AAD). Limits can instead be expressed in terms of the number of bytes, or blocks, of plaintext and maybe AAD in total.

To aid in translating between message-based and byte/block-based limits, a formulation of limits that includes a maximum message size (l) and the AEAD schemes' block length in bits (n) is provided.

All limits are based on the total number of messages, either the number of protected messages (q) or the number of forgery attempts (v); which correspond to CL and IL respectively.

Limits are then derived from those bounds using a target attacker probability. For example, given an integrity advantage of $IA = v * (81 / 2^{106})$ and a targeted maximum attacker success probability of $IA = p$, the algorithm remains secure, i.e., the adversary's advantage does not exceed the targeted probability of success, provided that $v \leq (p * 2^{106}) / 81$. In turn, this implies that $v \leq (p * 2^{103}) / 1$ is the corresponding limit.

To apply these limits, implementations can count the number of messages that are protected or rejected against the determined limits (q and v respectively). This requires that messages cannot exceed the maximum message size (l) that is chosen.

This analysis assumes a message-based approach to setting limits. Implementations that use byte counting rather than message counting could use a maximum message size (l) of one to determine a limit for q that can be applied with byte counting. This results in attributing per-message overheads to every byte, so the resulting limit could be significantly lower than necessary. Actions, like rekeying, that are taken to avoid the limit might occur more often as a result.

5. Single-Key AEAD Limits

This section summarizes the confidentiality and integrity bounds and limits for modern AEAD algorithms used in IETF protocols, including: AEAD_AES_128_GCM [[RFC5116](#)], AEAD_AES_256_GCM [[RFC5116](#)], AEAD_AES_128_CCM [[RFC5116](#)], AEAD_CHACHA20_POLY1305 [[RFC8439](#)], AEAD_AES_128_CCM_8 [[RFC6655](#)].

The CL and IL values bound the total number of encryption and forgery queries (q and v). Alongside each value, we also specify these bounds.

5.1. AEAD_AES_128_GCM and AEAD_AES_256_GCM

The CL and IL values for AES-GCM are derived in [[AEBounds](#)] and summarized below. For this AEAD, $n = 128$ and $t = 128$ [[GCM](#)]. In this example, the length s is the sum of AAD and plaintext, as described in [[GCMProofs](#)].

5.1.1. Confidentiality Limit

$$CA \leq ((s + q + 1)^2) / 2^{129}$$

This implies the following usage limit:

$$q + s \leq p^{(1/2)} * 2^{(129/2)} - 1$$

Which, for a message-based protocol with $s \leq q * l$, if we assume that every packet is size l , produces the limit:

$$q \leq (p^{(1/2)} * 2^{(129/2)} - 1) / (l + 1)$$

5.1.2. Integrity Limit

$$IA \leq 2 * (v * (l + 1)) / 2^{128}$$

This implies the following limit:

$$v \leq (p * 2^{127}) / (l + 1)$$

5.2. AEAD_CHACHA20_POLY1305

The only known analysis for AEAD_CHACHA20_POLY1305 [[ChaCha20Poly1305Bounds](#)] combines the confidentiality and integrity limits into a single expression, covered below:

$$CA \leq v * ((8 * l) / 2^{106})$$

$$IA \leq v * ((8 * l) / 2^{106})$$

This advantage is a tight reduction based on the underlying Poly1305 PRF [[Poly1305](#)]. It implies the following limit:

$$v \leq (p * 2^{103}) / l$$

5.3. AEAD_AES_128_CCM

The CL and IL values for AEAD_AES_128_CCM are derived from [[CCM-ANALYSIS](#)] and specified in the QUIC-TLS mapping specification [[I-D.ietf-quic-tls](#)]. This analysis uses the total number of underlying block cipher operations to derive its bound. For CCM, this number is the sum of: the length of the associated data in blocks, the length

of the ciphertext in blocks, the length of the plaintext in blocks, plus 1.

In the following limits, this is simplified to a value of twice the length of the packet in blocks, i.e., $2l$ represents the effective length, in number of block cipher operations, of a message with l blocks. This simplification is based on the observation that common applications of this AEAD carry only a small amount of associated data compared to ciphertext. For example, QUIC has 1 to 3 blocks of AAD.

For this AEAD, $n = 128$ and $t = 128$.

