

From Stateless to Stateful: Generic Authentication and Authenticated Encryption Constructions with Application to TLS

Colin Boyd¹ Britta Hale¹ Stig Frode Mjølsnes¹ Douglas Stebila² †

¹ Norwegian University of Science and Technology, NTNU, Trondheim, Norway
`{colin.boyd,britta.hale,stig.mjolsnes}@item.ntnu.no`
² Queensland University of Technology, Brisbane, Australia
`stebila@qut.edu.au`

Abstract. Authentication and authenticated encryption with associated data (AEAD) are applied in cryptographic protocols to provide message integrity. The definitions in the literature and the constructions used in practice all protect against forgeries, but offer varying levels of protection against replays, reordering, and drops. As a result of the lack of a systematic hierarchy of authentication and AEAD security notions, gaps have arisen in the literature, specifically in the provable security analysis of the Transport Layer Security (TLS) protocol. We present a hierarchy of authentication and AEAD security notions, interpolating between the lowest level of protection (against forgeries) and the highest level (against forgeries, replays, reordering, and drops). We show generically how to construct higher level schemes from a basic scheme and appropriate use of sequence numbers, and apply that to close the gap in the analysis of TLS record layer encryption.

Keywords: authentication, authenticated encryption with associated data (AEAD), Transport Layer Security (TLS) protocol, secure channels

1 Introduction

Message integrity is a vital security service demanded of cryptographic protocols, and is usually provided either by a message authentication code (MAC) or by a combined authenticated encryption scheme. The standard security property for a MAC is existential unforgeability under a chosen message attack.

There has been an extensive line of research on security notions and constructions for authenticated encryption schemes, with initial definitions given by Katz and Yung [14], Bellare and Namprempre [4], and Krawczyk [17]. For message confidentiality, an authenticated encryption scheme could achieve indistinguishability under either an adaptive chosen plaintext (IND-CPA) or an adaptive chosen ciphertext (IND-CCA a.k.a. IND-CCA2) attack. For message integrity, an authenticated encryption scheme could achieve either integrity of plaintexts

†Supported by the Australian Research Council (ARC) Discovery Project, grant DP130104304

(INT-PTXT) or of ciphertexts (INT-CTXT). Shrimpton [29] combined the separate INT-CTXT and IND-CCA experiments into a single experiment which he called IND-CCA3.

Bellare and Namprempre [4] and Krawczyk [17] also investigated how to construct authenticated encryption schemes from MACs and symmetric encryption, evaluating three construction paradigms: encrypt-and-MAC, MAC-then-encrypt, and encrypt-then-MAC.

Rogaway [26] defined the notion of *authenticated encryption with associated data (AEAD)*, to capture the common real-world scenario in which some data (such as packet headers) needs to be sent authentically alongside a ciphertext, but need not be encrypted, and AEAD has taken prominence over plain authenticated encryption in recent years.

Despite the utility of authenticated encryption and AEAD, it is not enough to realize the secure channel property expected of cryptographic protocols for two reasons. First, secure channel protocols are often expected to perform an initial establishment of the encryption key using a key exchange protocol; see for example the original paper on secure channels by Canetti and Krawczyk [6] (and the follow-up by Namprempre [22]) as well as recent realizations such as the authenticated and confidential channel establishment (ACCE) model of Jager et al. [13]. (In this paper, we will not focus on the key exchange establishment phase of secure channels.) Second, and more important for this paper, applications often expect reliable delivery of a sequence of messages: that no attacker can replay messages, deliver them in a different order in which they were sent, or drop some messages without later detection.

To capture the notion of delivery of a sequence of messages, Bellare et al. [3] introduced *stateful authenticated encryption*, with two security properties: stateful integrity of ciphertexts (INT-SFCTXT) and stateful indistinguishability of ciphertexts (IND-SFCCA). Kohno et al. [16] extended the statefulness to AEAD schemes, and gave a hierarchy of 5 integrity notions: type 1) security against forgeries; type 2) type 1 plus security against replays; type 3) type 2 plus security against reordering; type 4) type 3 plus detection of previous drops but still accepting subsequent messages; type 5) type 4 plus but not accepting subsequent messages. The type 5 notion of Kohno et al. [16] is equivalent to the stateful authenticated encryption notion of Bellare et al. [3].

Paterson et al. [23] revisit AEAD definitions in the context of the Transport Layer Security (TLS) protocol. They present a combined AEAD security notion called *length-hiding authenticated encryption (LHAE)*, which provides message integrity and confidentiality similar to the type-5 security of Kohno et al. [16], even for messages of different length (hence “length-hiding”), and in a single combined security property (following Shrimpton [29]). Paterson et al. then go on to show that, under appropriate length conditions on the message authentication tag, a simplified form of the encode-then-MAC-then-encrypt form of encryption in the TLS record layer in ciphersuites that use a block cipher in CBC mode is a secure length-hiding authenticated encryption scheme. The simplification is that the statefulness aspects (sequence numbers) are not considered.

Jager et al. [13] and Krawczyk et al. [18], in their provable security analyses of the full TLS protocol (covering both the authenticated key exchange in the TLS handshake and the TLS record layer), rely on an extension of the work of Paterson et al. [23], namely a form of *stateful length-hiding authenticated encryption (sLHAE)*. Unfortunately, the work of Paterson et al. did not show that TLS encode-then-MAC-then-encrypt satisfies sLHAE, only LHAE. To our knowledge, this gap remains in the literature until now.

1.1 Our contributions

In this work, we construct a hierarchy of authentication and AEAD security notions, show how to construct schemes with higher levels of security from a scheme with the lowest level of security combined with sequence numbers, and apply these techniques to TLS record layer encryption to bridge the gap between LHAE [23] and sLHAE [13].

First, we construct a hierarchy of authentication levels:

1. protection against forgeries,
2. protection against forgeries and replays,
3. protection against forgeries, replays, and reordering of messages, and
4. protection against forgeries, replays, reordering of messages, and dropped messages.

We give a similar hierarchy of definitions for AEAD, with single-experiment AEAD notions that combine integrity and indistinguishability, following Shrimpton [29]. In both cases, these hierarchy levels can be viewed as interpolating between existing *stateless* notions at our level 1 and existing *stateful* notions at our level 4.

Continuing, we show how to construct level 2, 3, and 4 schemes from level 1 schemes. The constructions are not surprising: by appropriate incorporation and checking of sequence numbers, the receiver can ensure it is receiving a valid sequence of sent messages. However, our constructions incorporate a degree of generality: rather than fixing how the sequence numbers are incorporated, we allow an *encoding scheme* to include them either *implicitly* or *explicitly*. For example, in an explicit encoding scheme, the sequence number might be authenticated and then transmitted alongside the ciphertext, in the manner of DTLS. Alternatively, in an implicit encoding scheme, the sequence number might be incorporated into the authentication calculation but not actually transmitted across the wire (since the receiving party ought to know what packet number to expect); this is how TLS works, for example.

