

Rethinking Zero-Knowledge TLS: Proof of Protocol Execution and the Reclaim Protocol

Andrey Bozhko[†] and Kirill Kutsenok^{*}

[†]University of Luxembourg

andrey.bozhko@uni.lu

^{*}Reclaim Protocol and King's College London

kirill@reclaimprotocol.org

September 23, 2026

Abstract

In a zkTLS protocol, a user proves to a third party a property of data obtained from an online service over TLS, revealing nothing beyond the statement proven. Each existing scheme, however, comes with its own security model, written for that scheme alone, so the schemes are not proven secure against a common definition. The root cause is a mismatch of levels: existing models are stated at the level of TLS and therefore reflect how each construction handles it. Yet TLS is only the transport, and a zkTLS proof is about the application-layer protocol executed over it. We therefore introduce the first transport-agnostic security model for zkTLS, stated at the application layer and mentioning neither TLS nor any particular construction. The resulting notion, *Zero-Knowledge Proof of Protocol Execution* (zkPoPE), defines what it means for a party that has executed an arbitrary two-party protocol with a counterparty to prove, in zero knowledge, that its private input and the reply it received satisfy a chosen property, and, crucially, that the reply comes from a real execution with that counterparty rather than being fabricated. This proof-of-execution requirement separates zkPoPE from a NIZK over the input–reply pair: knowing a satisfying pair does not by itself show that an execution occurred. We formalize zkPoPE as a UC ideal functionality $\mathcal{F}_{\text{zkPoPE}}$. We then introduce the *Reclaim Protocol*, a zkTLS scheme that UC-realizes $\mathcal{F}_{\text{zkPoPE}}$ under a deployment-aligned trust model and underlies a production system. Reclaim builds on *Distributed-AEAD*, a new three-party authenticated-encryption primitive of independent interest. On a lower-mid-range phone, Reclaim attests to a 1200 kB response in 7.28 s, two orders of magnitude beyond any size previously benchmarked for zkTLS.

Contents

1	Introduction	3
2	Technical Overview	5
2.1	Zero-Knowledge Proof of Protocol Execution	5
2.2	Reclaim Protocol	7
2.3	Open Problems	11
3	Zero-Knowledge Proof of Protocol Execution	11
3.1	System Model and Parties	11
3.2	Ideal Functionality	12
3.3	Relation to Existing Models	14
4	Distributed AEAD	14
4.1	Syntax	15
4.2	Security	16
5	The Reclaim Protocol	16
5.1	Hybrid model and interface	16
5.2	Admissible Protocols and Predicates	16
5.3	Protocol $\mathcal{RP}$	18
5.4	Security	20
6	Benchmarks	21
A	Message Flow of the Reclaim Protocol	26
B	Preliminaries	26
B.1	Notation	26
B.2	Request–Reply Protocols and Request Extractability	27
B.3	Ideal Functionalities	29
C	Security of Distributed-AEAD	31
C.1	Security Properties	32
C.2	Assumptions on the underlying primitives	34
C.3	Main theorem	35
C.4	Concrete schemes	37
D	Reclaim Protocol Security Proof	37
D.1	Model and conventions	38
D.2	Simulator	38
D.3	Indistinguishability: Case B (soundness)	40
D.4	Indistinguishability: Cases A, C, D (zero-knowledge)	41
D.5	Wrap-up	43

1 Introduction

Zero-knowledge TLS (zkTLS) protocols, also known as TLS oracles [ZMM⁺20], allow a User acting as the TLS client to convince a third party of statements about the data exchanged over a TLS connection. A zkTLS proof attests that the session was established with a specific server, identified by its TLS certificate, and that the application data exchanged over it satisfies a public predicate Φ , while revealing nothing further about the connection or its transcript. In the canonical example, a proof of solvency, a User logs into their bank’s web portal over TLS and proves that their balance exceeds a threshold T , revealing neither their credentials, nor their account number, nor the exact balance. Any verifier, be it a lender, an exchange, or a KYC provider, can accept this proof without involving the User or the bank in the verification. The zero-knowledge component is not optional: TLS-protected traffic typically carries credentials, session cookies, and unrelated personal data alongside the value being proved, so every realistic use of zkTLS must hide nearly the entire session.

The TLS transcript is not itself such a proof. TLS [RD08, Res18] authenticates the server to the client and protects application data against a network adversary, but its authentication binds the two endpoints alone. After the handshake, the User holds the symmetric keys for *both* directions and can produce a syntactically valid TLS record stream that decrypts to any plaintext of its choice. A transcript captured by the User is therefore not, by itself, evidence of what the server sent. One way to close this gap is to have the server sign its responses, and a standard for it exists in HTTP Message Signatures [BRS24]. Mechanisms in this spirit are deployed: Stripe authenticates the webhook events it delivers [Str], but what such mechanisms call a signature is a MAC under a secret shared between the two endpoints, and it convinces the counterparty alone, since anyone able to verify it could equally have produced it. Third-party-verifiable response signing has seen no broad deployment, and none as a general mechanism for arbitrary websites. The obstacle is usually the provider, which has no incentive to make the data it serves verifiable by third parties, and equally none to run any other modified software for the benefit of Users or third parties. zkTLS designs therefore impose *server-obliviousness*: to the server, the execution must look like an ordinary TLS session.

With the server unavailable and the client able to forge the transcript on its own, evidence must come from elsewhere: every known zkTLS protocol introduces auxiliary parties, which we call *attestors*, to ensure that the User cannot emit a proof inconsistent with what was exchanged. The *Reclaim Protocol*, which we introduce and analyze, has two attestors, each running in a trusted execution environment (TEE).

A security analysis of Reclaim needs a security definition, and none in the literature serves. Since DECO [ZMM⁺20], zkTLS constructions come with their own ideal functionalities or bespoke models, each written for one construction and bound to the specifics of TLS, a habit we call *protocol-inseparability*. DECO’s definition, for instance, requires the server side of any realizing protocol to run standard TLS, and ORIGO’s security theorem [ELW⁺24] is stated for a single TLS 1.3 cipher suite. Two zkTLS protocols are thus hard to compare, since they realize distinct functionalities, and a construction not built over TLS cannot be evaluated against this literature at all.

The root of this habit is a mismatch of levels. TLS is, in the end, only the transport layer over which a User and a server execute some application-layer protocol¹, such as a bank’s balance-inquiry flow. The property a User wishes to prove about a TLS session is a property of *that* application-layer execution rather than of the TLS bytes carrying it. Stated at the level of the application-layer protocol, the definition does not mention the transport, and the same problem statement covers other secure transports (e.g., IPsec [FK11] or QUIC [IT21]), exchanges that no transport signs for third parties, such as a conversation with an AI service, and even deniable channels such as end-to-end encrypted messaging [CGCD⁺17]. The latter are not a counterexample but the extreme case: the transcript proves nothing by design, so all evidentiary weight falls on the proof-of-execution machinery. The common object is a party’s participation in a protocol and

¹We use *application-layer* for everything above the transport: the message syntax carrying the exchange, such as HTTP, *together with* the business logic offered through it. The two are conceptually separable, and the IETF’s guidance treats the latter, which has no settled name of its own, as an application built on HTTP as a “substrate” [Not22]. We take them as a single object, since the messages of such a protocol are by definition the byte strings handed to the transport, and the properties we prove of them are stated over those bytes.

the wish to prove a property of it in zero knowledge, whatever the transport. To our knowledge, no zkTLS security definition has been stated at this level.

Stating the definition at this level is the central move of this paper. We call the resulting notion *Zero-Knowledge Proof of Protocol Execution* (zkPoPE), and the conventional zkTLS problem is then its instance in which the proven protocol runs over TLS.² zkPoPE is closely related to the problem NIZKs solve, and Section 2 obtains its functionality as a minimal adaptation of the NIZK one. It is nonetheless not a NIZK: its language cannot be checked without the counterparty’s private state, it admits no meaningful witness extraction, and a proof must attest to an *interaction* rather than to a computation the prover could run alone. Each difference forces a departure from the NIZK template. Beyond that, the definition restores comparability, since constructions for all the settings above become candidate realizations of a single notion, and it is forward-compatible: whatever transport eventually displaces TLS, and however deeply the post-quantum migration reshapes the cryptography beneath it [Con26], the model stays as it is.

Reclaim’s trust assumptions expose a second gap in the zkTLS literature. Existing constructions form two families. *Cryptographic-attestor* designs use no trusted hardware: *proxy* variants observe TLS ciphertexts on the wire and commit to them before the User releases the key material [LEFS23, ELW⁺24], while *MPC* variants compute the handshake and record encryption jointly with the User, so that neither party alone holds the traffic keys [ZMM⁺20, KK24, XYWY23, CDH⁺25]. Here theorems and deployments describe different situations: the theorems are stated for n attestors under a k -of- n honesty assumption, while the protocols are deployed with two or three of them. Closing this gap requires an argument that the security analysis does not supply, namely that the named attestors are independently operated, have an economic stake in honest behavior, and so on. The substantive trust assumption is thus the behavioral honesty of a specific small committee, justified outside the protocol. The *TEE-based* family, exemplified by Town Crier [ZCC⁺16], runs the TLS client inside a TEE such as Intel SGX [Int15] and emits signed transcripts whose authenticity reduces, via remote attestation, to the integrity of the TEE. It is proven and deployed in the same configuration, but relies on the integrity *and confidentiality* of an entire TEE-resident TLS client. Confidentiality is required since the transported data is itself sensitive, the only case in which zkTLS is of interest at all. The long line of side-channel and microarchitectural attacks on commodity TEEs [VBMW⁺18, vSKGY20, CSB⁺26] has made this confidentiality assumption hard to maintain, and the approach has been left dormant. Each family nonetheless has an element worth preserving: cryptographic-attestor designs spread trust across several parties, so that no single party holds material catastrophic if compromised, while the TEE-based approach is proven and deployed in the same configuration. What is missing is a construction that keeps both elements without the problematic assumptions just discussed and without the cost the cryptographic-attestor branch imposes on the User, whose device runs multi-party computation or zero-knowledge proof generation at a cost that grows sharply with the amount of data proven.

This work makes two contributions: the security model and a zkTLS construction that realizes it in the TLS setting. More specifically:

- (1) **The first transport-agnostic security model for zkTLS: zkPoPE.** The model is stated for two-party protocol executions, independently of construction, transport, and application-layer protocol. We define, in the UC framework [Can01], an ideal functionality $\mathcal{F}_{\text{zkPoPE}}$ for the User, Service, and Verifier roles, obtained from the NIZK functionality through the departures that proof of execution forces. We regard the introduction of zkPoPE and of $\mathcal{F}_{\text{zkPoPE}}$ as the principal conceptual contribution of this work, because the object it isolates, proof of execution, is one that no proof system for a locally checkable relation can express.
- (2) **The Reclaim Protocol.** We present the *Reclaim Protocol*,³ a zkTLS protocol that builds on a new three-party authenticated-encryption primitive, *Distributed-AEAD*. The latter splits the encryption of a record under unmodified AES-GCM or ChaCha20-Poly1305 among three roles, so that the key holder never sees the bytes the data owner marks secret, while the data owner never holds the key. We prove it secure

²Strictly applied, the naming logic of zkPoPE would have zkTLS denote a scheme proving execution of TLS itself, whereas existing zkTLS schemes, Reclaim included, prove execution of an application-layer protocol that runs over TLS. We retain zkTLS in line with established usage.

³An informal description of an early version of the Reclaim Protocol appeared in a whitepaper [Pro25] on the deploying company’s website. It contains none of the security model, Distributed-AEAD, functionality, or proof presented here.

under standard assumptions on those primitives. Reclaim runs against an unmodified TLS server, and we prove it UC-realizes $\mathcal{F}_{\text{zkPoPE}}$ under stated assumptions on its two attestors and for an explicitly limited class of protocols and predicates: request-reply protocols from which the User’s input can be extracted, as holds for public request formats such as HTTP, with predicates over the request and reply bytes the User has chosen not to hide. It keeps both elements identified above: trust spread across two attestors, and proof and deployment in the same configuration under narrowly scoped TEE guarantees. The User in the Reclaim Protocol runs neither multi-party computation nor proof generation, and its cost scales only with the size of the response rather than with the amount of data proven. On a lower-mid-range phone, proving responses of up to 25 kB takes about 3 s end to end, and a 1200 kB response 7.28 s, against a previous zkTLS maximum of 8 kB in 6.2 s [LEFS23]. Reclaim underlies a deployed production system, but we claim no software contribution.

Organization. Section 2 gives a technical overview of the contributions. Section 3 formalizes zkPoPE and $\mathcal{F}_{\text{zkPoPE}}$ and relates them to existing models (Section 3.3). Section 4 introduces Distributed-AEAD, Section 5 specifies Reclaim, and Section 5.4 states the theorem with a proof sketch. Section 6 reports the benchmarks. Appendices hold the preliminaries (Appendix B), the Distributed-AEAD security analysis (Appendix C), and the full proof (Appendix D).

2 Technical Overview

This section gives an informal account of the two contributions. Section 2.1 derives the zkPoPE definition: we fix the setting, state what the User claims, identify the two requirements this claim places on a proof system, and obtain the ideal functionality $\mathcal{F}_{\text{zkPoPE}}$ from the NIZK one through the departures those requirements force. Section 2.2 describes Reclaim, our realization of $\mathcal{F}_{\text{zkPoPE}}$ over TLS, and the Distributed-AEAD primitive at its core.

2.1 Zero-Knowledge Proof of Protocol Execution

2.1.1 Setting

As argued in Section 1, the object of a zkTLS proof is the execution of an application-layer protocol, and the transport carrying it does not enter the statement. We model the participants at that level and restrict to two-party protocols, which covers the zkTLS class, and we call the parties U and $Service$.⁴

Let $\mathcal{US} = \{\text{Service}_{\text{id}}\}_{\text{id} \in \mathcal{ID}}$ be a universe of Services. Each $\text{Service}_{\text{id}}$ is initialized by some setup procedure that produces a triplet $(\text{id}, \text{sec}_{\text{id}}, \text{Prot}_{\text{id}})$: a unique identifier, the Service’s private state (internal databases, per-user records), and a public description of the protocol a User can run with it. We write $\text{Service}_{\text{id}}$ for both the Service and its triplet. Besides the Services, there are a User U and a Verifier V .

The User runs Prot_{id} with $\text{Service}_{\text{id}}$ on a private input data of the format Prot_{id} expects⁵ and obtains

$$\text{reply} = \text{out}_U(\text{Prot}_{\text{id}}(\text{Service}_{\text{id}}(\text{sec}_{\text{id}}), U(\text{data}))),$$

its local output at the end of the interaction. The Service obtains a local output $\text{out}_{\text{Service}}$, which may update sec_{id} . The User then wishes to prove a property of $(\text{data}, \text{reply})$ to the Verifier. The two agree, possibly out of band, on a predicate Φ over the input-reply space of Prot_{id} .⁶ Φ may depend on parts of data the User does not wish to hide, such as the public part of a query, while credentials stay secret. The User’s goal is to convince the Verifier that the protocol was executed and that $\Phi(\text{data}, \text{reply}) = 1$.

⁴We write “Service” rather than “server” since the object of interest is the application-layer functionality, not the host running it. Protocols Prot_{id} with more than two parties are left to future work (Section 2.3).

⁵We make no assumption about how the input is chosen. Inputs of incorrect format are rejected by the Service as per the specification of Prot_{id} .

⁶ Φ may be specific to one User-Verifier pair or standardized across many sessions of Prot_{id} , for example, a predicate fixed for a class of governmental web services with a common interface. We make no assumption either way.

2.1.2 Two requirements

In proof-system terms, the User proves knowledge of an input $\mathbf{data}$ for which Φ holds on the reply of an execution of Prot_{id} with $\text{Service}_{\text{id}}$. The statement is (id, Φ) , and we write $\mathcal{L}_{\mathcal{US}}$ for the language of such statements (Section 3.1). The witness is $\mathbf{data}$ alone: the reply is determined by the execution, not supplied by the prover. The relation cannot be evaluated without sec_{id} , so membership is relative to the deployment $\mathcal{US}$ rather than to the public descriptions.

The first requirement is zero-knowledge in the standard sense: the proof reveals nothing about $(\mathbf{data}, \text{reply})$ beyond Φ and the fact that the pair satisfies it. This fixes what “knowledge” means in the “zk” of zkTLS and zkPoPE.

The second requirement is that the proof attest to an execution of Prot_{id} between the User and $\text{Service}_{\text{id}}$, for an *arbitrary* application-layer protocol. No proof system can check the relation of $\mathcal{L}_{\mathcal{US}}$ directly, since that needs sec_{id} . The fallback, a NIZK for the locally checkable relation “ $\Phi(\mathbf{data}, \text{reply}) = 1$ ” on a User-supplied pair, certifies knowledge of the pair but not execution. The two coincide for some pairs: if reply carries a Service signature over part of $\mathbf{data}$ (as in [BRS24]) and Φ verifies it, knowledge of $(\mathbf{data}, \text{reply})$ is evidence of execution. For most protocols of interest they do not: a reply may be no more than 200 OK, and a User can satisfy “ $\mathbf{data}$ contains the recipient’s address and reply is OK” by writing the pair down, whether or not a transfer took place. This separates zkPoPE from every proof system for a locally checkable relation, and is why every zkTLS scheme adds an execution-binding mechanism (attestors, MPC, or TEEs). We leave the characterization of the pairs for which a NIZK suffices to future work (Section 2.3).

2.1.3 The ideal functionality

zkPoPE thus asks for the zero-knowledge of a NIZK over $(\mathbf{data}, \text{reply})$ together with a guarantee that the pair was produced by an execution of Prot_{id} with the registered $\text{Service}_{\text{id}}$. We adapt the NIZK formalization to the second requirement in the UC framework [Can01]. The out-of-band parts of the setting map onto the environment, and the resulting ideal functionality (IF) $\mathcal{F}_{\text{zkPoPE}}$ is a building block for applications consuming zkPoPE proofs.

$\mathcal{F}_{\text{zkPoPE}}$ keeps the interface of the NIZK IF: **Prove** and **Verify** with the semantics of [GOS06], and a **Setup** message as in formulations that place the scheme’s algorithms inside the functionality [LR25]. In the NIZK IF:

- **Setup** initializes the functionality, after which all parties hold a CRS.
- **Prove** takes a statement and a witness from a party and checks them against the relation. On success, it obtains a proof from the simulator Sim , returns it to the party, and records it. Drawing the proof from Sim captures zero-knowledge.
- **Verify** takes a statement and a proof and checks them against the record. If they are not recorded, it asks Sim for a witness and checks the relation on it. This captures knowledge soundness.

$\mathcal{F}_{\text{zkPoPE}}$ departs from this template at three points, one for each difference from NIZK named in Section 1. The formal specification is given in Figure 2.

Setup is relative to a deployment. Since membership in $\mathcal{L}_{\mathcal{US}}$ depends on the Services’ private states, **Setup** is invoked twice. A Service first registers its triplet $(\text{id}, \text{sec}_{\text{id}}, \text{Prot}_{\text{id}})$ and is added to the universe $\mathcal{US}_{\text{sid}}$. Only afterwards do the remaining parties invoke **Setup**, supplying the set of Services they expect to interact with. The IF checks it against $\mathcal{US}_{\text{sid}}$ and asks the simulator for the scheme’s algorithms and the CRS.

Prove runs the protocol. This is the central modification and carries both requirements. A User-supplied $(\mathbf{data}, \text{reply})$ does not witness execution, so $\mathcal{F}_{\text{zkPoPE}}$ accepts no reply from the User: on input $(\text{sid}, \text{Prove}, \text{id}, \Phi, \mathbf{data})$ it runs Prot_{id} itself between $\text{Service}_{\text{id}}(\text{sec}_{\text{id}})$ and $\text{U}(\mathbf{data})$, hands the genuine reply to the User and $\text{out}_{\text{Service}}$ to the Service, evaluates $\Phi(\mathbf{data}, \text{reply})$, and only on success asks the simulator for a proof. The simulator learns (P, Φ, id) , the verdict, and the length of reply , but never $\mathbf{data}$ or reply : this is the zero-knowledge content. The reply comes from an interaction with the registered $\text{Service}_{\text{id}}$, whose ideal-world view thus matches a real

run: this is the proof-of-execution content. A corrupted prover may name Φ only after receiving **reply**. Since it holds the pair and could evaluate any predicate on it anyway, nothing false becomes provable.

Verify is record-based. $\mathcal{F}_{\text{zkPoPE}}$ accepts $(P, \Phi, \text{id}, \pi)$ if and only if this tuple appears in the list Π_{id} of proofs issued via **Prove**. Soundness holds by construction: every accepted proof was produced from a real execution in which $\Phi(\text{data}, \text{reply}) = 1$. There is no extractor: the relation of $\mathcal{L}_{\text{US}}$ cannot be checked without sec_{id} , and an extracted witness would not certify execution. A proof is valid because $\mathcal{F}_{\text{zkPoPE}}$ issued it.

2.2 Reclaim Protocol

Reclaim realizes $\mathcal{F}_{\text{zkPoPE}}$ when the transport is TLS, against a Service S_{id} running an unmodified TLS server. Figure 1 shows the parties and their connections.




Figure 1: Parties of the Reclaim protocol. U , A_K and A_T are pairwise connected by secure channels $\mathcal{F}_{\text{SC}}$ (double lines); U talks to S_{id} over the public network; V receives only the proof. Figure 5 in Appendix A gives the message flow.

Supported protocols and predicates. Reclaim supports a restricted class of protocols and predicates: *request-reply* protocols with *extractable* requests. In a request-reply protocol, the User sends one message and the Service answers with one, as in a web-API call. A request is extractable if from it one can recover an input on which the protocol regenerates it. On the predicate side, Reclaim supports conditions on the bytes of the request and the response at positions Φ leaves disclosed, subject to three requirements. The condition must depend on the User’s input only through the request that input produces, it must not be satisfied by the reply a Service returns to a malformed request, and it must fix the byte at every position it does not hide, since the Verifier sees all of them.


Extractability is assumed rather than enforced because of Reclaim’s cost target: keeping the User’s work small rules out having it prove anything about the request it sends. Section 5.2 states the conditions informally alongside the protocol, and Appendix B.2 states them formally.

Distributing trust between two attestors. A single attestor holding the traffic key, decrypting everything, and signing the result would make proof of execution easy and destroy zero-knowledge: it would see **data** and **reply** in full. Reclaim splits the job between two attestors, A_K and A_T , and partitions the roles so that no single view suffices to recover the User’s session data. Both attestors are semi-honest, and at most one of them is corrupted in any execution, or alternatively the User is malicious and both attestors are honest. A_K holds the traffic key but sees the request only with the private bytes one-time-padded out, while A_T sees the wire ciphertexts and the pad but never the key. The User holds neither the key nor any MAC material, and so cannot forge a record under the session.

Distributed-AEAD. The mechanism behind this separation, and the technical core of Reclaim, is a new primitive, *Distributed-AEAD*. Producing an AES-GCM or ChaCha20-Poly1305 record⁷ has two stages. First, a keystream and a MAC subkey are derived from the traffic key and the nonce, independently of the message. Encryption and tagging are then a public computation on these and the plaintext. Schematically, with the

⁷The mandatory-to-implement and recommended cipher suites of TLS1.3 [Res18] and the dominant choices in deployed TLS1.2 [RD08].

nonce and associated data left implicit:



The key and the plaintext therefore need not co-reside. Distributed-AEAD distinguishes three roles: a *key party*, which holds the traffic key and emits the derived material, a *data party*, which owns the message, and a *tag party*, which combines the two into the ciphertext-and-tag pair.

The data party (the User) one-time-pads the bytes it wants private and sends the padded message to the key party, which thus encrypts a message equal to the original on the public bytes and uniformly random on the private ones. The data party sends the pad to the tag party. Encryption under both cipher suites is an XOR with the keystream, so XOR'ing the pad out of the resulting ciphertext yields the encryption of the *original* message, bit for bit, without exposing the private bytes to the tag party. With K_s denoting the keystream derived from the key and the nonce, the cancellation is explicit:

$$\begin{aligned}
 & \underbrace{(P \oplus \text{mask}) \oplus K_s}_{\text{key party encrypts the masked plaintext}} \oplus \underbrace{\text{mask}}_{\text{tag party removes the mask}} \\
 &= P \oplus K_s \oplus \underbrace{(\text{mask} \oplus \text{mask})}_0 = C, \\
 & \text{mask}[i] = \begin{cases} \text{an independent uniform byte,} & i \text{ is private,} \\ 0, & \text{otherwise.} \end{cases}
 \end{aligned}$$

The tag party then tags the recovered ciphertext with the MAC subkey. To the Service, the record looks like one an ordinary TLS client produced.

