

A Formal Treatment of Accountable Proxying over TLS

Karthikeyan Bhargavan¹, Ioana Boureanu², Antoine Delignat-Lavaud³, Pierre-Alain Fouque⁴, and Cristina Onete⁵

¹ Inria de Paris, ² University of Surrey, SCCS ³ Microsoft Research, ⁴ Université de Rennes 1, IRISA,

⁵ Université de Limoges, XLIM, CNRS 7252

Email: karthikeyan.bhargavan@inria.fr; i.boureanu@surrey.ac.uk; antld@microsoft.com,
pa.fouque@gmail.com, cristina.onete@gmail.com

Abstract—Much of Internet traffic nowadays passes through active proxies, whose role is to inspect, filter, cache, or transform data exchanged between two endpoints. To perform their tasks, such proxies modify channel-securing protocols, like TLS, resulting in serious vulnerabilities. Such problems are exacerbated by the fact that middleboxes are often invisible to one or both endpoints, leading to a lack of *accountability*. A recent protocol, called mcTLS, pioneered accountability for proxies, which are authorized by the endpoints and given limited read/write permissions to application traffic.

Unfortunately, we show that mcTLS is insecure: the protocol modifies the TLS protocol, exposing it to a new class of *middlebox-confusion attacks*. Such attacks went unnoticed mainly because mcTLS lacked a formal analysis and security proofs. Hence, our second contribution is to formalize the goal of accountable proxying over secure channels. Third, we propose a provably-secure alternative to soon-to-be-standardized mcTLS: a generic and modular protocol-design that carefully composes generic secure channel-establishment protocols, which we prove secure. Finally, we present a proof-of-concept implementation of our design, instantiated with unmodified TLS 1.3 draft 23, and evaluate its overheads.

Keywords-mcTLS, TLS 1.3, provable security

I. INTRODUCTION

Internet protocols are largely designed around the end-to-end principle, which says that all application logic, except for the mundane activity of forwarding packets, should reside at the endpoints. However, a good portion of Internet traffic today passes through one or several active proxies or *middleboxes* that inspect, filter, and transform packets, based on dynamically configurable policies. These middleboxes include content delivery networks, personal and enterprise-level firewalls, compression proxies, malware scanners, parental-control content filters, and many other in-network functionalities that are considered desirable by client institutions, web servers, and network operators.

All of these functionalities require read access and/or write access to client-server traffic, and are consequently hindered by the prevalence of end-to-end encryption protocols like Transport Layer Security (TLS). To continue working, middlebox providers have resorted to a variety of *ad-hoc* techniques that enable them to decrypt TLS traffic between clients and servers. These techniques necessarily contradict the end-to-end security goals of TLS, by turning a well-studied two-party secure channel into a non-standard

multi-party cryptographic protocol with unknown security properties. In fact, middleboxes often introduce new threats, since adversaries aiming to break secure connections can now attack not just endpoints but also the middleboxes. Our goal in this paper is to precisely identify these threats and to formally define the security goals of proxied TLS connections, in order to enable a rigorous evaluation of existing and new middlebox-based designs.

Content Delivery Networks. CDNs like Akamai and Cloudflare are good examples of widely-used active proxies. For example, Akamai owns 233,000 HTTP and HTTPS caching proxy servers in over 130 countries and within more than 1,600 networks around the world, serving 27% of the world’s network traffic on behalf of many of the top websites. In October 2015, they estimated that 45% of their traffic was TLS-encrypted traffic [29] and this number has undoubtedly significantly increased since.

In order to serve TLS-traffic on behalf of a website, CDNs like Akamai need the website-owner to allow them to hold a valid X.509 certificate for the website’s domain(s) and the associated private key on their behalf, essentially licensing the CDN to impersonate the website to connecting web-browsers. This form of delegation endows a high degree of trust in the CDN infrastructure, since a bug or attack on any of the CDN’s “edge” servers around the world may allow an attacker to steal the website’s private key or other sensitive user data (*e.g.* see the recent Cloudbleed bug in Cloudflare [35]). Websites can try to mitigate this risk by delegating only certain subdomains to the CDN, rather than the full web-applications, but this has limited effectiveness on websites where JavaScript from one subdomain is routinely loaded on another.