5.3.1. Confidentiality Limit

$$\begin{aligned} CA &\leq (2l * q)^2 / 2^n \\ &\leq (2l * q)^2 / 2^{128} \end{aligned}$$

This implies the following limit:

$$q \leq \sqrt{(p * 2^{126}) / l^2}$$

5.3.2. Integrity Limit

$$\begin{aligned} IA &\leq v / 2^t + (2l * (v + q))^2 / 2^n \\ &\leq v / 2^{128} + (2l * (v + q))^2 / 2^{128} \end{aligned}$$

This implies the following limit:

$$v + (2l * (v + q))^2 \leq p * 2^{128}$$

In a setting where v or q is sufficiently large, v is negligible compared to $(2l * (v + q))^2$, so this can be simplified to:

$$v + q \leq p^{(1/2)} * 2^{63} / l$$

5.4. AEAD_AES_128_CCM_8

The analysis in [\[CCM-ANALYSIS\]](#) also applies to this AEAD, but the reduced tag length of 64 bits changes the integrity limit calculation considerably.

$$\begin{aligned} IA &\leq v / 2^t + (2l * (v + q))^2 / 2^n \\ &\leq v / 2^{64} + (2l * (v + q))^2 / 2^{128} \end{aligned}$$

This results in reducing the limit on v by a factor of 2^{64} .

$$v * 2^{64} + (2l * (v + q))^2 \leq p * 2^{128}$$

5.5. Single-Key Examples

An example protocol might choose to aim for a single-key CA and IA that is at most 2^{50} . If the messages exchanged in the protocol are at most a common Internet MTU of around 1500 bytes, then a value for l might be set to 2^7 . The values in [Table 2](#) show values of q and v that might be chosen under these conditions.

AEAD	Maximum q	Maximum v
AEAD_AES_128_GCM	$2^{32.5}$	2^{71}
AEAD_AES_256_GCM	$2^{32.5}$	2^{71}
AEAD_CHACHA20_POLY1305	n/a	2^{46}
AEAD_AES_128_CCM	2^{30}	2^{30}
AEAD_AES_128_CCM_8	$2^{30.9}$	2^{13}

Table 2: Example limits

AEAD_CHACHA20_POLY1305 provides no limit to q based on the provided analysis.

The limit for q on AEAD_AES_128_CCM and AEAD_AES_128_CCM_8 is reduced due to a need to reduce the value of q to ensure that IA does not exceed the target. This assumes equal proportions for q and v for AEAD_AES_128_CCM. AEAD_AES_128_CCM_8 in permits a much smaller value of v due to the shorter tag, which permits a higher limit for q .

Some protocols naturally limit v to 1, such as TCP-based variants of TLS, which terminate sessions on decryption failure. If v is limited to 1, q can be increased to 2^{31} for both CCM AEADs.

6. Multi-Key AEAD Limits

In the multi-key setting, each user is assumed to have an independent and identically distributed key, though nonces may be re-used across users with some very small probability. The success probability in attacking one of these many independent keys can be generically bounded by the success probability of attacking a single key multiplied by the number of keys present [[MUSecurity](#)], [[GCM-MU](#)]. Absent concrete multi-key bounds, this means the attacker advantage in the multi-key setting is the product of the single-key advantage and the number of keys.

This section summarizes the confidentiality and integrity bounds and limits for the same algorithms as in [Section 5](#) for the multi-key setting. The CL and IL values bound the total number of encryption and forgery queries (q and v). Alongside each value, we also specify these bounds.

6.1. AEAD_AES_128_GCM and AEAD_AES_256_GCM

Concrete multi-key bounds for AEAD_AES_128_GCM and AEAD_AES_256_GCM exist due to Theorem 4.3 in [GCM-MU2], which covers protocols with nonce randomization, like TLS 1.3 [TLS] and QUIC [RFC9001].

Results for AES-GCM without nonce randomization are captured by Theorem 3.1 in [GCM-MU2], which apply to protocols such as TLS 1.2 [RFC5246]. This produces similar limits under most conditions.

For this AEAD, $n = 128$, $t = 128$, and $r = 96$; the key length is $k = 128$ or $k = 256$ for AEAD_AES_128_GCM and AEAD_AES_256_GCM respectively.