We use this generic construction to close the gap in the provable security analysis of TLS record layer encryption. Paterson et al.’s analysis of a simplified form of TLS encode-then-MAC-then-encrypt (Π_{PRS}) shows that it satisfies the LHAE notion, equivalent to our level 1. We can formulate TLS’s use of sequence numbers as an encoding scheme in our generic construction, and then see that the full form of TLS encode-then-MAC-then-encrypt (Π_{TLS}) is equivalent to our level-4 generic construction applied to Π_{PRS} , and thus Π_{TLS} achieves level-4 AEAD security, equivalent to sLHAE. Fig. 1 illustrates the connection between

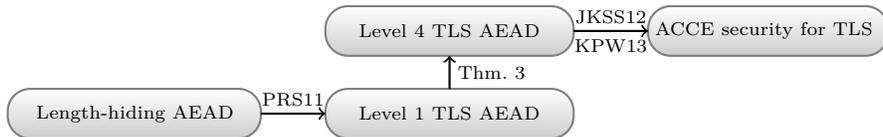


Fig. 1: TLS channel analysis.

our work and that of Paterson et al., Jager et al., and Krawczyk et al., depicting how the construction from level-1 AEAD to level-4 AEAD builds a missing and necessary bridge in the analysis of TLS.

Relation with existing work. The work most closely related to ours is the manuscript of Kohno et al. [16], who gave a hierarchy of AEAD notions. Our AEAD hierarchy maps on to theirs: our levels 1, 2, 3, and 4 correspond to their types 1, 2, 3, and 5, respectively. There are several differences with our work. They give constructions of higher level schemes directly from encryption and MAC schemes in the encrypt-and-MAC, MAC-then-encrypt, and encrypt-then-MAC paradigms, whereas we show how to construct higher levels generically from lower level schemes. Their AEAD hierarchy uses separate integrity and indistinguishability experiments at each level, whereas we use a single combined experiment at each level. We also give a hierarchy of authentication notions, not just AEAD notions, and thereby expand applicability to schemes outside of the AEAD context. Finally, we connect the hierarchy and our generic constructions with TLS record layer encryption.

Connection with secure channel definitions. One motivation of our work was to understand the difference between the original CK01 secure channel definition of Canetti and Krawczyk [6] and the ACCE model of Jager et al. [13]. The confidentiality and integrity notions in CK01 and their NetAut protocol correspond with level 1 of our AEAD hierarchy – stateless authenticated encryption. A comment in their paper does require that the receiver “check for uniqueness of the incoming message”, which would upgrade to level 2 in our hierarchy, and this is the notion that was used in a subsequent work by Namprempre [22]. In contrast, Jager et al.’s ACCE notion maps to level 4 of our AEAD hierarchy – sLHAE.

Application to real-world protocols. Each level of our AEAD hierarchy maps to the requirements expected in some real-world protocols:

- Level 1: DTLS [24, 25]: Datagram TLS provides basic authentication, allows packets to be dropped, and will receive packets out of order, queuing them for future processing.
- Level 2: IPsec Authentication Header (AH) [15]: IPsec Authentication Header protocol provides similar replay detection using a window of recently received packets combined with dropping packets that are “too old”.
- Level 2: DTLS with optional replay detection: Datagram TLS does allow optional replay detection [24, 25, §3.3] using a similar technique to IPsec AH.
- Level 3: 802.11 [12] is designed to preventing reordering and to detect replays but allows for packet dropping.

- Level 4: TLS [7] is designed to receive a message sequence strictly as a sent, and will be discussed at greater length in Section 4.

A recent analysis [20] of the QUIC protocol [19] employed an AEAD level comparable to our level 1 AEAD; however, the replay-detection abilities of QUIC suggest that a higher authentication level should be achievable.

1.2 Additional related work

There are several additional lines of work on authenticated encryption.

One line of research views data “as a stream”, rather than a discrete sequence of messages; practical implementations receive data byte-by-byte rather than as atomic messages in security definitions. Albrecht et al. [1] showed how to carry out a plaintext recovery attack against the Secure Shell (SSH) protocol as a result of byte-by-byte processing. This motivated the need for non-atomic authenticated encryption definitions [5, 8]. The work of Fischlin et al. [8] in particular is motivated by protocols such as TLS, SSH, and QUIC, and describes checks that can again be correlated with our level-4 AEAD notion. It would be interesting to expand stream-based analysis in the direction of our hierarchical levels for protocols that allow packet dropping. For example, the QUIC protocol [19] runs over UDP and tolerates a degree of packet loss, making analysis under a level-4 stream-based notion inappropriate.

Another line of research focuses on the use of nonces in authenticated encryption [26, 27], and more recently for the specific purposes of protecting implementations that misuse counters or nonces [28, 9, 11]. Meanwhile, Hoang et al. [10] define a notion of robust authenticated encryption which incorporates padding properties similar to the stateless form of LHAE of Paterson et al. [23]. Finally, additional recent work focuses on defining authenticated encryption results in the constructive cryptography framework [21, 2].

2 Authentication Hierarchy

In this section, we formalize our 4-tier hierarchy of authentication notions, each level building on the previous, and show how to achieve higher level notions from level-1 combined with appropriate checks on sequence numbers.

2.1 Definitions

Definition 1. A stateful authentication scheme Π for a message space \mathcal{M} , a key space \mathcal{K} , and an output space \mathcal{C} is a tuple of algorithms:

- $\text{Kgn}() \xrightarrow{\$} k$: A probabilistic key generation algorithm that outputs a key k .
- $\text{Snd}(k, m, st_E) \xrightarrow{\$} (c, st_E)$: A probabilistic authentication algorithm that takes as input a key $k \in \mathcal{K}$, a message $m \in \mathcal{M}$, and an authentication state st_E , and outputs a tagged message $c \in \mathcal{C}$ and updated state st_E .
- $\text{Rcv}(k, c, st_D) \rightarrow (m, \alpha, st_D)$: A deterministic verification algorithm that takes as input a key $k \in \mathcal{K}$, a tagged message $c \in \mathcal{C}$, and a verification state st_D , and outputs either a message $m \in \mathcal{M}$ or an error symbol \perp , a bit $\alpha \in \{0, 1\}$, and an updated state st_D .