An alternative to the CDN-architecture above, denoted Keyless SSL, is offered as a premium service by the CDN called Cloudflare. Cloudflare only gets the X.509 certificate for a CDN-ed domain (and the public key within the cert), and Cloudflare’s customers securely store the associated private keys on their own servers, only granting a limited key-usage API to the Cloudflare’s proxy servers. This design requires a minor refactoring of the TLS handshake, but even this small change results in a new three-party protocol whose security guarantees are subtle. Indeed, recently, Bhargavan et al. [4]

demonstrated several vulnerabilities of Keyless SSL; they used a provable security approach to formally analyze it and proposed alternative designs that achieve stronger security goals, albeit with reduced performance.

Client-side firewalls and content filters. Firewalls are commonly used in enterprises and educational institutions to protect computers from malware and age-inappropriate content. To be able to inspect TLS traffic between machines inside the firewall and the open Web, these firewall proxies effectively mount a man-in-the-middle attack on the TLS connection by asking all clients to install a CA certificate that they can then use to issue certificates for any web server. This gross misuse of the public key infrastructure completely bypasses the end-to-end guarantees of TLS, and hence requires complete trust in the design and implementation of the proxy itself.

Worryingly, recent studies of such proxies uncovered a plethora of serious security issues, from inappropriate or inexistent certification validation and ciphersuite downgrades, to the execution of completely invalid TLS handshakes [15], [34]. A typical example of what can go wrong is the Superfish scandal of 2015 [18] where the private key of the CA certificate used by a client-side proxy was leaked, allowing attackers to impersonate any website to any client who used the proxy. More generally, client-side proxies that install root CA certificates are completely invisible to both the client and server, making it impossible for web browsers or websites to impose their own security policies for sensitive transactions, say between a client and a bank.

Fine-grained access control with mcTLS. A recent protocol called mcTLS [33] proposes a different design where all middleboxes are fully visible to both the client and the server, both of whom must agree on the read/write privileges of each middlebox before the middleboxes get to intercept the TLS connection. Hence, mcTLS offers fine-grained access control to both the client and server over all proxied TLS connections. To achieve this efficiently, the protocol significantly modifies the TLS handshake protocol, aggressively sharing messages and key material across multiple hops to perform several TLS handshakes at once.

Because of its excellent performance and support for fine-grained policies, mcTLS is being standardized within the ETSI standards organization as the default middlebox security protocol for network operators.¹ Consequently, the protocol is poised to have a big impact on practical middlebox designs. However, before the protocol is standardized and widely deployed, we believe it is important to formalize its security goals and verify that mcTLS meets them.

The authors of mcTLS do offer an informal security analysis [33] of the protocol, but without a formal proof. Indeed, as we show in the next section, certain intended use-

cases of mcTLS are demonstrably insecure, which motivates the need for formal analysis, and provides a stark warning against tampering with existing protocols like TLS without a formal proof. This is particularly important for TLS 1.3, the new version of TLS which was designed to be provably secure and has received considerable attention from the academic community, but whose strong guarantees can still be completely broken by badly designed proxies.

Our approach and contributions. In this paper, we adopt a *provable security* approach to proxied TLS connections:

- We show that existing proxying mechanisms, including mcTLS, fail to ensure intuitive security notions of authentication, confidentiality, and integrity, even in common proxying scenarios.
- We provide (to our knowledge) the first fully formalized security definition for proxied TLS. We call this definition *authenticated and confidential channel establishment with accountable proxies* (ACCE-AP).
- We provide a modular protocol construction that only allows proxying with the full, explicit knowledge and consent of both partners (akin to mcTLS). Unlike mcTLS, our construction is provably secure and can be instantiated with any authenticated key-exchange (AKE) protocols which respect some reasonable conditions². In particular, our design composes TLS connections in a way that allows us to rely on existing proofs of TLS, rather than prove our protocol's security from scratch.
- We describe a proof-of-concept implementation, showing how our new design can be deployed modularly on top of miTLS, a high-assurance TLS 1.3 library.

We provide a more extended review of related literature in Appendix A.

II. ACTIVE PROXY ARCHITECTURES AND THEIR FLAWS

The goal of an active proxy is to provide in-network functionality between a client and a server, while preserving (as much as possible) end-to-end confidentiality, authentication, and data integrity. Figure 1 depicts a typical scenario with one proxy. We note that the proxy may offer its services to many clients and servers, some of whom may be malicious or otherwise controlled by an adversary. Furthermore, we assume that the TLS connections between the client and middlebox, and between the middlebox and server are subject to standard network attacks.

A. Proxying Architectures

Existing proxying architectures can be broadly categorized in terms of which participants are aware of the proxying. We discuss two common setups below, before focusing on mcTLS, one of the most sophisticated and convincing proxying protocols designed to date.

