

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: May 4, 2017

A. Bittau
Google
D. Boneh
D. Giffin
Stanford University
M. Handley
University College London
D. Mazieres
Stanford University
E. Smith
Kestrel Institute
October 31, 2016

TCP-ENO: Encryption Negotiation Option
draft-ietf-tcpinc-tcpeno-06

Abstract

Despite growing adoption of TLS [[RFC5246](#)], a significant fraction of TCP traffic on the Internet remains unencrypted. The persistence of unencrypted traffic can be attributed to at least two factors. First, some legacy protocols lack a signaling mechanism (such as a "STARTTLS" command) by which to convey support for encryption, making incremental deployment impossible. Second, legacy applications themselves cannot always be upgraded, requiring a way to implement encryption transparently entirely within the transport layer. The TCP Encryption Negotiation Option (TCP-ENO) addresses both of these problems through a new TCP option kind providing out-of-band, fully backward-compatible negotiation of encryption.

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 <http://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 May 4, 2017.

Internet-Draft

tcpeno

October 2016

Copyright Notice

Copyright (c) 2016 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 (<http://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.	Requirements language	3
2.	Introduction	3
2.1.	Design goals	3
3.	Terminology	4
4.	TCP-ENO specification	5
4.1.	ENO option	6
4.2.	The global suboption	9
4.3.	TCP-ENO roles	10
4.4.	Specifying suboption data length	10
4.5.	The negotiated TEP	12
4.6.	TCP-ENO handshake	12
4.7.	Data in SYN segments	13
4.8.	Negotiation transcript	15
5.	Requirements for TEPs	15
5.1.	Session IDs	16
6.	Examples	18
7.	Design rationale	19
7.1.	Future developments	19
7.2.	Handshake robustness	20
7.3.	Suboption data	21
7.4.	Passive role bit	21
7.5.	Option kind sharing	21
8.	Experiments	22
9.	Security considerations	22
10.	IANA Considerations	23
11.	Acknowledgments	24

12.	References	25
12.1.	Normative References	25
12.2.	Informative References	25
Authors'	Addresses	26

Internet-Draft

tcpeno

October 2016

[1.](#) Requirements language

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 [[RFC2119](#)].

[2.](#) Introduction

Many applications and protocols running on top of TCP today do not encrypt traffic. This failure to encrypt lowers the bar for certain attacks, harming both user privacy and system security. Counteracting the problem demands a minimally intrusive, backward-compatible mechanism for incrementally deploying encryption. The TCP Encryption Negotiation Option (TCP-ENO) specified in this document provides such a mechanism.

Introducing TCP options, extending operating system interfaces to support TCP-level encryption, and extending applications to take advantage of TCP-level encryption all require effort. To the greatest extent possible, the effort invested in realizing TCP-level encryption today needs to remain applicable in the future should the need arise to change encryption strategies. To this end, it is useful to consider two questions separately:

1. How to negotiate the use of encryption at the TCP layer, and
2. How to perform encryption at the TCP layer.

This document addresses question 1 with a new TCP option, ENO. TCP-ENO provides a framework in which two endpoints can agree on one among multiple possible TCP encryption protocols or `_TEPs_`. For future compatibility, TEPs can vary widely in terms of wire format, use of TCP option space, and integration with the TCP header and segmentation. However, ENO abstracts these differences to ensure the introduction of new TEPs can be transparent to applications taking advantage of TCP-level encryption.

Question 2 is addressed by one or more companion TEP specification documents. While current TEPs enable TCP-level traffic encryption today, TCP-ENO ensures that the effort invested to deploy today's TEPs will additionally benefit future ones.

[2.1.](#) Design goals

TCP-ENO was designed to achieve the following goals:

1. Enable endpoints to negotiate the use of a separately specified TCP encryption protocol or `_TEP_`.

Bittau, et al.

Expires May 4, 2017

[Page 3]

Internet-Draft

tcpeno

October 2016

2. Transparently fall back to unencrypted TCP when not supported by both endpoints.
3. Provide out-of-band signaling through which applications can better take advantage of TCP-level encryption (for instance, by improving authentication mechanisms in the presence of TCP-level encryption).
4. Provide a standard negotiation transcript through which TEPs can defend against tampering with TCP-ENO.
5. Make parsimonious use of TCP option space.
6. Define roles for the two ends of a TCP connection, so as to name each end of a connection for encryption or authentication purposes even following a symmetric simultaneous open.

[3.](#) Terminology

We define the following terms, which are used throughout this document:

SYN segment

A TCP segment in which the SYN flag is set

ACK segment

A TCP segment in which the ACK flag is set (which includes most segments other than an initial SYN segment)

non-SYN segment

A TCP segment in which the SYN flag is clear

SYN-only segment

A TCP segment in which the SYN flag is set but the ACK flag is clear

SYN-ACK segment

A TCP segment in which the SYN and ACK flags are both set

Active opener

A host that initiates a connection by sending a SYN-only segment. With the BSD socket API, an active opener calls "connect". In client-server configurations, active openers are typically clients.

Passive opener

A host that does not send a SYN-only segment, but responds to one with a SYN-ACK segment. With the BSD socket API, passive openers

call "listen" and "accept", rather than "connect". In client-server configurations, passive openers are typically servers.

Simultaneous open

The act of symmetrically establishing a TCP connection between two active openers (both of which call "connect" with BSD sockets). Each host of a simultaneous open sends both a SYN-only and a SYN-ACK segment. Simultaneous open is less common than asymmetric open with one active and one passive opener, but can be used for NAT traversal by peer-to-peer applications [[RFC5382](#)].