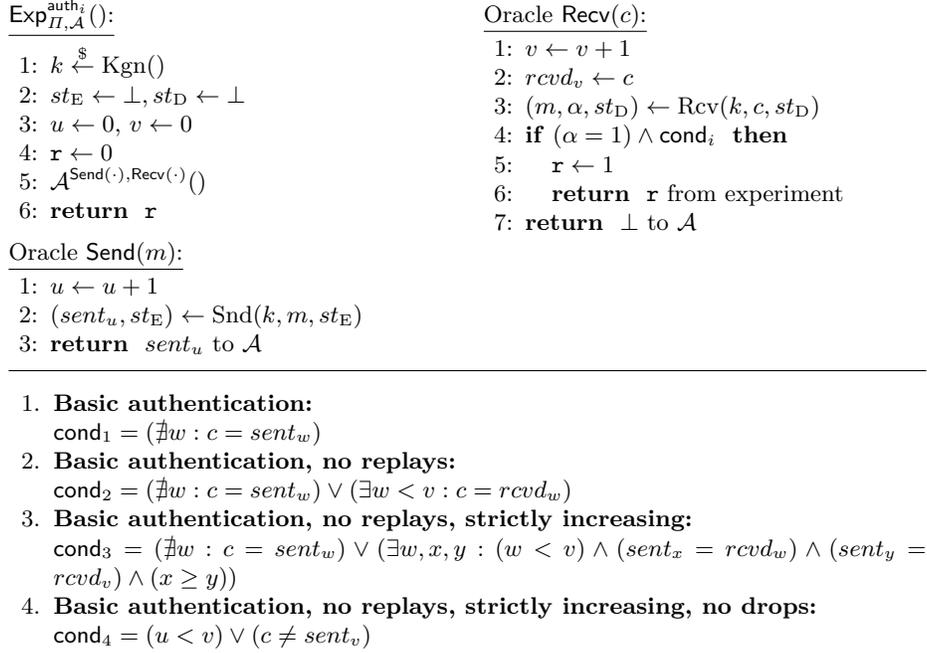


Fig. 2: Stateful authentication experiment auth_i with authentication condition cond_i for stateful authentication scheme $\Pi = (\text{Kgn}, \text{Snd}, \text{Rcv})$ and adversary \mathcal{A} .

On first use, st_E and st_D are initialized to \perp .

Correctness is defined in the natural way: for all $m \in \mathcal{M}$, all $k \xleftarrow{\$} \text{Kgn}()$, all st_E and st_D defined in any sequence of encryptions and decryptions respectively, and all c such that $(c, st'_E) \leftarrow \text{Snd}(k, m, st_E)$, we have that $\text{Rcv}(k, c, st_D) = (m, 1, st'_D)$.

Note that in the case of a Rcv (message authentication check) failure, the receive algorithm outputs a failure symbol \perp , $\alpha = 0$ to denote a failed receipt, and an updated state $st_D: (\perp, 0, st_D) \leftarrow \text{Rcv}(k, c, st_D)$. Otherwise, the algorithm outputs the correctly received message m , $\alpha = 1$ to denote successful receipt, and an updated state $st_D: (m, 1, st_D) \leftarrow \text{Rcv}(k, c, st_D)$.

Formally we define a stateful authentication security experiment that can be *parameterized* with different authentication conditions to capture various levels of authentication. Four graded levels of authentication are defined for the experiment, correlated to different conditions, cond_i , under which an adversary \mathcal{A} wins, as shown in Fig. 2. Note that cond_4 is strongly linked to authentication demands in analyses of TLS [23, 13], a protocol with strict authentication requirements.

Definition 2. Let Π be a stateful authentication scheme and let \mathcal{A} be an adversary algorithm. Let $i \in \{1, \dots, 4\}$. The stateful authentication experiment for Π with authentication condition cond_i is given by $\text{Exp}_{\Pi, \mathcal{A}}^{\text{auth}_i}$ in Fig. 2. We define $\text{Adv}_{\Pi}^{\text{auth}_i}(\mathcal{A}) = \Pr \left[\text{Exp}_{\Pi}^{\text{auth}_i}(\mathcal{A}) = 1 \right]$.

Remark 1. If the authenticated message c takes the form of a ciphertext, then level-1 authentication is equivalent to INT-CTXT. If c is such that $c = (m, \text{MAC}(m))$, where MAC is a message authentication code, then level-1 authentication is equivalent to SUF-CMA. In order to maximize the application potential of our results, we provide the generality for either application.

2.2 Relations among authentication notions

Each of the authentication notions sequentially implies the security of the levels below it. In the following theorem, the security implications between levels are formalized, with security at Level 2 implying security at Level 1, etc.

Theorem 1 (Level- $(i+1)$ authentication implies level- i authentication). *Let $\Pi = (\text{Kgn}, \text{Snd}, \text{Rcv})$ be an authentication scheme and let $i \in \{1, 2, 3\}$. For any adversary \mathcal{A} , $\text{Adv}_{\Pi}^{\text{auth}_i}(\mathcal{A}) \leq \text{Adv}_{\Pi}^{\text{auth}_{i+1}}(\mathcal{A})$.*

The proof of Theorem 1 can be found in the full version and is omitted here due to space restrictions.

2.3 Constructing higher level authentication schemes

In this section, we generically show how to build higher level authentication schemes based on lower level authentication schemes and the inclusion of sequence numbers with appropriate checks. Since currently implemented protocols use both implicit and explicit sequence numbers, we generalize our model for an arbitrary *encoding scheme* which captures both implicit and explicit sequence numbers.

Definition 3 (Authentication encoding scheme). *An (authentication) encoding scheme Coding for a sequence number space \mathbf{S} and message space \mathcal{M} is a pair of algorithms:*

- $\text{Ecd}(\text{sqn}, m) \rightarrow m_{\text{ecd}}$: *A deterministic encoding algorithm that takes as input a sequence number $\text{sqn} \in \mathbf{S}$ and a message $m \in \mathcal{M}$, and outputs an encoded message $m_{\text{ecd}} \in \mathcal{M}_{\text{ecd}}$, where \mathcal{M}_{ecd} is the encoded version of \mathcal{M} .*
- $\text{Dcd}(\text{sqnlist}, m_{\text{ecd}}) \rightarrow (\text{sqn}, m, \alpha)$: *A deterministic decoding algorithm that takes as input a sequence number list $\text{sqnlist} \subset \mathbf{S}$ and an encoded message $m_{\text{ecd}} \in \mathcal{M}_{\text{ecd}}$, and outputs a sequence number $\text{sqn} \in \mathbf{S}$, a message $m \in \mathcal{M}$ or an error symbol \perp , and a status variable $\alpha = 1$ if decoding was successful or $\alpha = 0$ otherwise.*

In our construction of higher level authentications, we will require that Ecd is collision-resistant.

We can construct schemes that use either implicit or explicit sequence numbers using Definition 3. For example, the scheme with $\text{Ecd}(\text{sqn}, m) := \text{sqn} \| m$ has an explicit sequence number, and may be very applicable in practice since sqn is sent explicitly with the message. An alternative scheme with implicit sequence numbers would be $\text{Ecd}(\text{sqn}, m) := m \| \text{MAC}(\text{sqn})$. Thus elements of the space \mathcal{M}_{ecd} may take various forms, contingent on the properties desirable for

Coding. We will see in Section 4.2 that the TLS record layer protocol uses an encoding scheme based on the second example above. We formally distinguish explicit and implicit sequence numbers as follows:

Definition 4. *We say that authentication encoding scheme **Coding** uses explicit sequence numbers if $\text{Dcd}(\emptyset, \text{Ecd}(\text{sqn}, m)) = (\text{sqn}, m, 1)$ for all sqn and all m , and that **Coding** uses implicit sequence numbers otherwise.*

We now present our generic constructions of level- i authentication schemes from a level-1 authentication scheme. The heart of our construction is a sequence number check TEST_i that will correspond to the authentication condition cond_i . Our constructions can accommodate any collision-resistant encoding scheme **Coding**, with either implicit or explicit sequence numbers; this requirement is specifically important in implicit authentication where the sequence number is not physically present on receipt. For conciseness, the notation Π'_i for $P(\Pi, \text{Ecd}, \text{TEST}_i)$ will be generally employed.