6.1.1. Authenticated Encryption Security Limit

Protocols with nonce randomization have a limit of:

$$AEA \leq ((q+v)*l*B / 2^{127}) + (1 / 2^{48})$$

This implies the following limit:

$$q + v \leq (p * 2^{127} - 2^{79}) / (l * B)$$

This assumes that B is much larger than 100; that is, each user enciphers significantly more than 1600 bytes of data. Otherwise, B should be increased by 161 for AEAD_AES_128_GCM and by 97 for AEAD_AES_256_GCM.

Protocols without nonce randomization have limits that are essentially the same provided that p is not less than 2^{-48} , as the simplified expression for AEA does not include the 2^{-48} term:

$$q + v \leq p * 2^{127} / (l * B)$$

Without nonce randomization, B should be increased by an additional 0.5.

6.1.2. Confidentiality Limit

The confidentiality advantage is essentially dominated by the same terms as the AE advantage for protocols with nonce randomization:

$$CA \leq (q*l*B / 2^{127}) + (1 / 2^{48})$$

This implies the following limit:

$$q \leq (p * 2^{127} - 2^{79}) / (l * B)$$

As before, the limit without nonce randomization is:

$$q \leq (p * 2^{127}) / (1 * B)$$

6.1.3. Integrity Limit

There is currently no dedicated integrity multi-key bound available for AEAD_AES_128_GCM and AEAD_AES_256_GCM. The AE limit can be used to derive an integrity limit as:

$$IA \leq AEA$$

[Section 6.1.1](#) therefore contains the integrity limits.

6.2. AEAD_CHACHA20_POLY1305, AEAD_AES_128_CCM, and AEAD_AES_128_CCM_8

There are currently no concrete multi-key bounds for AEAD_CHACHA20_POLY1305, AEAD_AES_128_CCM, or AEAD_AES_128_CCM_8. Thus, to account for the additional factor u , i.e., the number of keys, each p term in the confidentiality and integrity limits is replaced with p / u .

6.2.1. AEAD_CHACHA20_POLY1305

The combined confidentiality and integrity limit for AEAD_CHACHA20_POLY1305 is as follows.

$$\begin{aligned} v &\leq ((p / u) * 2^{106}) / 81 \\ &\leq (p * 2^{103}) / (1 * u) \end{aligned}$$

6.2.2. AEAD_AES_128_CCM and AEAD_AES_128_CCM_8

The integrity limit for AEAD_AES_128_CCM is as follows.

$$v + q \leq (p / u)^{(1/2)} * 2^{63} / 1$$

Likewise, the integrity limit for AEAD_AES_128_CCM_8 is as follows.

$$v * 2^{64} + (21 * (v + q))^2 \leq (p / u) * 2^{128}$$

7. Security Considerations

The different analyses of AEAD functions that this work is based upon generally assume that the underlying primitives are ideal. For example, that the pseudorandom function (PRF) or pseudorandom permutation (PRP) the AEAD builds upon is indistinguishable from a truly random function. Thus, the advantage estimates assume that the attacker is not able to exploit a weakness in an underlying primitive.

Many of the formulae in this document depend on simplifying assumptions, from differing models, which means that results are not

universally applicable. When using this document to set limits, it is necessary to validate all these assumptions for the setting in which the limits might apply. In most cases, the goal is to use assumptions that result in setting a more conservative limit, but this is not always the case. As an example of one such simplification, this document defines v as the total number of failed decryption queries (that is, failed forgery attempts), whereas models usually count in v all forgery attempts.

The CA and IL values defined in this document are upper bounds based on existing cryptographic research. Future analysis may introduce tighter bounds. Applications SHOULD NOT assume these bounds are rigid, and SHOULD accommodate changes. In particular, in two-party communication, one participant cannot regard apparent overuse of a key by other participants as being in error, when it could be that the other participant has better information about bounds.

Note that the limits in this document apply to the adversary's ability to conduct a single successful forgery. For some algorithms and in some cases, an adversary's success probability in repeating forgeries may be noticeably larger than that of the first forgery. As an example, [MF05] describes such multiple forgery attacks in the context of AES-GCM in more detail.

8. IANA Considerations

This document does not make any request of IANA.

9. References

9.1. Normative References

[AEAD] Rogaway, P., "Authenticated-Encryption with Associated-Data", September 2002, <<https://cseweb.ucsd.edu/~mihir/papers/musu.pdf>>.

[AEBounds] Luykx, A. and K. Paterson, "Limits on Authenticated Encryption Use in TLS", 8 March 2016, <<http://www.isg.rhul.ac.uk/~kp/TLS-AEbounds.pdf>>.