²On the one hand, this means that we can use TLS 1.3 as our building-block. On the other hand, we do not use any specific property of a given version of TLS 1.3, not in the construction nor in the proofs.

¹https://portal.etsi.org/webapp/WorkProgram/Report_WorkItem.asp?WKI_ID=52930

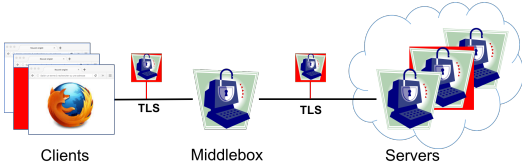


Figure 1: Active proxying

Invisible Proxies. Before encrypted connections became the norm, a variety of in-network proxies were used to provide compression, caching, load-balancing, and other bandwidth-saving or latency-reducing services. Once they were installed and configured, these proxies were *fully invisible* to both the client and server, since they could operate directly on unencrypted data. To continue working over TLS, these proxies need to be able to impersonate the client and the server to each other, a “feature” called *SSL Interception*. On the Web, clients are usually anonymous, but servers are identified by X.509 public-key certificates, and so such proxies need to have access to a valid certificate (and private key) for the server. This means that the proxy cannot be perfectly invisible to both parties. Either the proxy needs to be fully trusted by the server, who provisions it with a certificate, or it needs to be even more trusted by the client, who allows it to install a CA certificate that the proxy can use to impersonate any web server. We call the former design *client invisible* and the latter design *server invisible*.

Invisible proxies are attractive because they require little change to existing client-server deployment, but they are vulnerable to a number of attacks because the security of the TLS ecosystem relies on a careful collaboration between mechanisms implemented by both clients (web browsers) and servers (websites), and these mechanisms can be bypassed if one or both endpoints are unaware of the proxying. For example, modern browsers remember the certificates of important websites, through a technique called public-key pinning, and are able to enforce sophisticated certificate revocation policies through protocols like OCSP, but these protections no longer apply if the proxy uses its own certificate. Similarly, a browser or website may require a modern version of TLS with a strong ciphersuite, but the proxy may downgrade the connection to use a legacy version of the protocol. In addition, the high level of trust needed in the proxying software itself has been undermined by a series of flaws and attacks on middleboxes [15], [34], [18].

Accountable Proxies. The many vulnerabilities of invisible proxies stem from their adoption of ad hoc mechanisms based on misusing the public key infrastructure. An alternative is to design proxies that are not only visible, but request explicit authorization from one or both endpoints, and do not interfere with connections for which they do not have authorization. This would allow sensitive websites, like those for online banking, to forbid middleboxes and ensure that their clients’ account details are not visible to any party other than the user’s web browser. Such *accountable* designs may even

become mandatory because of privacy laws like the recently adopted Data Protection Regulation (EU regulation 2016/679), which requires that personal information pertaining to any citizen may only be sent, handled, and stored with that individual’s express consent.

Accountability allows network operators to reintroduce in-network functionality like caching and compression proxies over TLS without requiring full trust from clients and servers. For example, the Keyless SSL proposal from Cloudflare [4], and the related Lurk draft standard [30], enable server-visible caching proxies that require the proxy to contact a key server on every TLS session, allowing the key server to easily disable proxying based on a variety of server-side policies. These designs are not without their flaws, and alternative designs can offer even stronger and provably secure notions of accountability [4].

A new and radical proxy architecture, called mcTLS, offers even more fine-grained control over proxies’ behaviors, by requiring both the client and the server to negotiate and agree upon the read-write access rights given to each intermediate proxy. In the rest of this section, we describe this important new protocol, currently undergoing standardization by ETSI, and evaluate its security. In particular, we describe several *middlebox confusion attacks* against mcTLS and use them to motivate a rigorous formal analysis of active proxying.

B. The mcTLS protocol

The mcTLS protocol uses the same message formats and cryptographic constructions as TLS 1.2 (hence the name) but is a completely new multi-party protocol that modifies TLS in significant ways. The protocol features a client C , a server S , and a number of middleboxes situated between them. The server and the middleboxes are provisioned with public-key certificates. Each middlebox is given read and/or write permission to various mcTLS *contexts*, which are portions of an application-data stream, such as HTTP bodies/headers. For example, a firewall middlebox may only have read access for a particular HTTP header, whereas a caching middlebox may have write access to the full HTTP response. To enforce these access rights, the client and server generate read and write keys for each context and deliver them to each authorized middlebox over point-to-point TLS-like channels.