Definition 5 (P construction). *Let Π be a (level-1) authentication scheme, **Coding** be an encoding scheme, and let TEST_i be one of the conditions specified in Fig. 3. Define $\Pi'_i := P(\Pi, \text{Coding}, \text{TEST}_i)$ as the authentication scheme resulting from apply construction P in Fig. 3.*

In this construction, the check TEST_2 corresponds to the condition for level-2 authentication. Basic level-1 authentication is assumed, so TEST_2 's protection against replays implies replay protection for condition cond_2 . Namely, if $\exists w < v : c = \text{rcvd}_w$ then $\exists j : \text{sqn} = \text{st}_D.\text{sqnlist}_j$, since identical authenticated messages must contain identical sequence numbers. Similar connections exist between TEST_3 and cond_3 and TEST_4 and cond_4 . Note that to check TEST_2 it is necessary to maintain a record of all previously received sqn ; thus $\text{st}_D.\text{sqnlist}$ must be a complete record. However, for TEST_3 and TEST_4 , it is strictly only necessary for $\text{st}_D.\text{sqnlist}$ to contain the last received sqn .

The following theorem shows that the P construction with TEST_i achieves level- i authentication. Notably Theorem 2 depends on the collision-resistance of **Ecd**. For many encoding schemes, this follows immediately. For example, the simple concatenation scheme $\text{Ecd}(\text{ctr}, m) = \text{ctr}||m$ is clearly collision-resistant when assuming unambiguous concatenation. When such a scheme is used, the advantage of \mathcal{A} is then directly reducible to the advantage of \mathcal{F} . Due to space restrictions, the proof of Theorem 2 can be found in the full version of this paper.

Theorem 2. *Let Π be a secure level-1 authentication scheme and **Coding** be an authentication encoding scheme with collision-resistant encoding. Let $i \in \{2, 3, 4\}$. Then $\Pi'_i = P(\Pi, \text{Coding}, \text{TEST}_i)$, constructed as in Fig. 3, is a secure level- i authentication scheme. Specifically, let \mathcal{A} be an adversary algorithm that runs in time t and asks q_s **Send** queries and q_r **Recv** queries, and let $q = q_s + q_r$. Then there exists an adversary \mathcal{B} that runs in time $t_{\mathcal{B}} \approx t$ and asks no more than $q_{\mathcal{B}} = \frac{1}{2}q_s(q_s - 1)$ queries, and an adversary \mathcal{F} that runs in time $t_{\mathcal{F}} \approx t$ and asks $q_{\mathcal{F}} = q$ queries, such that $\text{Adv}_{P(\Pi, \text{Coding}, \text{TEST}_i)}^{\text{auth}_i}(\mathcal{A}) \leq \text{Adv}_{\Pi}^{\text{auth}_1}(\mathcal{F}) + \text{Adv}_{\text{Ecd}}^{\text{collision}}(\mathcal{B})$.*

$\Pi'_i.\text{Kgn}()$ 1: return $\Pi.\text{Kgn}()$	$\Pi'_i.\text{Rcv}(k, c, st'_D)$ 1: if $st'_D.\text{status} = \text{failed}$ then 2: return $(\perp, 0, st'_D)$ 3: $(m_\Pi, \alpha, st'_D.\text{subst})$ $\leftarrow \Pi.\text{Rcv}(k, c, st'_D.\text{subst})$ 4: if $\alpha = 1$ then 5: $(\text{sqn}, m, \alpha) \leftarrow \text{Dcd}(st'_D.\text{sqnlist}, m_\Pi)$ 6: if $(\alpha = 0) \vee \text{TEST}_i$ then 7: $st'_D.\text{status} = \text{failed}$ 8: return $(\perp, 0, st'_D)$ 9: $st'_D.\text{sqnlist} = st'_D.\text{sqnlist} \parallel \text{sqn}$ 10: return (m, α, st'_D)
$\Pi'_i.\text{Snd}(k, m, st'_E)$ 1: $(c, st'_E.\text{subst})$ $\leftarrow \Pi.\text{Snd}(k, \text{Ecd}(st'_E.\text{ctr}, m), st'_E.\text{subst})$ 2: $st'_E.\text{ctr} \leftarrow st'_E.\text{ctr} + 1$ 3: return (c, st'_E)	

Sequence number tests for building Π' , correlated to authentication levels:

- **Basic authentication, no replays:**
 $\text{TEST}_2 = (\exists j : \text{sqn} = st'_D.\text{sqnlist}_j)$
- **Basic authentication, no replays, strictly increasing:**
 $\text{TEST}_3 = (\exists j : \text{sqn} \succ st'_D.\text{sqnlist}_j)$
- **Basic authentication, no replays, strictly increasing, no drops:**
 $\text{TEST}_4 = (\exists j : \text{sqn} \succ st'_D.\text{sqnlist}_j) \vee (\text{sqn} \neq \max\{st'_D.\text{sqnlist}_j\} + 1)$

Description of states st'_E and st'_D :

- $st'_E.\text{subst} := st_E$, where st_E is the state in Π
- $st'_E.\text{ctr}$. When $\Pi'.\text{Snd}$ is initialized, $st'_E.\text{ctr} \leftarrow 0$.
- $st'_D.\text{subst} := st_D$, where st_D is the state in Π
- $st'_D.\text{status}$. Once $st'_D.\text{status} = \text{failed}$ it is not reset and all subsequently received messages are also immediately aborted.
- $st'_D.\text{sqnlist}$, an ordered list of sequence numbers previously received. It is required that $|st'_D.\text{sqnlist}| \geq 1$ after the first received sqn ; i.e. the size of the ordered set is maintained at 1 or greater. When $\Pi'.\text{Snd}$ is initialized, $st'_D.\text{sqnlist} \leftarrow \perp$.

Fig. 3: Construction P of a level- i authentication scheme Π'_i from a level-1 authentication scheme Π and encoding scheme $\text{Coding} = (\text{Ecd}, \text{Dcd})$.

The time-cost for checking using implicit sequence numbers could be considerable when using a Level 2 or Level 3 authentication notion due to the need to check against all previously received messages. However, to our knowledge, there are no real-world implementations using implicit sequence numbers at these levels. Implicit sequence numbers have been used in instances where Level 4 authentication is desired, but explicit sequence numbers are usually employed at the lower levels. Logically, this also corresponds to desirable real-world instantiation formats; if a protocol allows packets to be dropped then it would be inconvenient to base authentication upon information that is not explicitly sent in each packet. Alternatively, if no drops are allowed, authentication can be checked against explicit or implicit information.