TEP

A TCP encryption protocol intended for use with TCP-ENO and specified in a separate document.

TEP identifier

A unique 7-bit value in the range 0x20-0x7f that IANA has assigned to a TEP.

Negotiated TEP

The single TEP governing a TCP connection, determined by use of the TCP ENO option specified in this document.

[4.](#) TCP-ENO specification

TCP-ENO extends TCP connection establishment to enable encryption opportunistically. It uses a new TCP option kind to negotiate one among multiple possible TCP encryption protocols or TEPs. The negotiation involves hosts exchanging sets of supported TEPs, where each TEP is represented by a `_suboption_` within a larger TCP ENO option in the offering host's SYN segment.

If TCP-ENO succeeds, it yields the following information:

- o A negotiated TEP, represented by a unique 7-bit TEP identifier,
- o A few extra bytes of suboption data from each host, if needed by the TEP,
- o A negotiation transcript with which to mitigate attacks on the negotiation itself,
- o Role assignments designating one endpoint "host A" and the other endpoint "host B", and
- o A bit indicating whether or not the application at each end knows it is using TCP-ENO.

If TCP-ENO fails, encryption is disabled and the connection falls back to traditional unencrypted TCP.

The remainder of this section provides the normative description of the TCP ENO option and handshake protocol.

[4.1.](#) ENO option

TCP-ENO employs an option in the TCP header [[RFC0793](#)]. There are two equivalent kinds of ENO option, shown in Figure 1. [Section 10](#) specifies which of the two kinds is permissible and/or preferred.

byte	0	1	2		N+1	(N+2 bytes total)
	+-----+-----+-----+---...---+-----+					
	Kind= Len=					

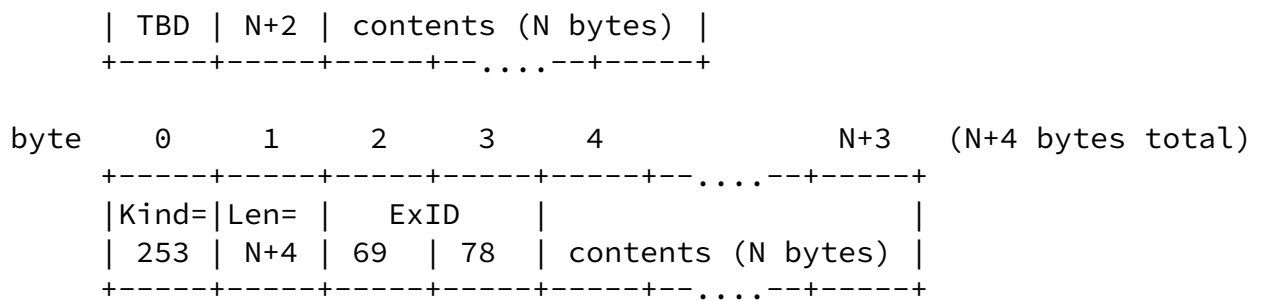



Figure 1: Two equivalent kinds of TCP-ENO option

The contents of an ENO option can take one of two forms. A SYN form, illustrated in Figure 2, appears only in SYN segments. A non-SYN form, illustrated in Figure 3, appears only in non-SYN segments. The SYN form of ENO acts as a container for zero or more suboptions, labeled "Opt_0", "Opt_1", ... in Figure 2. The non-SYN form, by its presence, acts as a one-bit acknowledgment, with the actual contents ignored by ENO. Particular TEPs MAY assign additional meaning to the contents of non-SYN ENO options. When a negotiated TEP does not assign such meaning, the contents of a non-SYN ENO option MUST be zero bytes.






Figure 2: SYN form of ENO




Figure 3: Non-SYN form of ENO, where N MAY be 0

Every suboption starts with a byte of the form illustrated in Figure 4. The high bit "v", when set, introduces suboptions with variable-length data. When "v = 0", the byte itself constitutes the entirety of the suboption. The 7-bit value "glt" expresses one of:

- o Global configuration data (discussed in [Section 4.2](#)),
- o Suboption data length for the next suboption (discussed in [Section 4.4](#)), or
- o An offer to use a particular TEP defined in a separate TEP specification document.

4.2. The global suboption

Suboptions 0x00-0x1f are used for global configuration that applies regardless of the negotiated TEP. A TCP SYN segment **MUST** include at most one ENO suboption in this range. A receiver **MUST** ignore all but the first suboption in this range so as to anticipate updates to ENO that assign new meaning to bits in subsequent global suboptions. The value of a global suboption byte is interpreted as a bitmask, illustrated in Figure 5.

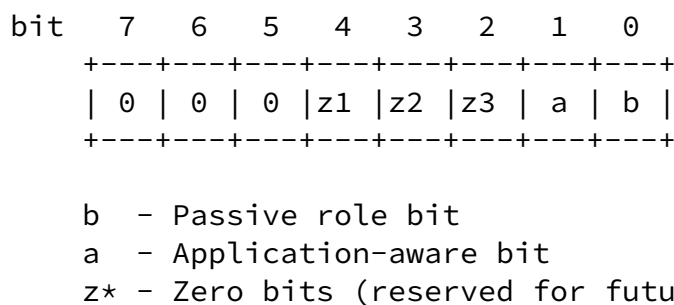


Figure 5: Format of the global suboption byte

The fields of the bitmask are interpreted as follows:

b

The passive role bit **MUST** be 1 for all passive openers. For active openers, it **MUST** default to 0, but implementations **MUST** provide an API through which an application can set "b = 1" before initiating an active open. (Manual configuration of "b" is necessary for simultaneous open.)

a

The application-aware bit "a" is an out-of-band signal indicating that the application on the sending host is aware of TCP-ENO and has been extended to alter its behavior in the presence of encrypted TCP. Implementations **MUST** set this bit to 0 by default, and **SHOULD** provide an API through which applications can change the value of the bit as well as examine the value of the bit sent by the remote host. Implementations **SHOULD** furthermore support a `_mandatory_` application-aware mode in which TCP-ENO is automatically disabled if the remote host does not set "a = 1".

z1, z2, z3

The "z" bits are reserved for future updates to TCP-ENO. They **MUST** be set to zero in sent segments and **MUST** be ignored in received segments.