Figure 2 depicts the mcTLS protocol when run for a single middlebox, denoted MW . The handshake consists of three interleaved TLS-DHE (ephemeral Diffie-Hellman) handshakes that share messages and key material: the primary handshake between the client and the server, and two intermediate handshakes between the client and middlebox, and middlebox and server. The server and middlebox authenticate themselves on all connections, but the client is unauthenticated. The client and server exchange nonces N_C, N_S and Diffie-Hellman key shares g^x, g^y ; the middlebox adds its own nonce N_{MW} and separate key shares g^z, g^t for use with the client and server.

At the end of the handshake, three sets of channel keys are computed using the TLS 1.2 key-derivation mechanism:

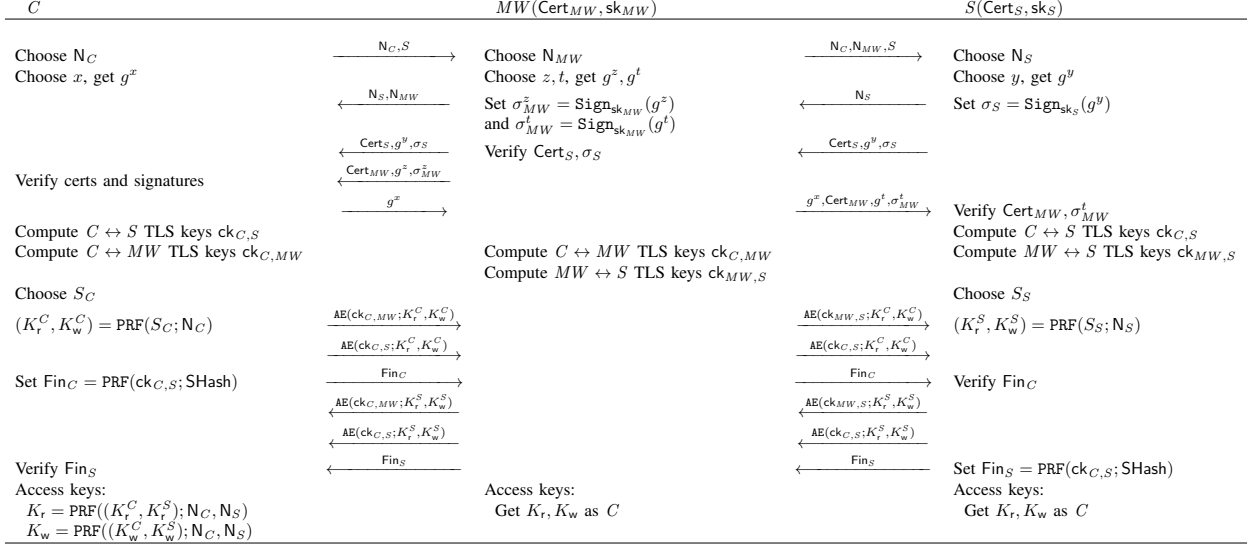


Figure 2: The mcTLS Handshake.

the client and server derive keys from g^{xy} , N_C , and N_S , the client and middlebox derive keys from g^{xz} , N_C , and N_{MW} , the middlebox and server derive keys from g^{yt} , N_{MW} , and N_S . These channel keys are then used to distribute context-specific read and write keys. The client and server generate independent context keys, say by applying a key derivation function to local temporary secrets (S_C, S_S), and deliver these keys to each other and to each authorized middlebox. After delivering the context keys, the two endpoints complete the handshake by exchanging Finished messages that contain MACs over the *entire transcript* as seen by each endpoints over the two sessions it runs, to ensure that none of the handshake messages has been tampered with.

Next, the client and server start exchanging application data; the data stream is divided into fragments, each fragment is labeled with a specific context and encrypted with the corresponding context key. Middleboxes that are authorized to read or write a particular context can then use the corresponding context keys to decrypt and/or reencrypt these packets, but not others.

The mcTLS protocol extends and deviates from the standard TLS 1.2 handshake in many ways. First, it adds a middlebox negotiation extension to the ClientHello and ServerHello messages, so that the client and server can agree upon the sequence of middleboxes they wish to use. Second, mcTLS reuses nonces and ephemeral Diffie-Hellman key shares between multiple connections, which is not recommended but is relatively harmless. Third, mcTLS only completes a full TLS handshake between the client and server; the client-middlebox and middlebox-server handshakes do not include any Finished messages, and so the middlebox cannot know whether the handshake has been tampered with. This is a more significant change to TLS, and as we shall

see below, it leads to serious attacks on mcTLS. Fourth, mcTLS modifies the record-layer protocol to use context keys instead of channel keys for encryption. This is again an important change that exposes the protocol to record-layer attacks. Finally, to reduce server overhead, the mcTLS authors [33] propose a *client key distribution mode* where the client unilaterally generates and distributes context keys to the server and all middleboxes. As we shall see, this mode also weakens the protocol and enables attacks.