Decryption of the response runs in reverse, with one ordering constraint: the tag party receives the ciphertext-tag pair, obtains the MAC material from the key party, and verifies the tag, and only then does the key party release the keystream to the User, who decrypts in private. Thus, at the User,

$$\underbrace{Resp}_{\text{response plaintext}} = \underbrace{Resp_{\text{enc}}}_{\text{verified response ciphertext}} \oplus \underbrace{K_{\text{resp}}}_{\text{keystream released only after tag verification}}.$$

When the tag is verified, the User holds no key-derived material, so the verified tag certifies the ciphertext the Service emitted under keys from a server-authenticated handshake. This chain, from certificate through handshake keys to response tag, links the response the User observes to an execution with the registered $\text{Service}_{\text{id}}$, which is what Prove of $\mathcal{F}_{\text{zkPoPE}}$ requires of reply .

The four-phase flow. Reclaim wraps these two Distributed-AEAD calls, one for the request and one for the response,⁸ in a four-phase flow. In Phase A, A_K runs the TLS handshake against S_{id} with the User mediating the network. A_K is the cryptographic peer, while the User owns the endpoint S_{id} observes and forwards messages between them. Thus S_{id} sees a vanilla TLS handshake with a vanilla client, and server-obliviousness holds at the level of the program S_{id} executes. Phase B encrypts the request via Distributed-AEAD and forwards the record to S_{id} , and Phase C handles the response by the tag-then-keystream procedure.

The proof transcript. Phase D produces the proof. The redaction sets are those fixed by Φ . The User has named the request set in Phase B and names the response set now. Each attestor signs the slice of

⁸We describe a single request-response exchange: the User obtains session state such as a cookie by ordinary use of the Service, then issues one authenticated request through Reclaim. Multi-round exchanges are a routine extension we omit.

the joint view its role exposes: A_K signs its view of the request, zeroed at the redacted positions, together with the response keystream restricted to the non-redacted positions, and A_T signs the response ciphertext it verified in Phase C, likewise zeroed. Both signatures cover the redaction set and a per-session tag fixed in the handshake, which binds them to one session and one redaction. Neither attester signs the response plaintext, which neither of them held. Moreover, A_T checks in Phase B that the pad touches only redacted positions, so that A_K 's view agrees with the request actually sent everywhere else.

The zeroing matters: A_K 's view of the request is the padded plaintext and A_T holds the pad, so a proof carrying the padded bytes at redacted positions would let a corrupted A_T unmask them. The Verifier reads the request-side inputs of Φ off A_K 's signed view, which equals the plaintext at the non-redacted positions, and obtains the response-side inputs by XOR'ing A_T 's signed ciphertext with A_K 's signed keystream there. On the response side, this reconstruction is position-wise:

$$\begin{aligned} \underbrace{\widehat{Resp}[i]}_{\text{verifier's reconstructed response}} &= \underbrace{\widehat{Resp}_{\text{enc}}[i]}_{\text{ciphertext signed by the tag party}} \oplus \underbrace{\tilde{K}_{\text{resp}}[i]}_{\text{keystream signed by the key party}} \\ &= \begin{cases} Resp_{\text{enc}}[i] \oplus K_{\text{resp}}[i] = Resp[i], & i \text{ is disclosed,} \\ 0 \oplus 0 = 0, & i \text{ is hidden.} \end{cases} \end{aligned}$$

Φ is then evaluated on the resulting partial views.

Efficiency. Across all parties, Reclaim's computation is that of one TLS execution between A_K and S_{id} , plus the redactions, two signatures, and a few off-path messages. On the User's device, it is a subset of a TLS client's work, namely the application logic and message forwarding, plus sampling a pad, XOR'ing it in, and choosing the redaction set: no multi-party computation and no proof generation, which dominate the User-side cost of every cryptographic-attestor construction, MPC [ZMM⁺20, KK24, XYWY23] and proxy [LEFS23, ELW⁺24] alike. The User's cost scales with the size of the response, not with the size of the data proven, so attesting to an entire response costs no more than attesting to one byte of it. On a lower-mid-range phone, Reclaim attests to a 1200 kB response in 7.28 s, whereas the prior zkTLS maximum is 8 kB in 6.7 s [LEFS23]. Full results are in Section 6.

Extension: hiding predicate inputs from the Verifier. The Verifier in Reclaim reconstructs the non-redacted bytes and evaluates Φ on them, so Reclaim provides *selective disclosure*: the Verifier learns the disclosed bytes and the two lengths, and nothing else of the session. Predicates whose inputs must stay hidden, such as thresholds, set membership, or the proof of solvency of Section 1, need a zero-knowledge proof layered between User and Verifier on top of Reclaim. The choice of that layer is orthogonal to our contributions, and we do not pursue it.

We specify Reclaim in Section 5 in a hybrid UC model, where the TLS handshake, certificate verification, and the links among U , A_K , and A_T are ideal functionalities ($\mathcal{F}_{\text{KE}}$, $\mathcal{F}_{\text{CERT}}$, $\mathcal{F}_{\text{SC}}$ [CK02, Can04]), and we analyze Distributed-AEAD on its own in Section 4. This separates Reclaim's cryptographic content from TLS message handling, treated elsewhere [GMP⁺08a, KMO⁺15].

2.2.1 Security of the Reclaim Protocol

We prove in Section 5.4 that Reclaim UC-realizes $\mathcal{F}_{\text{zkPoPE}}$ in the hybrid model with $\mathcal{F}_{\text{SC}}$, $\mathcal{F}_{\text{KE}}$, $\mathcal{F}_{\text{CERT}}$, and $\mathcal{F}_{\text{US}}$. The first three functionalities abstract the TLS handshake, certificate verification, and the secure links among U , A_K , and A_T , as in the previous subsection. The fourth, $\mathcal{F}_{\text{US}}$, models the U - S_{id} wire, and its leakage depends on A_K 's corruption status. With A_K honest, it is an adversarial network, leaking every wire byte and permitting messages to be modified, reordered, and dropped. With A_K corrupted, it is a secure channel, leaking only message lengths. A corrupted A_K already holds the session keys, so an adversary that could read the wire as well would simply decrypt the records.

Corruption pattern. Corruption is static. Writing FC for the fully corrupted parties and HCC for the semi-honestly corrupted ones, a choice (FC, HCC) is *admissible* if

$$FC \subseteq \{U\}, \quad HCC \subseteq \{A_K, A_T\}, \quad |HCC| \leq 1, \quad FC \neq \emptyset \implies HCC = \emptyset.$$

Services are honest, and $\mathcal{V}$ is omitted from FC because proofs are visible to the adversary by default, so a corrupted $\mathcal{V}$ sees nothing new. The bound $|HCC| \leq 1$ is needed because each attestor holds one half of what is required to recover the User’s session data, and two corruptions would put both halves in one view. Semi-honesty is needed because a fully malicious A_K holds the traffic key and could manipulate the AEAD keystream, breaking soundness. This is the single-party threat model of Distributed-AEAD (Section 4.2).

Theorem 2.1 (Security of Reclaim; informal statement of Theorem 5.1) Let D-AEAD satisfy the confidentiality and ciphertext-integrity guarantees of Section 4.2, let every Prot_{id} in the deployment be request-reply and request-extractable, and let every predicate be supported for its protocol. Then $\mathcal{RP}$ UC-realizes $\mathcal{F}_{\text{zkPoPE}}$ in the $(\mathcal{F}_{\text{SC}}, \mathcal{F}_{\text{KE}}, \mathcal{F}_{\text{CERT}}, \mathcal{F}_{\text{US}})$ -hybrid model under the admissible corruption pattern above.

The statement carries three restrictions, on the protocols, on the predicates, and on the corruptions. Each is forced by a point at which the proof would otherwise fail:

- Request extractability is needed because against a malicious U only a request is observable, and $\mathcal{F}_{\text{zkPoPE}}$ must nonetheless be supplied with an input.
- The conditions on the predicate are needed because $\mathcal{F}_{\text{zkPoPE}}$ evaluates Φ on the User’s input while $\mathcal{V}$ evaluates it on (a part of) the transcript, so the two verdicts have to agree, and because $\mathcal{V}$ sees every byte the predicate does not withhold, so those bytes must be fixed by the predicate.
- The corruption pattern is needed for the reason just given.

Beyond the assumptions on Distributed-AEAD and the modeling conventions of Section 5.4, the list is complete, since a UC realization must reproduce $\mathcal{F}_{\text{zkPoPE}}$ in every execution the environment can arrange: whatever the construction needs appears either as a hypothesis or as a step of the proof. A construction-specific security model may leave such conditions implicit.

Only request extractability limits which application-layer protocols a deployment may target. Where the request format is a public serialization, its parser inverts it and the condition holds outright. It must be assumed only for formats that are not publicly invertible, such as those embedding a signature under a U -held key, a commitment, or a zero-knowledge proof. For those, the assumption has content only in an idealized model: in the plain model, the environment may hand the adversary a ready-made request, so the extractor must be an inverter, and the condition collapses to invertibility (Appendix B.2).

Matching the corruption pattern to the deployment. The deployment places each attestor in a TEE under an independent operator. A TEE guarantees *integrity*, that the platform runs the attested code, and *confidentiality*, that enclave state stays hidden. We rely only on the first. Integrity is enforced by remote attestation, whereas attacks extracting enclave secrets [VBMW⁺18, vSKGY20, CSB⁺26] are practical, even though each targets specific hardware generations and is either closed by vendor updates or requires physical access to the machine. Integrity ensures that each attestor runs its program, so the only gain against it is reading its state, which is a semi-honest corruption. A breach of one enclave then yields at most one half of Distributed-AEAD’s split, the traffic key without the private bytes or the private bytes without the key. The attestors run on distinct architectures under distinct operators, so compromising both takes two attacks, and we treat them as non-colluding. The U is the system’s actual user, with an incentive to forge records, so it is fully malicious. Colluding with an attestor would require breaking that attestor’s integrity, which is why U is never corrupted alongside one. Town Crier’s analysis [ZCC⁺16] demands the same integrity of an enclave running a full TLS client, and, in addition, confidentiality of its whole view.

From $\mathcal{F}_{\text{CERT}}$ to remote attestation. $\mathcal{F}_{\text{CERT}}[A_K]$ and $\mathcal{F}_{\text{CERT}}[A_T]$ are instantiated by signing keys that exist only inside the enclaves. At launch, each attestor samples a key pair and places the verification key in a remote-attestation report, which the platform signs over a measurement of the code it runs. A verifier checks, once per attestor, that the report verifies under the vendor’s attestation infrastructure and that its measurement equals the published program, and, per proof, that σ_K and σ_T verify under the keys the reports carry. The signing key is the only enclave state whose secrecy the design requires.

2.3 Open Problems

We list directions naturally suggested by the present work.

Characterizing the zkPoPE landscape. Each construction in this space – Reclaim and the zkTLS protocols surveyed in Section 1 and Section 3.3 – rests on a particular combination of assumptions on the auxiliary parties, on the transport, and on the protocol–predicate pair $(\text{Prot}_{\text{id}}, \Phi)$. A fundamental open direction is to characterize the landscape of zkPoPE solvability as a whole: which choices of $(\text{Prot}_{\text{id}}, \Phi)$ and corruption model admit a zkPoPE realization, and at what minimum communication and computational costs in auxiliary parties. Some attester-free regimes are partially understood – e.g., protocols whose replies the Service signs, as in RFC 9421 [BRS24] – but the boundary between these regimes and the attester-bearing regime in which Reclaim and the rest of the existing zkTLS literature operate remains uncharted. Both constructive results that extend the attester-free regime and impossibility results that delimit it from above would clarify the structure of the problem class.

A universal zkPoPE protocol. The characterization above has a constructive counterpart, which we regard as the central open problem posed by this work: a *universal* zkPoPE protocol, i.e., a single construction serving arbitrary protocol–predicate pairs over arbitrary transports, playing for interactions the role that zkVMs – proof systems for the correct execution of arbitrary programs [BSCG⁺13] – play for local computation. Reclaim shows that the request–reply TLS instance admits an efficient realization under a deployment-aligned trust model. Whether one construction can serve the whole class, and under which minimal assumptions, remains open.

Beyond two-party protocols. Our formalization of zkPoPE restricts Prot_{id} to a two-party interaction between U and $\text{Service}_{\text{id}}$. Extending the model and the surrounding constructions to multi-party protocols, in which any party can prove a property of its participation in a protocol with several other parties, would broaden the class of attestable executions considerably and is a natural complement to the present work.

Increasing adversarial capabilities. The corruption pattern under which Reclaim is proved secure (at most one of A_K , A_T corrupted semi-honestly, or alternatively a malicious U with both attestors honest) was derived by ruling out attacks against stronger adversaries. Reducing the gap between the corruption model under which security can be proved and the natural all-powerful adversary – in particular, accommodating fully malicious attestors – remains open and is likely to require additional cryptographic machinery on top of Distributed-AEAD.

3 Zero-Knowledge Proof of Protocol Execution

3.1 System Model and Parties

We fix the notation for the setting of Section 2.1.

Universe of services: A *universe of services* is a finite set

$$\mathcal{US} = \{ (\text{id}, \text{Prot}_{\text{id}}, \text{sec}_{\text{id}}) : \text{id} \in \mathcal{ID} \},$$

each entry comprising a public identifier id , a publicly known two-party application-layer protocol Prot_{id} that the service runs with a user, and the service’s private state sec_{id} .⁹ $\mathcal{US}$ is fixed by an out-of-band setup prior to any zkPoPE session and is treated as a session-level parameter; we write $\text{Service}_{\text{id}}$ for the triple in $\mathcal{US}$ indexed by id .

Parties: A zkPoPE deployment over $\mathcal{US}$ involves a user U , a verifier V , and the services in $\mathcal{US}$. Concrete realizations may involve auxiliary parties, such as the two attestors of Section 5; these do not appear in the abstract interface.

⁹We keep sec_{id} fixed for the lifetime of the deployment for simplicity. A real service’s private state might evolve between sessions; membership is then read relative to the state at the session in question, exactly as the service’s coins are in Remark 3.1.

Predicate and language: A *predicate* Φ is a boolean function over input–reply pairs $(\text{data}, \text{reply})$ of Prot_{id} , agreed upon out-of-band between U and V . The user’s task is to convince the verifier of membership of (id, Φ) in the language

$$\mathcal{L}_{\mathcal{US}} = \left\{ (\text{id}, \Phi) \mid \exists \text{data} : \Phi(\text{data}, \text{out}_{\mathsf{U}}(\text{Prot}_{\text{id}}(\text{Service}_{\text{id}}(\text{sec}_{\text{id}}), \mathsf{U}(\text{data})))) = 1 \right\}.$$

Membership in $\mathcal{L}_{\mathcal{US}}$ is defined relative to the specific deployment $\mathcal{US}$, not to the public descriptions $\{(\text{id}, \text{Prot}_{\text{id}})\}_{\text{id}}$ alone, since it depends on the services’ private states sec_{id} . Φ is a public parameter; data is the user’s private witness, and comprises exactly what U needs in order to run Prot_{id} — credentials, cookies, query parameters — and nothing the protocol does not consume. Membership asks nothing further of Φ ; a realization may.¹⁰

Remark 3.1 We have treated Prot_{id} as a deterministic function of $(\text{sec}_{\text{id}}, \text{data})$. The protocols of interest are randomized, and the two sides’ coins enter the language differently. The user’s tape R_{U} joins data in the witness: the user samples it and can be said to know it. A service draws fresh coins each time it is invoked. Writing R_{Service} for the coins of a single session, membership asserts that the predicate is satisfied for some coins the service could supply:

$$\mathcal{L}_{\mathcal{US}} = \left\{ (\text{id}, \Phi) \mid \exists \text{data}, R_{\mathsf{U}} : \begin{array}{c} \Phi(\text{data}, \text{out}_{\mathsf{U}}(\text{Prot}_{\text{id}}(\text{Service}_{\text{id}}(\text{sec}_{\text{id}}; R_{\text{Service}}), \\ \mathsf{U}(\text{data}; R_{\mathsf{U}})))) = 1 \end{array} \right\}.$$

Membership so read is a statement about knowledge alone: satisfying coins exist, but nothing says a run produced them. That gap is the execution requirement of Section 2.1, which the next subsection imposes on top of membership. We drop the tapes from the notation in the rest of the paper.

Interface: A zkPoPE scheme exposes its parties to the environment $\mathcal{Z}$ through four input messages:

- $(\text{sid}, \text{Setup}, (\text{id}, \text{Prot}, \text{sec}))$: service registration, submitted by a service.
- $(\text{sid}, \text{Setup}, \{(\text{pid}_i, \text{id}_i, \text{Prot}_i)\}_{i=1}^l)$: client-side setup, declaring the set of services a non-service party expects to interact with.
- $(\text{sid}, \text{Prove}, \text{id}, \Phi, \text{data}, R_{\mathsf{U}})$: proving, submitted by the user; R_{U} is the user-side random tape discussed in Remark 3.1.
- $(\text{sid}, \text{Verify}, \text{id}, P_i, \Phi, \pi)$: non-interactive verification, submitted by the verifier.

The four tags are the four input handlers of $\mathcal{F}_{\text{zkPoPE}}$ below.

3.2 Ideal Functionality

We capture the security requirements of zkPoPE through the ideal functionality $\mathcal{F}_{\text{zkPoPE}}$ defined in Figure 2. $\mathcal{F}_{\text{zkPoPE}}$ runs in the environment $\mathcal{Z}$, in the presence of a simulator Sim , and a set of machines $\mathsf{P} = \{P_1, \dots, P_N\}$ partitioned into users (which we group together with the verifier and any auxiliary parties a concrete realization may require) and services. It is parameterized by the security parameter 1^λ and by statically chosen sets FC of fully corrupted and HCC of honest-but-curious corrupted parties.

Notation. Inside the handlers we use the following semantic instructions:

- **Break:** abort the current activation without output.
- **Parse:** extract well-formed messages matching a given pattern.

¹⁰Reclaim’s verifier reconstructs part of the transcript and evaluates Φ on it, which restricts the predicates it supports; Section 5.2 fixes that restriction. It is a feature of how Reclaim verifies rather than of zkPoPE: a realization whose verifier checks a proof instead of re-evaluating the predicate would carry a different restriction, or none.

<u>Functionality $\mathcal{F}_{\text{zkPoPE}}$</u>
On input $(\text{sid}, \text{Setup}, (\text{id}, \text{Prot}, \text{sec}))$ from a service P: 1. Check there was no user-side Setup for sid, Break otherwise. 2. Check P is not in the universe of services $\mathcal{US}_{\text{sid}}$, and there is no service with id id in the universe, Break otherwise. 3. Add P as $\text{Service}_{\text{id}} = (\text{id}, \text{sec}, \text{Prot})$ to the universe of services for sid: $\mathcal{US}_{\text{sid}} \leftarrow \mathcal{US}_{\text{sid}} \cup (\text{pid}, \text{Service}_{\text{id}}).$ 4. Send $(\text{sid}, \text{ServiceAdded}, \text{pid}, \text{id}, \text{Prot})$ to Sim.
On input $(\text{sid}, \text{Setup}, \{\text{pid}_i, \text{id}_i, \text{Prot}_i\}_{i=1}^l)$ from honest users: 1. Check there was no user-side Setup for sid, Break otherwise. 2. Check that all input services are consistent with $\mathcal{US}_{\text{sid}}$. 3. Send $(\text{sid}, \text{Setup})$ to Sim, receive in reply crs and a tuple $\text{Algs} = (\text{Setup}, \text{Prove}, \text{Verify}).$ 4. Store $(\text{sid}, \text{crs}, \text{Algs})$. 5. Send $(\text{sid}, \text{Setup}, \text{crs})$ to the calling parties.
On input $(\text{sid}, \text{Prove}, \text{id}, \Phi, \text{data}, R_U)$ from a party P: 1. Check there was Setup for sid. 2. Retrieve $(\text{sid}, \text{crs}, \text{Algs})$ and $(\text{pid}, \text{Service}_{\text{id}} = (\text{id}, \text{sec}, \text{Prot})) \in \mathcal{US}_{\text{sid}}$. 3. If P is honest and $\Phi(\text{data}, \text{reply}) = 0$ for every reply, send $(\text{sid}, \text{Prove}, \perp)$ to P and Break. <i>Simulate the protocol between user and service.</i> 4. Compute $(\text{reply}, \text{out}_{\text{Service}}) \leftarrow \text{Prot}_{\text{id}}(\text{Service}_{\text{id}}(\text{sec}_{\text{id}}; R_{\text{Service}}), \text{U}(\text{data}; R_U))$, with R_{Service} the coins $\text{Service}_{\text{id}}$'s tape supplies for this activation. 5. Send $(\text{sid}, \text{Executed}, \text{reply})$ to P and $(\text{sid}, \text{Executed}, \text{out}_{\text{Service}})$ to P_{pid}. 6. If P is corrupted, send $(\text{sid}, \text{Predicate}, P)$ to Sim, receive $(\text{sid}, \text{Predicate}, \Phi')$ in reply, and set $\Phi \leftarrow \Phi'$. 7. Compute the binary value $\text{res} = \Phi(\text{data}, \text{reply})$. <i>Issue a proof if the predicate holds.</i> 8. Send $(\text{sid}, \text{Prove}, P, \Phi, \text{id}, \text{res}, \text{reply})$ to Sim. 9. If res is equal to 0, send $(\text{sid}, \text{Prove}, \perp)$ to P, and Break. 10. Wait for a reply (Proof, π) from Sim. 11. Check $\text{Verify}(\text{crs}, \pi) = 1$, otherwise send $(\text{sid}, \text{Prove}, \perp)$ to P, and Break. 12. Send $(\text{sid}, \text{Prove}, P, \Phi, \text{id}, \pi)$ to P. 13. Store $(P, \Phi, \text{id}, \pi)$ to list Π_{sid}, and Break.
On input $(\text{sid}, \text{Verify}, \text{id}, P_i, \Phi, \pi)$ from P: 1. Check there was Setup for sid. 2. Check if there is an entry $(P_i, \Phi, \text{id}, \pi)$ in Π_{sid}: • If so, send $(\text{sid}, \text{Verify}, (\text{id}, P_i, \Phi, \pi), 1)$ to P, • Otherwise, send $(\text{sid}, \text{Verify}, (\text{id}, P_i, \Phi, \pi), 0)$ to P, and Break.

Figure 2: Ideal Functionality $\mathcal{F}_{\text{zkPoPE}}$ for zkPoPE.

Every message $\mathcal{F}_{\text{zkPoPE}}$ sends to a party is a *private delayed output* in the sense of [Can01]: Sim is told that a message to that party is pending and decides when it is delivered; its contents reach Sim only when the recipient is corrupted. We write $\text{Algs} = (\text{Setup}, \text{Prove}, \text{Verify})$ for the bundle of scheme algorithms supplied by Sim at setup time, together with the common reference string crs ; inside $\mathcal{F}_{\text{zkPoPE}}$ these have their standard NIZK meanings, the protocol-level Verify checking a proof against crs . Sim produces crs by its own simulated setup and keeps the trapdoor, so $\mathcal{F}_{\text{zkPoPE}}$ holds none. On a successful Prove activation $\mathcal{F}_{\text{zkPoPE}}$ obtains the proof from Sim, sending it $(\text{sid}, \text{Prove}, P, \Phi, \text{id}, \text{res}, |\text{reply}|)$ and awaiting a reply, as in the ideal functionality for non-interactive zero-knowledge of [GOS06]. The message is sent on every activation and not only on a satisfied predicate, since a party running a Prove session is observable whatever the predicate evaluates to, and it carries the length of the reply, as a channel carrying the reply would leak it; $\mathcal{F}_{\text{zkPoPE}}$ awaits a proof only when $\text{res} = 1$. A corrupted prover fixes its predicate after the run, through Sim, and may submit $\Phi = \perp$ on input: such a prover holds the witness $(\text{data}, \text{reply})$ once the run is over and could evaluate any predicate on it, so nothing false becomes provable and the honest prover’s guarantees are unaffected. An honest prover whose input no reply can complete to a satisfying pair is turned away before the run and without notice to Sim, as in the functionality of [GOS06], which ignores a Prove on $(x, w) \notin R$; the test is part of Φ ’s specification.