A SYN segment without an explicit global suboption has an implicit global suboption of 0x00. Because passive openers **MUST** always set "b

Internet-Draft

tcpeno

October 2016

= 1", they cannot rely on this implicit 0x00 byte and **MUST** include an explicit global suboption in their SYN-ACK segments.

[4.3.](#) TCP-ENO roles

TCP-ENO uses abstract roles to distinguish the two ends of a TCP connection. These roles are determined by the "b" bit in the global suboption. The host that sent an implicit or explicit suboption with "b = 0" plays the "A" role. The host that sent "b = 1" plays the "B" role.

If both sides of a connection set "b = 1" (which can happen if the active opener misconfigures "b" before calling "connect"), or both sides set "b = 0" (which can happen with simultaneous open), then TCP-ENO **MUST** be disabled and the connection **MUST** fall back to unencrypted TCP.


TEP specifications **SHOULD** refer to TCP-ENO's A and B roles to specify asymmetric behavior by the two hosts. For the remainder of this document, we will use the terms "host A" and "host B" to designate the hosts with roles A and B, respectively, in a connection.

[4.4.](#) Specifying suboption data length

A TEP **MAY** optionally make use of one or more bytes of suboption data. The presence of such data is indicated by setting "v = 1" in the initial suboption byte (see Figure 4). By default, suboption data extends to the end of the TCP option. Hence, if only one suboption requires data, the most compact way to encode it is to place it last in the ENO option, after all other suboptions. As an example, in Figure 2, the last suboption, "Opt_i", has suboption data and thus requires "v = 1"; however, the suboption data length can be inferred from the total length of the TCP option.

When a suboption with data is not last in an ENO option, the sender **MUST** explicitly specify the suboption data length for the receiver to know where the next suboption starts. The sender does so by preceding the suboption with a length byte, depicted in Figure 6. The length byte encodes a 5-bit value "nnnnn". Adding one to "nnnnn"


yields the length of the suboption data (not including the length byte or the TEP identifier). Hence, a length byte can designate anywhere from 1 to 32 bytes of suboption data (inclusive).



nnnnn - 5-bit value encoding (length - 1)

Figure 6: Format of a length byte

A suboption preceded by a length byte MUST be a TEP identifier ("glt >= 0x20") and MUST have "v = 1". Figure 7 shows an example of such a suboption.



length byte - specifies nnnnn
 TEP identifier - MUST have v = 1 and glt >= 0x20
 suboption data - length specified by nnnnn+1

Figure 7: Suboption with length byte

A host MUST ignore an ENO option in a SYN segment and MUST disable encryption if either:

1. A length byte indicates that suboption data would extend beyond the end of the TCP ENO option, or
2. A length byte is followed by an octet in the range 0x00-0x9f

(meaning the following byte has "v = 0" or "glt < 0x20").

Because the last suboption in an ENO option is special-cased to have its length inferred from the 8-bit TCP option length, it MAY contain more than 32 bytes of suboption data. Other suboptions are limited to 32 bytes by the length byte format. The TCP header itself can only accommodate a maximum of 40 bytes of options per segment, however, so regardless of the length byte could not fit more than one suboption over 32 bytes. That said, TEPs MAY define the use of multiple suboptions with the same TEP identifier in the same SYN segment, providing another way to convey over 32 bytes of suboption data even with length bytes.

[4.5.](#) The negotiated TEP

A TEP identifier "glt" (with "glt >= 0x20") is `_valid_` for a connection when:

1. Each side has sent a suboption for "glt" in its SYN-form ENO option,
2. Any suboption data in these "glt" suboptions is valid according to the TEP specification and satisfies any runtime constraints, and
3. If an ENO option contains multiple suboptions with "glt", then such repetition is well-defined by the TEP specification.

The `_negotiated TEP_` is the last valid TEP identifier in host B's SYN-form ENO option. This definition means host B specifies TEP suboptions in order of increasing priority, while host A does not influence TEP priority.

A passive opener (which is always host B) sees the remote host's SYN segment before constructing its own SYN-ACK. Hence, a passive opener SHOULD include only one TEP identifier in SYN-ACK segments and SHOULD ensure this TEP identifier is valid. However, simultaneous open or implementation considerations can prevent host B from offering only

one TEP.

4.6. TCP-ENO handshake

A host employing TCP-ENO for a connection MUST include an ENO option in every TCP segment sent until either encryption is disabled or the host receives a non-SYN segment.

A host MUST disable encryption, refrain from sending any further ENO options, and fall back to unencrypted TCP if any of the following occurs:

1. Any segment it receives up to and including the first received ACK segment does not contain a ENO option (or contains an ill-formed SYN-form ENO option),
2. The SYN segment it receives does not contain a valid TEP identifier, or
3. It receives a SYN segment with an incompatible global suboption. (Specifically, incompatible means the two hosts set the same "b" value or the connection is in mandatory application-aware mode and the remote host set "a = 0".)