C. Middlebox Confusion Attacks on mcTLS

To better understand the security guarantees of mcTLS, let us first consider some of its intended use case scenarios.

In the first scenario, a client C uses a firewall middlebox MW to protect it from malware on the web. It requires MW to filter all its connections, including those to a trusted server S , and those to an attacker-controlled website \mathcal{A} . MW has read access to all incoming data and it imposes different filtering policies for different websites; in particular, it has stricter rules for \mathcal{A} than for S . The adversary's goal is to bypass MW and deliver malware to C .

In the second scenario, a client C uses a caching proxy MW to speed up its accesses to the web. The proxy retrieves and caches static pages for a trusted server S as well as an untrusted server \mathcal{A} . When C connects to S via MW and asks for a specific resource, e.g. GET/login.html, MW looks in the cache it holds for S and immediately delivers the page if it is found. The attacker's goal is to read or write private user data (e.g. her password) that is sent between C and S .

Note that in both these scenarios, the middlebox MW makes important decisions based on the identity of the server. Consequently, if an attacker can confuse the middlebox about the identity of the server, it can fool the middlebox into making incorrect decisions. While mcTLS seeks to

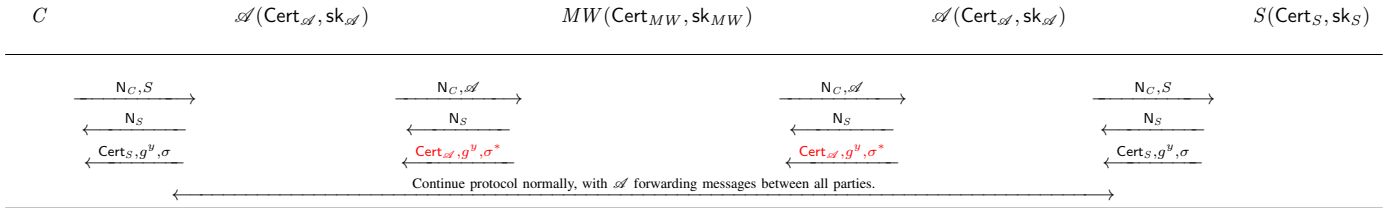


Figure 3: Unknown Key Share Attack on mcTLS Middleboxes.

protect endpoints against malicious middleboxes, the security analysis of the protocol appears to overlook attacks against middleboxes themselves, a category of attacks that we call *middlebox confusion attacks*.

Poisoning Caches with an Unknown Key Share Attack.

In our first attack, a client C wants to communicate with a trusted server S through an honest middlebox MW . However, the network between C and MW is controlled by the adversary, which interferes with the mcTLS protocol, as depicted in Figure 3. Importantly, the attacker only modifies messages that are seen by MW ; it then restores the correct messages and forwards them to C and S , so that they cannot detect this tampering.

When C sends its connection request (ClientHello) to MW , the attacker replaces the identity of S (typically indicated in the TLS-SNI extension) with its own identity \mathcal{A} . MW then initiates a connection with the attacker’s server, and \mathcal{A} forwards these handshake messages to S . When S returns its certificate and signature over its key share g^y , the attacker replaces them with its own certificate and signature over the same server key share. When MW forwards these messages to C , the attacker intercepts and replaces the certificate and signature with the original messages sent by S . The mcTLS handshake proceeds normally to conclusion, since the transcripts at the endpoints match. The transcript at the middlebox is quite different, but MW does not send or receive any Finished MACs in mcTLS, and so it cannot detect the attacker’s tampering.

At the end of this handshake, the client and server correctly believe that they have a connection to each other via MW , but MW believes that it is proxying a connection between C and the attacker-controlled website \mathcal{A} . This is a form of *unknown key share* attack on the middlebox, and even though the attacker does not know any of the context keys, it can still mount a serious attack. If the middlebox is a caching proxy, it will now deliver to C content that was previously retrieved from the attacker’s website, thereby allowing the attacker to inject data (and hence JavaScript) into connections between C and S . Hence, a network attacker can completely break the server authentication guarantees of caching proxies in mcTLS.