3.3 Relation to Existing Models

DECO [ZMM⁺20] introduced the oracle ideal functionality $\mathcal{F}_{\text{Oracle}}$, a UC abstraction of a TLS query–response interaction with three roles, prover, verifier and server, and a query template instantiated by a private prover input. The verifier participates online, so non-interactive proofs are not covered; the definition binds the server side to standard TLS (Definition 3.1 in [ZMM⁺20]); and no proof object exists, since the designated verifier receives the evaluated statement $\text{Stmt}(R)$ from the functionality and nothing else, so attestations are not transferable and verification is not a separate operation. Its descendants keep this shape. Garble-then-Prove [XYWY23] defines $\mathcal{F}_{\text{AuthData}}$ on top of it, with the three roles, the query template and the commit-then-prove flow. ACTS [MVVZ26] keeps DECO’s design and replaces its final zero-knowledge stage with a predicate blind signature, so that attestations are standard RSA-PSS signatures; its formal analysis is confined to that component. ORIGO [ELW⁺24] defines $\mathcal{F}_{DP}$ for a proxy verifier in a malicious-network setting, with the protocol fixed to TLS 1.3 and AES-GCM, which ties its theorem to that AEAD. DiStefano [CDH⁺25] rejects monolithic DECO-style proofs, since any change of primitive or ciphersuite forces a new one, and analyzes its handshake, query and commitment phases separately around the *designated-commitment TLS* primitive; this restores modularity but gives up any single object against which schemes could be compared, and stays inside the transport, with phases specific to TLS 1.3 and a record layer targeting AES-GCM. Ticket to Hide [LRV26] builds a bespoke functionality for proving a connection to one of several servers without revealing which; TLSNotary [KK24] is documented informally; Janus [LEFS23] formalizes asymmetric privacy through an HVZK proof system; Town Crier [ZCC⁺16] casts security as datagram authenticity for a feed delivered by an SGX-resident TLS client.

In every one of these works TLS is essential to the definition or the analysis, and no functionality in the line has been reused unchanged by a successor: each modified, re-tailored or abandoned its predecessor’s, which is the incomparability of Section 1 playing out in the literature. $\mathcal{F}_{\text{zkPoPE}}$ takes the other exit from protocol-inseparability: one functionality for the whole class, with the transport removed from it.

4 Distributed AEAD

The Reclaim Protocol relies on a three-party variant of authenticated encryption. The encryption and authentication of a message are jointly produced by two non-colluding parties, both distinct from the data party. The primitive must guarantee that no single party learns the inputs of the others, and that the data party alone cannot forge a valid ciphertext-tag pair. We define it directly over the class of stream-cipher AEAD designs and do not pursue further generalization in this work. Concrete instantiations based on AES-GCM and ChaCha20-Poly1305, the AEAD schemes underlying TLS, appear in Appendix C.4.

4.1 Syntax

Fix associated spaces $\mathcal{K} = \mathcal{K}(\lambda)$ (keys), $\mathcal{N} = \mathcal{N}(\lambda)$ (nonces), $\mathcal{T} = \mathcal{T}(\lambda)$ (tags), $\mathcal{TM} = \mathcal{TM}(\lambda)$ (tag material), $\mathcal{AD} \subseteq \{0, 1\}^*$ (associated data), $\mathcal{M} \subseteq \{0, 1\}^*$ (messages), and $\mathcal{C} \subseteq \{0, 1\}^*$ (ciphertexts).

Definition 4.1 (Distributed stream-cipher AEAD) A *distributed stream-cipher AEAD scheme* is a tuple

$$\text{D-AEAD} = (\text{KGen}, \text{GenEncMat}, \text{GenTagMat}, \text{ComputeTag}, \text{VerifyTag})$$

of PPT algorithms which are defined as follows:

- $\text{KGen}(1^\lambda) \rightarrow k \in \mathcal{K}$;
- $\text{GenEncMat} : \mathcal{K} \times \mathcal{N} \times \mathbb{N} \rightarrow \{0, 1\}^*$ outputs an ℓ -bit keystream on (k, n, ℓ) ;
- $\text{GenTagMat} : \mathcal{K} \times \mathcal{N} \rightarrow \mathcal{TM}$ outputs tag material $\tau \in \mathcal{TM}$ on (k, n) ;
- $\text{ComputeTag} : \mathcal{TM} \times \mathcal{AD} \times \mathcal{C} \rightarrow \mathcal{T}$ outputs a tag $T \in \mathcal{T}$ on (τ, AD, C) ;
- $\text{VerifyTag} : \mathcal{TM} \times \mathcal{AD} \times \mathcal{C} \times \mathcal{T} \rightarrow \{0, 1\}$ outputs a bit on (τ, AD, C, T) , with $\text{VerifyTag}(\tau, AD, C, \text{ComputeTag}(\tau, AD, C)) = 1$ for every $(\tau, AD, C) \in \mathcal{TM} \times \mathcal{AD} \times \mathcal{C}$.

The scheme defines two three-party interactive protocols Π_{AEnc} and Π_{ADec} (Figures 3–4), executed by the *key party* $\mathcal{P}_K$ holding k ; the *data party* $\mathcal{P}_D$ holding message P and a *secret-index set* $\text{SR} \subseteq [|P|]$ (designating positions of P whose contents must remain hidden from $\mathcal{P}_K$ and $\mathcal{P}_T$) in Π_{AEnc} and the ciphertext-tag pair (C, T) in Π_{ADec} ; and the *tag party* $\mathcal{P}_T$ holding no private input.

We write $\text{embed}_{\text{SR}, \ell}(M) \in \{0, 1\}^\ell$ for the embedding of $M \in \{0, 1\}^{|\text{SR}|}$ at the positions indexed by SR , with zeros elsewhere.

Common input (sid, n, AD, P). Private inputs: $\mathcal{P}_K: k$; $\mathcal{P}_D: P, \text{SR}$; $\mathcal{P}_T: \perp$. 1. $\mathcal{P}_D$: sample $M \xleftarrow{\\$} \{0, 1\}^{ \text{SR} }$; set $\tilde{P} \leftarrow P \oplus \text{embed}_{\text{SR}, P }(M)$; send $\tilde{P}$ to $\mathcal{P}_K$ and $\text{embed}_{\text{SR}, P }(M)$ to $\mathcal{P}_T$. 2. $\mathcal{P}_K$: compute $K_s \leftarrow \text{GenEncMat}(k, n, P)$ and $\tau \leftarrow \text{GenTagMat}(k, n)$; set $\tilde{C} \leftarrow \tilde{P} \oplus K_s$; send $(\tilde{C}, \tau)$ to $\mathcal{P}_T$; output ok. 3. $\mathcal{P}_T$: set $C \leftarrow \tilde{C} \oplus \text{embed}_{\text{SR}, P }(M)$ and $T \leftarrow \text{ComputeTag}(\tau, AD, C)$; send (C, T) to $\mathcal{P}_D$; output ok. 4. $\mathcal{P}_D$: upon receipt of (C, T), output (C, T).

Figure 3: Π_{AEnc} protocol.

Common input (sid, n, AD). Private inputs: $\mathcal{P}_K: k$; $\mathcal{P}_D: (C, T)$; $\mathcal{P}_T: \perp$. 1. $\mathcal{P}_D$: send (C, T) to $\mathcal{P}_T$. 2. $\mathcal{P}_T$: send C to $\mathcal{P}_K$; 3. $\mathcal{P}_K$: compute $\tau \leftarrow \text{GenTagMat}(k, n)$; send τ to $\mathcal{P}_T$. 4. $\mathcal{P}_T$: compute $T' \leftarrow \text{ComputeTag}(\tau, AD, C)$; if $T' \neq T$ send $\perp$ to $\mathcal{P}_K$ and $\mathcal{P}_D$ and output $\perp$; otherwise, send ok to $\mathcal{P}_K$ and $\mathcal{P}_D$ and output ok. 5. $\mathcal{P}_K$: upon ok from $\mathcal{P}_T$, compute $K_s \leftarrow \text{GenEncMat}(k, n, C)$, send K_s to $\mathcal{P}_D$, and output ok; upon $\perp$ from $\mathcal{P}_T$, output $\perp$. 6. $\mathcal{P}_D$: upon K_s from $\mathcal{P}_K$, output $\tilde{P} \leftarrow C \oplus K_s$; upon $\perp$ from $\mathcal{P}_T$, output $\perp$.
--

Figure 4: Π_{ADec} protocol.

Correctness. An honest execution of Π_{AEnc} delivers to $\mathcal{P}_D$ the pair (C, T) with $C = P \oplus \text{GenEncMat}(k, n, |P|)$ and $T = \text{ComputeTag}(\text{GenTagMat}(k, n), AD, C)$, the AEAD record of P under (k, n) ; an honest execution of Π_{ADec} on (C, T) delivers P to $\mathcal{P}_D$ when T is the valid tag for (AD, C) under (k, n) , and $\perp$ otherwise.

4.2 Security

We adopt a *static, single-party* corruption model. The adversary $\mathcal{A}$ is *semi-honest* when corrupting $\mathcal{P}_K$ or $\mathcal{P}_T$, observing the party’s view but following the prescribed protocol, and *malicious* when corrupting $\mathcal{P}_D$, free to deviate arbitrarily from Π_{AEnc} and Π_{ADec} . This asymmetry is essential for Reclaim’s setting, as discussed in Section 2.2.1.

We require two security properties of D-AEAD; Appendix C gives the games and the reductions to standard assumptions on the primitives. Integrity is stronger than standard AEAD INT-CTXT, since $\mathcal{P}_D$ ’s view contains the plaintext and adaptive protocol-internal state. Confidentiality is restricted to the SR positions: in Reclaim, the positions outside SR are the bytes the predicate discloses to the verifier, so their indistinguishability is not required.

Confidentiality. $\mathcal{P}_K$ or $\mathcal{P}_T$, semi-honestly running adaptively many concurrent honest sessions of Π_{AEnc} and Π_{ADec} , learns nothing about $\mathcal{P}_D$ ’s plaintext at the SR positions in Π_{AEnc} and nothing about the plaintext at any position in Π_{ADec} .

Ciphertext integrity. $\mathcal{P}_D$, arbitrarily deviating from Π_{AEnc} and Π_{ADec} across adaptively many concurrent sessions with honest $\mathcal{P}_K$ and $\mathcal{P}_T$, cannot produce a tuple (n, AD, C, T) that an honest run of Π_{ADec} accepts and that was not delivered to $\mathcal{P}_D$ in any prior Π_{AEnc} session.

Joint guarantee. Together the two properties say that every record an honest Π_{ADec} accepts was produced with knowledge of k , hence not by $\mathcal{P}_D$ alone, while the views of $\mathcal{P}_K$ and $\mathcal{P}_T$ reveal nothing about the plaintext at the SR positions in Π_{AEnc} and nothing about it at all in Π_{ADec} ; and neither protocol completes without both $\mathcal{P}_K$ and $\mathcal{P}_T$. Reclaim uses this to obtain an authentic, redacted record of $\mathcal{P}_D$ ’s payloads held jointly by two independent parties.

5 The Reclaim Protocol

5.1 Hybrid model and interface

We specify the Reclaim Protocol $\mathcal{RP}$ in the $(\mathcal{F}_{\text{SC}}, \mathcal{F}_{\text{CERT}}, \mathcal{F}_{\text{KE}}, \mathcal{F}_{\text{US}})$ -hybrid UC model [Can01], with A_K and A_T semi-honest. $\mathcal{F}_{\text{SC}}$ [CK02] carries every message among $\{U, A_K, A_T\}$, $\mathcal{F}_{\text{KE}}$ [CK02] produces the Phase A session key, $\mathcal{F}_{\text{CERT}}$ [Can04] provides three signing facilities, and $\mathcal{F}_{\text{US}}$ models the U - S_{id} wire: every Phase A-C message sent to or received from S_{id} “over the network” below travels over $\mathcal{F}_{\text{US}}$, whose leakage regime is motivated in Section 2.2.1. Appendix B.3 gives the message syntax of each functionality.

The three $\mathcal{F}_{\text{CERT}}$ instances are pre-registered with fixed roles. $\mathcal{F}_{\text{CERT}}[S_{\text{id}}]$ models S_{id} ’s long-term TLS certificate and authenticates its key-exchange contribution in Phase A, following [GMP⁺08b]; $\mathcal{F}_{\text{CERT}}[A_K]$ and $\mathcal{F}_{\text{CERT}}[A_T]$ sign the attestors’ per-session transcript views in Phase D, and theirs are the only signatures $\mathcal{V}$ checks. All Prove sessions share the three instances and the three pairwise $\mathcal{F}_{\text{SC}}$ channels sid_{UA_K} , sid_{UA_T} , $\text{sid}_{A_KA_T}$; concurrent sessions are disambiguated by a per-session tag ssid , fixed in Phase A and carried by every later message of the session, which both Phase-D signatures cover.

Two routing decisions keep S_{id} unmodified while authentication of the TLS session resides with the attestors. Although A_K initiates the key exchange, the network endpoint S_{id} observes belongs to U : every non- $\mathcal{F}_{\text{KE}}$ Phase A message between A_K and S_{id} is relayed by U over the network. And the call to $\mathcal{F}_{\text{KE}}$ is made by A_K and S_{id} directly, with U not a party to it. S_{id} ’s view is thus a standard TLS handshake with U , while the output of $\mathcal{F}_{\text{KE}}$ and the verdict of $\mathcal{F}_{\text{CERT}}[S_{\text{id}}]$ are held by A_K and never appear on the network.

5.2 Admissible Protocols and Predicates

$\mathcal{RP}$ is analyzed for a restricted class of application-layer protocols and of predicates over them. This subsection fixes that class. We state the conditions informally here; their formal counterparts are Definitions B.1, B.2 and B.5 of Appendix B.2.

The protocol. We require Prot_{id} to be *request-reply*: the user sends one message Req , computed from its input data and its coins by a map we write $\text{Prot}_{\text{id}}.\text{U}_1$; the service answers with one message Resp computed from sec_{id} and Req ; and the user’s reply is that response. We require further that Prot_{id} be *request-extractable*: from a request, one can recover an input and coins that regenerate exactly that request.

The latter requirement is the price of a design choice: U is never asked to show how it built Req from data . Demanding that would admit every protocol, since the demonstration would itself yield the extraction, but it would put a proof about the request format on the user’s critical path, with cost growing in the format and the predicate. $\mathcal{RP}$ keeps the user’s cost low, so a malicious U may put an arbitrary string on the wire, and the request is the only handle the analysis has on the session; request extractability is what turns it back into an input. Little is lost, because $\text{Prot}_{\text{id}}.\text{U}_1$ is public and, in the common cases, invertible.

- Where it is invertible the condition holds with no assumption beyond the shape of the protocol: requests that are the input verbatim, or a public serialization of fields of data , possibly with a nonce or timestamp drawn from the coins, whose parser is the inverse.
- Where it is not — a request carrying, for example, a signature under a U -held key, a commitment, or a zero-knowledge proof — an idealized model widens the class: a local random-oracle or algebraic-group analysis hands the extractor the adversary’s oracle queries or group-element representations, which for many such formats is what it needs. The widening is partial.

Appendix B.2 treats the notion in full; widening the class further we leave to future work.

The predicate. A predicate Φ is a per-session parameter agreed between U and $\mathcal{V}$ in advance. $\mathcal{F}_{\text{zkPoPE}}$ evaluates it on the user’s input and the reply; $\mathcal{V}$ holds neither data nor sec_{id} and evaluates only on the request and response bytes it reconstructs. We therefore require Φ to be *admissible* for Prot_{id} (Definition B.5): there exists a predicate $\hat{\Phi}$ over request–response pairs that agrees with Φ on every execution, and no rejection reply satisfies $\hat{\Phi}$. The first condition makes $\mathcal{V}$ ’s verdict the one the language speaks about; the second makes $\mathcal{V}$ reject the proof of a session whose request no data generates.

We require moreover that $\hat{\Phi}$ have the following form. Its *format* names the map $\text{Prot}_{\text{id}}.\text{U}_1$, specialized to the request layout in question, together with the lengths $|\text{Req}|, |\text{Resp}|$ and two position sets $\text{SR}_{\text{req}} \subseteq \{1, \dots, |\text{Req}|\}$ and $\text{SR}_{\text{resp}} \subseteq \{1, \dots, |\text{Resp}|\}$ marking the bytes hidden from $\mathcal{V}$ and from both attestors; both sets depend on $\hat{\Phi}$ and Prot_{id} alone. Since Prot_{id} has fluid but deterministic structure, the format is a small program rather than a constant table. Its *condition* is a propositional combination over $\{\wedge, \vee, \neg\}$ of atoms “the byte at position p equals the constant c ”, with p outside SR_{req} on the request side and outside SR_{resp} on the response side, given as an offset or as a parse-tree location under Prot_{id} ; bytes are not interpreted beyond constant equality and not compared with one another. We write $\hat{\Phi}(\text{Req}, \text{Resp}) \in \{0, 1\}$ for the result; since no atom reads a withheld position, it is the same on the zero-filled pair $\mathcal{V}$ reconstructs. Rejection soundness then amounts to the response-side atoms being inconsistent with the response head returned on a malformed request. We call predicates of this shape *equivalence predicates*.

Finally, $\hat{\Phi}$ must *determine what it discloses*: any two pairs satisfying it agree at every position outside SR_{req} and SR_{resp} . $\mathcal{V}$ receives every such byte, so otherwise a proof could reveal which of two satisfying transcripts took place, which the statement (id, Φ) does not. Consequently the disclosed bytes are a function of Φ and Prot_{id} , and a predicate of this form is fixed by its position sets and the bytes it discloses; we take that pair as its representation wherever predicates are compared, as in $\mathcal{F}_{\text{zkPoPE}}$ ’s record. Disjunction and negation remain available where they admit no second value at a disclosed position. An admissible predicate of this form we call *supported*.

The description of $\hat{\Phi}$ is that of Φ : the format and the condition are what U and $\mathcal{V}$ agree on, and Φ is the same condition read on the user’s input and reply. We therefore write Φ for both and let the arguments disambiguate; the appendix keeps them apart.

Notation and conventions. For a position set $R \subseteq \{1, \dots, \ell\}$ and $M \in \{0, 1\}^{|R|}$, $\text{embed}_{R, \ell}(M) \in \{0, 1\}^\ell$ places the i -th symbol of M at the i -th smallest position of R and zero elsewhere, so $X \mapsto X \oplus \text{embed}_{R, \ell}(M)$ is the identity outside R and, for uniform M , randomizes X on R ; $\text{zero}_R(X)$ is X with the positions in R set to zero. The algorithms GenEncMat , GenTagMat , ComputeTag , VerifyTag are those of D-AEAD (Definition 4.1),

instantiated for the session's cipher suite as in Appendix C.4; we omit their associated-data argument, which $\mathcal{RP}$ does not rely on. KDF [Kra10] expands the $\mathcal{F}_{KE}$ key κ into two directional keys and two initialization vectors $(k_{enc}, k_{dec}, IV_{enc}, IV_{dec})$, the vectors serving as the D-AEAD nonces of the request and response records. We describe $\mathcal{RP}$ for one AEAD record per direction; several records are handled by deriving each record's nonce from the direction's IV and the record sequence number, as in TLS 1.3 [Res18].

5.3 Protocol $\mathcal{RP}$

Figure 5 in Appendix A summarizes the message flow described below.

Session state. All Prove sessions share the three $\mathcal{F}_{SC}$ channels, so U , A_K and A_T each keep, per sid , a table of active sessions indexed by $ssid$, recording the session's position in Phases A–D and every value generated or received in it. A_K creates the entry when it generates $ssid$ in Phase A; A_T on the first Phase-B message, from U or A_K , bearing a new $ssid$; U on the $(sid, \text{Prove}, \dots)$ input, binding it to $ssid$ on receipt of $(\text{handshake-done}, ssid)$ after checking $ssid = (sid, U, id, r)$ against the sid_{KE} it relayed. A party processes a message under $ssid$ only if it is the one its state for $ssid$ expects next from that sender on that channel, and drops it silently otherwise.

Setup. On $(sid, \text{Setup}, (id, \text{Prot}, \text{sec}))$ from $\mathcal{Z}$, S_{id} records $(id, \text{Prot}, \text{sec})$ as its service descriptor for sid and spawns $\mathcal{F}_{CERT}[S_{id}]$ by issuing $(\text{Sign}, sid_{S_{id}}, \perp)$ and discarding the response. On $(sid, \text{Setup}, \{(pid_i, id_i, \text{Prot}_i)\}_{i=1}^l)$ from $\mathcal{Z}$, each of U , A_K , A_T , $\mathcal{V}$ records the list as its admissible services for sid ; A_K and A_T spawn $\mathcal{F}_{CERT}[A_K]$ and $\mathcal{F}_{CERT}[A_T]$ in the same way, and the three channels sid_{UA_K} , sid_{UA_T} , $sid_{A_KA_T}$ are established via $\mathcal{F}_{SC}$.

Protocol $\mathcal{RP}$: Prove

Only U is driven by $\mathcal{Z}$; A_K , A_T and S_{id} react to the messages of Phases A–D, and every party silently drops a message that fails its checks or the session-state convention. On input $(sid, \text{Prove}, id, \Phi, \text{data}, R_U)$ from $\mathcal{Z}$, U :

1. If id is not in its admissible-service list, outputs $(sid, \text{Prove}, \perp)$ and halts; otherwise retrieves $(pid_{S_{id}}, \text{Prot})$.
2. Computes $Req \leftarrow \text{Prot}_{id}.U_1(\text{data}; R_U)$ under the format named by Φ and extracts SR_{req}, SR_{resp} . If Req differs at a position outside SR_{req} from the value Φ fixes there, so that no response can satisfy Φ , outputs $(sid, \text{Prove}, \perp)$ and halts.
3. Runs Phases A–D. On recovering $Resp$ at the end of Phase C, outputs $(sid, \text{Executed}, Resp)$ and evaluates $\Phi(\text{data}, Resp)$; if 0, outputs $(sid, \text{Prove}, \perp)$ and halts without entering Phase D.
4. If a phase aborts, outputs $(sid, \text{Prove}, \perp)$; on completing Phase D with transcript π , outputs $(sid, \text{Prove}, U, \Phi, id, \pi)$.

Phase A — Handshake (*participants* U , A_K , S_{id})

- On the Prove input, U sends id to A_K via $\mathcal{F}_{SC}[sid_{UA_K}]$.
- A_K checks id against its list, retrieving $(pid_{S_{id}}, \text{Prot})$, and aborts otherwise; samples $r \xleftarrow{\$} \{0, 1\}^\lambda$; sets $sid_{KE} \leftarrow (A_K, pid_{S_{id}}, id, r)$ and $ssid \leftarrow (sid, U, id, r)$; and sends $(\text{HandshakeRequest}, sid_{KE})$ to U , who forwards it to S_{id} over the network.
- On $(\text{HandshakeRequest}, sid_{KE})$ with $sid_{KE} = (A_K, pid_{S_{id}}, id, r)$ matching its PID and a registered id , S_{id} obtains $\sigma_{S_{id}}$ on (sid_{KE}, A_K, S_{id}) from $\mathcal{F}_{CERT}[S_{id}]$, sends it to U over the network, and sends $(\text{Establish-session}, sid_{KE}, pid_{S_{id}}, A_K, \text{responder})$ to $\mathcal{F}_{KE}$. U forwards $\sigma_{S_{id}}$ to A_K .
- A_K verifies $\sigma_{S_{id}}$ on (sid_{KE}, A_K, S_{id}) at $\mathcal{F}_{CERT}[S_{id}]$, aborting on 0, and sends $(\text{Establish-session}, sid_{KE}, A_K, pid_{S_{id}}, \text{initiator})$ to $\mathcal{F}_{KE}$.
- On (key, sid_{KE}, κ) from $\mathcal{F}_{KE}$, A_K and S_{id} each derive $(k_{enc}, k_{dec}, IV_{enc}, IV_{dec}) \leftarrow \text{KDF}(\kappa, sid_{KE})$. A_K sends $(\text{handshake-done}, ssid)$ to U , who records $ssid$.