Hosts MUST NOT alter SYN-form ENO options in retransmitted segments, or between the SYN and SYN-ACK segments of a simultaneous open, with two exceptions for an active opener. First, an active opener MAY unilaterally disable ENO (and thus remove the ENO option) between retransmissions of a SYN-only segment. (Such removal could be useful if middleboxes are dropping segments with the ENO option.) Second, an active opener performing simultaneous open MAY include no TCP-ENO option in its SYN-ACK if the received SYN caused it to disable encryption according to the above rules (for instance because role negotiation failed).

Once a host has both sent and received an ACK segment containing an ENO option, encryption MUST be enabled. Once encryption is enabled, hosts MUST follow the specification of the negotiated TEP and MUST NOT present raw TCP payload data to the application. In particular, data segments MUST NOT contain plaintext application data, but rather ciphertext, key negotiation parameters, or other messages as determined by the negotiated TEP.

[4.7.](#) Data in SYN segments

TEPs MAY specify the use of data in SYN segments so as to reduce the number of round trips required for connection setup. The meaning of data in a SYN segment with an ENO option (a SYN+ENO segment) is determined by the last TEP identifier in the ENO option, which we term the segment's `_SYN TEP_`.

A host sending a SYN+ENO segment MUST NOT include data in the segment unless the SYN TEP's specification defines the use of such data. Furthermore, to avoid conflicting interpretations of SYN data, a SYN+ENO option MUST NOT include a non-empty TCP Fast Open (TFO) option [[RFC7413](#)].

Because a host can send SYN data before knowing which if any TEP will govern a connection, hosts implementing ENO are REQUIRED to discard data from SYN+ENO segments when the SYN TEP does not govern the connection or when there is any ambiguity over the meaning of the SYN data. This requirement applies to hosts that implement ENO even when ENO has been disabled by configuration. However, note that discarding SYN data is already common practice [[RFC4987](#)] and the new requirement applies only to segments containing ENO options.

More specifically, a host that implements ENO MUST discard the data in a received SYN+ENO segment if any of the following applies:

- o ENO fails and TEP-indicated encryption is disabled for the connection,

- o The received segment's SYN TEP is not the negotiated TEP,
- o The negotiated TEP does not define the use of SYN data, or
- o The SYN segment contains a non-empty TFO option or any other TCP option implying a conflicting definition of SYN data.

A host discarding SYN data in compliance with the above requirement MUST NOT acknowledge the sequence number of the discarded data, but rather MUST acknowledge the other host's initial sequence number as if the received SYN segment contained no data. Furthermore, after

discarding SYN data, such a host MUST NOT assume the SYN data will be identically retransmitted, and MUST process data only from non-SYN segments.

If a host sends a SYN+ENO segment with data and receives acknowledgment for the data, but the SYN TEP governing the data is not the negotiated TEP (either because a different TEP was negotiated or because ENO failed to negotiate encryption), then the host MUST reset the TCP connection. Proceeding in any other fashion risks misinterpreted SYN data.

If a host sends a SYN-only SYN+ENO segment bearing data and subsequently receives a SYN-ACK segment without an ENO option, that host MUST reset the connection even if the SYN-ACK segment does not acknowledge the SYN data. The issue is that unacknowledged data may nonetheless have been cached by the receiver; later retransmissions intended to supersede this unacknowledged data could fail to do so if the receiver gives precedence to the cached original data. Implementations MAY provide an API call for a non-default mode in which unacknowledged SYN data does not cause a connection reset, but applications MUST only use this mode when a higher-layer integrity check would anyway terminate a garbled connection.

To avoid unexpected connection resets, ENO implementations MUST disable the use of data in SYN-only segments by default. Such data MAY be enabled by an API command. In particular, implementations MAY provide a per-connection mandatory encryption mode that automatically resets a connection if ENO fails, and MAY enable SYN data in this mode.

To satisfy the requirement of the previous paragraph, all TEPs SHOULD support a normal mode of operation that avoids data in SYN-only segments. An exception is TEPs intended to be disabled by default.

[4.8.](#) Negotiation transcript

To defend against attacks on encryption negotiation itself, TEPs need a way to reference a transcript of TCP-ENO's negotiation. In

particular, a TEP MUST with high probability fail to reach key agreement between two honest endpoints if the TEP's selection resulted from tampering with the contents of SYN-form ENO options. (Of course, in the absence of endpoint authentication, two honest endpoints can still each end up talking to a man-in-the-middle attacker rather than to each other.)

TCP-ENO defines its negotiation transcript as a packed data structure consisting of two TCP-ENO options exactly as they appeared in the TCP header (including the TCP option kind, TCP option length byte, and, for option kind 253, the bytes 69 and 78 as illustrated in Figure 1). The transcript is constructed from the following, in order:

1. The TCP-ENO option in host A's SYN segment, including the kind and length bytes.
2. The TCP-ENO option in host B's SYN segment, including the kind and length bytes.

Note that because the ENO options in the transcript contain length bytes as specified by TCP, the transcript unambiguously delimits A's and B's ENO options.

[5.](#) Requirements for TEPs

TCP-ENO affords TEP specifications a large amount of design flexibility. However, to abstract TEP differences away from applications requires fitting them all into a coherent framework. As such, any TEP claiming an ENO TEP identifier MUST satisfy the following normative list of properties.