3 Authenticated Encryption Hierarchy

In this section, we build equivalent notions for *authenticated encryption with associated data (AEAD) schemes*. AEAD security is typically defined by extending the authentication notion with a type of *left-or-right* encryption game.

3.1 Definitions

Definition 6. A stateful AEAD scheme Π for a message space \mathcal{M} , an associated data space \mathcal{AD} , a key space \mathcal{K} , and a ciphertext space \mathcal{C} , is a tuple of algorithms:

- $\text{Kgn}() \xrightarrow{\mathbb{S}} k$: A probabilistic key generation algorithm that outputs a key k
- $\text{E}(k, \ell, \mathbf{ad}, m, st_{\text{E}}) \xrightarrow{\mathbb{S}} (c, st'_{\text{E}})$: A probabilistic encryption algorithm that takes as input a key $k \in \mathcal{K}$, a length $\ell \in \mathbb{Z}$, associated data $\mathbf{ad} \in \mathcal{AD}$, a message $m \in \mathcal{M}$, and an encryption state st_{E} , and outputs a ciphertext $c \in \mathcal{C}$ and updated state st'_{E} .
- $\text{D}(k, \mathbf{ad}, c, st_{\text{D}}) \rightarrow (\mathbf{ad}, m, \alpha, st'_{\text{D}})$: A deterministic decryption algorithm that takes as input a key $k \in \mathcal{K}$, associated data $\mathbf{ad} \in \mathcal{AD}$, a ciphertext c , and a decryption state st_{D} , and outputs either associated data \mathbf{ad} or an error symbol \perp , a message $m \in \mathcal{M}$ or an error symbol \perp , a bit $\alpha \in \{0, 1\}$, and an updated state st'_{D} .

Compared with stateful authentication schemes in Definition 1, AEAD schemes utilize two further fields: \mathbf{ad} , which is for associated data (such as authenticated but unencrypted header data), and an optional length field ℓ .

Correctness is defined in an analogous manner to that of stateful authentication schemes. Correspondingly we define 4 levels of stateful AEAD security.

Definition 7. Let Π be a stateful AEAD scheme and let \mathcal{A} be an PPT adversarial algorithm. Let $i \in \{1, \dots, 4\}$ and let $b \in \{0, 1\}$. The stateful AEAD experiment for Π with condition cond_i and bit b is given by $\text{Exp}_{\Pi}^{\text{aead}_i - b}(\mathcal{A})$ in Fig. 4. We define $\text{Adv}_{\Pi}^{\text{aead}_i}(\mathcal{A}) = \left| \Pr \left[\text{Exp}_{\Pi}^{\text{aead}_i - 1}(\mathcal{A}) = 1 \right] - \Pr \left[\text{Exp}_{\Pi}^{\text{aead}_i - 0}(\mathcal{A}) = 1 \right] \right|$.

The **Encrypt** and **Decrypt** oracles in Fig. 4 work together to provide both an authentication experiment and ciphertext indistinguishability experiment. When $b = 0$, the adversary always gets m_0 encrypted and never receives any decryption information. When $b = 1$, the adversary always gets m_1 encrypted and potentially receives decryption information. If the adversary makes an attempt to forge ciphertexts or violate the sequencing condition (modelled by the **out-of-sync** flag), then a secure stateful AEAD scheme should return \perp in all subsequent decryption queries. If the adversary has caused the encryptor and decryptor to get out of sync (by forging a ciphertext or violating the sequencing condition) and ever receives non- \perp from **Decrypt**, the adversary learns $b = 1$.

When ℓ is not used, the level-1 notion aead_1 corresponds to IND-CCA and INT-CTXT security of a stateless AEAD scheme.

When ℓ is used for length, the level-4 notion aead_4 corresponds to the stateful length-hiding authenticated encryption security notion of Krawczyk et al. [18] which is a slight modification of that of Jager et al. [13].

<p>$\text{Exp}_{\Pi, \mathcal{A}}^{\text{aead}_i - b}()$:</p> <ol style="list-style-type: none"> 1: $k \xleftarrow{\\$} \text{Kgn}()$ 2: $st_E \leftarrow \perp, st_D \leftarrow \perp$ 3: $u \leftarrow 0, v \leftarrow 0$ 4: out-of-sync $\leftarrow 0$ 5: $b' \xleftarrow{\\$} \mathcal{A}^{\text{Encrypt}(\cdot), \text{Decrypt}(\cdot)}()$ 6: return b' <p>Oracle Encrypt($\ell, \text{ad}, m_0, m_1$):</p> <ol style="list-style-type: none"> 1: $u \leftarrow u + 1$ 2: $(\text{sent}.c^{(0)}, st_E^{(0)}) \leftarrow E(k, \ell, \text{ad}, m_0, st_E)$ 3: $(\text{sent}.c^{(1)}, st_E^{(1)}) \leftarrow E(k, \ell, \text{ad}, m_1, st_E)$ 4: if $\text{sent}.c^{(0)} = \perp$ or $\text{sent}.c^{(1)} = \perp$ then 5: return \perp 6: $(\text{sent}.ad_u, \text{sent}.c_u, st_E) := (\text{ad}, \text{sent}.c^{(b)}, st_E^{(b)})$ 7: return $\text{sent}.c_u$ 	<p>Oracle Decrypt(ad, c):</p> <ol style="list-style-type: none"> 1: if $b = 0$ then 2: return \perp 3: $v \leftarrow v + 1$ 4: $\text{rcvd}.c_v \leftarrow c$ 5: $(\text{ad}, m, \alpha, st_D) \leftarrow D(k, \text{ad}, c, st_D)$ 6: if $(\alpha = 1) \wedge \text{cond}_i$ then 7: out-of-sync $\leftarrow 1$ 8: if out-of-sync = 1 then 9: return m 10: return \perp
<ol style="list-style-type: none"> 1. Basic authenticated encryption: $\text{cond}_1 = (\nexists w : (c = \text{sent}.c_w) \wedge (\text{ad} = \text{sent}.ad_w))$ 2. Basic authenticated encryption, no replays: $\text{cond}_2 = (\nexists w : (c = \text{sent}.c_w) \wedge (\text{ad} = \text{sent}.ad_w)) \vee (\exists w < v : c = \text{rcvd}.c_w)$ 3. Basic authenticated encryption, no replays, strictly increasing: $\text{cond}_3 = (\nexists w : (c = \text{sent}.c_w) \wedge (\text{ad} = \text{sent}.ad_w)) \vee (\exists w, x, y : (w < v) \wedge (\text{sent}.c_x = \text{rcvd}.c_w) \wedge (\text{sent}.c_y = \text{rcvd}.c_v) \wedge (x \geq y))$ 4. Basic authenticated encryption, no replays, strictly increasing, no drops: $\text{cond}_4 = (u < v) \vee (c \neq \text{sent}.c_v) \vee (\text{ad} \neq \text{sent}.ad_v)$ 	

Fig. 4: Stateful AEAD experiment aead_i with authentication condition cond_i for stateful AEAD scheme $\Pi = (\text{Kgn}, E, D)$ and adversary \mathcal{A} .