[AEComposition] Bellare, M. and C. Namprempre, "Authenticated Encryption: Relations among notions and analysis of the generic composition paradigm", July 2007, <<http://cseweb.ucsd.edu/~mihir/papers/oem.pdf>>.

[CCM-ANALYSIS] Jonsson, J., "On the Security of CTR + CBC-MAC", Selected Areas in Cryptography pp. 76-93, DOI 10.1007/3-540-36492-7_7, 2003, <https://doi.org/10.1007/3-540-36492-7_7>.

[ChaCha20Poly1305Bounds]

Procter, G., "A Security Analysis of the Composition of ChaCha20 and Poly1305", 11 August 2014, <<https://eprint.iacr.org/2014/613.pdf>>.

[GCM]

Dworkin, M., "Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC", NIST Special Publication 800-38D, November 2007.

[GCM-MU]

Bellare, M. and B. Tackmann, "The Multi-User Security of Authenticated Encryption: AES-GCM in TLS 1.3", 27 November 2017, <<https://eprint.iacr.org/2016/564.pdf>>.

[GCM-MU2]

Hoang, V.T., Tessaro, S., and A. Thiruvengadam, "The Multi-user Security of GCM, Revisited: Tight Bounds for Nonce Randomization", 15 October 2018, <<https://eprint.iacr.org/2018/993.pdf>>.

[GCMProofs]

Iwata, T., Ohashi, K., and K. Minematsu, "Breaking and Repairing GCM Security Proofs", 1 August 2012, <<https://eprint.iacr.org/2012/438.pdf>>.

[MUSecurity]

Bellare, M., Boldyreva, A., and S. Micali, "Public-Key Encryption in a Multi-user Setting: Security Proofs and Improvements", May 2000, <<https://cseweb.ucsd.edu/~mihir/papers/musu.pdf>>.

[Poly1305]

Bernstein, D., "The Poly1305-AES Message-Authentication Code", Fast Software Encryption pp. 32-49, DOI 10.1007/11502760_3, 2005, <https://doi.org/10.1007/11502760_3>.

[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/rfc/rfc2119>>.

[RFC5116]

McGrew, D., "An Interface and Algorithms for Authenticated Encryption", RFC 5116, DOI 10.17487/RFC5116, January 2008, <<https://www.rfc-editor.org/rfc/rfc5116>>.

[RFC6655]

McGrew, D. and D. Bailey, "AES-CCM Cipher Suites for Transport Layer Security (TLS)", RFC 6655, DOI 10.17487/RFC6655, July 2012, <<https://www.rfc-editor.org/rfc/rfc6655>>.

[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/rfc/rfc8174>>.

[RFC8439]

Nir, Y. and A. Langley, "ChaCha20 and Poly1305 for IETF Protocols", RFC 8439, DOI 10.17487/RFC8439, June 2018, <<https://www.rfc-editor.org/rfc/rfc8439>>.

9.2. Informative References

[I-D.ietf-quic-tls] Thomson, M. and S. Turner, "Using TLS to Secure QUIC", Work in Progress, Internet-Draft, draft-ietf-quic-tls-34, 14 January 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-quic-tls-34>>.

[MF05] McGrew, D.A. and S.R. Fluhrer, "Multiple forgery attacks against Message Authentication Codes", 31 May 2005, <<https://csrc.nist.gov/CSRC/media/Projects/Block-Cipher-Techniques/documents/BCM/Comments/CWC-GCM/multi-forge-01.pdf>>.

[NonceDisrespecting] Bock, H., Zauner, A., Devlin, S., Somorovsky, J., and P. Jovanovic, "Nonce-Disrespecting Adversaries -- Practical Forgery Attacks on GCM in TLS", 17 May 2016, <<https://eprint.iacr.org/2016/475.pdf>>.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/rfc/rfc5246>>.

[RFC9001] Thomson, M., Ed. and S. Turner, Ed., "Using TLS to Secure QUIC", RFC 9001, DOI 10.17487/RFC9001, May 2021, <<https://www.rfc-editor.org/rfc/rfc9001>>.

[TLS] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/rfc/rfc8446>>.

Authors' Addresses

Felix Günther
ETH Zurich

Email: mail@felixguenther.info

Martin Thomson
Mozilla

Email: mt@lowentropy.net

Christopher A. Wood
Cloudflare

Email: caw@heapingbits.net