- o TEPs MUST protect TCP data streams with authenticated encryption.
- o TEPs MUST define a session ID whose value identifies the TCP connection and, with overwhelming probability, is unique over all time if either host correctly obeys the TEP. [Section 5.1](#) describes the requirements of the session ID in more detail.
- o TEPs MUST NOT permit the negotiation of any encryption algorithms with significantly less than 128-bit security.
- o TEPs MUST NOT allow the negotiation of null cipher suites, even for debugging purposes. (Implementations MAY support debugging modes that allow applications to extract their own session keys.)

- o TEPs MUST NOT depend on long-lived secrets for data confidentiality, as implementations SHOULD provide forward secrecy some bounded, short time after the close of a TCP connection. (Exceptions to forward secrecy are permissible only at the implementation level, and only in response to hardware or architectural constraints--e.g., storage that cannot be securely erased.)
- o TEPs MUST protect and authenticate the end-of-file marker conveyed by TCP's FIN flag. In particular, a receiver MUST with high probability detect a FIN flag that was set or cleared in transit and does not match the sender's intent. A TEP MAY discard a segment with such a corrupted FIN bit, or may abort the connection in response to such a segment. However, any such abort MUST raise an error condition distinct from an authentic end-of-file condition.
- o TEPs MUST prevent corrupted packets from causing urgent data to be delivered when none has been sent. A TEP MAY do so by cryptographically protecting the URG flag and urgent pointer alongside ordinary payload data. Alternatively, a TEP MAY disable urgent data functionality by clearing the URG flag on all received segments and returning errors in response to sender-side urgent-data API calls. Implementations SHOULD avoid negotiating TEPs that disable urgent data by default. The exception is when applications and protocols are known never to send urgent data.

[5.1.](#) Session IDs

Each TEP MUST define a session ID that is computable by both endpoints and uniquely identifies each encrypted TCP connection. Implementations MUST expose the session ID to applications via an API extension. Applications that are aware of TCP-ENO SHOULD, when practical, authenticate the TCP endpoints by incorporating the values of the session ID and TCP-ENO role (A or B) into higher-layer authentication mechanisms.

In order to avoid replay attacks and prevent authenticated session IDs from being used out of context, session IDs MUST be unique over all time with high probability. This uniqueness property MUST hold even if one end of a connection maliciously manipulates the protocol in an effort to create duplicate session IDs. In other words, it MUST be infeasible for a host, even by violating the TEP specification, to establish two TCP connections with the same session ID to remote hosts properly implementing the TEP.

To prevent session IDs from being confused across TEPs, all session

TEP identifier in host B's SYN segment. If the "v" bit was 1 in host B's SYN segment, then it is also 1 in the session ID. However, only the first byte is included, not the suboption data. Figure 8 shows the resulting format. This format is designed for TEPs to compute unique identifiers; it is not intended for application authors to pick apart session IDs. Applications SHOULD treat session IDs as monolithic opaque values and SHOULD NOT discard the first byte to shorten identifiers. (An exception is for non-security-relevant purposes, such as gathering statistics about negotiated TEPs.)




Figure 8: Format of a session ID

Though TEP specifications retain considerable flexibility in their definitions of the session ID, all session IDs MUST meet the following normative list of requirements:

- o The session ID MUST be at least 33 bytes (including the one-byte suboption), though TEPs may choose longer session IDs.
- o The session ID MUST depend in a collision-resistant way on all of the following (meaning it is computationally infeasible to produce collisions of the session ID derivation function unless all of the following quantities are identical):
 - * Fresh data contributed by both sides of the connection,
 - * Any public keys, public Diffie-Hellman parameters, or other public asymmetric cryptographic parameters that are employed by the TEP and have corresponding private data that is known by only one side of the connection, and
 - * The negotiation transcript specified in [Section 4.8](#).
- o Unless and until applications disclose information about the

session ID, all but the first byte MUST be computationally indistinguishable from random bytes to a network eavesdropper.

- o Applications MAY choose to make session IDs public. Therefore, TEPs MUST NOT place any confidential data in the session ID (such as data permitting the derivation of session keys).

[6.](#) Examples

This subsection illustrates the TCP-ENO handshake with a few non-normative examples.

```
(1) A -> B:  SYN      ENO<X,Y>
(2) B -> A:  SYN-ACK  ENO<b=1,Y>
(3) A -> B:  ACK      ENO<>
[rest of connection encrypted according to TEP Y]
```

Figure 9: Three-way handshake with successful TCP-ENO negotiation

Figure 9 shows a three-way handshake with a successful TCP-ENO negotiation. The two sides agree to follow the TEP identified by suboption Y.

```
(1) A -> B:  SYN      ENO<X,Y>
(2) B -> A:  SYN-ACK
(3) A -> B:  ACK
[rest of connection unencrypted legacy TCP]
```

Figure 10: Three-way handshake with failed TCP-ENO negotiation

Figure 10 shows a failed TCP-ENO negotiation. The active opener (A) indicates support for TEPs corresponding to suboptions X and Y. Unfortunately, at this point one of several things occurs:

1. The passive opener (B) does not support TCP-ENO,
2. B supports TCP-ENO, but supports neither of TEPs X and Y, and so does not reply with an ENO option,
3. B supports TCP-ENO, but has the connection configured in

mandatory application-aware mode and thus disables ENO because A's SYN segment does not set the application-aware bit, or

4. The network stripped the ENO option out of A's SYN segment, so B did not receive it.

Whichever of the above applies, the connection transparently falls back to unencrypted TCP.

```
(1) A -> B:  SYN      ENO<X,Y>
(2) B -> A:  SYN-ACK  ENO<b=1,X> [ENO stripped by middlebox]
(3) A -> B:  ACK
[rest of connection unencrypted legacy TCP]
```