Analogously to Section 2.2, level- $(i+1)$ AEAD security implies level- i AEAD security. The details are omitted due to space restrictions.

3.2 Constructing higher level AEAD schemes

Similarly to Section 3, we can construct higher level AEAD schemes based on a level-1 AEAD scheme with the inclusion of sequence numbers with appropriate checks. We again generalize the approach using an encoding scheme that captures both implicit and explicit sequence numbers.

Definition 8 (AEAD encoding scheme). *An AEAD encoding scheme Coding for a sequence number space S , a message space \mathcal{M} , and an associated data space \mathcal{AD} is a pair of algorithms:*

- $\text{Ecd}(\text{sqn}, \text{ad}, m) \rightarrow (\text{ad}_{\text{ecd}}, m_{\text{ecd}})$: *A deterministic encoding algorithm that takes as input a sequence number $\text{sqn} \in S$, associated data $\text{ad} \in \mathcal{AD}$, and a message $m \in \mathcal{M}$, and outputs an encoded associated data value $\text{ad}_{\text{ecd}} \in \mathcal{AD}_{\text{ecd}}$ and message $m_{\text{ecd}} \in \mathcal{M}_{\text{ecd}}$, where $\mathcal{AD}_{\text{ecd}}$ and \mathcal{M}_{ecd} are the encoded versions of associated data space \mathcal{AD} and message space \mathcal{M} , respectively.*

$\Pi'_i.\text{Kgn}():$ 1: return $\Pi.\text{Kgn}()$	$\Pi'_i.\text{D}(k, \text{ad}, c, st'_D):$ 1: if $st'_D.\text{status} = \text{failed}$ then 2: return $(\perp, 0, st_D)$ 3: $(\text{ad}_\Pi, m_\Pi, \alpha, st'_D.\text{subst})$ $\leftarrow \Pi.\text{D}(k, \text{ad}, c, st'_D.\text{subst})$ 4: if $\alpha = 1$ then 5: $(\text{sqn}, \text{ad}, m, \alpha)$ $\leftarrow \text{Dcd}(st'_D.\text{sqnlist}, \text{ad}_\Pi, m_\Pi)$ 6: if $(\alpha = 0) \vee \text{TEST}i$ then 7: $st'_D.\text{status} = \text{failed}$ 8: return $(\perp, 0, st'_D)$ 9: $st'_D.\text{sqnlist} = st'_D.\text{sqnlist} \parallel \text{sqn}$ 10: return (m, α, st'_D)
$\Pi'_i.\text{E}(k, \ell, \text{ad}, m, st'_E):$ 1: (ad_Π, m_Π) $\leftarrow \text{Ecd}(st'_E.\text{ctr}, \text{ad}, m)$ 2: $(c, st'_E.\text{subst})$ $\leftarrow \Pi.\text{E}(k, m_\Pi, \text{ad}_\Pi, \ell, st'_E.\text{subst})$ 3: $st'_E.\text{ctr} \leftarrow st'_E.\text{ctr} + 1$ 4: return (c, st'_E)	

Description of states st'_E and st'_D :

- $st'_E.\text{subst} := st_E$, where st_E is the state in Π
- $st'_E.\text{ctr}$. When $\Pi'.\text{E}$ is initialized, $st'_E.\text{ctr} \leftarrow 0$.
- $st'_D.\text{subst} := st_D$, where st_D is the state in Π
- $st'_D.\text{status}$. Once $st'_D.\text{status} = \text{failed}$ it is not reset and all subsequently received messages are also immediately aborted.
- $st'_D.\text{sqnlist}$, an ordered list of sequence numbers previously received. It is required that $|st'_D.\text{sqnlist}| \geq 1$ after the first received sqn ; i.e. the size of the ordered set is maintained at 1 or greater. When $\Pi'.\text{E}$ is initialized, $st'_D.\text{sqnlist} \leftarrow \perp$.

Fig. 5: Construction P_{AEAD} of a level- i AEAD scheme Π'_i from a level-1 AEAD scheme Π and AEAD encoding scheme $\text{Coding} = (\text{Ecd}, \text{Dcd})$, with $\text{TEST}i$ as shown in Fig. 3.

- $\text{Dcd}(\text{sqnlist}, \text{ad}_{\text{ecd}}, m_{\text{ecd}}) \rightarrow (\text{sqn}, \text{ad}, m, \alpha)$: A deterministic decoding algorithm that takes as input a sequence number list $\text{sqnlist} \subset \mathcal{S}$, an encoded associated data value ad_{ecd} , and an encoded message $m_{\text{ecd}} \in \mathcal{M}_{\text{ecd}}$, and outputs a sequence number $\text{sqn} \in \mathcal{S}$, associated data $\text{ad} \in \mathcal{AD}$ or an error symbol \perp , a message $m \in \mathcal{M}$ or an error symbol \perp , and a status variable $\alpha = 1$ if decoding was successful or $\alpha = 0$ otherwise.

Definition 9. We say that AEAD encoding scheme Coding uses explicit sequence numbers if, for all sqn , ad , and m , when $\text{Ecd}(\text{sqn}, \text{ad}, m) = (\text{ad}_{\text{ecd}}, m_{\text{ecd}})$, we have that $\text{Dcd}(\perp, \text{ad}_{\text{ecd}}, m_{\text{ecd}}) = (\text{sqn}, \text{ad}, m, 1)$. Otherwise, we say that Coding uses implicit sequence numbers.

Definition 10 (P_{AEAD} construction). Let Π be a (level-1) AEAD scheme, Coding be an AEAD encoding scheme, and let $\text{TEST}i$ be a condition specified in Fig. 3. Define $\Pi'_i := P_{\text{AEAD}}(\Pi, \text{Ecd}, \text{TEST}i)$ as the AEAD scheme resulting from applying construction P_{AEAD} in Fig. 5.

Theorem 3. Let Π be a secure level-1 AEAD scheme and Coding be an AEAD encoding scheme with collision-resistant encoding. Let $\text{TEST}i$ be defined as in Fig. 3 and $i \in \{2, 3, 4\}$. Then $\Pi'_i = P_{\text{AEAD}}(\Pi, \text{Coding}, \text{TEST}i)$, constructed as in Fig. 5, is a secure level- i AEAD scheme. Specifically, let \mathcal{A} be an adversary

algorithm that runs in time t and asks q_e **Encrypt** queries and q_d **Decrypt** queries, and let $q = q_e + q_d$. Then there exists an adversary \mathcal{B} that runs in time $t_{\mathcal{B}} \approx t$ and asks no more than $q_{\mathcal{B}} = \frac{1}{2}q_e(q_e - 1)$ queries, and an adversary \mathcal{F} that runs in time $t_{\mathcal{F}} \approx t$ and asks $q_{\mathcal{F}} = q$ queries, such that $\mathbf{Adv}_{P_{\text{AEAD}}(\Pi, \text{Coding}, \text{TEST}i)}^{\text{aeadi}}(\mathcal{A}) \leq \mathbf{Adv}_{\Pi}^{\text{aeadi}_1}(\mathcal{F}) + \mathbf{Adv}_{\text{Ecd}}^{\text{collision}}(\mathcal{B})$.