Phase B — Request Handling (*participants* U, A_K, A_T , and S_{id} as a TLS endpoint; instantiates $D\text{-AEAD}.\Pi_{A_{Enc}}$)

- On (handshake-done, ssid), U samples $M_{req} \xleftarrow{\$} \{0, 1\}^{|\text{SR}_{req}|}$, sets $\widetilde{Req} \leftarrow Req \oplus \text{embed}_{\text{SR}_{req}, |Req|}(M_{req})$, and sends (ssid, $\widetilde{Req}$, SR_{req}) to A_K via $\mathcal{F}_{SC}[\text{sid}_{UA_K}]$ and (ssid, $\text{embed}_{\text{SR}_{req}, |Req|}(M_{req})$) to A_T via $\mathcal{F}_{SC}[\text{sid}_{UA_T}]$.
- On (ssid, $\widetilde{Req}$, SR_{req}), A_K records SR_{req} , computes $K_{req} \leftarrow \text{GenEncMat}(k_{enc}, IV_{enc}, |Req|)$, $\tau_{req} \leftarrow \text{GenTagMat}(k_{enc}, IV_{enc})$ and $\widetilde{Req}_{enc} \leftarrow \widetilde{Req} \oplus K_{req}$, and sends (ssid, $\widetilde{Req}_{enc}$, τ_{req} , SR_{req}) to A_T via $\mathcal{F}_{SC}[\text{sid}_{A_K A_T}]$.
- On (ssid, Y) from U and (ssid, $\widetilde{Req}_{enc}$, τ_{req} , SR_{req}) from A_K , A_T checks that $|Y| = |\widetilde{Req}_{enc}|$ and that Y is zero outside SR_{req} , aborting otherwise; records (ssid, SR_{req}); computes $Req_{enc} \leftarrow \widetilde{Req}_{enc} \oplus Y$ and $T_{req} \leftarrow \text{ComputeTag}(\tau_{req}, Req_{enc})$; and sends (ssid, Req_{enc} , T_{req}) to U via $\mathcal{F}_{SC}[\text{sid}_{UA_T}]$.
- U forwards (Req_{enc} , T_{req}) to S_{id} over the network; ssid never travels on the wire.
- S_{id} checks $\text{VerifyTag}(\text{GenTagMat}(k_{enc}, IV_{enc}), Req_{enc}, T_{req}) = 1$, aborting otherwise; recovers $Req \leftarrow Req_{enc} \oplus \text{GenEncMat}(k_{enc}, IV_{enc}, |Req_{enc}|)$; computes $(Resp, \text{out}_{Service}) \leftarrow \text{Prot}_{id}.\text{Service}(\text{sec}, Req; R_{Service})$ on fresh coins; outputs (sid, Executed, $\text{out}_{Service}$) to $\mathcal{Z}$; and sends ($Resp_{enc}$, T_{resp}) to U over the network, with $Resp_{enc} \leftarrow Resp \oplus \text{GenEncMat}(k_{dec}, IV_{dec}, |Resp|)$ and $T_{resp} \leftarrow \text{ComputeTag}(\text{GenTagMat}(k_{dec}, IV_{dec}), Resp_{enc})$. This is standard TLS record processing; S_{id} runs no code of $\mathcal{RP}$.

Phase C — Response Handling (*participants* U, A_T, A_K ; instantiates $D\text{-AEAD}.\Pi_{A_{Dec}}$)

- On ($Resp_{enc}$, T_{resp}) from the network, U sends (ssid, $Resp_{enc}$, T_{resp}) to A_T via $\mathcal{F}_{SC}[\text{sid}_{UA_T}]$; A_T sends (ssid, $|Resp_{enc}|$) to A_K via $\mathcal{F}_{SC}[\text{sid}_{A_K A_T}]$; A_K replies with (ssid, τ_{resp}), $\tau_{resp} \leftarrow \text{GenTagMat}(k_{dec}, IV_{dec})$.
- A_T computes $b \leftarrow \text{VerifyTag}(\tau_{resp}, Resp_{enc}, T_{resp})$ and sends (ssid, ok) if $b = 1$, and (ssid, $\perp$) otherwise, to both A_K and U , outputting the same.
- On (ssid, ok), A_K computes $K_{resp} \leftarrow \text{GenEncMat}(k_{dec}, IV_{dec}, |Resp_{enc}|)$ and sends (ssid, K_{resp}) to U via $\mathcal{F}_{SC}[\text{sid}_{UA_K}]$; on (ssid, $\perp$) it outputs $\perp$.
- On (ssid, ok) from A_T and (ssid, K_{resp}) from A_K , U outputs $Resp \leftarrow Resp_{enc} \oplus K_{resp}$; on (ssid, $\perp$) it outputs $\perp$.

Phase D — Transcript Generation (*participants* U, A_K, A_T)

- On completing Phase C, U sends (begin-transcript, ssid, SR_{resp}) to A_K via $\mathcal{F}_{SC}[\text{sid}_{UA_K}]$.
- A_K retrieves the SR_{req} recorded for ssid, computes $Req_{null} \leftarrow \text{zero}_{\text{SR}_{req}}(\widetilde{Req})$ and $\tilde{K}_{resp} \leftarrow \text{zero}_{\text{SR}_{resp}}(K_{resp})$, obtains σ_K on $(U, id, ssid, \text{SR}_{req}, \text{SR}_{resp}, Req_{null}, \tilde{K}_{resp})$ from $\mathcal{F}_{CERT}[A_K]$, and sends (begin-transcript, ssid, id, SR_{resp}) to A_T via $\mathcal{F}_{SC}[\text{sid}_{A_K A_T}]$ and (ssid, σ_K , Req_{null} , $\tilde{K}_{resp}$) to U via $\mathcal{F}_{SC}[\text{sid}_{UA_K}]$.
- A_T retrieves the SR_{req} recorded for ssid, computes $\widetilde{Resp}_{enc} \leftarrow \text{zero}_{\text{SR}_{resp}}(Resp_{enc})$, obtains σ_T on $(U, id, ssid, \text{SR}_{req}, \text{SR}_{resp}, \widetilde{Resp}_{enc})$ from $\mathcal{F}_{CERT}[A_T]$, and sends (ssid, σ_T , $\widetilde{Resp}_{enc}$) to U via $\mathcal{F}_{SC}[\text{sid}_{UA_T}]$.
- U assembles $\pi \leftarrow (\sigma_K, \sigma_T, id, ssid, \text{SR}_{req}, \text{SR}_{resp}, Req_{null}, \tilde{K}_{resp}, \widetilde{Resp}_{enc})$ and outputs (sid, Prove, U, Φ, id, π) to $\mathcal{Z}$.

Protocol $\mathcal{RP}$: Verify (*non-interactive; $\mathcal{V}$ holds A_K, A_T and $\text{sid}_K, \text{sid}_T$ as protocol parameters*)

On input (sid, Verify, id, P_i, Φ, π) from $\mathcal{Z}$, $\mathcal{V}$ outputs (sid, Verify, (id, P_i, Φ, π), b), where $b = 0$ if check 1 or 2 fails and $b = b_1 \wedge b_2 \wedge b_3$ otherwise.

1. id is in its admissible-service list for sid.
2. π parses as $(\sigma_K, \sigma_T, id_\pi, ssid, \text{SR}_{req}, \text{SR}_{resp}, Req_{null}, \tilde{K}_{resp}, \widetilde{Resp}_{enc})$ with $ssid = (sid', P', id', r)$, $id_\pi = id' = id$, $sid' = sid$, $P' = P_i$, and $(\text{SR}_{req}, \text{SR}_{resp})$ the position sets fixed by the format of Φ .
3. b_1 and b_2 are the verdicts of $\mathcal{F}_{CERT}[A_K]$ on σ_K over $(P_i, id, ssid, \text{SR}_{req}, \text{SR}_{resp}, Req_{null}, \tilde{K}_{resp})$ and of $\mathcal{F}_{CERT}[A_T]$ on σ_T over $(P_i, id, ssid, \text{SR}_{req}, \text{SR}_{resp}, \widetilde{Resp}_{enc})$.
4. $b_3 \leftarrow \Phi(Req_{null}, \widetilde{Resp})$ with $\widetilde{Resp} \leftarrow \widetilde{Resp}_{enc} \oplus \tilde{K}_{resp}$.

5.4 Security

Theorem 5.1 Let D-AEAD be a distributed stream-cipher AEAD scheme satisfying the confidentiality and ciphertext-integrity guarantees of Section 4.2, let every Prot_{id} in the deployment be request-reply and request-extractable, and let every predicate be supported for its protocol, in the sense of Section 5.2. Then the Reclaim protocol $\mathcal{RP}$ of Section 5 UC-realizes the ideal functionality $\mathcal{F}_{\text{zkPoPE}}$ of Section 3 in the $(\mathcal{F}_{\text{SC}}, \mathcal{F}_{\text{KE}}, \mathcal{F}_{\text{CERT}}, \mathcal{F}_{\text{US}})$ -hybrid model, under the admissible corruption pattern recalled in Section 5.4.1.

5.4.1 Model and Corruption Pattern

$\mathcal{F}_{\text{US}}$ and the admissible corruption pattern are those of Section 2.2.1. The pattern has four cases, which the proof treats separately: Case A (passive), $FC = HCC = \emptyset$; Case B (malicious U), $FC = \{\text{U}\}$, $HCC = \emptyset$; Case C, $HCC = \{\text{A}_K\}$; and Case D, $HCC = \{\text{A}_T\}$, with $FC = \emptyset$ in the last two. Phases B and C of $\mathcal{RP}$ instantiate D-AEAD's Π_{AEnc} and Π_{ADec} under $\text{U} \leftrightarrow \mathcal{P}_D$, $\text{A}_K \leftrightarrow \mathcal{P}_K$, $\text{A}_T \leftrightarrow \mathcal{P}_T$, with SR_{req} and SR_{resp} as the secret-index sets. At most one of the three is corrupted in any execution, so Case B is D-AEAD's malicious- $\mathcal{P}_D$ regime and Cases C, D its semi-honest $\mathcal{P}_K$ and $\mathcal{P}_T$ regimes.

5.4.2 Simulator

Sim runs the hybrid functionalities and an in-head copy of every uncorrupted party, forwarding a semi-honest attestor's view to $\mathcal{A}$. On $\mathcal{F}_{\text{zkPoPE}}$.Setup it commits the algorithms (Setup, Prove, Verify) of Section 5, with Verify checking σ_K, σ_T under the verification algorithms it commits to $\mathcal{F}_{\text{CERT}}[\text{A}_K], \mathcal{F}_{\text{CERT}}[\text{A}_T]$. It releases each delayed output of $\mathcal{F}_{\text{zkPoPE}}$ when the corresponding party of $\mathcal{RP}$ would produce it.

Honest User (Cases A, C, D). $\mathcal{F}_{\text{zkPoPE}}$'s announcement $(P, \Phi, \text{id}, \text{res}, |\text{reply}|)$ tells Sim that a session has begun. A supported Φ fixes, through its format and its disclosure-determinacy, the position sets $(\text{SR}_{\text{req}}, \text{SR}_{\text{resp}})$, the lengths, and the bytes at every disclosed position; let $\text{Req}^*, \text{Resp}^*$ carry those bytes and zeros on the two sets. Sim runs $\mathcal{RP}$ with its in-head parties on the request Req^* and, in place of the service's computation, the response Resp^* (a zero string of the announced length when $\text{res} = 0$), stopping after Phase C when $\text{res} = 0$. The proof it returns is the π its in-head U assembles. The simulation differs from a real execution only in the plaintexts, which agree with the real ones outside the secret positions whenever anything derived from them is signed.

Malicious User (Case B). $\mathcal{A}$ drives U's channels; the in-head attestors and service respond as prescribed.

When the in-head S_{id} accepts a request record, Sim recovers the request $\text{Req}' = \widetilde{\text{Req}}^{\mathcal{A}} \oplus Y^{\mathcal{A}}$ from $\mathcal{A}$'s two Phase-B messages — the same reconstruction the honest A_T performs, whose check that $Y^{\mathcal{A}}$ is zero outside $\text{SR}_{\text{req}}^{\mathcal{A}}$ makes Req' agree with the signed Req_{null} at every disclosed position — and applies the extractor to obtain an input and tape $(\text{data}', R'_{\text{U}})$ regenerating Req' . It submits these to $\mathcal{F}_{\text{zkPoPE}}$.Prove on U's behalf, receives the reply as the corrupted caller's output, encrypts it as the response record, and runs Phases C and D against $\mathcal{A}$. It then names the predicate: the supported one with $\mathcal{A}$'s position sets asserting the disclosed bytes of the transcript, or the identically-false one if the reply is a rejection. On $\text{res} = 1$ it answers the proof request with the π its attestors signed, so that $\mathcal{F}_{\text{zkPoPE}}$'s record matches what $\mathcal{A}$ holds.

5.4.3 Indistinguishability

Zero-knowledge (Cases A, C, D). The simulation replaces the request and response, on the positions in SR_{req} and SR_{resp} (and everywhere in a response that fails Φ , which is never signed), by fixed bytes. In Cases A and D, $\mathcal{A}$ controls the wire, and ciphertext integrity lets the simulated S_{id} and A_T accept only each other's records; $\mathcal{A}$'s view is then $\mathcal{P}_T$'s view of Π_{AEnc} and Π_{ADec} together with the wire, so a hybrid over sessions and directions reduces the change to D-AEAD's $\mathcal{P}_T$ -confidentiality ($\text{IND-CCA-}\mathcal{P}_T$). In Case C the wire is hidden by $\mathcal{F}_{\text{US}}$ and $\mathcal{A}$'s view is $\mathcal{P}_K$'s: everything in it is a function of the session key and of $\widetilde{\text{Req}}$, which is one-time-padded on SR_{req} and equal to the public bytes elsewhere, so the change is invisible ($\text{IND-CPA-}\mathcal{P}_K$, advantage 0).

Soundness (Case B). $\mathcal{F}_{\text{zkPoPE}}.\text{Verify}$ accepts only proofs it recorded, so indistinguishability from the real execution says that a malicious $\mathcal{U}$ obtains no accepted π beyond those, and that those are accepted. Three hops. First, every record that S_{id} or A_T accepts is one the other honest side produced in that session, by D-AEAD ciphertext integrity. Second, $\mathcal{F}_{\text{CERT}}$ ’s sign tables pin an accepted π^* to a single session: both signatures cover the session tag, the position sets and the redacted views, and a mismatch between the sessions behind σ_K^* and σ_T^* needs a nonce collision, of probability at most $q^2 2^{-\lambda-1}$ over q sessions. $\mathcal{V}$ then evaluates Φ on the disclosed bytes of that session’s transcript, and by disclosure-determinacy exactly one supported predicate with those position sets accepts it — the one Sim names. Third, the extractor’s input regenerates the session’s request, so $\mathcal{F}_{\text{zkPoPE}}$ ’s in-head run returns the response $\mathcal{A}$ received and records π under that same predicate; this costs the extraction error over q sessions.

6 Benchmarks

Setup. We benchmark the Reclaim Protocol with the prover running on a mobile phone equipped with a Qualcomm Snapdragon 888 5G processor and 8 GB of RAM, released in 2021, representative of the lower-mid range of current consumer mobile hardware. The two TEE attestors run on Google Cloud confidential-computing virtual machines in the **europe-north1-a** region, backed by AMD SEV-SNP. The prover communicates with the attestors over Wi-Fi with a typical round-trip time of 75 ms to the region. Every reported value is the mean of ten independent executions, and we record the wall-clock interval between the start of the TLS handshake on the prover and the moment the verifier accepts the resulting attestation.

Scope of the measurement. End-to-end latency in our protocol is dominated by communication between the prover and the attestors, and this cost scales with the length of the server response rather than with the size of the proven substring. A modest implementation-level optimization that ships only the proven bytes in a small number of network rounds saves approximately 0.5 s, which is immaterial at the timescales reported below. We therefore fix the benchmark to the most demanding configuration and prove the *entire* server response in every run.

Comparison with prior work. Existing zkTLS protocols are not evaluated on response sizes comparable to ours. The largest published evaluation is Janus [LEFS23], which proves 8 kB of sensitive TLS data in 6.2 s under TLS 1.2 and 6.7 s under TLS 1.3. All other known constructions evaluate at smaller sizes, with most below 4 kB and none above the 8 kB ceiling set by Janus. Since we prove the entire response, our protocol attests to approximately 500 kB of data in the same time budget, roughly 60 times larger. A direct comparison at the sizes of Table 1 is therefore not possible, and we instead place each row next to the wall-clock time of a plain HTTPS fetch of the same resource on the same device, executed with `curl`. This isolates the protocol overhead from the irreducible network and transport cost.

Table 1: End-to-end latency of the Reclaim Protocol against a plain HTTPS fetch of the same resource, executed with `curl` on a lower-mid range mobile device. Each value is the mean of ten runs. Sizes report the HTTPS response in binary kilobytes (1 kB = 1024 B). Δ is the difference of the two timings and isolates the cost of the protocol.

Site	Size (kB)	TLS (s)	Reclaim (s)	Δ (s)
example.com	1.23	0.39	2.90	2.51
iana.org	6.00	0.62	3.11	2.49
jquery.com	25.12	0.45	3.21	2.76
wikipedia.org	117.50	0.52	4.67	4.15
wikipedia.org/Main_Page	225.35	0.60	4.79	4.19
wikipedia.org/Software_eng.	408.73	0.64	5.67	5.03
github.com	557.39	0.68	6.66	5.98
wikipedia.org/World_War_II	1201.66	0.84	7.28	6.44

Results. Table 1 reports the timings. Across response sizes spanning three orders of magnitude, from 1.23 kB to 1201.66 kB, the protocol completes in under 7.3 s. The additive overhead over the HTTPS baseline grows from 2.5 s to 6.4 s, a 2.5-fold increase across a roughly 1000-fold increase in response size. These absolute timings, reached at proven data sizes up to two orders of magnitude beyond any previously reported, easily place the Reclaim Protocol within deployable range on lower-mid range mobile hardware.

Use of AI tools. We used Anthropic’s Claude (Table 5.1) to proofread the technical details and proofs, check the text for consistency, improve its style, readability and editorial quality, and composing TikZ schemes. The authors reviewed every change and take responsibility for the content.

References

- [BCPR14] Nir Bitansky, Ran Canetti, Omer Paneth, and Alon Rosen. On the existence of extractable one-way functions. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 505–514. ACM Press, May / June 2014.
- [Ber05] Daniel J. Bernstein. The poly1305-aes message-authentication code. In *Proceedings of the 12th International Conference on Fast Software Encryption, FSE’05*, page 32–49, Berlin, Heidelberg, 2005. Springer-Verlag.
- [BP15] Elette Boyle and Rafael Pass. Limits of extractability assumptions with distributional auxiliary input. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology – ASIACRYPT 2015, Part II*, volume 9453 of *Lecture Notes in Computer Science*, pages 236–261. Springer, Berlin, Heidelberg, November / December 2015.
- [BRS24] Annabelle Backman, Justin Richer, and Manu Sporny. HTTP Message Signatures. RFC 9421, February 2024.
- [BSCG⁺13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, Eran Tromer, and Madars Virza. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In *Advances in Cryptology – CRYPTO 2013*, volume 8043 of *LNCS*, pages 90–108. Springer, 2013.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science (FOCS 2001)*, pages 136–145. IEEE, 2001.
- [Can04] Ran Canetti. Universally composable signature, certification, and authentication. In *Proceedings of the 17th IEEE Computer Security Foundations Workshop (CSFW)*, pages 219–233. IEEE, 2004.
- [CDG⁺18] Jan Camenisch, Manu Drijvers, Tommaso Gagliardoni, Anja Lehmann, and Gregory Neven. The wonderful world of global random oracles. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 280–312. Springer, Cham, April / May 2018.
- [CDH⁺25] Sofia Celi, Alex Davidson, Hamed Haddadi, Gonçalo Pestana, and Joe Rowell. DiStefano: Decentralized infrastructure for sharing trusted encrypted facts and nothing more. In *Network and Distributed System Security Symposium (NDSS) 2025*, 2025.
- [CGCD⁺17] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. In *2017 IEEE European Symposium on Security and Privacy*, pages 451–466, 2017.

- [CJS14] Ran Canetti, Abhishek Jain, and Alessandra Scafuro. Practical UC security with a global random oracle. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 2014: 21st Conference on Computer and Communications Security*, pages 597–608. ACM Press, November 2014.
- [CK02] Ran Canetti and Hugo Krawczyk. Universally composable notions of key exchange and secure channels. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 337–351. Springer, 2002.
- [Con26] Deirdre Connolly. ML-KEM Post-Quantum Key Agreement for TLS 1.3. Internet-Draft draft-ietf-tls-mlkem-07, Internet Engineering Task Force, February 2026. Work in Progress.
- [CSB⁺26] Jalen Chuang, Alex Seto, Nicolas Berrios, Stephan van Schaik, Christina Garman, and Daniel Genkin. Tee.fail: Breaking trusted execution environments via ddr5 memory bus interposition. In *47th IEEE Symposium on Security and Privacy (IEEE S&P ’26)*. IEEE Computer Society, 2026.
- [Dam92] Ivan Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In Joan Feigenbaum, editor, *Advances in Cryptology – CRYPTO’91*, volume 576 of *Lecture Notes in Computer Science*, pages 445–456. Springer, Berlin, Heidelberg, August 1992.
- [Dwo07] Morris J. Dworkin. Recommendation for block cipher modes of operation: Galois/Counter mode (GCM) and GMAC. NIST Special Publication 800-38D, National Institute of Standards and Technology, Gaithersburg, MD, November 2007.
- [ELW⁺24] Jens Ernstberger, Jan Lauinger, Yinnan Wu, Arthur Gervais, and Sebastian Steinhorst. ORIGO: Proving provenance of sensitive data with constant communication. Cryptology ePrint Archive, Paper 2024/447, 2024.
- [FK11] Sheila Frankel and Suresh Krishnan. IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap. RFC 6071, February 2011.
- [FKL18] Georg Fuchsbauer, Eike Kiltz, and Julian Loss. The algebraic group model and its applications. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part II*, volume 10992 of *Lecture Notes in Computer Science*, pages 33–62. Springer, Cham, August 2018.
- [FNR22] Roy T. Fielding, Mark Nottingham, and Julian Reschke. HTTP/1.1. RFC 9112, June 2022.
- [GMP⁺08a] Sebastian Gajek, Mark Manulis, Olivier Pereira, Ahmad-Reza Sadeghi, and Jörg Schwenk. Universally composable security analysis of tls. In *International Conference on Provable Security*, pages 313–327. Springer, 2008.
- [GMP⁺08b] Sebastian Gajek, Mark Manulis, Olivier Pereira, Ahmad-Reza Sadeghi, and Jörg Schwenk. Universally composable security analysis of TLS. In *Provable Security (ProvSec 2008)*, volume 5324 of *LNCS*, pages 313–327. Springer, 2008.
- [GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for NP. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 339–358. Springer, Berlin, Heidelberg, May / June 2006.
- [GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd Annual ACM Symposium on Theory of Computing*, pages 99–108. ACM Press, June 2011.

- [Int15] Intel Corporation. *Intel® Software Guard Extensions (Intel® SGX)*, 2015. Official Intel overview document.
- [IT21] Jana Iyengar and Martin Thomson. QUIC: A UDP-Based Multiplexed and Secure Transport. RFC 9000, May 2021.
- [KK24] Maciej Kalka and Marek Kirejczyk. A comprehensive review of tlsnotary protocol, 2024.
- [KMO⁺15] Markulf Kohlweiss, Ueli Maurer, Cristina Onete, Björn Tackmann, and Daniele Venturi. (de-) constructing tls 1.3. In *International Conference on Cryptology in India*, pages 85–102. Springer, 2015.
- [Kra10] Hugo Krawczyk. Cryptographic extraction and key derivation: The HKDF scheme. In *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *LNCS*, pages 631–648. Springer, 2010.
- [LEFS23] Jan Lauinger, Jens Ernstberger, Andreas Finkenzeller, and Sebastian Steinhorst. Janus: Fast privacy-preserving data provenance for TLS. Cryptology ePrint Archive, Paper 2023/1377, 2023.
- [LR25] Anna Lysyanskaya and Leah Namisa Rosenbloom. Adaptive uc nizk for practical applications. In *International Conference on Cryptology and Information Security in Latin America*, pages 76–108. Springer, 2025.
- [LRV26] Ryan Little, Daniel S. Roche, and Mayank Varia. Ticket to hide: Private, practical proofs of provenance for TLS. Cryptology ePrint Archive, Paper 2026/540, 2026.
- [MV04] David A. McGrew and John Viega. The security and performance of the galois/counter mode (gcm) of operation. In *Proceedings of the 5th International Conference on Cryptology in India, INDOCRYPT’04*, page 343–355, Berlin, Heidelberg, 2004. Springer-Verlag.
- [MVVZ26] Pierpaolo Della Monica, Ivan Visconti, Andrea Vitaletti, and Marco Zecchini. ACTS: Attestations of contents in TLS sessions, 2026. Cryptology ePrint Archive, Paper 2026/755.
- [Nao03] Moni Naor. On cryptographic assumptions and challenges (invited talk). In Dan Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 96–109. Springer, Berlin, Heidelberg, August 2003.
- [NL15] Yoav Nir and Adam Langley. ChaCha20 and Poly1305 for IETF Protocols. RFC 7539, May 2015.
- [Not22] Mark Nottingham. Building Protocols with HTTP. RFC 9205, June 2022. BCP 56.
- [Pro25] Reclaim Protocol. Reclaim Protocol Whitepaper (v2), 2025.
- [RD08] Eric Rescorla and Tim Dierks. The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246, August 2008.
- [Res18] Eric Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446, August 2018.
- [Str] Stripe, Inc. Resolve webhook signature verification errors. <https://docs.stripe.com/webhooks/signature>.
- [VBMW⁺18] Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F. Wenisch, Yuval Yarom, and Raoul Strackx. Foreshadow: extracting the keys to the intel sgx kingdom with transient out-of-order execution. In *Proceedings of the 27th USENIX Conference on Security Symposium, SEC’18*, page 991–1008, USA, 2018. USENIX Association.

