HTTP Working Group Internet-Draft Intended status: Standards Track Expires: May 4, 2017

# Encrypted Content-Encoding for HTTP draft-ietf-httpbis-encryption-encoding-04

### Abstract

This memo introduces a content coding for HTTP that allows message payloads to be encrypted.

Note to Readers

Discussion of this draft takes place on the HTTP working group mailing list (ietf-http-wg@w3.org), which is archived at <a href="https://lists.w3.org/Archives/Public/ietf-http-wg/">https://lists.w3.org/Archives/Public/ietf-http-wg/</a> .

Working Group information can be found at <a href="http://httpwg.github.io/">http://httpwg.github.io/</a>; source code and issues list for this draft can be found at <a href="https://github.com/httpwg/http-extensions/labels/encryption">https://github.com/httpwg/http-extensions/labels/encryption</a> .

#### Status of This Memo

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

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 <u>http://datatracker.ietf.org/drafts/current/</u>.

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 May 4, 2017.

## Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents

(<u>http://trustee.ietf.org/license-info</u>) 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

$\underline{1}$ . Introduction	<u>2</u>
<u>1.1</u> . Notational Conventions	<u>3</u>
2. The "aes128gcm" HTTP Content Coding	<u>3</u>
2.1. Encryption Content Coding Header	<u>5</u>
2.2. Content Encryption Key Derivation	<u>6</u>
<u>2.3</u> . Nonce Derivation	<u>7</u>
<u>3</u> . Crypto-Key Header Field	<u>7</u>
<u>4</u> . Examples	<u>8</u>
<u>4.1</u> . Encryption of a Response	<u>8</u>
<u>4.2</u> . Encryption with Multiple Records	<u>9</u>
5. Security Considerations	<u>9</u>
5.1. Key and Nonce Reuse	<u>10</u>
5.2. Data Encryption Limits	<u>10</u>
<u>5.3</u> . Content Integrity	<u>10</u>
5.4. Leaking Information in Headers	<u>11</u>
<u>5.5</u> . Poisoning Storage	<u>11</u>
5.6. Sizing and Timing Attacks	<u>12</u>
<u>6</u> . IANA Considerations	<u>12</u>
6.1. The "aes128gcm" HTTP Content Coding	<u>12</u>
<u>6.2</u> . Crypto-Key Header Field	<u>12</u>
<u>6.3</u> . The HTTP Crypto-Key Parameter Registry	<u>12</u>
<u>6.3.1</u> . keyid	<u>13</u>
<u>6.3.2</u> . aes128gcm	<u>13</u>
<u>7</u> . References	<u>13</u>
<u>7.1</u> . Normative References	<u>13</u>
<u>7.2</u> . Informative References	<u>14</u>
Appendix A. JWE Mapping	<u>15</u>
Appendix B. Acknowledgements	<u>16</u>
Author's Address	<u>16</u>

## **1**. Introduction

It is sometimes desirable to encrypt the contents of a HTTP message (request or response) so that when the payload is stored (e.g., with a HTTP PUT), only someone with the appropriate key can read it.

For example, it might be necessary to store a file on a server without exposing its contents to that server. Furthermore, that same

Internet-Draft

HTTP encryption coding

file could be replicated to other servers (to make it more resistant to server or network failure), downloaded by clients (to make it available offline), etc. without exposing its contents.

These uses are not met by the use of TLS [<u>RFC5246</u>], since it only encrypts the channel between the client and server.

This document specifies a content coding (<u>Section 3.1.2 of [RFC7231]</u>) for HTTP to serve these and other use cases.

This content coding is not a direct adaptation of message-based encryption formats - such as those that are described by [<u>RFC4880</u>], [<u>RFC5652</u>], [<u>RFC7516</u>], and [<u>XMLENC</u>] - which are not suited to stream processing, which is necessary for HTTP. The format described here cleaves more closely to the lower level constructs described in [<u>RFC5116</u>].

To the extent that message-based encryption formats use the same primitives, the format can be considered as sequence of encrypted messages with a particular profile. For instance, <u>Appendix A</u> explains how the format is congruent with a sequence of JSON Web Encryption [<u>RFC7516</u>] values with a fixed header.

This mechanism is likely only a small part of a larger design that uses content encryption. How clients and servers acquire and identify keys will depend on the use case. Though a complete key management system is not described, this document defines an Crypto-Key header field that can be used to convey keying material.

## **<u>1.1</u>**. Notational Conventions

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 [<u>RFC2119</u>].