Figure 11: Failed TCP-ENO negotiation because of network filtering

Figure 11 Shows another handshake with a failed encryption negotiation. In this case, the passive opener B receives an ENO option from A and replies. However, the reverse network path from B to A strips ENO options. Hence, A does not receive an ENO option from B, disables ENO, and does not include a non-SYN-form ENO option when ACKing B's SYN segment. The lack of ENO in A's ACK segment signals to B that the connection will not be encrypted. At this point, the two hosts proceed with an unencrypted TCP connection.

```
(1) A -> B:  SYN      ENO<Y,X>
(2) B -> A:  SYN      ENO<b=1,X,Y,Z>
(3) A -> B:  SYN-ACK  ENO<Y,X>
(4) B -> A:  SYN-ACK  ENO<b=1,X,Y,Z>
[rest of connection encrypted according to TEP Y]
```

Figure 12: Simultaneous open with successful TCP-ENO negotiation

Figure 12 shows a successful TCP-ENO negotiation with simultaneous open. Here the first four segments MUST contain a SYN-form ENO option, as each side sends both a SYN-only and a SYN-ACK segment. The ENO option in each host's SYN-ACK is identical to the ENO option in its SYN-only segment, as otherwise connection establishment could not recover from the loss of a SYN segment. The last valid TEP in host B's ENO option is Y, so Y is the negotiated TEP.

This section describes some of the design rationale behind TCP-ENO.

[7.1.](#) Future developments

TCP-ENO is designed to capitalize on future developments that could alter trade-offs and change the best approach to TCP-level encryption (beyond introducing new cipher suites). By way of example, we discuss a few such possible developments.

Various proposals exist to increase option space in TCP [I-D.ietf-tcpm-tcp-edo][[I-D.briscoe-tcpm-inspace-mode-tcpbis](#)][I-D.touch-tcpm-tcp-syn-ext-opt]. If SYN segments gain large options, it becomes possible to fit public keys or Diffie-Hellman parameters into SYN segments. Future TEPs can take advantage of this by performing key agreement directly within suboption data, both simplifying protocols and reducing the number of round trips required for connection setup.

If TCP gains large SYN option support, the 32-byte limit on length bytes may prove problematic. This draft intentionally aborts TCP-ENO if a length byte is followed by an octet in the range 0x00-0x9f. Any document updating TCP's option size limit can also enable larger

suboptions by updating this draft to assign meaning to such currently undefined byte sequences.

New revisions to socket interfaces [[RFC3493](#)] could involve library calls that simultaneously have access to hostname information and an underlying TCP connection. Such an API enables the possibility of authenticating servers transparently to the application, particularly in conjunction with technologies such as DANE [[RFC6394](#)]. An update to TCP-ENO can adopt one of the "z" bits in the global suboption to negotiate the use of an endpoint authentication protocol before any application use of the TCP connection. Over time, the consequences of failed or missing endpoint authentication can gradually be increased from issuing log messages to aborting the connection if some as yet unspecified DNS record indicates authentication is mandatory. Through shared library updates, such endpoint authentication can potentially be added transparently to legacy applications without recompilation.

TLS can currently only be added to legacy applications whose

protocols accommodate a STARTTLS command or equivalent. TCP-ENO, because it provides out-of-band signaling, opens the possibility of future TLS revisions being generically applicable to any TCP application.

[7.2.](#) Handshake robustness

Incremental deployment of TCP-ENO depends critically on failure cases devolving to unencrypted TCP rather than causing the entire TCP connection to fail.

Because a network path may drop ENO options in one direction only, a host must know not just that the peer supports encryption, but that the peer has received an ENO option. To this end, ENO disables encryption unless it receives an ACK segment bearing an ENO option. To stay robust in the face of dropped segments, hosts must continue to include non-SYN form ENO options in segments until such point as they have received a non-SYN segment from the other side.

One particularly pernicious middlebox behavior found in the wild is load balancers that echo unknown TCP options found in SYN segments back to an active opener. The passive role bit "b" in global suboptions ensures encryption will always be disabled under such circumstances, as sending back a verbatim copy of an active opener's SYN-form ENO option always causes role negotiation to fail.

[7.3.](#) Suboption data

TEPs can employ suboption data for session caching, cipher suite negotiation, or other purposes. However, TCP currently limits total option space consumed by all options to only 40 bytes, making it impractical to have many suboptions with data. For this reason, ENO optimizes the case of a single suboption with data by inferring the length of the last suboption from the TCP option length. Doing so saves one byte.

[7.4.](#) Passive role bit

TCP-ENO, TEPs, and applications all have asymmetries that require an unambiguous way to identify one of the two connection endpoints. As an example, [Section 4.8](#) specifies that host A's ENO option comes before host B's in the negotiation transcript. As another example, an application might need to authenticate one end of a TCP connection with a digital signature. To ensure the signed message cannot not be interpreted out of context to authenticate the other end, the signed message would need to include both the session ID and the local role, A or B.

A normal TCP three-way handshake involves one active and one passive opener. This asymmetry is captured by the default configuration of the "b" bit in the global suboption. With simultaneous open, both hosts are active openers, so TCP-ENO requires that one host manually configure "b = 1". An alternate design might automatically break the symmetry to avoid this need for manual configuration. However, all such designs we considered either lacked robustness or consumed precious bytes of SYN option space even in the absence of simultaneous open. (One complicating factor is that TCP does not know it is participating in a simultaneous open until after it has sent a SYN segment. Moreover, with packet loss, one host might never learn it has participated in a simultaneous open.)