- [vSKGY20] Stephan van Schaik, Andrew Kwong, Daniel Genkin, and Yuval Yarom. SGAXe: How SGX fails in practice. <https://sgaxeattack.com/>, 2020.
- [XYWY23] Xiang Xie, Kang Yang, Xiao Wang, and Yu Yu. Lightweight authentication of web data via garble-then-prove. Cryptology ePrint Archive, Paper 2023/964, 2023.
- [ZCC⁺16] Fan Zhang, Ethan Cecchetti, Kyle Croman, Ari Juels, and Elaine Shi. Town crier: An authenticated data feed for smart contracts. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, page 270–282, New York, NY, USA, 2016. Association for Computing Machinery.
- [ZMM⁺20] Fan Zhang, Deepak Maram, Harjasleen Malvai, Steven Goldfeder, and Ari Juels. Deco: Liberating web data using decentralized oracles for tls. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, CCS '20, page 1919–1938. ACM, October 2020.

A Message Flow of the Reclaim Protocol




Figure 5: Message flow of the Reclaim protocol $\mathcal{RP}$: what each party holds and the messages of Phases A–D (Section 5.3). Messages among U, A_K and A_T travel over secure channels $\mathcal{F}_{SC}$; those between U and S_{id} over the public network as ordinary TLS records. The traffic keys live only at A_K , the ciphertexts and the User’s one-time pad only at A_T , so neither view alone recovers the session data. The proof $\pi = (\sigma_K, \sigma_T, id, ssid, SR_{req}, SR_{resp}, Req_{null}, \widetilde{K}_{resp}, \widetilde{Resp}_{enc})$ carries the disclosed positions and the request and response with private bytes zeroed, both signed by the attestors.

B Preliminaries

B.1 Notation

We write n for the security parameter and 1^n for its unary encoding; PPT abbreviates probabilistic polynomial-time, and $\text{negl}(n)$ denotes a negligible function. For $n \in \mathbb{N}$ we write $[n] := \{1, \dots, n\}$. We use $|\mathcal{X}|$ for the cardinality of a set and $|x|$ for the bit-length of a string x ; $\perp$ denotes a distinguished failure (“no value”) symbol.

We write $x \xleftarrow{\$} \mathcal{X}$ for sampling x uniformly from a finite set $\mathcal{X}$, and $y \leftarrow A(x)$ for assigning to y the output of a (possibly randomized) algorithm A on input x . To make randomness explicit we write $y \leftarrow A(x; \rho)$ for running A on input x with random tape ρ , and $(y, \rho) \leftarrow A(x)$ for saving the random tape used into ρ ; with ρ fixed, A is deterministic, which lets a relation or a proof refer to the exact coins an algorithm used. For a probabilistic experiment we write $\Pr[E : \text{exp}]$ for the probability of event E after running exp . Two distribution ensembles X, Y are *identically distributed*, written $X \equiv Y$, if they coincide, and *computationally indistinguishable*, written $X \approx_c Y$, if no PPT distinguisher separates them with non-negligible advantage.

For a two-party protocol Π we write

$$\Pi\langle P_1(\text{in}_1), P_2(\text{in}_2) \rangle$$

for the joint execution in which party P_i is initialized with private input in_i . We write $\text{out}_{P_i}(\Pi\langle P_1(\text{in}_1), P_2(\text{in}_2) \rangle)$ for the local output produced by P_i at the end of the execution, and collect both outputs as

$$(\text{out}_{P_1}, \text{out}_{P_2}) \leftarrow \Pi\langle P_1(\text{in}_1), P_2(\text{in}_2) \rangle.$$

When the execution is randomized, each party may additionally receive a random tape; we write $\Pi\langle P_1(\text{in}_1; r_1), P_2(\text{in}_2; r_2) \rangle$ to make those tapes explicit.

B.2 Request–Reply Protocols and Request Extractability

We define the class of two-party protocols we call *request–reply*, a notion of *request extractability* for protocols in that class, and a matching notion of *admissibility* for predicates over them. Both concern a protocol and a predicate alone. Section 5.2 explains informally why the analysis of $\mathcal{RP}$ appeals to them.

The request–reply class. We treat protocols carrying one message in each direction.

Definition B.1 (Request–reply protocol) A two-party protocol Prot between $\mathbf{U}$ and Service is *request–reply* if it is specified by two efficiently computable maps

$$\begin{aligned} \text{Prot.U}_1 &: (\text{data}, R_{\mathbf{U}}) \mapsto \text{Req}, \\ \text{Prot.Service} &: (\text{sec}, \text{Req}, R_{\text{Service}}) \mapsto (\text{Resp}, \text{out}_{\text{Service}}), \end{aligned}$$

and its execution consists of $\mathbf{U}$ computing and sending $\text{Req} \leftarrow \text{Prot.U}_1(\text{data}; R_{\mathbf{U}})$; Service computing $(\text{Resp}, \text{out}_{\text{Service}}) \leftarrow \text{Prot.Service}(\text{sec}, \text{Req}; R_{\text{Service}})$, sending Resp to $\mathbf{U}$, and locally outputting $\text{out}_{\text{Service}}$; and $\mathbf{U}$ locally outputting the response it received, $\text{reply} = \text{Resp}$. We write

$$\mathcal{R}_{\text{Prot}} = \{ \text{Prot.U}_1(\text{data}; R) : \text{data} \in \{0, 1\}^*, R \in \{0, 1\}^* \}$$

for the set of *well-formed* requests of Prot , and posit as part of Prot 's specification a publicly recognizable set $\mathcal{E}_{\text{Prot}} \subseteq \{0, 1\}^*$ of *rejection replies*, such that for every $\text{Req} \notin \mathcal{R}_{\text{Prot}}$ and every $(\text{sec}, R_{\text{Service}})$ the service replies with some $\text{Resp} \in \mathcal{E}_{\text{Prot}}$ and produces the null local output $\text{out}_{\text{Service}} = \perp$.

Taking the user's output to be the response verbatim costs nothing over the predicates admitted below (Remark B.6). The rejection set formalizes the assumption of Section 2.1 that the Service rejects inputs of incorrect format; that it is publicly recognizable makes a rejection reply as informative as no reply to any party inspecting the transcript. Over HTTP, the 4xx status line in the response head identifies a malformed request to a party holding neither data nor sec .

Request extractability. The notion asks that a request identify an input producing it, in the following sense.

Definition B.2 (Request extractability) A request–reply protocol Prot is *request-extractable* if for every PPT algorithm $\mathcal{A}$ there exists a PPT algorithm $\text{E}_{\mathcal{A}}$ such that, for every auxiliary input $\text{aux} \in \{0, 1\}^{\text{poly}(n)}$,

$$\varepsilon_{\text{Prot}, \mathcal{A}}^{\text{ext}}(n) := \Pr[\text{Req}^* \in \mathcal{R}_{\text{Prot}} \wedge \text{Prot.U}_1(\text{data}'; R') \neq \text{Req}^*] \leq \text{negl}(n),$$

where $\varepsilon_{\text{Prot}, \mathcal{A}}^{\text{ext}}(n)$ denotes the supremum over aux of the probability, taken over the experiment

$$\begin{aligned} R_{\mathcal{A}} &\xleftarrow{\$} \{0, 1\}^{\text{poly}(n)}; \\ \text{Req}^* &\leftarrow \mathcal{A}(1^n, \text{aux}; R_{\mathcal{A}}); \\ (\text{data}', R') &\leftarrow \text{E}_{\mathcal{A}}(1^n, \text{aux}, R_{\mathcal{A}}, \text{Req}^*). \end{aligned}$$

The extractor receives $\mathcal{A}$'s auxiliary input and random tape and is invoked once, on the request $\mathcal{A}$ has produced: it is *straight-line*, never rewinding $\mathcal{A}$ or restarting it on fresh coins. This is the most demanding of the usual extractor shapes, and the one our application requires. The extractor need only produce *some* preimage of Req^* , not the one $\mathcal{A}$ had in mind; the two may differ whenever **data** carries components the request does not transmit — fields the format discards, credentials that authenticate without appearing verbatim, slack in an encoding. Nothing is lost: by Definition B.1 everything after the first message is a function of Req^* and the service's state, so any two preimages elicit the same response and, by Definition B.5 below, the same predicate value.

The following unconditional variant is the form in which the protocols surveyed at the end of this subsection meet the condition.

Definition B.3 (Request invertibility) A request-reply protocol Prot is *request-invertible* if there is a single deterministic polynomial-time algorithm E such that $\text{Prot.U}_1(\text{data}'; R') = Req$ for every $Req \in \mathcal{R}_{\text{Prot}}$, where $(\text{data}', R') \leftarrow E(Req)$.

Fact B.4 Request invertibility implies request extractability with $\varepsilon_{\text{Prot}, \mathcal{A}}^{\text{ext}}(n) = 0$ for every $\mathcal{A}$: take $E_{\mathcal{A}} = E$, which ignores $(\text{aux}, R_{\mathcal{A}})$.

Admissible predicates. Request extractability recovers an input regenerating the request; it does not relate a predicate's value on that input to its value on the transcript. The second notion does.

Definition B.5 (Admissible predicate) Let Prot be a request-reply protocol. A predicate Φ over input-reply pairs of Prot is *admissible for* Prot if there exists a predicate $\hat{\Phi}$ over request-response pairs such that:

- (i) (*Transcript determinacy.*) For all data , R_U and $Resp$,

$$\Phi(\text{data}, Resp) = \hat{\Phi}(\text{Prot.U}_1(\text{data}; R_U), Resp).$$

- (ii) (*Rejection soundness.*) $\hat{\Phi}(Req, Resp) = 0$ for every Req and every $Resp \in \mathcal{E}_{\text{Prot}}$.

Clause (i) says that Φ depends on **data** only through the request the execution puts on the wire; clause (ii) rules out the predicates a malformed request could satisfy. The equivalence predicates of Section 5.2 are specified from the transcript side: one writes down $\hat{\Phi}$, and Φ is that condition read on the user's input and reply. Clause (i) is then immediate, and clause (ii) holds whenever the response-side atoms are inconsistent with a **4xx** response head.

Remark B.6 (Dropping the user's post-processing) Definition B.1 fixes $\text{reply} = Resp$. One might allow the user a post-processing map Prot.U_2 , taking $(\text{data}, R_U, Resp)$ to reply — decrypting under a key held in **data**, say. Over admissible predicates this is no restriction. Clause (i) would then read

$$\Phi(\text{data}, \text{Prot.U}_2(\text{data}, R_U, Resp)) = \hat{\Phi}(\text{Prot.U}_1(\text{data}; R_U), Resp),$$

so the predicate's value depends on nothing Prot.U_2 contributes beyond the transcript, and replacing Prot.U_2 by the identity, with $\hat{\Phi}$ as the predicate, leaves requests, responses and predicate values unchanged on every input. This is a property of the class treated here, not of the notion of Section 3.1, where the user's output is an arbitrary local output of an arbitrary two-party protocol.

The two notions combine as follows.

Lemma 1 (Extraction consistency) Let Prot be request-extractable, let Φ be admissible for Prot , let $\mathcal{A}$ be PPT with extractor $E_{\mathcal{A}}$, and fix $(\text{sec}, R_{\text{Service}})$. Let $Req^* \leftarrow \mathcal{A}(1^n, \text{aux}; R_{\mathcal{A}})$, let $(\text{data}', R') \leftarrow E_{\mathcal{A}}(1^n, \text{aux}, R_{\mathcal{A}}, Req^*)$, and let $(Resp^*, \text{out}_{\text{Service}}^*) \leftarrow \text{Prot.Service}(\text{sec}, Req^*; R_{\text{Service}})$. Then, except with probability $\varepsilon_{\text{Prot}, \mathcal{A}}^{\text{ext}}(n)$, exactly one of the following holds.

- (i) $Req^* \in \mathcal{R}_{\text{Prot}}$, and the execution $\text{Prot}(\text{Service}(\text{sec}; R_{\text{Service}}), U(\text{data}'; R'))$ places exactly Req^* on the wire. It therefore runs to $(Resp^*, \text{out}_{\text{Service}}^*)$, and its user-side output satisfies $\Phi(\text{data}', Resp^*) = \hat{\Phi}(Req^*, Resp^*)$.

(ii) $Req^* \notin \mathcal{R}_{\text{Prot}}$, and $\hat{\Phi}(Req^*, Resp^*) = 0$.

Proof. The two cases are exhaustive and mutually exclusive by definition of $\mathcal{R}_{\text{Prot}}$. In case (i), Definition B.2 gives $\text{Prot.U}_1(\text{data}'; R') = Req^*$ except with probability $\varepsilon_{\text{Prot}, \mathcal{A}}^{\text{ext}}(n)$; the first message of the execution is this value by Definition B.1, and the service’s computation is a function of $(\text{sec}, Req^*, R_{\text{Service}})$ alone, so it returns $(Resp^*, \text{out}_{\text{Service}}^*)$. The claim about Φ is clause (i) of Definition B.5 applied to $(\text{data}', R', Resp^*)$. In case (ii), Definition B.1 gives $Resp^* \in \mathcal{E}_{\text{Prot}}$, and clause (ii) of Definition B.5 gives $\hat{\Phi}(Req^*, Resp^*) = 0$.

Achievability. Request extractability concerns Prot alone. It holds trivially when $\text{data} = Req$ and Prot.U_1 is the identity, the shape of the oracle functionality of DECO [ZMM⁺20] and its descendants, and it holds by Fact B.4 whenever Prot.U_1 is a public injective serialization, the extractor being its parser. An HTTP request is of this kind: the message syntax of [FNR22] fixes the encoding of request line, header fields and body and specifies the parse recovering each, so a Prot.U_1 assembling a request from fields of data is inverted by that parse. Invertibility survives a nonce, timestamp or padding drawn from R_U : the parser reads them off Req and returns them as R' , which is where letting the extractor choose R' does its work.

Where Req carries a value the party producing it cannot open — a signature under a U-held key, a commitment, a zero-knowledge proof, for example — Prot.U_1 is not efficiently invertible, and Definition B.2 asks instead for an extractor that reads the signing key or the opening off the adversary’s code and coins. Invertibility is a property of the function Prot.U_1 , settled by examining Prot ; extractability asserts an extractor for *every* adversary, in the $\forall \mathcal{A} \exists E_{\mathcal{A}}$ shape of knowledge-of-exponent [Dam92] and of the extraction assumptions behind succinct arguments, and like them it is not falsifiable in the sense of [Nao03, GW11]. Discharging it means assuming it, or working in a model that hands the extractor a lever on the adversary: its random-oracle queries, or the representations an algebraic adversary supplies with its group elements [FKL18]. The lever helps only where the adversary itself produced the values to be recovered, so the widening is partial.¹¹ Quantifying over every auxiliary input has a consequence worth stating. Let $\mathcal{D}$ be the adversary that outputs its auxiliary input. Its extractor $E_{\mathcal{D}}$ must, for every $\text{aux} \in \mathcal{R}_{\text{Prot}}$, return a preimage of aux except with negligible probability over its own coins; it is a randomized inverter of Prot.U_1 on $\mathcal{R}_{\text{Prot}}$. In the plain model, request extractability is therefore request invertibility up to a negligible error, and the non-invertible formats above are not extractable at all: the auxiliary input may carry a signed request whose signing key no extractor recovers. This is the barrier of [BCPR14, BP15] in our setting. The notion has content beyond invertibility only in idealized models in which a well-formed request cannot be handed over but must be built with oracle queries: a local random oracle, whose queries the extractor sees, or an algebraic group, whose representations it is given. The same observation is what lets a simulation-based proof use the notion. Against the dummy adversary, whose requests are its auxiliary input, the extractor is $E_{\mathcal{D}}$, fixed once, and the simulator runs it on the environment’s instructions without knowing the environment. We adopt Definition B.2 as stated so as to cover these models uniformly; the invertible cases above need none of this.

Section 2.1 raises the complementary question of the protocol–predicate pairs for which knowledge of $(\text{data}, \text{reply})$ already witnesses execution. Both are conditions on (Prot, Φ) ; a systematic account is left to future work (Section 2.3).

B.3 Ideal Functionalities

We work in the $(\mathcal{F}_{\text{SC}}, \mathcal{F}_{\text{CERT}}, \mathcal{F}_{\text{KE}}, \mathcal{F}_{\text{US}})$ -hybrid model. The first three are standard UC functionalities lifted from prior work [CK02, Can04], with one structural modification to $\mathcal{F}_{\text{CERT}}$ (described below) that we adopt so that the simulator can commit to its certification algorithms at setup time rather than be queried per signature. $\mathcal{F}_{\text{US}}$ is new to this work; it was motivated informally in the Technical Overview, and we give its formal definition here. All four functionalities are stated below.

¹¹Which values the lever reaches depends on how the oracle is modeled. If it is local to the execution under analysis, every query put to it by or for the adversary is visible and extraction goes through. A *global* random oracle [CJS14, CDG⁺18] is shared with everything outside that execution, which may query it unobserved; the lever is then lost, and only the restricted observable and programmable variants recover part of it. A value an outside party obtained from the oracle and handed to $\mathcal{A}$ is, in this respect, auxiliary input.

Secure Channel $\mathcal{F}_{\text{SC}}$

Functionality $\mathcal{F}_{\text{SC}}$ (secure channel between two parties P_i, P_j).

Implicit state: per session identifier sid , an activation record (initially empty).

- **Activation.** On $(\text{Establish-session}, \text{sid}, P', \text{role})$ from a party P , where $\text{role} \in \{\text{initiator}, \text{responder}\}$: record (P, P', role) and notify Sim. Mark the session active once both P_i and P_j have submitted matching activations.
- **Transmission.** On $(\text{Send}, \text{sid}, m)$ from an active partner P : leak $(\text{sid}, |m|, P)$ to Sim, then output $(\text{Received}, \text{sid}, m)$ to the other partner.
- **Deactivation.** On $(\text{Expire-session}, \text{sid})$ from an active partner: mark the session inactive and notify Sim.

Figure 6: Functionality $\mathcal{F}_{\text{SC}}$ [CK02].

Guarantee. Between two uncorrupted partners, $\mathcal{F}_{\text{SC}}$ leaks only the lengths of transmitted messages to Sim; in particular, $\mathcal{A}$ learns no contents of $(\text{Send}, \cdot, m)$ messages, nor can it modify, reorder, or drop them.

Certification $\mathcal{F}_{\text{CERT}}$

This is an adaptation of the UC certification functionality of [Can04]: the simulator commits a $(\text{Sign}, \text{Vfy})$ algorithm pair at the first activation of each instance, and $\mathcal{F}_{\text{CERT}}$ subsequently runs these algorithms locally on Sign and Verify calls. Unforgeability when the owner is uncorrupted is enforced by $\mathcal{F}_{\text{CERT}}$'s internal sign table $\mathcal{T}$, independently of the Vfy algorithm supplied by Sim.

Functionality $\mathcal{F}_{\text{CERT}}[\text{pid}]$ (certification, owned by signer identity pid ; instances indexed by $\text{sid} = (\text{pid}, \text{sid}')$).

State: an algorithm pair $(\text{Sign}, \text{Vfy})$, initially undefined; a table $\mathcal{T} \subseteq \{0, 1\}^* \times \{0, 1\}^*$ of signed pairs, initially empty.

- **Initialization** (triggered on the first activation under sid). Notify Sim of the activation; on receipt of $(\text{Algorithms}, \text{sid}, \text{Sign}, \text{Vfy})$ from Sim, record $(\text{Sign}, \text{Vfy})$ as this instance's algorithms. Then proceed with the current activation below.
- **Sign.** On $(\text{Sign}, \text{sid}, m)$ from pid : compute $\sigma \leftarrow \text{Sign}(m)$, insert (m, σ) into $\mathcal{T}$, and output $(\text{Signature}, \text{sid}, m, \sigma)$ to pid . Notify Sim of the activation.
- **Verify.** On $(\text{Verify}, \text{sid}, m, \sigma)$ from any party P :
 - ◊ if pid is uncorrupted, output $(\text{Verified}, \text{sid}, m, 1)$ if $(m, \sigma) \in \mathcal{T}$, and $(\text{Verified}, \text{sid}, m, 0)$ otherwise;
 - ◊ if pid is corrupted, output $(\text{Verified}, \text{sid}, m, \text{Vfy}(m, \sigma))$.

Figure 7: Functionality $\mathcal{F}_{\text{CERT}}[\text{pid}]$, adapted from [Can04] with simulator-provided algorithms.

Guarantee (unforgeability). If the owner pid is uncorrupted, $\mathcal{F}_{\text{CERT}}[\text{pid}].\text{Verify}$ outputs 1 on (m, σ) if and only if pid previously requested $\mathcal{F}_{\text{CERT}}[\text{pid}].\text{Sign}(m)$ and received σ in reply. In particular, no third party, including one that holds the committed Sign and Vfy algorithms via Sim's setup-time disclosure, can produce a fresh $(m^*, \sigma^*) \notin \mathcal{T}$ that $\mathcal{F}_{\text{CERT}}[\text{pid}]$ accepts.

Key Exchange $\mathcal{F}_{\text{KE}}$

Guarantee. When both endpoints are uncorrupted, κ is uniformly distributed in $\{0, 1\}^\lambda$ and is delivered identically to both endpoints; Sim learns only the activation and not κ .

Functionality $\mathcal{F}_{KE}$ (key exchange between two parties P_i, P_j). <i>State:</i> per session identifier sid, an activation record (initially empty). – Activation. On $(\text{Establish-session}, sid, P_i, P_j, \text{role})$ from a party $P \in \{P_i, P_j\}$, with $\text{role} \in \{\text{initiator}, \text{responder}\}$: record $(P, P_i, P_j, \text{role})$ and notify Sim. – Key delivery. Once both P_i and P_j have submitted matching activations under sid: ◊ if both endpoints are uncorrupted, sample $\kappa \xleftarrow{\\$} \{0, 1\}^\lambda$, output $(\text{key}, sid, \kappa)$ to each of P_i, P_j, and notify Sim of the delivery; ◊ if either endpoint is corrupted, request κ from Sim and output $(\text{key}, sid, \kappa)$ to each of P_i, P_j.

Figure 8: Functionality $\mathcal{F}_{KE}$ [CK02].

User–Service Channel $\mathcal{F}_{US}$

The new functionality $\mathcal{F}_{US}$ is parameterized by the corruption status of A_K : it behaves as an adversarial network when A_K is uncorrupted, and as a UC secure channel when A_K is corrupted. The motivation for this dichotomy is laid out in Subsection 2.2.1 of the Technical Overview.

Functionality $\mathcal{F}_{US}$ (User–Service channel; transmission mode determined by the corruption status of A_K). <i>State:</i> per session identifier sid, an activation record (initially empty). – Activation. On $(\text{Establish-session}, sid, P', \text{role})$ from U and from S_{id} with matching roles: mark the session active and notify Sim. – Transmission. On (Send, sid, m) from an active partner P: ◊ if A_K is uncorrupted (<i>adversarial-network mode</i>): leak (sid, m, P) to Sim, and on Sim’s instruction $(\text{Deliver}, sid, m')$, output $(\text{Received}, sid, m')$ to the other partner. Sim may also instruct $\mathcal{F}_{US}$ to drop the message. ◊ if A_K is corrupted (<i>secure-channel mode</i>): leak (sid, m , P) to Sim, then output $(\text{Received}, sid, m)$ to the other partner.
--

Figure 9: Functionality $\mathcal{F}_{US}$ (new).