The proof of Theorem 3 is omitted due to space restrictions.

4 Authenticated Encryption in TLS

The work of Paterson et al. [18] showed that the MAC-then-encode-then-encrypt mode of CBC encryption in TLS 1.2 (with sufficiently long MAC tags) is a secure length-hiding authenticated encryption (LHAE) scheme, assuming the encryption function is a strong pseudorandom permutation and the MAC is a pseudorandom function. Their definition corresponds to level 1 of our AEAD hierarchy. Several subsequent work on the provable security of TLS, such as that of Jager et al. [13] and Krawczyk et al. [18], assume that the TLS record layer is a secure *stateful* length-hiding authenticated encryption (sLHAE) scheme, corresponding to level 4 of our AEAD hierarchy. To our knowledge, there has as of yet been no formal connection between the LHAE result of Paterson et al. and the sLHAE requirement of subsequent works; we address that gap in this section by bringing sequence numbers into the modeling using the framework in the previous sections.

4.1 TLS sequence numbers and authentication level

The TLS record layer utilizes sequence numbers to ensure detection of deleted or reordered records [7, p. 94]. Being 64-bits long, sequence number exhaustion for any given connection is unlikely and the specification demands renegotiation should it occur. Sequence numbers are sent implicitly by inclusion under the MAC (or AEAD). When instantiated, “the first record transmitted under a particular connection state MUST use sequence number 0” [7, §6.1] and each subsequent record increments the sequence number. Sequence numbers are continuous across record types (application and alert).

When the ciphersuite uses MAC-then-encode-then-encrypt, the MAC tag is computed as follows, where k is the MAC key (either `MAC_write_key` or `MAC_read_key`, depending on the direction), `sqn` is the 64-bit sequence number, and `m` is the (possibly compressed) TLS plaintext object (called `TLSCompressed`) [7]: `MAC(k, sqn || m.type || m.version || m.length || m.fragment)`. Since the sequence number is implicit, a receiver will check the MAC verification using the expected sequence number. If the check fails, a `bad_record_mac` alert (type 20) will be generated – an alert that is always fatal [7, §7.2.2].

When the ciphersuite is uses a combined AEAD scheme, the sequence number, as well as several other values, are included in the additional data field [7]: `ad = sqn || m.type || m.version || m.length`. The ciphertext is then $c \leftarrow \text{Encrypt}(k, m.length, ad, m.fragment, st_E)$. The sequence number is not transmitted in the ciphertext. AEAD decryption is applied using the expected sequence number. Decryption failure must also result in a `bad_record_mac` fatal alert [7, §6.2.3.3].

$\Pi_{PRS}.E(k, \ell, \mathbf{ad}, m, \perp)$: 1: $(k_m, k_e) \leftarrow k$ 2: $t \leftarrow \text{MAC}(k_m, \mathbf{ad}, m)$ 3: $c \leftarrow E(k_e, \ell, m, t)$ 4: return (c, \perp)	$\Pi_{PRS}.D(k, \mathbf{ad}, c, \perp)$: 1: $(k_m, k_e) \leftarrow k$ 2: $(m, t, \alpha) \leftarrow D(k_e, c)$ 3: if $\text{MAC}(k_m, \mathbf{ad}, m) \neq t$ then 4: return $(\perp, 0, \perp)$ 5: return (m, α, \perp)
$\Pi_{TLS}.E(k, \ell, \mathbf{ad}, m, st_E)$: 1: $(k_m, k_e) \leftarrow k$ 2: $t \leftarrow \text{MAC}(k_m, st_E.\text{ctr} \parallel \mathbf{ad}, m)$ 3: $c \leftarrow E(k_e, \ell, m, t)$ 4: $st_E.\text{ctr} \leftarrow st_E.\text{ctr} + 1$ 5: return (c, st_E)	$\Pi_{TLS}.D(k, \mathbf{ad}, c, st_D)$: 1: $(k_m, k_e) \leftarrow k$ 2: if $st_D.\text{status} = \text{failed}$ then 3: return $(\perp, 0, st_D)$ 4: $(m, t, \alpha) \leftarrow D(k_e, c)$ 5: if $\text{MAC}(k_m, st_D.\text{ctr} \parallel \mathbf{ad}, m) \neq t$ then 6: $\alpha \leftarrow 0$ 7: if $\alpha = 0$ then 8: $st_D.\text{status} \leftarrow \text{failed}$ 9: return $(\perp, 0, st_D)$ 10: $st_D.\text{ctr} \leftarrow st_D.\text{ctr} + 1$ 11: return (m, α, st_D)

Fig. 6: Construction of AEAD schemes Π_{PRS} (Paterson et al. [23] variant of TLS MAC-then-encode-then-encrypt) and Π_{TLS} (TLS MAC-then-encode-then-encrypt) from encode-then-encrypt scheme (E, D).

4.2 From TLS Level-1 AEAD to Level-4 AEAD

Paterson et al. [23] show that a simplified version of TLS MAC-then-encode-then-encrypt, which we call Π_{PRS} and describe in the top half of Fig. 6, satisfies level-1 AEAD security. By design, Π_{PRS} includes the sequence number field in the \mathbf{ad} , but never initializes it as Π_{PRS} is not stateful. However, the TLS record layer protocol as actually used is stateful and, as such, ought to achieve a higher level of AEAD; namely, it should satisfy level-4 AEAD. The bottom half of Fig. 6 shows the TLS MAC-then-encode-then-encrypt record layer with the use of sequence numbers as specified in the standard.

Our framework allows us to immediately show that Π_{TLS} satisfies level-4 AEAD security: we incorporate the sequence numbers in an implicit AEAD encoding scheme Coding_{TLS} , and then view Π_{TLS} as the result of applying the P_{AEAD} construction to Π_{PRS} and Coding_{TLS} .

Define AEAD encoding scheme $\text{Coding}_{TLS} = (\text{Ecd}_{TLS}, \text{Dcd}_{TLS})$ as follows:

- $\text{Ecd}_{TLS}(\text{sqn}, \mathbf{ad}, m) = (\text{sqn} \parallel \mathbf{ad}, m)$
- $\text{Dcd}_{TLS}(\text{sqnlist}, \text{sqn} \parallel \mathbf{ad}, m) = (\text{sqn}, \mathbf{ad}, m, \alpha)$

where $\alpha = 1$ if and only if sqn and sqnlist satisfy TEST4 in Fig. 3, $\mathbf{ad} \neq \perp$, and $m \neq \perp$.

Theorem 4. $\Pi_{TLS} = P_{AEAD}(\Pi_{PRS}, \text{Coding}_{TLS}, \text{TEST4})$.