[7.5.](#) Option kind sharing

This draft does not specify the use of ENO options beyond the first few segments of a connection. Moreover, it does not specify the content of ENO options in non-SYN segments, only their presence. As a result, any use of option kind TBD (or option kind 253 with ExID 0x454E) after the SYN exchange does not conflict with this document. Because in addition ENO guarantees at most one negotiated TEP per connection, TEPs will not conflict with one another or ENO if they use ENO's option kind for out-of-band signaling in non-SYN segments.

[8.](#) Experiments

This document has experimental status because TCP-ENO's viability depends on middlebox behavior that can only be determined *a posteriori*. Specifically, we must determine to what extent

middleboxes will permit the use of TCP-ENO. Once TCP-ENO is deployed, we will be in a better position to gather data on two types of failure:

1. Middleboxes downgrading TCP-ENO connections to unencrypted TCP. This can happen if middleboxes strip unknown TCP options or if they terminate TCP connections and relay data back and forth.
2. Middleboxes causing TCP-ENO connections to fail completely. This can happen if applications perform deep packet inspection and start dropping segments that unexpectedly contain ciphertext.

The first type of failure is tolerable since TCP-ENO is designed for incremental deployment anyway. The second type of failure is more problematic, and, if prevalent, will require the development of techniques to avoid and recover from such failures.

9. Security considerations

An obvious use case for TCP-ENO is opportunistic encryption--that is, encrypting some connections, but only where supported and without any kind of endpoint authentication. Opportunistic encryption protects against undetectable large-scale eavesdropping. However, it does not protect against detectable large-scale eavesdropping (for instance, if ISPs terminate TCP connections and proxy them, or simply downgrade connections to unencrypted). Moreover, opportunistic encryption emphatically does not protect against targeted attacks that employ trivial spoofing to redirect a specific high-value connection to a man-in-the-middle attacker.

Achieving stronger security with TCP-ENO requires verifying session IDs. Any application relying on ENO for communications security **MUST** incorporate session IDs into its endpoint authentication. By way of example, an authentication mechanism based on keyed digests (such as Digest Access Authentication [[RFC7616](#)]) can be extended to include the role and session ID in the input of the keyed digest. To preserve backwards compatibility, applications **MAY** use the application-aware bit to negotiate the inclusion of session IDs in authentication.

Because TCP-ENO enables multiple different TEPs to coexist, security could potentially be only as strong as the weakest available TEP. In particular, if session IDs do not depend on the TCP-ENO transcript in

a strong way, an attacker can undetectably tamper with ENO options to force negotiation of a deprecated and vulnerable TEP. To avoid such problems, TEPs MUST compute session IDs using only well-studied and conservative hash functions. That way, even if other parts of a TEP are vulnerable, it is still intractable for an attacker to induce identical session IDs at both ends after tampering with ENO contents in SYN segments.

Implementations MUST NOT send ENO options unless they have access to an adequate source of randomness [RFC4086]. Without secret unpredictable data at both ends of a connection, it is impossible for TEPs to achieve confidentiality and forward secrecy. Because systems typically have very little entropy on bootup, implementations might need to disable TCP-ENO until after system initialization.

With a regular three-way handshake (meaning no simultaneous open), the non-SYN form ENO option in an active opener's first ACK segment MAY contain $N > 0$ bytes of TEP-specific data, as shown in Figure 3. Such data is not part of the TCP-ENO negotiation transcript, and hence MUST be separately authenticated by the TEP.

10. IANA Considerations

This document defines a new TCP option kind for TCP-ENO, assigned a value of TBD from the TCP option space. This value is defined as:

Kind	Length	Meaning	Reference
TBD	N	Encryption Negotiation (TCP-ENO)	[RFC-TBD]

TCP Option Kind Numbers

Early implementations of TCP-ENO and a predecessor TCP encryption protocol made unauthorized use of TCP option kind 69.

[RFC-editor: please glue the following text to the previous paragraph iff TBD == 69, otherwise delete it.] These earlier uses of option 69 are not compatible with TCP-ENO and could disable encryption or suffer complete connection failure when interoperating with TCP-ENO-compliant hosts. Hence, legacy use of option 69 MUST be disabled on hosts that cannot be upgraded to TCP-ENO.

[RFC-editor: please glue this to the previous paragraph regardless of the value of TBD.] More recent implementations used experimental option 253 per [RFC6994] with 16-bit ExID 0x454E, and SHOULD migrate to option TBD by default.

Internet-Draft

tcpeno

October 2016

This document defines a 7-bit "glt" field in the range of 0x20-0x7f for which IANA shall maintain a new sub-registry entitled "TCP encryption protocol identifiers" under the "Transmission Control Protocol (TCP) Parameters" registry. The description of this registry should be interpreted with respect to the terminology defined in [\[RFC5226\]](#).

The intention is for IANA to grant registration requests for TEP identifiers in anticipation of a published RFC. Hence, a Specification is Required. However, to allow for implementation experience, identifiers should be allocated prior to the RFC being approved for publication. A Designated Expert appointed by the IESG area director shall approve allocations once it seems more likely than not that an RFC will eventually be published. The Designated Expert shall post a request to the TCPINC WG mailing list (or a successor designated by the Area Director) for comment and review, including an Internet-Draft. Before a period of 30 days has passed, the Designated Expert will either approve or deny the registration request and publish a notice of the decision to the TCPINC WG mailing list or its successor, as well as informing IANA. A denial notice must be justified by an explanation, and in the cases where it is possible, concrete suggestions on how the request can be modified so as to become acceptable should be provided.

The initial values of the TCP-ENO encryption protocol identifier registry are shown in Table 2.