Guarantee. If A_K is uncorrupted, $\mathcal{A}$ sees every byte on the U – S_{id} wire and may modify, reorder, or drop messages at will. If A_K is corrupted, $\mathcal{F}_{US}$ acts as $\mathcal{F}_{SC}$ restricted to the U – S_{id} link: $\mathcal{A}$ learns only message lengths and cannot affect delivery.

C Security of Distributed-AEAD

This section gives the formal counterpart to the informal security discussion of Section 4.2. We provide game-based definitions of the two security requirements claimed for D-AEAD – splitting confidentiality into role-specific games $\text{IND-CPA-}\mathcal{P}_K$ and $\text{IND-CCA-}\mathcal{P}_T$ to match the asymmetric corruption model, and formalizing ciphertext integrity as a single game INT-CTXT against a malicious data party – and prove in Theorem C.6 that the construction of Section 4.1 satisfies all three under two assumptions on its underlying primitives: joint pseudorandomness of $(\text{GenEncMat}, \text{GenTagMat})$ and one-time tag unforgeability of $(\text{ComputeTag}, \text{VerifyTag})$. For the specific instantiations of Section C.4, both assumptions reduce in a single step to standard properties of the underlying primitives – PRF security of AES (resp. the ChaCha20 core) and unforgeability of the GHASH (resp. Poly1305), so the security of D-AEAD rests on the same foundations as the AEAD schemes we are interested in this work. Note that the two confidentiality games are not symmetric: $\text{IND-CPA-}\mathcal{P}_K$ holds information-theoretically, because $\mathcal{P}_K$ is shielded from the plaintext at SR by a one-time pad applied by $\mathcal{P}_D$

before $\mathcal{P}_K$ enters the protocol, whereas $\text{IND-CCA-}\mathcal{P}_T$ is computational and reduces to the pseudorandomness of the keystream and tag-material primitives.

Notation. For a corrupt role $X \in \{K, C\}$, we write view_X for the view of $\mathcal{P}_X$ in a single honest execution of Π_{AEnc} or Π_{ADec} : that party's private input, internal randomness, and incoming messages on its channels; outgoing messages and internal computations are deterministic functions of these. We write $\overline{\text{SR}} = [\ell] \setminus \text{SR}$ for the public positions of an ℓ -bit plaintext. The confidentiality games below treat SR as part of the public common input; only the message contents at the positions indexed by SR are required to be hidden from $\mathcal{P}_K$ and $\mathcal{P}_T$.

We formalize below the two security requirements of Section 4.2. Confidentiality is split per honest-but-curious role into $\text{IND-CPA-}\mathcal{P}_K$ (Definition C.1) and $\text{IND-CCA-}\mathcal{P}_T$ (Definition C.2); ciphertext integrity against a malicious data party is captured by INT-CTXT (Definition C.3). The reduction to assumptions on the underlying primitives is stated in Theorem C.6.

C.1 Security Properties

Confidentiality against $\mathcal{P}_K$. A scheme D-AEAD provides confidentiality against $\mathcal{P}_K$ if no PPT adversary, given the long-term key k and the view view_K that $\mathcal{P}_K$ would obtain in an honest execution of Π_{AEnc} run on one of two adversarially chosen, equal-length plaintexts P_0, P_1 agreeing on $\overline{\text{SR}}$, can determine which of the two plaintexts was the input to that execution. The mechanism that protects $P_b|_{\text{SR}}$ from $\mathcal{P}_K$ is the one-time-pad masking $\tilde{P} = P_b \oplus \text{embed}_{\text{SR}, \ell}(M)$ that $\mathcal{P}_D$ applies before $\mathcal{P}_K$ enters the protocol. Accordingly, the bound on the $\text{IND-CPA-}\mathcal{P}_K$ advantage is information-theoretic and uses none of the assumptions on D-AEAD's underlying primitives. Restriction to a single challenge with no oracle access is without loss of generality: since $\mathcal{A}$ holds k and the common inputs are public, every honest action of $\mathcal{P}_K$ and $\mathcal{P}_T$ in Π_{AEnc} or Π_{ADec} is a deterministic function of public quantities together with k , so $\mathcal{A}$ may simulate any number of further encryption and decryption sessions locally. In Π_{ADec} specifically, $\mathcal{P}_K$'s incoming messages reduce to $|C|$ and the verdict bit v : the length $|C|$ leaks only the plaintext length, which is already pinned down by the equal-length constraint $|P_0| = |P_1|$, and the verdict $v = \text{VerifyTag}(\text{GenTagMat}(k, n), AD, C, T)$ is a deterministic function of k and the queried tuple and hence computable by $\mathcal{A}$. More formally:

Definition C.1 (Confidentiality against $\mathcal{P}_K$) A scheme D-AEAD provides confidentiality against $\mathcal{P}_K$ if for every PPT $\mathcal{A}$,

$$\begin{aligned} \text{Adv}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CPA-}\mathcal{P}_K}(\lambda) &= |\Pr[\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CPA-}\mathcal{P}_K-1}(\lambda) = 1] - \Pr[\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CPA-}\mathcal{P}_K-0}(\lambda) = 1]| \\ &\leq \text{negl}(\lambda), \end{aligned}$$

where $\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CPA-}\mathcal{P}_K-b}$ is defined in figure 10.

Experiment $\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CPA-}\mathcal{P}_K-b}(\lambda)$:

1. $k \xleftarrow{\$} \text{KGen}(1^\lambda)$;
2. $(\text{sid}, n, AD, \ell, \text{SR}, P_0, P_1, \text{st}) \leftarrow \mathcal{A}(1^\lambda, k)$;
3. if $|P_0| \neq \ell$ or $|P_1| \neq \ell$ or $P_0|_{\overline{\text{SR}}} \neq P_1|_{\overline{\text{SR}}}$, then abort;
4. $M \xleftarrow{\$} \{0, 1\}^{|\text{SR}|}$; $\tilde{P} \leftarrow P_b \oplus \text{embed}_{\text{SR}, \ell}(M)$;
5. $b' \leftarrow \mathcal{A}(\text{st}, \tilde{P})$; return b' .

Figure 10: Experiment $\text{IND-CPA-}\mathcal{P}_K$

Note that the view view_K of $\mathcal{P}_K$ in the honest Π_{AEnc} run on common input $(\text{sid}, n, AD, \ell, \text{SR})$ and private input k is reduced to the single incoming message $\tilde{P}$: the remaining elements $K_s = \text{GenEncMat}(k, n, \ell)$, $\tau = \text{GenTagMat}(k, n)$, and $\tilde{C} = \tilde{P} \oplus K_s$ are deterministic functions of $(k, n, \ell, \tilde{P})$, already in $\mathcal{A}$'s state.

Confidentiality against $\mathcal{P}_T$. A scheme D-AEAD provides confidentiality against $\mathcal{P}_T$ if no PPT adversary, given the view view_C that $\mathcal{P}_T$ would obtain in a single challenge Π_{AEnc} execution on one of two adversarially chosen, equal-length plaintexts P_0, P_1 agreeing on $\overline{\text{SR}}$, and adaptive oracle access to honest executions of Π_{AEnc} (the encryption oracle $\mathcal{O}_{\text{Enc}}$, nonce-respecting, with the challenge nonce excluded from later queries) and Π_{ADec} (the decryption oracle $\mathcal{O}_{\text{Dec}}$, with the challenge ciphertext-tag tuple excluded from later queries), can determine which of the two plaintexts was the input to the challenge execution made through the challenge oracle $\mathcal{O}_{\text{Chal-}b}$. Unlike in $\text{IND-CPA-}\mathcal{P}_K$, the relevant protection here is the AEAD encryption step itself: $\mathcal{P}_T$ sees the ciphertext $C = P_b \oplus K_s$ but never sees the keystream K_s or the masked plaintext $\tilde{P}$ from which K_s could be recovered, so the $\text{IND-CCA-}\mathcal{P}_T$ advantage is bounded computationally and the reduction goes through pseudorandomness of $(\text{GenEncMat}, \text{GenTagMat})$. More formally:

Definition C.2 (Confidentiality against $\mathcal{P}_T$) A scheme D-AEAD provides confidentiality against $\mathcal{P}_T$ if for every PPT $\mathcal{A}$,

$$\begin{aligned} \text{Adv}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T}(\lambda) &= |\Pr[\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T-1}(\lambda) = 1] - \Pr[\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T-0}(\lambda) = 1]| \\ &\leq \text{negl}(\lambda), \end{aligned}$$

where $\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T-b}$ together with its oracles $\mathcal{O}_{\text{Enc}}, \mathcal{O}_{\text{Dec}}, \mathcal{O}_{\text{Chal-}b}$ is defined in figure 11.

Experiment $\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T-b}(\lambda)$:

1. $k \xleftarrow{\$} \text{KGen}(1^\lambda); \mathcal{N}_{\text{used}} \leftarrow \emptyset; \text{ch} \leftarrow \perp;$
2. $b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Enc}}, \mathcal{O}_{\text{Dec}}, \mathcal{O}_{\text{Chal-}b}}(1^\lambda);$
3. return b' .

Oracle $\mathcal{O}_{\text{Enc}}(n, AD, P, \text{SR})$:

1. if $n \in \mathcal{N}_{\text{used}}$ return $\perp$; otherwise $\mathcal{N}_{\text{used}} \leftarrow \mathcal{N}_{\text{used}} \cup \{n\};$
2. $M \xleftarrow{\$} \{0, 1\}^{|\text{SR}|};$
3. $\tilde{P} \leftarrow P \oplus \text{embed}_{\text{SR}, |P|}(M);$
4. $K_s \leftarrow \text{GenEncMat}(k, n, |P|);$
5. $\tau \leftarrow \text{GenTagMat}(k, n);$
6. $\tilde{C} \leftarrow \tilde{P} \oplus K_s; C \leftarrow \tilde{C} \oplus \text{embed}_{\text{SR}, |P|}(M);$
7. $T \leftarrow \text{ComputeTag}(\tau, AD, C);$
8. return $(M, \tilde{C}, \tau, C, T).$

Oracle $\mathcal{O}_{\text{Dec}}(n, AD, C, T)$:

1. if $(n, AD, C, T) = \text{ch}$ return $\perp$;
2. $\tau \leftarrow \text{GenTagMat}(k, n);$
3. $v \leftarrow \text{VerifyTag}(\tau, AD, C, T);$
4. return $(\tau, v).$

Oracle $\mathcal{O}_{\text{Chal-}b}(n^*, AD^*, \text{SR}^*, P_0, P_1)$:

1. if $|P_0| \neq |P_1|$, or $P_0|_{\overline{\text{SR}}^*} \neq P_1|_{\overline{\text{SR}}^*}$, or $n^* \in \mathcal{N}_{\text{used}}$, or $\text{ch} \neq \perp$, return $\perp$;
2. $\mathcal{N}_{\text{used}} \leftarrow \mathcal{N}_{\text{used}} \cup \{n^*\};$
3. compute $(M^*, \tilde{C}^*, \tau^*, C^*, T^*)$ by running the body of $\mathcal{O}_{\text{Enc}}$ on $(n^*, AD^*, P_b, \text{SR}^*)$, omitting its freshness check;
4. $\text{ch} \leftarrow (n^*, AD^*, C^*, T^*);$
5. return $(M^*, \tilde{C}^*, \tau^*, C^*, T^*).$

Figure 11: Experiment $\text{IND-CCA-}\mathcal{P}_T$ and oracles $\mathcal{O}_{\text{Enc}}, \mathcal{O}_{\text{Dec}}, \mathcal{O}_{\text{Chal-}b}$

Note that the returned tuples coincide with view_C in an honest run: in Π_{AEnc} , $\mathcal{P}_T$ receives M from $\mathcal{P}_D$ and $(\tilde{C}, \tau)$ from $\mathcal{P}_K$ and emits (C, T) to $\mathcal{P}_D$; in Π_{ADec} , $\mathcal{P}_T$ receives (C, T) from $\mathcal{P}_D$ (already in the query) and τ from $\mathcal{P}_K$, and emits the verdict v .

Ciphertext integrity. A scheme D-AEAD provides ciphertext integrity if no PPT adversary, given adaptive access to an encryption oracle $\mathcal{O}_{\text{Enc}}$ (returning the ciphertext-tag pair an honest Π_{AEnc} run would produce on an adversarially chosen plaintext, nonce-respecting on the encryption side) and a decryption oracle $\mathcal{O}_{\text{Dec}}$ (returning the verdict an honest Π_{ADec} run would produce on an adversarially chosen tuple, with no nonce restriction), can submit to $\mathcal{O}_{\text{Dec}}$ a tuple (n^*, AD^*, C^*, T^*) that was not previously delivered by $\mathcal{O}_{\text{Enc}}$ and that Π_{ADec} accepts. The two oracles capture exactly the leverage available to a malicious $\mathcal{P}_D$ when interacting with honest $\mathcal{P}_K$ and $\mathcal{P}_T$: every outgoing message $\mathcal{P}_D$ can place on its Π_{AEnc} or Π_{ADec} channels induces the same honest-party output as a single $\mathcal{O}_{\text{Enc}}$ or $\mathcal{O}_{\text{Dec}}$ call, respectively; the formal protocol-to-game reduction is given inside the proof of Theorem C.6.

Definition C.3 (Ciphertext integrity) A scheme D-AEAD provides ciphertext integrity if for every PPT $\mathcal{A}$, the advantage

$$\text{Adv}_{\text{D-AEAD}, \mathcal{A}}^{\text{INT-CTXT}}(\lambda) = \Pr[\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{INT-CTXT}}(\lambda) = 1] \leq \text{negl}(\lambda),$$

where the experiment $\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{INT-CTXT}}$ together with its oracles $\mathcal{O}_{\text{Enc}}$ and $\mathcal{O}_{\text{Dec}}$ is defined in figure 12.




Figure 12: Experiment INT-CTXT and oracles $\mathcal{O}_{\text{Enc}}, \mathcal{O}_{\text{Dec}}$.

Note that the $\mathcal{O}_{\text{Enc}}$ output (C, T) coincides with what a (possibly malicious) $\mathcal{P}_D$ receives from $\mathcal{P}_T$ at the end of an honest Π_{AEnc} run on the effective plaintext induced by $\mathcal{P}_D$'s outgoing messages; the $\mathcal{O}_{\text{Dec}}$ verdict v coincides with the verdict $\mathcal{P}_D$ obtains from honest $\mathcal{P}_K$ and $\mathcal{P}_T$ at the end of Π_{ADec} on (n, AD, C, T) .

C.2 Assumptions on the underlying primitives

The reduction in Theorem C.6 uses two assumptions on the underlying primitives: joint pseudorandomness of $(\text{GenEncMat}, \text{GenTagMat})$ and one-time tag unforgeability of $(\text{ComputeTag}, \text{VerifyTag})$. Both are properties of the components in isolation, and both reduce in a single step to standard primitive properties for the concrete schemes of Section C.4.

Joint pseudorandomness. A pair $(\text{GenEncMat}, \text{GenTagMat})$ is *jointly pseudorandom* if no PPT adversary distinguishes the pair of oracles $(\mathcal{O}_S, \mathcal{O}_T)$ instantiated with a fresh key $k \xleftarrow{\$} \text{KGen}(1^\lambda)$ from the same pair instantiated with uniformly sampled outputs. The stream oracle $\mathcal{O}_S$ is nonce-respecting and returns a keystream of the requested length on each fresh nonce; the tag-material oracle $\mathcal{O}_T$ accepts repeated queries on the same nonce and returns the same value on each repetition, capturing the fact that $\text{GenTagMat}(k, n)$ depends only on (k, n) and is executed both at encryption and at every decryption attempt on n . Without loss of generality we assume the tag-material space of form $\mathcal{TM} = \mathcal{SP} \times \mathcal{TM}_\tau$ for some sets $\mathcal{SP}$ (the *shared subkey* space) and $\mathcal{TM}_\tau$ (the *per-nonce material* space). We require that GenTagMat decomposes as $\text{GenTagMat}(k, n) = (\sigma(k), \tau(k, n))$ for some functions $\sigma : \mathcal{K} \rightarrow \mathcal{SP}$ and $\tau : \mathcal{K} \times \mathcal{N} \rightarrow \mathcal{TM}_\tau$, that is, the first coordinate of $\tau^{(n)}$ depends only on k . When $\mathcal{SP} = \{\perp\}$ the decomposition is trivial. More formally:

Definition C.4 (Joint pseudorandomness) A pair $(\text{GenEncMat}, \text{GenTagMat})$ is *jointly pseudorandom* if for every PPT $\mathcal{A}$,

$$\begin{aligned} \text{Adv}_{(\text{GenEncMat}, \text{GenTagMat}), \mathcal{A}}^{\text{jPR}}(\lambda) &= |\Pr[\text{Exp}_{\mathcal{A}}^{\text{jPR-1}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A}}^{\text{jPR-0}}(\lambda) = 1]| \\ &\leq \text{negl}(\lambda), \end{aligned}$$

where $\text{Exp}_{\mathcal{A}}^{\text{jPR-}b}$ is defined in figure 13.

Experiment $\text{Exp}_{(\text{GenEncMat}, \text{GenTagMat}), \mathcal{A}}^{\text{jPR}-b}(\lambda):$ 1. $\mathcal{N}_{\text{used}}^S \leftarrow \emptyset; \mathbf{T} \leftarrow [];$ 2. if $b = 1$: $k \xleftarrow{\\$} \text{KGen}(1^\lambda);$ 3. if $b = 0$: $\sigma \xleftarrow{\\$} \mathcal{SP};$ 4. $b' \leftarrow \mathcal{A}^{\mathcal{O}_S, \mathcal{O}_\mathbf{T}}(1^\lambda);$ return b'.	Oracle $\mathcal{O}_S(n, \ell):$ 1. if $n \in \mathcal{N}_{\text{used}}^S$ return $\perp$; otherwise $\mathcal{N}_{\text{used}}^S \leftarrow \mathcal{N}_{\text{used}}^S \cup \{n\};$ 2. if $b = 1$ return $\text{GenEncMat}(k, n, \ell)$ else return $K_s \xleftarrow{\\$} \{0, 1\}^\ell.$ Oracle $\mathcal{O}_\mathbf{T}(n):$ 1. if $\mathbf{T}[n] \neq \perp$ return $\mathbf{T}[n];$ 2. if $b = 1$ set $\mathbf{T}[n] \leftarrow \text{GenTagMat}(k, n);$ 3. else: $\tau_n \xleftarrow{\\$} \mathcal{TM}_\tau$, set $\mathbf{T}[n] \leftarrow (\sigma, \tau_n);$ 4. return $\mathbf{T}[n].$
--	--

Figure 13: Experiment jPR over the product structure $\mathcal{TM} = \mathcal{SP} \times \mathcal{TM}_\tau$.

One-time tag unforgeability. Sample a shared subkey $\sigma \xleftarrow{\$} \mathcal{SP}$ once and, on the first appearance of each nonce n , a per-nonce material $\tau_n \xleftarrow{\$} \mathcal{TM}_\tau$. The adversary is allowed at most one $\mathcal{O}_{\text{Tag}}$ call per nonce and arbitrarily many $\mathcal{O}_{\text{Ver}}$ calls on arbitrary nonces; no PPT adversary can produce a fresh tuple (n, AD, C, T) that VerifyTag accepts. More formally:

Definition C.5 (One-time tag unforgeability) A pair $(\text{ComputeTag}, \text{VerifyTag})$ is *one-time tag unforgeable* if for every PPT $\mathcal{A}$,

$$\text{Adv}_{(\text{ComputeTag}, \text{VerifyTag}), \mathcal{A}}^{\text{OTUF}}(\lambda) = \Pr[\text{Exp}_{\mathcal{A}}^{\text{OTUF}}(\lambda) = 1] \leq \text{negl}(\lambda),$$

where $\text{Exp}_{\mathcal{A}}^{\text{OTUF}}$ and its oracles $\mathcal{O}_{\text{Tag}}, \mathcal{O}_{\text{Ver}}$ are defined in figure 14.

Experiment $\text{Exp}_{(\text{ComputeTag}, \text{VerifyTag}), \mathcal{A}}^{\text{OTUF}}(\lambda):$ 1. $\sigma \xleftarrow{\\$} \mathcal{SP}; \mathbf{T} \leftarrow []; \mathcal{S} \leftarrow \emptyset; \text{Tagged} \leftarrow \emptyset; \text{win} \leftarrow 0;$ 2. run $\mathcal{A}^{\mathcal{O}_{\text{Tag}}, \mathcal{O}_{\text{Ver}}}(1^\lambda);$ return win.	Oracle $\mathcal{O}_{\text{Tag}}(n, AD, C):$ 1. if $n \in \text{Tagged}$ return $\perp$; otherwise $\text{Tagged} \leftarrow \text{Tagged} \cup \{n\};$ 2. if $\mathbf{T}[n] = \perp$: $\tau_n \xleftarrow{\\$} \mathcal{TM}_\tau$, $\mathbf{T}[n] \leftarrow (\sigma, \tau_n);$ 3. $T \leftarrow \text{ComputeTag}(\mathbf{T}[n], AD, C); \mathcal{S} \leftarrow \mathcal{S} \cup \{(n, AD, C, T)\};$ return $T.$
Oracle $\mathcal{O}_{\text{Ver}}(n, AD, C, T):$ 1. if $\mathbf{T}[n] = \perp$: $\tau_n \xleftarrow{\\$} \mathcal{TM}_\tau$, $\mathbf{T}[n] \leftarrow (\sigma, \tau_n);$ 2. $v \leftarrow \text{VerifyTag}(\mathbf{T}[n], AD, C, T);$ 3. if $v = 1$ and $(n, AD, C, T) \notin \mathcal{S}$, set win $\leftarrow 1$; return $v.$	

Figure 14: Experiment OTUF with shared subkey σ and per-nonce one-time tag access.

C.3 Main theorem

Theorem C.6 (Security of D-AEAD) Let $\mathcal{A}$ be a PPT adversary making at most q_e queries to $\mathcal{O}_{\text{Enc}}$ and q_d queries to $\mathcal{O}_{\text{Dec}}$ in the experiments where these oracles are present.

1. $\text{Adv}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CPA-PK}}(\lambda) = 0.$
2. There exists a PPT adversary $\mathcal{B}$ such that

$$\text{Adv}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-P}_T}(\lambda) \leq 2 \cdot \text{Adv}_{(\text{GenEncMat}, \text{GenTagMat}), \mathcal{B}}^{\text{jPR}}(\lambda).$$

3. There exist PPT adversaries $\mathcal{B}, \mathcal{C}$ such that

$$\text{Adv}_{\text{D-AEAD}, \mathcal{A}}^{\text{INT-CTXT}}(\lambda) \leq \text{Adv}_{(\text{GenEncMat}, \text{GenTagMat}), \mathcal{B}}^{\text{jPR}}(\lambda) + \text{Adv}_{(\text{ComputeTag}, \text{VerifyTag}), \mathcal{C}}^{\text{OTUF}}(\lambda).$$

Proof. We treat the three claims in turn.

Confidentiality against $\mathcal{P}_K$. The bit b enters $\text{Exp}^{\text{IND-CPA-}\mathcal{P}_K-b}$ only through $\tilde{P} = P_b \oplus \text{embed}_{\text{SR},\ell}(M)$, where $M \xleftarrow{\$} \{0,1\}^{|\text{SR}|}$ is uniform and independent of P_b . At positions in SR , $\tilde{P}|_{\text{SR}} = P_b|_{\text{SR}} \oplus M$ is a one-time-pad encryption of $P_b|_{\text{SR}}$ under the uniform key M , so $\tilde{P}|_{\text{SR}}$ is uniform on $\{0,1\}^{|\text{SR}|}$ irrespective of b . At positions in $\overline{\text{SR}}$ the embedded mask is zero, so $\tilde{P}|_{\overline{\text{SR}}} = P_b|_{\overline{\text{SR}}}$, which equals $P_0|_{\overline{\text{SR}}} = P_1|_{\overline{\text{SR}}}$. The distribution of $\tilde{P}$ is therefore independent of b , and the advantage is zero.

Confidentiality against $\mathcal{P}_T$. Define the hybrid experiment Hyb^b as $\text{Exp}^{\text{IND-CCA-}\mathcal{P}_T-b}$ with the following modification: on first occurrence of nonce n , sample $\tau^{(n)} \xleftarrow{\$} \mathcal{TM}$ and cache it; whenever $\mathcal{O}_{\text{Enc}}$ or $\mathcal{O}_{\text{Chal-}b}$ first uses n , additionally sample $K_s^{(n)} \xleftarrow{\$} \{0,1\}^{|P|}$ and cache it. All subsequent invocations of $\text{GenEncMat}(k, n, \cdot)$ or $\text{GenTagMat}(k, n)$ on n return the cached values. We bound the advantage in two steps.