Base64url encoding is defined in <u>Section 2 of [RFC7515]</u>.

### 2. The "aes128gcm" HTTP Content Coding

The "aes128gcm" HTTP content coding indicates that a payload has been encrypted using Advanced Encryption Standard (AES) in Galois/Counter Mode (GCM) as identified as AEAD\_AES\_128\_GCM in [RFC5116], Section 5.1. The AEAD\_AES\_128\_GCM algorithm uses a 128 bit content encryption key.

Using this content coding requires knowledge of a key. The Crypto-Key header field (<u>Section 3</u>) can be included to describe how the content encryption key is derived or retrieved. Keys might be

[Page 3]

provided in messages that are separate from those with encrypted content using Crypto-Key, or provided through external mechanisms.

The "aes128gcm" content coding uses a single fixed set of encryption primitives. Cipher suite agility is achieved by defining a new content coding scheme. This ensures that only the HTTP Accept-Encoding header field is necessary to negotiate the use of encryption.

The "aes128gcm" content coding uses a fixed record size. The final encoding consists of a header (see <u>Section 2.1</u>), zero or more fixed size encrypted records, and a partial record. The partial record MUST be shorter than the fixed record size.

++	content is rs octets minus padding
data	of between 2 and 65537 octets;
++	the last record is smaller
I	
V	
++	add padding to get rs octets;
pad   data	the last record contains
++	up to rs minus 1 octets
I	
V	
++	encrypt with AEAD_AES_128_GCM;
ciphertext	final size is rs plus 16 octets
++	the last record is smaller

The record size determines the length of each portion of plaintext that is enciphered, with the exception of the final record, which is necessarily smaller. The record size ("rs") is included in the content coding header (see <u>Section 2.1</u>).

AEAD\_AES\_128\_GCM produces ciphertext 16 octets longer than its input plaintext. Therefore, the length of each enciphered record other than the last is equal to the value of the "rs" parameter plus 16 octets. To prevent an attacker from truncating a stream, an encoder MUST append a record that contains only padding and is smaller than the full record size if the final record ends on a record boundary. A receiver MUST fail to decrypt if the final record ciphertext is less than 18 octets in size or equal to the record size plus 16 (that is, the size of a full encrypted record). Valid records always contain at least two octets of padding and a 16 octet authentication tag.

Each record contains between 2 and 65537 octets of padding, inserted into a record before the enciphered content. Padding consists of a

[Page 4]

Internet-Draft

two octet unsigned integer in network byte order, followed that number of zero-valued octets. A receiver MUST fail to decrypt if any padding octet other than the first two are non-zero, or a record has more padding than the record size can accommodate.

The nonce for each record is a 96-bit value constructed from the record sequence number and the input keying material. Nonce derivation is covered in Section 2.3.

The additional data passed to each invocation of AEAD\_AES\_128\_GCM is a zero-length octet sequence.

A consequence of this record structure is that range requests [RFC7233] and random access to encrypted payload bodies are possible at the granularity of the record size. Partial records at the ends of a range cannot be decrypted. Thus, it is best if range requests start and end on record boundaries.

Selecting the record size most appropriate for a given situation requires a trade-off. A smaller record size allows decrypted octets to be released more rapidly, which can be appropriate for applications that depend on responsiveness. Smaller records also reduce the additional data required if random access into the ciphertext is needed. Applications that depend on being able to pad by arbitrary amounts cannot increase the record size beyond 65537 octets.

Applications that don't depending on streaming, random access, or arbitrary padding can use larger records, or even a single record. A larger record size reduces the processing and data overheads.

### **2.1**. Encryption Content Coding Header

The content coding uses a header block that includes all parameters needed to decrypt the content (other than the key). The header block is placed in the body of a message ahead of the sequence of records.

+----+ | salt (16) | rs (4) | idlen (1) | keyid (idlen) | +----+

salt: The "salt" parameter comprises the first 16 octets of the "aes128gcm" content coding header. The same "salt" parameter value MUST NOT be reused for two different payload bodies that have the same input keying material; generating a random salt for every application of the content coding ensures that content encryption key reuse is highly unlikely.

[Page 5]

- rs: The "rs" or record size parameter contains an unsigned 32-bit integer in network byte order that describes the record size in octets. Note that it is therefore impossible to exceed the 2^36-31 limit on plaintext input to AEAD\_AES\_128\_GCM. Values smaller than 3 are invalid.
- keyid: The "keyid" parameter can be used to identify the keying material that is used. When the Crypto-Key header field is used, the "keyid" identifies a matching value in that field. The "keyid" parameter MUST be used if keying material included in an Crypto-Key header field is needed to derive the content encryption key. The "keyid" parameter can also be used to identify keys in an application-specific fashion.