Bypassing Firewalls using Client Key Distribution. Our second attack is similar in structure to the previous scenario, but is exploitable in a different way. This time, the client C wishes to connect to the attacker’s website \mathcal{A} via a firewall

proxy MW . The attacker intercepts the client’s request and tampers with it to make MW believe that the C wishes to connect to the trusted server S instead. The attacker then completes the handshake with the client (using its own key shares). The connection with S cannot be completed, since the transcripts at C and S do not match, but if the protocol uses the client-key-distribution mode, it then unilaterally generates and sends context keys to \mathcal{A} and MW .

At this point, the C correctly thinks it is connected to the attacker’s server \mathcal{A} via MW , but MW thinks it is proxying a connection between C and S . This is a *server impersonation* attack on the middlebox. It allows the attacker to inject malware to the client, which will not be filtered as strictly as it should be, because MW thinks that the message is being sent by the trusted server S . Hence, mcTLS with client key distribution cannot enforce site-specific firewalling policies.

Record-Layer Attacks. mcTLS modifies the TLS record protocol to enforce fine-grained access control, but the resulting protocol is too weak for many use-cases. For example, an attacker who controls the network between C and its firewall MW can always inject malware into C , bypassing the firewall. This is because mcTLS privileges the endpoints over the middleboxes, allowing them to read and write all contexts. So all the attacker has to do is to present an innocuous data stream to MW and then it can decrypt, modify, and reencrypt data between MW and C .

As a more subtle example, consider a scenario where the attacker has obtained the reading key for some context, by compromising some middlebox, for instance. Since it does not have write access, the attacker should not be able to tamper with the data. However, since a reading context-key can be used both to verify and to create new MACs, the attacker can fool honest middleboxes into accepting tampered data, and then restore the original data before the endpoint sees it. Hence, even with read-only access, an attacker can inject data into the caches for trusted servers, or it can cause firewall protections to be bypassed.

Towards Provably Secure Accountable Proxies. The mcTLS protocol represents the most robust and flexible proxying proposal to date, and we believe that its approach is commendable. The middlebox confusion attacks described above can be prevented by using stronger variants of TLS, such as TLS 1.3 or TLS 1.2 with the session-hash extension. The record-layer attacks are harder to prevent without significant redesign. The main lesson from our attacks, however,

is that the security guarantees of active proxying scenarios can be subtle and even well-designed proposals can benefit from rigorous formal analysis.

III. DEFINING ACCE-AP SECURITY

So, in this section, we formally capture the security guarantees we believe should be attained by secure and accountable proxying, in the shape of the following *authenticated and confidential channel establishment with accountable proxies* (ACCE-AP) notion.

Using ACCE. Traditional authenticated key-exchange (AKE) security models [2] aim to prove that obtained session keys are indistinguishable from random. This is a very strong, composable guarantee; pseudorandom keys may then be securely used by any symmetric-key primitive. By contrast, Authenticated and Confidential Channel Establishment (ACCE) security [20] relaxes this, focusing on the security of the channel obtained by using those keys. An ACCE-secure key-exchange protocol guarantees that the established keys are able to construct a secure channel. ACCE security was designed to capture the complexities of TLS 1.2, whose key-confirmation step acts as a distinguishing oracle for the pseudorandomness of the session keys, without damaging the security of the established channel [25], [8].

TLS 1.3 was designed to (provably) guarantee composable AKE security, unlike TLS 1.2. Indeed, the session keys in TLS 1.3 are indistinguishable from random, which, coupled with the security of the authenticated encryption algorithms, should yield strong, composable security. However, the composition of AKE and secure record-layer exchanges is not trivial [9]. As a result, proofs such as [13] cannot simply stop at the key-establishment step. We choose instead to use the ACCE definitions, which provides no composability, but –in turn– makes explicit the guarantees of confidentiality, authenticity, and integrity we may expect of the established channel.

We focus on accountable proxying over secure channels. A crucial aspect of this is indeed defining what security (in particular authenticity and integrity) can be guaranteed at the record layer. That is where the augmented risk of added middleboxes is most obvious. To highlight this, and to describe precisely what is lost by giving third parties access to encrypted traffic, we choose to rely on ACCE security. We review both AKE and ACCE terminology in the appendix.

On proxy visibility. One implicit assumption of mTLS and of our work is that both the client and the server are aware, when starting the handshake, of all the proxies that can be found between them. We believe this is a necessary property when sensitive communication is exchanged between the client and the server. Indeed, a client might *e.g.*, be more cautious with the data it requests from a server if it knew that the traffic is proxied by a middlebox.