Step 1: independence of the hybrid from b . The bit b enters Hyb^b only through the $\mathcal{O}_{\text{Chal-}b}$ call on P_b . Inside $\mathcal{O}_{\text{Chal-}b}$, the keystream K_s^* is uniform on $\{0,1\}^{|P_b|}$ and independent of M^* and τ^* , so $C^* = P_b \oplus K_s^*$ is uniform on $\{0,1\}^{|P_b|}$ and independent of (M^*, τ^*) . The remaining components of the challenge tuple are deterministic functions of (M^*, τ^*, C^*) and of $AD^*, \text{SR}^*, |P_b|$: $\tilde{C}^* = C^* \oplus \text{embed}_{\text{SR}^*, |P_b|}(M^*)$ and $T^* = \text{ComputeTag}(\tau^*, AD^*, C^*)$. Since $\mathcal{O}_{\text{Chal-}b}$ enforces $|P_0| = |P_1|$, the joint distribution of $(M^*, \tilde{C}^*, \tau^*, C^*, T^*)$ is identical for $b = 0$ and $b = 1$. The remaining oracles take inputs from $\mathcal{A}$ that are not functions of b and produce outputs by procedures that do not branch on b . Therefore the joint distribution of $\mathcal{A}$'s view in Hyb^b is identical for $b \in \{0,1\}$, so $\Pr[\text{Hyb}^0 = 1] = \Pr[\text{Hyb}^1 = 1]$.

Step 2: reduction to jPR. Define a PPT adversary $\mathcal{B}$ against jPR as follows.

1. Receive 1^λ and oracle access to $(\mathcal{O}_S, \mathcal{O}_T)$ of the jPR experiment, and sample $c \xleftarrow{\$} \{0,1\}$.
2. Run $\mathcal{A}$ and simulate $\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T-c}$, with each invocation of $\text{GenEncMat}(k, n, \ell)$ inside $\mathcal{O}_{\text{Enc}}$ or $\mathcal{O}_{\text{Chal-}b}$ replaced by a call to $\mathcal{O}_S(n, \ell)$ and each invocation of $\text{GenTagMat}(k, n)$ inside $\mathcal{O}_{\text{Enc}}$, $\mathcal{O}_{\text{Dec}}$, or $\mathcal{O}_{\text{Chal-}b}$ replaced by a call to $\mathcal{O}_T(n)$.
3. When $\mathcal{A}$ halts and outputs a bit b' , return $[b' = c]$.

When $(\mathcal{O}_S, \mathcal{O}_T)$ are the real jPR oracles, $\mathcal{B}$'s simulation is distributed exactly as $\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T-c}$, so $\Pr[\mathcal{B}=1 \mid \text{jPR}=1] = \frac{1}{2} + \frac{1}{2}(\Pr[\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T-1}=1] - \Pr[\text{Exp}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T-0}=1])$. When $(\mathcal{O}_S, \mathcal{O}_T)$ are the ideal jPR oracles, the simulation is distributed exactly as Hyb^c , and by Step 1 the output of $\mathcal{A}$ in Hyb^c is independent of c , so $\Pr[\mathcal{B}=1 \mid \text{jPR}=0] = \frac{1}{2}$. Taking the absolute difference, $\text{Adv}_{(\text{GenEncMat}, \text{GenTagMat}), \mathcal{B}}^{\text{jPR}}(\lambda) = \frac{1}{2} \text{Adv}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T}(\lambda)$, which gives the stated bound $\text{Adv}_{\text{D-AEAD}, \mathcal{A}}^{\text{IND-CCA-}\mathcal{P}_T}(\lambda) \leq 2 \cdot \text{Adv}_{(\text{GenEncMat}, \text{GenTagMat}), \mathcal{B}}^{\text{jPR}}(\lambda)$.

Ciphertext integrity. Define the hybrid experiment Hyb as $\text{Exp}^{\text{INT-CTXT}}$ with the same modification as above: each $(K_s^{(n)}, \tau^{(n)})$ is sampled uniformly on first use of n and cached. We bound the advantage in two steps.

Step 1: indistinguishability from the hybrid. Define a PPT adversary $\mathcal{B}$ against jPR by the same construction as in the IND-CCA- $\mathcal{P}_T$ proof, with the simulated experiment being $\text{Exp}^{\text{INT-CTXT}}$ and the returned value being the experiment's win bit. By the same argument, $|\Pr[\text{Exp}^{\text{INT-CTXT}} = 1] - \Pr[\text{Hyb} = 1]| = \text{Adv}_{\mathcal{B}}^{\text{jPR}}(\lambda)$.

Step 2: bounding the hybrid via OTUF. Define a PPT adversary $\mathcal{C}$ against OTUF as follows.

1. Receive 1^λ and oracle access to $(\mathcal{O}_{\text{Tag}}, \mathcal{O}_{\text{Ver}})$ of the OTUF experiment; initialize $\mathcal{N}_{\text{used}}, \mathcal{Q} \leftarrow \emptyset$.
2. Run $\mathcal{A}$ and simulate Hyb as follows.
3. On every $\mathcal{O}_{\text{Enc}}(n, AD, P)$ call: if $n \in \mathcal{N}_{\text{used}}$ return $\perp$, otherwise $\mathcal{N}_{\text{used}} \leftarrow \mathcal{N}_{\text{used}} \cup \{n\}$; sample $K_s \xleftarrow{\$} \{0,1\}^{|P|}$; set $C \leftarrow P \oplus K_s$; obtain $T \leftarrow \mathcal{O}_{\text{Tag}}(n, AD, C)$; $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(n, AD, C, T)\}$; return (C, T) .
4. On every $\mathcal{O}_{\text{Dec}}(n, AD, C, T)$ call: obtain $v \leftarrow \mathcal{O}_{\text{Ver}}(n, AD, C, T)$ and return v .
5. Halt when $\mathcal{A}$ halts.

The shared subkey σ and each per-nonce τ_n in $\mathcal{C}$'s simulation are sampled by the OTUF experiment exactly as Hyb would sample them, and each K_s is uniform on $\{0,1\}^{|P|}$, so $\mathcal{A}$'s view under $\mathcal{C}$ is distributed exactly

as in Hyb. The set $\mathcal{Q}$ maintained by $\mathcal{C}$ coincides with the OTUF experiment's $\mathcal{S}$, so $\text{win} = 1$ in OTUF exactly when $\mathcal{A}$'s $\mathcal{O}_{\text{Dec}}$ query returns 1 on a tuple absent from $\mathcal{Q}$, which is the win condition of Hyb. Hence $\text{Adv}_{(\text{ComputeTag}, \text{VerifyTag}), \mathcal{C}}^{\text{OTUF}}(\lambda) = \Pr[\text{Hyb} = 1]$.

$$\text{Combining. } \text{Adv}_{\text{D-AEAD}, \mathcal{A}}^{\text{INT-CTXT}}(\lambda) \leq |\Pr[\text{Exp}^{\text{INT-CTXT}} = 1] - \Pr[\text{Hyb} = 1]| + \Pr[\text{Hyb} = 1] = \text{Adv}_{(\text{GenEncMat}, \text{GenTagMat}), \mathcal{B}}^{\text{jPR}}(\lambda) + \text{Adv}_{(\text{ComputeTag}, \text{VerifyTag}), \mathcal{C}}^{\text{OTUF}}(\lambda).$$

C.4 Concrete schemes

AES-GCM [Dwo07]. Let $\kappa \in \{128, 256\}$. Set $|\mathcal{K}| = \kappa$, $|\mathcal{N}| = 96$, $|\mathcal{T}| = 128$, $\mathcal{SP} = \{0, 1\}^{128}$, $\mathcal{TM}_\tau = \{0, 1\}^{128}$, so $\mathcal{TM} = \mathcal{SP} \times \mathcal{TM}_\tau$. Let $\text{AES}_k : \{0, 1\}^{128} \rightarrow \{0, 1\}^{128}$ denote AES block encryption with κ -bit key, $\text{CTR}_k(n, \ell)$ the AES-CTR keystream of length ℓ with 32-bit block counter starting at 2 per [Dwo07], and GHASH_H the polynomial-evaluation hash over $\text{GF}(2^{128})$ with subkey H , applied to the standard GCM padding of (AD, C) . Define

$$\begin{aligned} \text{GenEncMat}(k, n, \ell) &= \text{CTR}_k(n, \ell), \\ \text{GenTagMat}(k, n) &= (H, S), \quad H = \text{AES}_k(0^{128}), \quad S = \text{AES}_k(n \parallel 0^{31} \parallel 1), \\ \text{ComputeTag}((H, S), AD, C) &= \text{GHASH}_H(AD, C) \oplus S, \\ \text{VerifyTag}((H, S), AD, C, T) &= [\text{GHASH}_H(AD, C) \oplus S = T], \end{aligned}$$

where H depends only on k and instantiates the shared subkey $\sigma \in \mathcal{SP}$, while S instantiates the per-nonce material $\tau_n \in \mathcal{TM}_\tau$.

All AES calls land on pairwise disjoint inputs: 0^{128} for H , $n \parallel 0^{31} \parallel 1$ for S , and $n \parallel 0^{31} \parallel c$ with $c \geq 2$ for $\text{CTR}_k(n, \cdot)$. PRP/PRF switching replaces AES_k by a uniformly random function, yielding a single uniform H shared across all nonces, an independent uniform $S^{(n)}$ per nonce, and uniform keystreams – the distribution of the jPR ideal experiment of figure 13 – so jPR reduces in a single step to PRP security of AES_κ . With $(\sigma, \tau_n) = (H, S_n)$ both uniform on $\{0, 1\}^{128}$ in OTUF, the ε -almost-XOR-universality of GHASH [MV04] bounds the forgery probability by $q_v \cdot \lceil \ell_{\max}/128 \rceil \cdot 2^{-128}$ over q_v verify queries on inputs of total length at most $\ell_{\max}$ bits, discharging OTUF unconditionally.

ChaCha20-Poly1305 [NL15]. Set $|\mathcal{K}| = 256$, $|\mathcal{N}| = 96$, $|\mathcal{T}| = 128$, $\mathcal{SP} = \{\perp\}$, $\mathcal{TM}_\tau = \{0, 1\}^{256}$, so $\mathcal{TM} \cong \mathcal{TM}_\tau$ and the shared component is trivial. Let $\text{ChaCha20}_k(n, c, \ell)$ denote the ChaCha20 keystream of ℓ bits under key k and 96-bit nonce n starting at 32-bit block counter c , and $\text{Poly1305}_{(r,s)}$ the Poly1305 MAC under one-time key (r, s) , applied to the standard RFC 8439 [NL15] padding of (AD, C) . Define

$$\begin{aligned} \text{GenEncMat}(k, n, \ell) &= \text{ChaCha20}_k(n, 1, \ell), \\ \text{GenTagMat}(k, n) &= (r, s), \quad r \parallel s = \text{ChaCha20}_k(n, 0, 256), \\ \text{ComputeTag}((r, s), AD, C) &= \text{Poly1305}_{(r,s)}(AD \parallel C), \\ \text{VerifyTag}((r, s), AD, C, T) &= [\text{Poly1305}_{(r,s)}(AD \parallel C) = T]. \end{aligned}$$

The counter-zero block consumed by GenTagMat and the counter- ≥ 1 blocks consumed by GenEncMat are produced by disjoint inputs to the ChaCha20 core, so PRF security of the core reduces jPR in a single step to a standard primitive assumption; with $\mathcal{SP}$ trivial there is no shared component to align. With each $\tau_n = (r_n, s_n) \xleftarrow{\$} \{0, 1\}^{256}$ on first use of n in OTUF, the ε -almost- Δ -universality of the Poly1305 polynomial hash [Ber05] bounds the forgery probability by $q_v \cdot 8 \lceil \ell_{\max}/16 \rceil \cdot 2^{-106}$ over q_v verify queries on inputs of total length at most $\ell_{\max}$ bytes, discharging OTUF unconditionally.

D Reclaim Protocol Security Proof

Theorem 5.1 (restated) Let D-AEAD be a distributed stream-cipher AEAD scheme satisfying the confidentiality and ciphertext-integrity guarantees of Section 4.2, let every Prot_{id} in the deployment be request-reply and request-extractable (Definitions B.1 and B.2), and let every predicate be supported for its protocol:

admissible (Definition B.5) and of the form fixed in Section 5.2. Then the Reclaim protocol $\mathcal{RP}$ of Section 5 UC-realizes the ideal functionality $\mathcal{F}_{\text{zkPoPE}}$ of Section 3 in the $(\mathcal{F}_{\text{SC}}, \mathcal{F}_{\text{KE}}, \mathcal{F}_{\text{CERT}}, \mathcal{F}_{\text{US}})$ -hybrid model, under the admissible corruption pattern recalled in Section D.1.

D.1 Model and conventions

The functionality $\mathcal{F}_{\text{US}}$ and the admissible corruption pattern were motivated in Subsection 2.2.1 of the Technical Overview. We restate them here for self-containment.

The User–Service channel $\mathcal{F}_{\text{US}}$. $\mathcal{F}_{\text{US}}$ is parameterized by the corruption status of A_K . When A_K is uncorrupted, $\mathcal{F}_{\text{US}}$ behaves as a standard adversarial network: every wire byte is leaked to the adversary, and the adversary may reorder or drop messages. When A_K is corrupted, $\mathcal{F}_{\text{US}}$ behaves as a UC secure channel and leaks only message lengths.

Admissible corruption pattern. A static corruption choice (FC, HCC) is *admissible* if $FC \subseteq \{U\}$, $HCC \subseteq \{A_K, A_T\}$, $FC \cap HCC = \emptyset$, $|HCC| \leq 1$, and $FC \neq \emptyset \implies HCC = \emptyset$. Services are honest by default; the Verifier is omitted from FC because proofs are treated as adversarial-visible by default and a corrupted $\mathcal{V}$ adds nothing beyond what the environment already has. The pattern reduces to four cases:

- **Case A (passive):** $FC = HCC = \emptyset$.
- **Case B (malicious U):** $FC = \{U\}$, $HCC = \emptyset$.
- **Case C (HCC A_K):** $FC = \emptyset$, $HCC = \{A_K\}$.
- **Case D (HCC A_T):** $FC = \emptyset$, $HCC = \{A_T\}$.

Mapping to D-AEAD’s roles. Phases B and C of $\mathcal{RP}$ instantiate D-AEAD’s sub-protocols Π_{AEnc} and Π_{ADec} with the role correspondence

$$U \longleftrightarrow \mathcal{P}_D, \quad A_K \longleftrightarrow \mathcal{P}_K, \quad A_T \longleftrightarrow \mathcal{P}_T,$$

and the secret-index sets SR_{req} , SR_{resp} on the request and response side. Under the admissible pattern, at most one of $\{U, A_K, A_T\}$ is corrupted in any single execution, which matches D-AEAD’s single-party threat model: $U \in FC$ corresponds to D-AEAD’s malicious- $\mathcal{P}_D$ regime, while $A_K \in HCC$ or $A_T \in HCC$ corresponds to its semi-honest $\mathcal{P}_K$ or $\mathcal{P}_T$ regime.

Conventions. We treat KDF as ideal: the tuple $(k_{\text{enc}}, k_{\text{dec}}, IV_{\text{enc}}, IV_{\text{dec}})$ derived from the uniform $\mathcal{F}_{\text{KE}}$ key is uniform with independent components, so the two directions of a session are keyed independently. A pseudorandom KDF adds one standard term to every bound below. By the session-state convention of Section 5.3, every honest party invokes $(\text{GenEncMat}, \text{GenTagMat})$ at most once per direction per session, so each nonce $IV_{\text{enc}}, IV_{\text{dec}}$ is used once.

Throughout, q denotes the number of Prove sessions $\mathcal{Z}$ instantiates, and every D-AEAD advantage is that of an adversary running in time comparable to $\mathcal{Z}$ and $\mathcal{A}$ together.

D.2 Simulator

We construct a simulator Sim such that $\text{REAL}_{\mathcal{RP}, \mathcal{A}, \mathcal{Z}} \approx_c \text{IDEAL}_{\mathcal{F}_{\text{zkPoPE}}, \text{Sim}, \mathcal{Z}}$ for every admissible corruption pattern. Sim runs the four hybrid functionalities of Appendix B.3 internally and, for every uncorrupted party of $\mathcal{RP}$, an in-head copy of that party; a corrupted U is $\mathcal{A}$ ’s, and a semi-honest attestor is run by Sim with its view forwarded to $\mathcal{A}$. Sim draws every value it samples from its random tape ρ , one slice per session, so that a session’s values are available wherever Sim needs them.

Setup. On $\mathcal{F}_{\text{zkPoPE}}.\text{Setup}$, Sim returns the empty string as crs and $\text{Algs} = (\text{Setup}, \text{Prove}, \text{Verify})$, with Setup and Prove the corresponding protocols of Section 5.3 and $\text{Verify}(\text{crs}, \pi)$ the algorithm that parses π and checks σ_K and σ_T , over the components of π as in $\mathcal{RP}$ ’s Verify , under the Vfy algorithms Sim commits to $\mathcal{F}_{\text{CERT}}[A_K]$

and $\mathcal{F}_{\text{CERT}}[A_T]$. At the first activation of each of these two instances Sim supplies a correct $(\text{Sign}, \text{Vfy})$ pair; every π it later returns to $\mathcal{F}_{\text{zkPoPE}}$ carries signatures produced by them, so the check at step 9 of **Prove** passes. Sim runs $\mathcal{F}_{\text{CERT}}[S_{\text{id}}]$ on behalf of the honest service, and configures $\mathcal{F}_{\text{US}}$ to adversarial-network mode when $A_K \notin HCC$ and to secure-channel mode when $A_K \in HCC$.

What Sim knows about a session. For a supported predicate Φ write $(\text{SR}_{\text{req}}, \text{SR}_{\text{resp}})$ for its position sets, $\ell_{\text{Req}}, \ell_{\text{Resp}}$ for the lengths its format fixes, and $\text{Req}^* \in \{0, 1\}^{\ell_{\text{Req}}}$, $\text{Resp}^* \in \{0, 1\}^{\ell_{\text{Resp}}}$ for the strings that carry, outside SR_{req} and SR_{resp} , the byte values every pair satisfying $\hat{\Phi}$ shares (Section 5.2), and zero on the two sets. All of these are functions of $(\Phi, \text{Prot}_{\text{id}})$. In the honest-User cases Sim learns $(P, \Phi, \text{id}, \text{res}, \ell)$ from $\mathcal{F}_{\text{zkPoPE}}$'s announcement, with ℓ the length of the reply, and nothing else about the session. When $\text{res} = 1$ the disclosed bytes of the real request and response are those of $(\text{Req}^*, \text{Resp}^*)$ and $\ell = \ell_{\text{Resp}}$. When $\text{res} = 0$ the real request still agrees with Req^* outside SR_{req} , by the check in step 2 of $\mathcal{RP}$'s **Prove**, while the response is unconstrained; we then let $\text{Resp}^* = 0^\ell$, and the session ends before anything derived from the response other than its encryption appears in any view. In Case B, Sim reads the request from $\mathcal{A}$'s traffic and receives the response from $\mathcal{F}_{\text{zkPoPE}}$.

Delayed outputs. Sim releases each pending output of $\mathcal{F}_{\text{zkPoPE}}$ at the moment the corresponding party of $\mathcal{RP}$ produces it: $(\text{Executed}, \text{out}_{\text{Service}})$ to the service when the in-head S_{id} accepts the request record; $(\text{Executed}, \text{reply})$ and, on $\text{res} = 0$, $(\text{Prove}, \perp)$ to the caller when the in-head U completes Phase C; the proof when Phase D completes. A pending output to a party for which no session was announced is the $\perp$ of an input no reply can satisfy, and is released at once. A session that $\mathcal{A}$ stalls leaves its outputs undelivered, as in **REAL**. When the in-head A_T rejects the response record of an honest-User session, the honest U of $\mathcal{RP}$ outputs $\perp$; Sim reproduces this by delivering $\mathcal{F}_{\text{zkPoPE}}$'s $\perp$ if $\text{res} = 0$, and by answering the proof request with $\perp$, which fails **Verify**, if $\text{res} = 1$.

Honest-User sessions (Cases A, C, D). $\mathcal{F}_{\text{zkPoPE}}$'s announcement $(\text{sid}, \text{Prove}, P, \Phi, \text{id}, \text{res}, \ell)$ is what tells Sim that a session of P for (Φ, id) has begun. Sim then runs $\mathcal{RP}$ with its in-head parties, letting the in-head U run on the request Req^* and the in-head S_{id} answer with the response Resp^* instead of evaluating Prot_{id} , while $\mathcal{A}$ acts on $\mathcal{F}_{\text{US}}$ and receives the view of the semi-honest attestor, if any. Concretely, with $\mathcal{A}$ delivering wire messages as it chooses:

- *Phase A.* The in-head A_K samples r , forms sid_{KE} and $\text{ssid} = (\text{sid}, P, \text{id}, r)$, and the handshake runs as specified over $\mathcal{F}_{\text{SC}}$, $\mathcal{F}_{\text{US}}$, $\mathcal{F}_{\text{CERT}}[S_{\text{id}}]$ and $\mathcal{F}_{\text{KE}}$; Sim 's $\mathcal{F}_{\text{KE}}$ hands both endpoints a κ drawn from ρ , from which they derive $(k_{\text{enc}}, k_{\text{dec}}, IV_{\text{enc}}, IV_{\text{dec}})$.
- *Phase B.* The in-head U samples M_{req} and sends $\widetilde{\text{Req}} = \text{Req}^* \oplus \text{embed}_{\text{SR}_{\text{req}}, \ell_{\text{Req}}}(M_{\text{req}})$ and SR_{req} to A_K , and $\text{embed}_{\text{SR}_{\text{req}}, \ell_{\text{Req}}}(M_{\text{req}})$ to A_T ; the in-head attestors run Π_{AEnc} honestly, A_K forwarding SR_{req} , and U forwards $(\text{Req}_{\text{enc}}, T_{\text{req}})$ over $\mathcal{F}_{\text{US}}$.
- *Phase C.* If the delivered request record is the one the in-head A_T produced, the in-head S_{id} encrypts Resp^* under $(k_{\text{dec}}, IV_{\text{dec}})$ and sends the record over $\mathcal{F}_{\text{US}}$; otherwise it rejects. The in-head U and attestors run Π_{ADec} on the record $\mathcal{A}$ delivers, A_T accepting it only if it is the one S_{id} sent. On completion the in-head U halts if $\text{res} = 0$.
- *Phase D.* If $\text{res} = 1$, the in-head parties run Phase D as specified, and U assembles

$$\pi = (\sigma_K, \sigma_T, \text{id}, \text{ssid}, \text{SR}_{\text{req}}, \text{SR}_{\text{resp}}, \text{Req}_{\text{null}}, \widetilde{K}_{\text{resp}}, \widetilde{\text{Resp}}_{\text{enc}}).$$

Sim answers the proof request with this π . In Case C it forwards to $\mathcal{A}$ the in-head A_K 's view: its handshake state, the received $\widetilde{\text{Req}}$ and SR_{req} and the material it computes in Phases B and C, and its Phase-D inputs and outputs. In Case D it forwards the in-head A_T 's view: the mask received from U , the messages received from A_K , the records and verdicts it computes, and its Phase-D outputs. The simulation differs from a real execution in two respects: the in-head S_{id} and A_T accept only records their honest counterpart produced, which Sim , running both, can check; and the plaintexts $(\text{Req}, \text{Resp})$ have been replaced by $(\text{Req}^*, \text{Resp}^*)$, which agree with them outside $(\text{SR}_{\text{req}}, \text{SR}_{\text{resp}})$ whenever $\text{res} = 1$.

Malicious-User sessions (Case B). $\mathcal{A}$ drives $\mathcal{U}$'s side of every channel, and Sim 's in-head A_K , A_T and S_{id} respond as $\mathcal{RP}$ prescribes, with one addition: the in-head S_{id} accepts a request record only if the in-head A_T produced it in that session, and the in-head A_T a response record only if the in-head S_{id} did, which Sim , running both, can check. $\mathcal{F}_{\text{CERT}}[S_{\text{id}}]$ accepts only pairs in its sign table, so a handshake context that $\mathcal{A}$ alters is rejected. The in-head S_{id} cannot evaluate Prot_{id} , holding neither sec_{id} nor the service's coins, so Sim obtains the response from $\mathcal{F}_{\text{zkPoPE}}$. It does so at the moment the in-head S_{id} accepts a request record $(\text{Req}_{\text{enc}}, T_{\text{req}})$ delivered over $\mathcal{F}_{\text{US}}$, the moment at which the real S_{id} would run Prot_{id} , in two stages.