### **2.2**. Content Encryption Key Derivation

In order to allow the reuse of keying material for multiple different HTTP messages, a content encryption key is derived for each message. The content encryption key is derived from the decoded value of the "salt" parameter using the HMAC-based key derivation function (HKDF) described in [<u>RFC5869</u>] using the SHA-256 hash algorithm [<u>FIPS180-4</u>].

The value of the "salt" parameter is the salt input to HKDF function. The keying material identified by the "keyid" parameter is the input keying material (IKM) to HKDF. Input keying material can either be prearranged, or can be described using the Crypto-Key header field (<u>Section 3</u>). The extract phase of HKDF therefore produces a pseudorandom key (PRK) as follows:

PRK = HMAC-SHA-256(salt, IKM)

The info parameter to HKDF is set to the ASCII-encoded string "Content-Encoding: aes128gcm" and a single zero octet:

cek\_info = "Content-Encoding: aes128gcm" || 0x00

Note: Concatenation of octet sequences is represented by the "||" operator.

AEAD\_AES\_128\_GCM requires a 16 octet (128 bit) content encryption key (CEK), so the length (L) parameter to HKDF is 16. The second step of HKDF can therefore be simplified to the first 16 octets of a single HMAC:

CEK = HMAC-SHA-256(PRK, cek\_info || 0x01)

[Page 6]

### **<u>2.3</u>**. Nonce Derivation

The nonce input to AEAD\_AES\_128\_GCM is constructed for each record. The nonce for each record is a 12 octet (96 bit) value that is produced from the record sequence number and a value derived from the input keying material.

The input keying material and salt values are input to HKDF with different info and length parameters.

The length (L) parameter is 12 octets. The info parameter for the nonce is the ASCII-encoded string "Content-Encoding: nonce", terminated by a a single zero octet:

```
nonce_info = "Content-Encoding: nonce" || 0x00
```

The result is combined with the record sequence number - using exclusive or - to produce the nonce. The record sequence number (SEQ) is a 96-bit unsigned integer in network byte order that starts at zero.

Thus, the final nonce for each record is a 12 octet value:

NONCE = HMAC-SHA-256(PRK, nonce\_info || 0x01) XOR SEQ

This nonce construction prevents removal or reordering of records. However, it permits truncation of the tail of the sequence (see <u>Section 2</u> for how this is avoided).

### **<u>3</u>**. Crypto-Key Header Field

A Crypto-Key header field can be used to describe the input keying material used by the "aes128gcm" content coding.

Ordinarily, this header field will not appear in the same message as the encrypted content. Including the encryption key with the encrypted payload reduces the value of using encryption to a somewhat complicated checksum. However, the Crypto-Key header field could be used in one message to provision keys for other messages.

The Crypto-Key header field uses the extended ABNF syntax defined in <u>Section 1.2 of [RFC7230]</u> and the "parameter" and "OWS" rules from [<u>RFC7231</u>].

Expires May 4, 2017 [Page 7]

Crypto-Key = #crypto-key-params crypto-key-params = [ parameter \*( OWS ";" OWS parameter ) ]

- keyid: The "keyid" parameter corresponds to the "keyid" parameter in the content coding.
- aes128gcm: The "aes128gcm" parameter contains the base64url-encoded octets [<u>RFC7515</u>] of the input keying material for the "aes128gcm" content coding.

Crypto-Key header field values with multiple instances of the same parameter name in a single crypto-key-params production are invalid.

The input keying material used by the key derivation (see <u>Section 2.2</u>) can be determined based on the information in the Crypto-Key header field.

The value or values provided in the Crypto-Key header field is valid only for the current HTTP message unless additional information indicates a greater scope.

Alternative methods for determining input keying material MAY be defined by specifications that use this content coding. This document only defines the use of the "aes128gcm" parameter which describes an explicit key.

The "aes128gcm" parameter MUST decode to at least 16 octets in order to be used as input keying material for "aes128gcm" content coding.

#### **<u>4</u>**. Examples

This section shows a few examples of the encrypted content coding.

Note: All binary values in the examples in this section use base64url encoding [<u>RFC7515</u>]. This includes the bodies of requests. Whitespace and line wrapping is added to fit formatting constraints.