In practice, the client and server might not be aware of all the proxies connected to them, and especially, to

their communication partners. However, it would suffice to ask those proxies to identify themselves (and request permissions) in a first protocol step, orthogonal to our design. No cryptographic requirements would be made of this step: *i.e.*, we authentication is needed. This will be taken care of during our protocol.

Notations. Following the approach of Bhargavan et al. [4], the two-party ACCE notion with mutual authentication will be denoted 2-ACCE, whereas for server-only authenticated handshakes we use the notation 2-SACCE. Like Bhargavan et al., we view the proxied handshake in the presence of (possibly multiple) middleboxes as a number of linked 2-party protocols that are executed in parallel.

A. An intuition of the ACCE-AP model

Tuples of 2-party protocols. Our protocol is run by several parties, which are either clients, or middleboxes, or servers. We view an n -party session (with one client, one server, and $n - 2$ middleboxes) as a set of smaller, 2-party sub-protocols, in which instances of one party play the client or the server of a traditional 2-party handshake. This is defined as the *role* that the party plays in a 2-party execution.

Each party may run multiple concurrent executions of a 2-ACCE or 2-SACCE protocol: each protocol session is executed by a party *instance*. As a consequence, in an n -party handshake, various instances of the same party maybe executed at the same time: in some of these the instance will be a client, in others, a server. The m -th instance of party P_i is denoted π_i^m .

Although this may seem restrictive, it is more a matter of notation: indeed we choose to separate the communication between, say, the middlebox and the client, and the middlebox and the server because it is easier to keep track of it. We do assume that those instances may share state, which is certainly the case for the TLS handshake.

Instance partnering. Two-party handshakes use the notion of *partnering* or *matching conversation* to define which instances of two parties execute the protocol together. When moving from 2- to multi-party handshakes, a crucial notion that extends partnering is that of *session binding*, defining which 2-party handshakes take place as part of the same multiparty handshake. The binding is established through stored *local* information on direct communication partners, but also *global* information describing the configuration of the handshake, the parties acting as endpoints, the precise proxies used, and their access rights. In our model, all the parties involved in the handshake need to keep track of some cumulative *attributes*. Since the endpoints control the handshake, they also additionally need to store *master* parameters, such as the access rights of each middlebox. We refer to the endpoint instances that authenticate the master parameters by *master instances*.

Configurations and contexts. Session binding is defined in

terms of a handshake *configuration*, which is a suite of party identities ordered from the client to the server. We always assume that the client remains anonymous in the handshake, *i.e.*, it always runs server-only authenticated handshakes. The middleboxes may play the role of either client or server, but always authenticate to their partner.

We divide each message into disjoint *contexts*, following mTLS terminology. The list of all possible contexts is denoted Γ . Contexts can be defined, *e.g.*, as the concatenation of the message *type* (either request or response) and a message *fragment* (the header, the body, or a fragment thereof). Thus, one context could be “*request | header*”.

Each middlebox will have some *read*, *write*, or *none* access *permissions* on the defined contexts. We assume that entities which have write permissions for a given context may also read those contexts.

ACCE-AP security. We capture the security of the accountably-proxied secure-channel establishment protocols (called ACCE-security with accountable proxies – in short, ACCE-AP) in terms of the following requirements:

The (*S*)ACCE security of each independent, 2-party handshake, when the two partners are honest;

The *soundness of the configuration* in these 3 cases:

Fully honest. If all parties are honest, then all the instances are properly coupled, and all parties finish the configuration successfully, with all the appropriate keys.
Malicious *MW*. No collusion of malicious proxies learn anything about keys they have no right to.

Malicious endpoint. If an endpoint behaves maliciously w.r.t. the established configuration, contexts, or permissions, this will be detected by the remaining participants.

B. The ACCE-AP model

In what follows, we formalize configurations, permissions, and contexts as described above.

Oracles in ACCE-AP. Let k be an arbitrarily fixed integer. We consider AKE protocols executed between a client, a server, and up to $k - 2$ middleboxes. More formally a *handshake configuration* $HConfig$ is a list of parties P_1, \dots, P_k , such that P_1 is a client, P_k is a server, and the remaining parties are middleboxes. The configuration is ordered so that conversation will be forwarded in increasing order of indices for request messages, and inversely for the response.

Moreover, in each handshake, each of the proxies is associated with a list of contexts and corresponding permissions, as formalized below. For any configuration $HConfig$ of size k , its corresponding *contextual access list* $acl_{HConfig}$ is a hash-table of size $k - 2$, indexed by *middleboxes* in that handshake. For a given proxy P_i , the entry in $acl_{HConfig}$ corresponding to P_i is a list of length $|\Gamma|$ which, for each context c , contains the permissions (*i.e.*, exactly one value amongst *none*, *read*, *orwrite*) given to P_i for that context c .