Value	Meaning	Reference
0x20	Experimental Use	
0x21	TCPCRYPT_ECDHE_P256	[I-D.ietf-tcpinc-tcpencrypt]
0x22	TCPCRYPT_ECDHE_P521	[I-D.ietf-tcpinc-tcpencrypt]
0x23	TCPCRYPT_ECDHE_Curve25519	[I-D.ietf-tcpinc-tcpencrypt]
0x24	TCPCRYPT_ECDHE_Curve448	[I-D.ietf-tcpinc-tcpencrypt]
0x30	TCP-Use-TLS	[I-D.ietf-tcpinc-use-tls]

Table 2: TCP encryption protocol identifiers

[11.](#) Acknowledgments

We are grateful for contributions, help, discussions, and feedback from the TCPINC working group, including Marcelo Bagnulo, David Black, Bob Briscoe, Jana Iyengar, Tero Kivinen, Mirja Kuhlewind, Yoav Nir, Christoph Paasch, Eric Rescorla, Kyle Rose, and Joe Touch. This work was partially funded by DARPA CRASH and the Stanford Secure Internet of Things Project.

Internet-Draft

tcpeno

October 2016

12. References

12.1. Normative References

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), DOI 10.17487/RFC0793, September 1981, <<http://www.rfc-editor.org/info/rfc793>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", [BCP 106](#), [RFC 4086](#), DOI 10.17487/RFC4086, June 2005, <<http://www.rfc-editor.org/info/rfc4086>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC6994] Touch, J., "Shared Use of Experimental TCP Options", [RFC 6994](#), DOI 10.17487/RFC6994, August 2013, <<http://www.rfc-editor.org/info/rfc6994>>.
- [RFC7413] Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", [RFC 7413](#), DOI 10.17487/RFC7413, December 2014, <<http://www.rfc-editor.org/info/rfc7413>>.

12.2. Informative References

- [I-D.briscoe-tcpm-inspace-mode-tcpbis]
Briscoe, B., "Inner Space for all TCP Options (Kitchen

Sink Draft - to be Split Up)", [draft-briscoe-tcpm-inspace-mode-tcpbis-00](#) (work in progress), March 2015.

[I-D.ietf-tcpinc-tcpcrypt]

Bittau, A., Boneh, D., Giffin, D., Hamburg, M., Handley, M., Mazieres, D., Slack, Q., and E. Smith, "Cryptographic protection of TCP Streams (tcpcrypt)", [draft-ietf-tcpinc-tcpcrypt-03](#) (work in progress), October 2016.

[I-D.ietf-tcpinc-use-tls]

Rescorla, E., "Using TLS to Protect TCP Streams", [draft-ietf-tcpinc-use-tls-01](#) (work in progress), May 2016.

Bittau, et al.

Expires May 4, 2017

[Page 25]

Internet-Draft

tcpeno

October 2016

[I-D.ietf-tcpm-tcp-edo]

Touch, J. and W. Eddy, "TCP Extended Data Offset Option", [draft-ietf-tcpm-tcp-edo-06](#) (work in progress), June 2016.

[I-D.touch-tcpm-tcp-syn-ext-opt]

Touch, J. and T. Faber, "TCP SYN Extended Option Space Using an Out-of-Band Segment", [draft-touch-tcpm-tcp-syn-ext-opt-05](#) (work in progress), October 2016.

[RFC3493] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W. Stevens, "Basic Socket Interface Extensions for IPv6", [RFC 3493](#), DOI 10.17487/RFC3493, February 2003, <<http://www.rfc-editor.org/info/rfc3493>>.

[RFC4987] Eddy, W., "TCP SYN Flooding Attacks and Common Mitigations", [RFC 4987](#), DOI 10.17487/RFC4987, August 2007, <<http://www.rfc-editor.org/info/rfc4987>>.

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

[RFC5382] Guha, S., Ed., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", [BCP 142](#), [RFC 5382](#), DOI 10.17487/RFC5382, October 2008, <<http://www.rfc-editor.org/info/rfc5382>>.

- [RFC6394] Barnes, R., "Use Cases and Requirements for DNS-Based Authentication of Named Entities (DANE)", [RFC 6394](https://www.rfc-editor.org/info/rfc6394), DOI 10.17487/RFC6394, October 2011, <<http://www.rfc-editor.org/info/rfc6394>>.
- [RFC7616] Shekh-Yusef, R., Ed., Ahrens, D., and S. Bremer, "HTTP Digest Access Authentication", [RFC 7616](https://www.rfc-editor.org/info/rfc7616), DOI 10.17487/RFC7616, September 2015, <<http://www.rfc-editor.org/info/rfc7616>>.

Authors' Addresses

Andrea Bittau
Google
345 Spear Street
San Francisco, CA 94105
US

Email: bittau@google.com

Bittau, et al.

Expires May 4, 2017

[Page 26]

Internet-Draft

tcpeno

October 2016

Dan Boneh
Stanford University
353 Serra Mall, Room 475
Stanford, CA 94305
US

Email: dabo@cs.stanford.edu

Daniel B. Giffin
Stanford University
353 Serra Mall, Room 288
Stanford, CA 94305
US

Email: dbg@scs.stanford.edu

Mark Handley
University College London
Gower St.

London WC1E 6BT
UK

Email: M.Handley@cs.ucl.ac.uk

David Mazieres
Stanford University
353 Serra Mall, Room 290
Stanford, CA 94305
US

Email: dm@uun.org

Eric W. Smith
Kestrel Institute
3260 Hillview Avenue
Palo Alto, CA 94304
US

Email: eric.smith@kestrel.edu