### <u>4.1</u>. Encryption of a Response

Here, a successful HTTP GET response has been encrypted using input keying material that is identified by the string "a1".

The encrypted data in this example is the UTF-8 encoded string "I am the walrus". The input keying material is included in the Crypto-Key header field. The content body contains a single record only and is shown here using base64url encoding for presentation reasons.

[Page 8]

Internet-Draft

HTTP/1.1 200 OK Content-Type: application/octet-stream Content-Length: 33 Content-Encoding: aes128gcm Crypto-Key: aes128gcm=B33e\_VeFr0yIHwFTIfmesA

9Y1iaZMzICC05D03y8dWiAAAopoAzpM9l8LHdpDa09C-UvT4kttTI\_edSsHv1o5b lWZ5mBYL

Note that the media type has been changed to "application/octetstream" to avoid exposing information about the content. Alternatively (and equivalently), the Content-Type header field can be omitted.

### 4.2. Encryption with Multiple Records

This example shows the same encrypted message, but split into records of 10 octets each. The first record includes a single additional octet of padding, which causes the end of the content to align with a record boundary, forcing the creation of a third record that contains only padding.

HTTP/1.1 200 OK Content-Length: 70 Content-Encoding: aes128gcm Crypto-Key: keyid="a1"; aes128gcm="B03ZVPxUlnLORbVGMpbT1Q"

\_lgOPHdbKmIaLnZC7\_8huQAAAAoCYTGkQWUSYylMKzMduBHDCFDwL2oODx8nkh0n uOTNrh48DaWSm02DiQPzQAOGe6xRAeBj588hH6jQRTh\_szFRS2Nwx9Aeuiic

## **<u>5</u>**. Security Considerations

This mechanism assumes the presence of a key management framework that is used to manage the distribution of keys between valid senders and receivers. Defining key management is part of composing this mechanism into a larger application, protocol, or framework.

Implementation of cryptography - and key management in particular can be difficult. For instance, implementations need to account for the potential for exposing keying material on side channels, such as might be exposed by the time it takes to perform a given operation. The requirements for a good implementation of cryptographic algorithms can change over time.

[Page 9]

#### 5.1. Key and Nonce Reuse

Encrypting different plaintext with the same content encryption key and nonce in AES-GCM is not safe [<u>RFC5116</u>]. The scheme defined here uses a fixed progression of nonce values. Thus, a new content encryption key is needed for every application of the content coding. Since input keying material can be reused, a unique "salt" parameter is needed to ensure a content encryption key is not reused.

If a content encryption key is reused - that is, if input keying material and salt are reused - this could expose the plaintext and the authentication key, nullifying the protection offered by encryption. Thus, if the same input keying material is reused, then the salt parameter MUST be unique each time. This ensures that the content encryption key is not reused. An implementation SHOULD generate a random salt parameter for every message; a counter could achieve the same result.

### 5.2. Data Encryption Limits

There are limits to the data that AEAD\_AES\_128\_GCM can encipher. The maximum value for the record size is limited by the size of the "rs" field in the header (see <u>Section 2.1</u>), which ensures that the 2^36-31 limit for a single application of AEAD\_AES\_128\_GCM is not reached [<u>RFC5116</u>]. In order to preserve a 2^-40 probability of indistinguishability under chosen plaintext attack (IND-CPA), the total amount of plaintext that can be enciphered MUST be less than 2^44.5 blocks of 16 octets [<u>AEBounds</u>].

If rs is a multiple of 16 octets, this means 398 terabytes can be encrypted safely, including padding. However, if the record size is not a multiple of 16 octets, the total amount of data that can be safely encrypted is reduced proportionally. The worst case is a record size of 3 octets, for which at most 74 terabytes of plaintext can be encrypted, of which at least two-thirds is padding.

### **<u>5.3</u>**. Content Integrity

This mechanism only provides content origin authentication. The authentication tag only ensures that an entity with access to the content encryption key produced the encrypted data.

Any entity with the content encryption key can therefore produce content that will be accepted as valid. This includes all recipients of the same HTTP message.

Furthermore, any entity that is able to modify both the Encryption header field and the HTTP message body can replace the contents.

Without the content encryption key or the input keying material, modifications to or replacement of parts of a payload body are not possible.

#### **5.4**. Leaking Information in Headers

Because only the payload body is encrypted, information exposed in header fields is visible to anyone who can read the HTTP message. This could expose side-channel information.

For example, the Content-Type header field can leak information about the payload body.