Stage 1 (request). Sim holds $\mathcal{A}$'s two Phase-B messages, $(\widetilde{\text{Req}}^{\mathcal{A}}, \text{SR}_{\text{req}}^{\mathcal{A}})$ to A_K and a mask $Y^{\mathcal{A}}$ to A_T that the in-head A_T has checked to be zero outside $\text{SR}_{\text{req}}^{\mathcal{A}}$, as forwarded by A_K , and sets $\text{Req}' \leftarrow \widetilde{\text{Req}}^{\mathcal{A}} \oplus Y^{\mathcal{A}}$. This is the plaintext the in-head A_T encrypted, hence, when the accepted record is the one A_T produced, the request the real S_{id} would decrypt.

Stage 2 (input). We take $\mathcal{A}$ to be the dummy adversary $\mathcal{D}$, without loss of generality [Can01], so every message of the corrupted $\mathcal{U}$ is an instruction of $\mathcal{Z}$. Let aux be the sequence of instructions $\mathcal{Z}$ has given $\mathcal{D}$ so far, all of which Sim has received; then $\text{Req}' = \mathcal{D}(1^\lambda, \text{aux})$. Sim runs the extractor $E_{\mathcal{D}}$ of Definition B.2, obtains $(\text{data}', R'_{\text{U}}) \leftarrow E_{\mathcal{D}}(1^\lambda, \text{aux}, \varepsilon, \text{Req}')$, with ε the empty tape since $\mathcal{D}$ uses no coins, and submits $(\text{id}, \text{Prove}, \text{id}, \perp, \text{data}', R'_{\text{U}})$ to $\mathcal{F}_{\text{zkPoPE}}$ on behalf of $\mathcal{U}$. $\mathcal{F}_{\text{zkPoPE}}$ runs Prot_{id} in head, delivers $\text{out}_{\text{Service}}$ to the service, which Sim releases at once, delivers the reply Resp' to Sim as the corrupted caller's output, and asks Sim for the predicate. Sim leaves that request pending, encrypts Resp' under $(k_{\text{dec}}, IV_{\text{dec}})$ as the in-head S_{id} 's response record, and continues Phases C and D against $\mathcal{A}$, its in-head attestors checking and signing whatever $\mathcal{A}$ sends, with the set $\text{SR}_{\text{req}}^{\mathcal{A}}$ from Phase B and the $\text{SR}_{\text{resp}}^{\mathcal{A}}$ that $\mathcal{A}$ supplies in Phase D.

Answering. When Phase D completes, Sim answers the pending predicate request. If $\text{Resp}' \in \mathcal{E}_{\text{Prot}_{\text{id}}}$ it returns the predicate that is identically 0. Otherwise it returns $\Phi^{\mathcal{A}}$, the supported predicate with position sets $(\text{SR}_{\text{req}}^{\mathcal{A}}, \text{SR}_{\text{resp}}^{\mathcal{A}})$ whose condition asserts, at every position outside them, the byte that Req' or Resp' carries there. $\mathcal{F}_{\text{zkPoPE}}$ then computes res , announces it and, on $\text{res} = 1$, requests a proof, which Sim answers with the π formed from the signatures its in-head attestors issued in Phase D and the values they signed. By construction $\widehat{\Phi}^{\mathcal{A}}(\text{Req}', \text{Resp}') = 1$; and when extraction succeeds, Lemma 1 gives that $\mathcal{F}_{\text{zkPoPE}}$'s in-head run put Req' on the wire, so its reply is Resp' and $\text{res} = \Phi^{\mathcal{A}}(\text{data}', \text{Resp}') = \widehat{\Phi}^{\mathcal{A}}(\text{Req}', \text{Resp}') = 1$, and $\mathcal{F}_{\text{zkPoPE}}$ records $(\mathcal{U}, \Phi^{\mathcal{A}}, \text{id}, \pi)$. If $\mathcal{A}$ never completes Phase D, the predicate request stays pending and nothing is recorded, as nothing is signed in REAL. If $\mathcal{A}$ never delivers an acceptable request record, $\mathcal{F}_{\text{zkPoPE}}$ is never invoked.

Two features of Definition B.2 make this construction possible inside a UC simulation: $E_{\mathcal{D}}$ is straight-line, so Sim never rewinds an execution whose schedule it does not control; and its auxiliary input is quantified universally, so it accepts the instruction stream $\mathcal{Z}$ produced, whatever it is. We write E_{ext} for the event that in some session $\text{Req}' \in \mathcal{R}_{\text{Prot}_{\text{id}}}$ but $\text{Prot}_{\text{id}}.\text{U}_1(\text{data}', R'_{\text{U}}) \neq \text{Req}'$. For each session the auxiliary input is fixed when the extractor runs, so Definition B.2 bounds the failure probability by $\varepsilon_{\text{Prot}_{\text{id}}, \mathcal{D}}^{\text{ext}}(\lambda)$, and over q sessions

$$\Pr[E_{\text{ext}}] \leq q \cdot \varepsilon_{\text{Prot}_{\text{id}}, \mathcal{D}}^{\text{ext}}(\lambda).$$

D.3 Indistinguishability: Case B (soundness)

In Case B, $\mathcal{A}$ controls $\mathcal{U}$ while $A_K, A_T, S_{\text{id}}, \mathcal{V}$ are honest, $\mathcal{F}_{\text{CERT}}[A_K]$ and $\mathcal{F}_{\text{CERT}}[A_T]$ run in their table-based honest-owner branch, and $\mathcal{F}_{\text{US}}$ is in adversarial-network mode. We show

$$\left| \Pr[\mathcal{Z} \rightarrow 1 \mid \text{REAL}_{\mathcal{RP}, \mathcal{A}, \mathcal{Z}}] - \Pr[\mathcal{Z} \rightarrow 1 \mid \text{IDEAL}_{\mathcal{F}_{\text{zkPoPE}}, \text{Sim}, \mathcal{Z}}] \right| \leq 2q \cdot \text{Adv}_{\text{D-AEAD}}^{\text{INT-CTXT}}(\lambda) + q^2 \cdot 2^{-\lambda-1} + q \cdot \varepsilon_{\text{Prot}_{\text{id}}, \mathcal{D}}^{\text{ext}}(\lambda), \quad (1)$$

through games $G_0 = \text{REAL}$, G_1 , G_2 , $G_3 = \text{IDEAL}$.

Hop $G_0 \rightarrow G_1$: wire integrity. G_1 is G_0 except that S_{id} rejects any request record $(\text{Req}_{\text{enc}}, T_{\text{req}})$ that A_T did not produce in that session, and A_T rejects any response record $(\text{Resp}_{\text{enc}}, T_{\text{resp}})$ that S_{id} did not produce in that session. The two games differ only on the event E_{wire} that G_0 's honest parties accept such a record in some session.

We bound $\Pr[E_{\text{wire}}]$ by an adversary $\mathcal{A}'_{\text{wire}}$ against $\text{Exp}_{\text{D-AEAD}}^{\text{INT-CTXT}}$. It guesses a session $s \in [q]$ and a direction, runs G_0 for $\mathcal{Z}$ and $\mathcal{A}$, simulating every honest party and functionality with keys it samples itself, except that on the guessed direction of session s it uses the game's key: the honest encryption of that direction — the A_K – A_T run of Π_{AEnc} on the request side, S_{id} 's AEAD encryption on the response side — is replaced by one $\mathcal{O}_{\text{Enc}}$ query, and the honest verification — S_{id} 's on the request side, A_T 's in Π_{ADec} on the response side — by $\mathcal{O}_{\text{Dec}}$ queries. Since each nonce is used once, the nonce-respecting $\mathcal{O}_{\text{Enc}}$ answers every redirected call. In Case B, $\mathcal{A}$ receives from the honest parties only what $\mathcal{P}_D$ receives in Π_{AEnc} and Π_{ADec} — the record (C, T) , the verdict, and the keystream K_{resp} released after a successful verdict, which $\mathcal{A}'_{\text{wire}}$ computes as $\text{Resp}_{\text{enc}} \oplus \text{Resp}$ since it knows the plaintext it encrypted — so the simulation is faithful, and a forgery in the guessed session and direction is an $\mathcal{O}_{\text{Dec}}$ query accepted on a tuple outside $\mathcal{Q}$. Hence

$$\Pr[E_{\text{wire}}] \leq 2q \cdot \text{Adv}_{\text{D-AEAD}, \mathcal{A}'_{\text{wire}}}^{\text{INT-CTXT}}(\lambda).$$

Hop $G_1 \rightarrow G_2$: record-based verification. For each session s that completes Phase D write π^s for the proof its signatures determine, $(\text{Req}^s, \text{Resp}^s)$ for the request S_{id} decrypted and the response it returned, and Φ^s for the supported predicate with s 's position sets asserting the disclosed bytes of $(\text{Req}^s, \text{Resp}^s)$, when $\text{Resp}^s \notin \mathcal{E}_{\text{Protid}}$. G_2 is G_1 except that $\mathcal{V}$ answers a query $(\text{id}, P_i, \Phi, \pi^*)$ by the rule: *accept iff $P_i = \text{U}$ and, for some session s with $\text{Resp}^s \notin \mathcal{E}_{\text{Protid}}$, $\pi^* = \pi^s$, $\text{id} = \text{id}^s$ and $\Phi = \Phi^s$.*

The two games differ only on the event E_{forge} that G_1 's Verify and the rule disagree on some query

$$\pi^* = (\sigma_K^*, \sigma_T^*, \text{id}^*, \text{ssid}^*, \text{SR}_{\text{req}}^*, \text{SR}_{\text{resp}}^*, \text{Req}_{\text{null}}^*, \widetilde{K}_{\text{resp}}^*, \widetilde{\text{Resp}}_{\text{enc}}^*).$$

Both $\mathcal{F}_{\text{CERT}}$ instances accept only pairs in their sign tables, and each honest attestor signs once per session, on a message containing that session's ssid. So σ_K^* pins $(P_i, \text{id}^*, \text{ssid}^*, \text{SR}_{\text{req}}^*, \text{SR}_{\text{resp}}^*, \text{Req}_{\text{null}}^*, \widetilde{K}_{\text{resp}}^*)$ to the message A_K signed in a unique session s_K , and σ_T^* pins $(P_i, \text{id}^*, \text{ssid}^*, \text{SR}_{\text{req}}^*, \text{SR}_{\text{resp}}^*, \widetilde{\text{Resp}}_{\text{enc}}^*)$ to the message A_T signed in a unique session s_T , with $\text{ssid}^{s_K} = \text{ssid}^* = \text{ssid}^{s_T}$. Each ssid carries a nonce $r \xleftarrow{\$} \{0, 1\}^\lambda$ that A_K samples fresh per session, so $s_K \neq s_T$ requires a collision among q nonces, of probability at most $q^2 \cdot 2^{-\lambda-1}$. When $s_K = s_T = s$, every component of π^* equals that of π^s , the signed messages carry $P_i = \text{U}$, and $\mathcal{V}$'s parse of ssid^* enforces $\text{id} = \text{id}^s$. It remains to compare the two verdicts on π^s . In G_1 , $\mathcal{V}$ checks that Φ 's position sets are $(\text{SR}_{\text{req}}^s, \text{SR}_{\text{resp}}^s)$ and evaluates Φ on $(\text{Req}_{\text{null}}^s, \widetilde{\text{Resp}}_{\text{enc}}^s \oplus \widetilde{K}_{\text{resp}}^s)$. In G_2 the record S_{id} decrypted is the one A_T produced, so $\text{Req}^s = \widetilde{\text{Req}}^s \oplus Y^s$ with Y^s zero outside SR_{req}^s by A_T 's check, and $\text{Req}_{\text{null}}^s = \text{zeros}_{\text{SR}_{\text{req}}^s}(\widetilde{\text{Req}}^s) = \text{zeros}_{\text{SR}_{\text{req}}^s}(\text{Req}^s)$; likewise the record A_T accepted is S_{id} 's, so $\widetilde{\text{Resp}}_{\text{enc}}^s \oplus \widetilde{K}_{\text{resp}}^s = \text{zeros}_{\text{SR}_{\text{resp}}^s}(\text{Resp}^s)$. $\mathcal{V}$ therefore evaluates $\widehat{\Phi}(\text{Req}^s, \text{Resp}^s)$. Since Φ is supported, this is 1 iff Φ 's sets are s 's and its condition fixes, at every disclosed position, the byte $(\text{Req}^s, \text{Resp}^s)$ carries there — that is, iff $\Phi = \Phi^s$ — and rejection soundness makes it 0 whenever $\text{Resp}^s \in \mathcal{E}_{\text{Protid}}$. This is the rule of G_2 . Hence

$$\Pr[E_{\text{forge}}] \leq q^2 \cdot 2^{-\lambda-1}.$$

Hop $G_2 \rightarrow G_3$. In IDEAL, $\mathcal{V}$ accepts $(\text{id}, P_i, \Phi, \pi^*)$ iff $(P_i, \Phi, \text{id}, \pi^*) \in \Pi_{\text{sid}}$, and Sim's in-head S_{id} and A_T apply the acceptance rule of G_1 . By the construction of Sim, Π_{sid} receives an entry for session s exactly when $\mathcal{A}$ completes Phase D and $\text{Resp}^s \notin \mathcal{E}_{\text{Protid}}$, and that entry is $(\text{U}, \Phi^s, \text{id}, \pi^s)$, provided the reply $\mathcal{F}_{\text{zkPoPE}}$ returned is the response Sim put on the wire; conditioned on $\overline{E_{\text{ext}}}$ it is, by Lemma 1. So the record of G_3 is the record of G_2 . The remaining views agree as well: the in-head parties follow $\mathcal{RP}$; the response Sim encrypts is $\text{Protid}.\text{Service}(\text{sec}_{\text{id}}, \text{Req}'; \cdot)$ on fresh coins, exactly the distribution of the real S_{id} 's response on the same request, independently across sessions; and the service's output $\text{outs}_{\text{service}}$ is released when the real S_{id} would produce it. Hence $G_2 \equiv G_3$ conditioned on $\overline{E_{\text{ext}}}$, and the bound follows by summing the three events.

D.4 Indistinguishability: Cases A, C, D (zero-knowledge)

In Cases A, C and D the User is honest, and Sim's simulation differs from REAL in the acceptance rule for wire records and in the plaintexts: the in-head U and S_{id} run on $(\text{Req}^*, \text{Resp}^*)$ in place of $(\text{Req}, \text{Resp})$.

The requests agree outside SR_{req} in every session, by step 2 of **Prove**. When $\text{res} = 1$ the responses agree outside SR_{resp} ; when $\text{res} = 0$ they may differ anywhere, but Phase D does not run, so nothing derived from the response other than its encryption appears in any view. In every game below the honest U's output (**Executed**, Resp) to $\mathcal{Z}$ is the real response, as $\mathcal{F}_{\text{zkPoPE}}$ delivers it in **IDEAL**. Everything else — the honest parties' code, the functionalities, the delivery of outputs — is identical, and the π that $\mathcal{F}_{\text{zkPoPE}}$ stores is the one the in-head U assembles. We show that $\mathcal{Z}$'s distinguishing advantage is at most

$$2q \cdot \text{Adv}_{\text{D-AEAD}}^{\text{INT-CTXT}}(\lambda) + 2q \cdot \text{Adv}_{\text{D-AEAD}}^{\text{IND-CCA-}\mathcal{P}_T}(\lambda)$$

in Cases A and D, and at most $q \cdot \text{Adv}_{\text{D-AEAD}}^{\text{IND-CPA-}\mathcal{P}_K}(\lambda)$ in Case C. By Theorem C.6 the first two terms are negligible under jPR and OTUF, and the third is zero.

Cases A and D. In Case A, $\mathcal{A}$ sees the $\mathcal{F}_{\text{US}}$ records $(\text{Req}_{\text{enc}}, T_{\text{req}})$ and $(\text{Resp}_{\text{enc}}, T_{\text{resp}})$, the $\mathcal{F}_{\text{SC}}$ length leaks, and the proofs output to $\mathcal{Z}$. In Case D it additionally sees $\mathcal{A}_T$'s view: the mask $\text{embed}_{\text{SR}_{\text{req}}, \ell_{\text{Req}}}(M_{\text{req}})$ from U, $(\widetilde{\text{Req}}_{\text{enc}}, \tau_{\text{req}}, \text{SR}_{\text{req}})$ from $\mathcal{A}_K$, the record $(\text{Req}_{\text{enc}}, T_{\text{req}})$ it computes, then $(\text{Resp}_{\text{enc}}, T_{\text{resp}})$, τ_{resp} and its verdict, and in Phase D $(\text{SR}_{\text{req}}, \text{SR}_{\text{resp}})$, $\widetilde{\text{Resp}}_{\text{enc}}$ and σ_T . Under $\mathcal{A}_T \leftrightarrow \mathcal{P}_T$ this is $\mathcal{P}_T$'s view in the session's Π_{AEnc} and Π_{ADec} runs, plus values that are functions of it and of public data; the Case-A view is a sub-view of it.

Let G_1 be **REAL** with the acceptance rule of the simulation: S_{id} and $\mathcal{A}_T$ reject any record their counterpart did not produce in that session. $\mathcal{F}_{\text{US}}$ is in adversarial-network mode, so $\mathcal{A}$ may deliver altered records, and **REAL** and G_1 differ only when one is accepted; the adversary $\mathcal{A}'_{\text{wire}}$ of Section D.3, now also simulating the honest U, bounds this by $2q \cdot \text{Adv}_{\text{D-AEAD}}^{\text{INT-CTXT}}(\lambda)$. From G_1 we hybridize over the $2q$ pairs (session, direction). Let $H_0 = G_1$, and let H_j run the first j pairs, in a fixed order, on the starred plaintext and the rest on the real one; $H_{2q} = \text{IDEAL}$. An adversary $\mathcal{A}'_{\text{IND-CCA-}\mathcal{P}_T}$ against **IND-CCA-}\mathcal{P}_T (Definition C.2) picks $j \xleftarrow{\$} [2q]$, runs H_{j-1} for $\mathcal{Z}$ and $\mathcal{A}$ with locally sampled keys everywhere except on pair j , and there queries**

$$\mathcal{O}_{\text{Chal-}b}(n, AD, \text{SR}, P_0, P_1)$$

with n the direction's nonce, P_0 the real plaintext and P_1 the starred one, $\text{SR} = \text{SR}_{\text{req}}$ on the request side, and on the response side $\text{SR} = \text{SR}_{\text{resp}}$ when $\text{res} = 1$ and $\text{SR} = [\ell_{\text{Resp}}]$ when $\text{res} = 0$; in every case P_0, P_1 have length ℓ and agree outside SR . The oracle returns $\mathcal{P}_T$'s view $(M, \widetilde{C}, \tau, C, T)$ of Π_{AEnc} on P_b . On the request side this is $(\text{embed}_{\text{SR}_{\text{req}}, \ell_{\text{Req}}}(M_{\text{req}}), \widetilde{\text{Req}}_{\text{enc}}, \tau_{\text{req}}, \text{Req}_{\text{enc}}, T_{\text{req}})$; on the response side (C, T) is S_{id} 's record and $\tau = \tau_{\text{resp}}$ is what $\mathcal{A}_T$ receives in Π_{ADec} , the other components being discarded. $\mathcal{A}'_{\text{IND-CCA-}\mathcal{P}_T}$ forwards to $\mathcal{A}$ the part of this view the case prescribes, and completes the session from it: the Phase-D values $\text{Req}_{\text{null}} = \text{zeros}_{\text{SR}_{\text{req}}}(P_b)$ and $\widetilde{\text{Resp}}_{\text{enc}} \oplus \widetilde{K}_{\text{resp}} = \text{zeros}_{\text{SR}_{\text{resp}}}(P_b)$ are the same for $b = 0, 1$, and $\widetilde{K}_{\text{resp}} = \text{zeros}_{\text{SR}_{\text{resp}}}(C \oplus P_b)$ is computable from the returned C and either plaintext. When $b = 0$ the view is that of H_{j-1} , when $b = 1$ that of H_j , and $\mathcal{A}'_{\text{IND-CCA-}\mathcal{P}_T}$ outputs $\mathcal{Z}$'s bit. Summing over j ,

$$|\Pr[\mathcal{Z} \rightarrow 1 \mid G_1] - \Pr[\mathcal{Z} \rightarrow 1 \mid \text{IDEAL}]| \leq 2q \cdot \text{Adv}_{\text{D-AEAD}, \mathcal{A}'_{\text{IND-CCA-}\mathcal{P}_T}}^{\text{IND-CCA-}\mathcal{P}_T}(\lambda).$$

Case C. $\mathcal{F}_{\text{US}}$ is in secure-channel mode, so $\mathcal{A}$ sees no wire bytes and delivers every record unaltered; the acceptance rule of the simulation changes nothing. Its view consists of the $\mathcal{F}_{\text{SC}}$ and $\mathcal{F}_{\text{US}}$ length leaks, the proofs output to $\mathcal{Z}$, and $\mathcal{A}_K$'s view: the handshake state $(r, \kappa, k_{\text{enc}}, k_{\text{dec}}, IV_{\text{enc}}, IV_{\text{dec}})$; $\widetilde{\text{Req}}$, SR_{req} and the material $K_{\text{req}}, \tau_{\text{req}}, \widetilde{\text{Req}}_{\text{enc}}$ it computes in Phase B; $|\text{Resp}_{\text{enc}}|$, τ_{resp} , the verdict and K_{resp} in Phase C; and $(\text{SR}_{\text{req}}, \text{SR}_{\text{resp}})$, $\text{Req}_{\text{null}}, \widetilde{K}_{\text{resp}}, \sigma_K$ in Phase D. Every component other than $\widetilde{\text{Req}}$ is a function of κ , of public data, and of $\widetilde{\text{Req}}$; the proof π is a function of $\widetilde{\text{Req}}$, κ , public data and $\widetilde{\text{Resp}}_{\text{enc}} = \text{zeros}_{\text{SR}_{\text{resp}}}(\text{Resp}_b \oplus K_{\text{resp}})$, which for $\text{res} = 1$ is the same for $\text{Resp}_0 = \text{Resp}$ and $\text{Resp}_1 = \text{Resp}^*$. So the two executions differ only in $\widetilde{\text{Req}} = \text{Req}_b \oplus \text{embed}_{\text{SR}_{\text{req}}, \ell_{\text{Req}}}(M_{\text{req}})$ with $\text{Req}_0 = \text{Req}$ and $\text{Req}_1 = \text{Req}^*$, per session. This is the **IND-CPA-}\mathcal{P}_K** experiment (Definition C.1) with $\text{SR} = \text{SR}_{\text{req}}$, and a hybrid over the q sessions gives the bound $q \cdot \text{Adv}_{\text{D-AEAD}}^{\text{IND-CPA-}\mathcal{P}_K}(\lambda)$, which is 0 by Theorem C.6: $\widetilde{\text{Req}}$ agrees with Req_b outside SR_{req} , where Req_0 and Req_1 coincide, and is uniform on SR_{req} under the one-time pad M_{req} , whichever b .

D.5 Wrap-up

For each admissible corruption pattern the distinguishing advantage of $\mathcal{Z}$ is bounded by the corresponding expression above: $2q \cdot \text{Adv}_{\text{D-AEAD}}^{\text{INT-CTXT}} + q^2 \cdot 2^{-\lambda-1} + q \cdot \varepsilon_{\text{Prot}_{\text{id}}, \mathcal{D}}^{\text{ext}}$ in Case B, $2q \cdot \text{Adv}_{\text{D-AEAD}}^{\text{INT-CTXT}} + 2q \cdot \text{Adv}_{\text{D-AEAD}}^{\text{IND-CCA-}\mathcal{P}_T}$ in Cases A and D, and 0 in Case C. By Theorem C.6, $\text{Adv}_{\text{D-AEAD}}^{\text{INT-CTXT}}$ and $\text{Adv}_{\text{D-AEAD}}^{\text{IND-CCA-}\mathcal{P}_T}$ are negligible under jPR and OTUF; $\varepsilon_{\text{Prot}_{\text{id}}, \mathcal{D}}^{\text{ext}}$ is negligible because Prot_{id} is request-extractable. Extraction is invoked only against a corrupted User. This establishes Theorem 5.1. $\square$