Theorem 4 follows semantically comparing Π_{TLS} and the scheme resulting from the construction $P_{AEAD}(\Pi_{PRS}, \text{Coding}_{TLS}, \text{TEST4})$.

Clearly, Ecd_{TLS} is collision-resistant due to the unambiguous parsing of sqn as a fixed-length 64-bit value. We can thus apply Theorem 3 to obtain Corollary 1.

Corollary 1. *The TLS record layer with MAC-then-encode-then-encrypt in CBC mode satisfies level-4 AEAD security. Specifically, let \mathcal{A} be an adversary algorithm that runs in time t against Π_{TLS} . Then there exists an adversary \mathcal{F} that runs in time $t_{\mathcal{F}} \approx t$ such that $\text{Adv}_{\Pi_{TLS}}^{\text{aead}_4}(\mathcal{A}) \leq \text{Adv}_{\Pi_{PRS}}^{\text{aead}_1}(\mathcal{F})$.*

From Paterson et al. [23] we know that the TLS record layer encryption in MAC-then-encode-then-encrypt CBC mode satisfies AEAD level-1 security when a secure cipher and message authentication code is used. Combined with Corollary 1, this means that the sLHAE security definition used by Jager et al. [13] and Krawczyk et al. [18] in their analyses of full TLS ciphersuites is achieved, and thus TLS is ACCE secure in this scenario.

References

1. Albrecht, M.R., Paterson, K.G., Watson, G.J.: Plaintext recovery attacks against SSH. In: 2009 IEEE Symposium on Security and Privacy. pp. 16–26. IEEE Computer Society Press (May 2009)
2. Badertscher, C., Matt, C., Maurer, U., Rogaway, P., Tackmann, B.: Augmented secure channels and the goal of the TLS 1.3 record layer. Cryptology ePrint Archive, Report 2015/394 (2015), <http://eprint.iacr.org/2015/394>
3. Bellare, M., Kohno, T., Namprempre, C.: Authenticated encryption in SSH: Provably fixing the SSH binary packet protocol. In: Atluri, V. (ed.) ACM CCS 02. pp. 1–11. ACM Press (Nov 2002)
4. Bellare, M., Namprempre, C.: Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 531–545. Springer, Heidelberg (Dec 2000)
5. Boldyreva, A., Degabriele, J.P., Paterson, K.G., Stam, M.: Security of symmetric encryption in the presence of ciphertext fragmentation. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 682–699. Springer, Heidelberg (Apr 2012)
6. Canetti, R., Krawczyk, H.: Analysis of key-exchange protocols and their use for building secure channels. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 453–474. Springer, Heidelberg (May 2001)
7. Dierks, T., Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.2 (2008), <https://tools.ietf.org/html/rfc5426>, RFC 5426
8. Fischlin, M., Günther, F., Marson, G.A., Paterson, K.G.: Data is a stream: Security of stream-based channels. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part II. LNCS, vol. 9216, pp. 545–564. Springer, Heidelberg (Aug 2015)
9. Fleischmann, E., Forler, C., Lucks, S.: McOE: A family of almost foolproof on-line authenticated encryption schemes. In: Canteaut, A. (ed.) FSE 2012. LNCS, vol. 7549, pp. 196–215. Springer, Heidelberg (Mar 2012)
10. Hoang, V.T., Krovetz, T., Rogaway, P.: Robust authenticated-encryption: AEZ and the problem that it solves. Cryptology ePrint Archive, Report 2014/793 (2014), <http://eprint.iacr.org/2014/793>
11. Hoang, V.T., Reyhanitabar, R., Rogaway, P., Vizár, D.: Online authenticated-encryption and its nonce-reuse misuse-resistance. Cryptology ePrint Archive, Report 2015/189 (2015), <http://eprint.iacr.org/2015/189>
12. IEEE 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications (2012), <http://dx.doi.org/10.1109/IEEESTD.2012.6178212>

13. Jager, T., Kohlar, F., Schäge, S., Schwenk, J.: On the security of TLS-DHE in the standard model. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 273–293. Springer, Heidelberg (Aug 2012)
14. Katz, J., Yung, M.: Unforgeable encryption and chosen ciphertext secure modes of operation. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 284–299. Springer, Heidelberg (Apr 2001)
15. Kent, S.: IP Authentication Header (2005), <https://tools.ietf.org/html/rfc4302>, RFC 4302
16. Kohno, T., Palacio, A., Black, J.: Building secure cryptographic transforms, or how to encrypt and MAC. Cryptology ePrint Archive, Report 2003/177 (2003), <http://eprint.iacr.org/2003/177>
17. Krawczyk, H.: The order of encryption and authentication for protecting communications (or: How secure is SSL?). In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 310–331. Springer, Heidelberg (Aug 2001)
18. Krawczyk, H., Paterson, K.G., Wee, H.: On the security of the TLS protocol: A systematic analysis. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 429–448. Springer, Heidelberg (Aug 2013)
19. Langley, A., Chang, W.T.: QUIC Wire Layout Specification (2015), https://docs.google.com/document/d/1WJvyZf1A02pq77y0Lbp9NsGjC1CHetAXV8I0fQe-B_U
20. Lychev, R., Jero, S., Boldyreva, A., Nita-Rotaru, C.: How secure and quick is QUIC? Provable security and performance analyses. In: 2015 IEEE Symposium on Security and Privacy. pp. 214–231. IEEE Computer Society Press (May 2015)
21. Maurer, U., Tackmann, B.: On the soundness of authenticate-then-encrypt: formalizing the malleability of symmetric encryption. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) ACM CCS 10. pp. 505–515. ACM Press (Oct 2010)
22. Namprempre, C.: Secure channels based on authenticated encryption schemes: A simple characterization. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 515–532. Springer, Heidelberg (Dec 2002)
23. Paterson, K.G., Ristenpart, T., Shrimpton, T.: Tag size does matter: Attacks and proofs for the TLS record protocol. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 372–389. Springer, Heidelberg (Dec 2011)
24. Rescorla, E., Modadugu, N.: Datagram Transport Layer Security (2006), <https://tools.ietf.org/html/rfc4347>, RFC 4347
25. Rescorla, E., Modadugu, N.: Datagram Transport Layer Security Version 1.2 (2012), <https://tools.ietf.org/html/rfc6347>, RFC 6347
26. Rogaway, P.: Authenticated-encryption with associated-data. In: Atluri, V. (ed.) ACM CCS 02. pp. 98–107. ACM Press (Nov 2002)
27. Rogaway, P., Bellare, M., Black, J., Krovetz, T.: OCB: A block-cipher mode of operation for efficient authenticated encryption. In: ACM CCS 01. pp. 196–205. ACM Press (Nov 2001)
28. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer, Heidelberg (May / Jun 2006)
29. Shrimpton, T.: A characterization of authenticated-encryption as a form of chosen-ciphertext security. Cryptology ePrint Archive, Report 2004/272 (2004), <http://eprint.iacr.org/2004/272>