There are a number of strategies available to mitigate this threat, depending upon the application's threat model and the users' tolerance for leaked information:

- Determine that it is not an issue. For example, if it is expected that all content stored will be "application/json", or another very common media type, exposing the Content-Type header field could be an acceptable risk.
- 2. If it is considered sensitive information and it is possible to determine it through other means (e.g., out of band, using hints in other representations, etc.), omit the relevant headers, and/ or normalize them. In the case of Content-Type, this could be accomplished by always sending Content-Type: application/octetstream (the most generic media type), or no Content-Type at all.
- 3. If it is considered sensitive information and it is not possible to convey it elsewhere, encapsulate the HTTP message using the application/http media type (<u>Section 8.3.2 of [RFC7230]</u>), encrypting that as the payload of the "outer" message.

## **<u>5.5</u>**. Poisoning Storage

This mechanism only offers encryption of content; it does not perform authentication or authorization, which still needs to be performed (e.g., by HTTP authentication [<u>RFC7235</u>]).

This is especially relevant when a HTTP PUT request is accepted by a server; if the request is unauthenticated, it becomes possible for a third party to deny service and/or poison the store.

Expires May 4, 2017 [Page 11]

## **<u>5.6</u>**. Sizing and Timing Attacks

Applications using this mechanism need to be aware that the size of encrypted messages, as well as their timing, HTTP methods, URIs and so on, may leak sensitive information.

This risk can be mitigated through the use of the padding that this mechanism provides. Alternatively, splitting up content into segments and storing the separately might reduce exposure. HTTP/2 [RFC7540] combined with TLS [RFC5246] might be used to hide the size of individual messages.

### <u>6</u>. IANA Considerations

### 6.1. The "aes128gcm" HTTP Content Coding

This memo registers the "aes128gcm" HTTP content coding in the HTTP Content Codings Registry, as detailed in <u>Section 2</u>.

- o Name: aes128gcm
- o Description: AES-GCM encryption with a 128-bit content encryption
  key
- o Reference: this specification

### 6.2. Crypto-Key Header Field

This memo registers the "Crypto-Key" HTTP header field in the Permanent Message Header Registry, as detailed in <u>Section 3</u>.

- o Field name: Crypto-Key
- o Protocol: HTTP
- o Status: Standard
- o Reference: this specification
- o Notes:

## 6.3. The HTTP Crypto-Key Parameter Registry

This memo establishes a registry for parameters used by the "Crypto-Key" header field under the "Hypertext Transfer Protocol (HTTP) Parameters" grouping. The "Hypertext Transfer Protocol (HTTP) Crypto-Key Parameters" operates under an "Specification Required" policy [RFC5226].

Expires May 4, 2017 [Page 12]

Entries in this registry are expected to include the following information:

- o Parameter Name: The name of the parameter.
- o Purpose: A brief description of the purpose of the parameter.
- o Reference: A reference to a specification that defines the semantics of the parameter.

The initial contents of this registry are:

### <u>6.3.1</u>. keyid

- o Parameter Name: keyid
- o Purpose: Identify the key that is in use.
- o Reference: this document

#### 6.3.2. aes128gcm

- o Parameter Name: aes128gcm
- Purpose: Provide an explicit input keying material value for the aes128gcm content coding.
- o Reference: this document

### 7. References

### 7.1. Normative References

[FIPS180-4]

Department of Commerce, National., "NIST FIPS 180-4, Secure Hash Standard", March 2012, <<u>http://csrc.nist.gov/publications/fips/fips180-4/</u> fips-180-4.pdf>.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <u>BCP 14</u>, <u>RFC 2119</u>, DOI 10.17487/RFC2119, March 1997, <<u>http://www.rfc-editor.org/info/rfc2119</u>>.
- [RFC5116] McGrew, D., "An Interface and Algorithms for Authenticated Encryption", <u>RFC 5116</u>, DOI 10.17487/RFC5116, January 2008, <<u>http://www.rfc-editor.org/info/rfc5116</u>>.

Expires May 4, 2017 [Page 13]

- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", <u>BCP 26</u>, <u>RFC 5226</u>, DOI 10.17487/RFC5226, May 2008, <<u>http://www.rfc-editor.org/info/rfc5226</u>>.
- [RFC5869] Krawczyk, H. and P. Eronen, "HMAC-based Extract-and-Expand Key Derivation Function (HKDF)", <u>RFC 5869</u>, DOI 10.17487/RFC5869, May 2010, <<u>http://www.rfc-editor.org/info/rfc5869</u>>.
- [RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", <u>RFC 7230</u>, DOI 10.17487/RFC7230, June 2014, <<u>http://www.rfc-editor.org/info/rfc7230</u>>.
- [RFC7231] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", <u>RFC 7231</u>, DOI 10.17487/RFC7231, June 2014, <<u>http://www.rfc-editor.org/info/rfc7231</u>>.
- [RFC7515] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Signature (JWS)", <u>RFC 7515</u>, DOI 10.17487/RFC7515, May 2015, <<u>http://www.rfc-editor.org/info/rfc7515</u>>.

## <u>7.2</u>. Informative References

#### [AEBounds]

- Luykx, A. and K. Paterson, "Limits on Authenticated Encryption Use in TLS", March 2016, <<u>http://www.isg.rhul.ac.uk/~kp/TLS-AEbounds.pdf</u>>.
- [RFC4880] Callas, J., Donnerhacke, L., Finney, H., Shaw, D., and R. Thayer, "OpenPGP Message Format", <u>RFC 4880</u>, DOI 10.17487/RFC4880, November 2007, <<u>http://www.rfc-editor.org/info/rfc4880</u>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", <u>RFC 5246</u>, DOI 10.17487/RFC5246, August 2008, <<u>http://www.rfc-editor.org/info/rfc5246</u>>.

- [RFC7233] Fielding, R., Ed., Lafon, Y., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Range Requests", <u>RFC 7233</u>, DOI 10.17487/RFC7233, June 2014, <<u>http://www.rfc-editor.org/info/rfc7233</u>>.
- [RFC7235] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Authentication", <u>RFC 7235</u>, DOI 10.17487/RFC7235, June 2014, <<u>http://www.rfc-editor.org/info/rfc7235</u>>.
- [RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", <u>RFC 7540</u>, DOI 10.17487/RFC7540, May 2015, <<u>http://www.rfc-editor.org/info/rfc7540</u>>.
- [XMLENC] Eastlake, D., Reagle, J., Hirsch, F., Roessler, T., Imamura, T., Dillaway, B., Simon, E., Yiu, K., and M. Nystroem, "XML Encryption Syntax and Processing", W3C Recommendation REC-xmlenc-core1-20130411 , January 2013, <<u>https://www.w3.org/TR/2013/REC-xmlenc-core1-20130411</u>>.

### <u>Appendix A</u>. JWE Mapping

The "aes128gcm" content coding can be considered as a sequence of JSON Web Encryption (JWE) objects [RFC7516], each corresponding to a single fixed size record that includes leading padding. The following transformations are applied to a JWE object that might be expressed using the JWE Compact Serialization:

- o The JWE Protected Header is fixed to the value { "alg": "dir", "enc": "A128GCM" }, describing direct encryption using AES-GCM with a 128-bit content encryption key. This header is not transmitted, it is instead implied by the value of the Content-Encoding header field.
- o The JWE Encrypted Key is empty, as stipulated by the direct encryption algorithm.
- o The JWE Initialization Vector ("iv") for each record is set to the exclusive or of the 96-bit record sequence number, starting at zero, and a value derived from the input keying material (see <u>Section 2.3</u>). This value is also not transmitted.

 The final value is the concatenated header, JWE Ciphertext, and JWE Authentication Tag, all expressed without base64url encoding. The "." separator is omitted, since the length of these fields is known.

Thus, the example in <u>Section 4.1</u> can be rendered using the JWE Compact Serialization as:

eyAiYWxnIjogImRpciIsICJlbmMiOiAiQTEyOEdDTSIgfQ..31iQYc1v4a36EgyJ. AM6TPZfCx3aQ2jvQvlL0-JLb.21Mj951Kwe\_WjluVZnmYFgs

Where the first line represents the fixed JWE Protected Header, an empty JWE Encrypted Key, and the algorithmically-determined JWE Initialization Vector. The second line contains the encoded body, split into JWE Ciphertext and JWE Authentication Tag.

### Appendix B. Acknowledgements

Mark Nottingham was an original author of this document.

The following people provided valuable input: Richard Barnes, David Benjamin, Peter Beverloo, JR Conlin, Mike Jones, Stephen Farrell, Adam Langley, John Mattsson, Julian Reschke, Eric Rescorla, Jim Schaad, and Magnus Westerlund.

Author's Address

Martin Thomson Mozilla

Email: martin.thomson@gmail.com

Expires May 4, 2017 [Page 16]