As explained, a proxied handshake consists of a number of parallel AKE sessions executed concomitantly. Each session is run between two instances, one for each of the two communicating parties. The number of executions per proxied handshake is protocol-specific: mTLS for instance ran three sessions for the 3-party case, and six for the 4-party case. The multiple parallel sessions are also taken into account by our modification of the traditional NewSession oracle (presented in the Appendix).

Party attributes. Following the ACCE model, *parties* keep track of values such as their long-term secret keys, while party *instances* store session-specific values, *e.g.*, session and partner identifiers, the secret bits used by the Encrypt and Decrypt oracles, and established keys. These attributes are described in detail in Appendix C.

Instance-local ACCE-AP attributes. In addition to typical ACCE attributes, presented in the Appendix, the following attributes are specific to sole party instances:

π_i^m .bid. The *binding identifier* π_i^m .bid of an instance π_i^m is a special session-identifier that ties π_i^m to a proxied handshake. Its value is set upon configuration acceptance. Note that π_i^m .bid could differ from the same instance’s *session* identifier π_i^m .sid³.

π_i^m . β . The *configuration-acceptance* bit π_i^m . β of instance π_i^m takes values in $\{0, 1, \perp\}$. Initially, π_i^m . β is set to \perp . If π_i^m accepts the configuration input to NewSession, then π_i^m . β is set to 1; in case of rejection, it is set to 0.

Cumulative ACCE-AP attributes. Cumulative attributes store the partnering, session, and binding information of all the instances of a given party taking part in a single handshake (which we call *siblings*).

π_i^m .siblings. The *siblings* π_i^m .siblings of instance π_i^m are all the instances of P_i output by the NewSession oracle that generated π_i^m , including π_i^m itself.

Cumulative attributes are inherent to k -party handshakes; they are needed to established accountability.

π_i^m .c.pid. The *cumulative partner identifier* π_i^m .c.pid of an instance π_i^m is a table of entries of the type $(\pi_i^n, \pi_i^n$.pid), for each $\pi_i^n \in \pi_i^m$.siblings.

π_i^m .c.bid. The *cumulative binding identifier* π_i^m .c.bid of an instance is a tuple of binding identifiers π_i^ℓ .bid. Each identifier π_i^ℓ .bid corresponds to one instance $\pi_i^\ell \in \pi_i^m$.siblings, *i.e.*, π_i^m .c.bid is $(\pi_i^\ell$.bid) $_{\pi_i^\ell \in \pi_i^m$.siblings}.

π_i^m .c.config. The *cumulative configuration* π_i^m .c.config of a proxy instance π_i^m stores the configuration it has accepted for a given handshake. This attribute is *middlebox*-specific. Endpoints will store a *master* configuration attribute π_i^m .m.config (see below).

π_i^m .c.perm. The *cumulative permission attribute* π_i^m .c.perm is a list of $|\Gamma|$ elements, each taking a single value in the set

³Both the partner and session identifiers are protocol-specific, and typically include public and private information that makes each session unique.

{none, read, write}. This value represents the permission P_i has on the corresponding context. If $P_i \in \mathcal{C} \cup \mathcal{S}$, then all permissions are set to write.

$\pi_i^m.c.txt.keys$. The *cumulative context-key list* $\pi_i^m.c.txt.keys$ consists of $|\Gamma|$ elements, each of them a tuple of values corresponding to the read and write keys for a given context (or \perp if the value is unavailable). All sibling instances of a configuration-accepting party P_i share the same keys; if moreover $P_i \in \mathcal{C} \cup \mathcal{S}$, all the keys are non- \perp .

Master ACCE-AP attributes. Only stored by the handshake's endpoints, these attributes keep track of information on the handshake configuration and all the access-rights pertaining to the proxies. Thus, endpoints will be able to enforce accountability for all proxies.

$\pi_i^m.m.config$. The *master configuration attribute* is the endpoint equivalent to the middleboxes' $\pi.c.config$. For an endpoint instance π_i^m accepting the configuration HConfig, $\pi_i^m.m.config$ stores HConfig if it is the configuration for which NewSession was run.

$\pi_i^m.m.perm$. The *master permission-list attribute* $\pi_i^m.m.perm$ stores, for each of the $k - 2$ middleboxes input to NewSession, their permissions per context if they correspond to the $acl_{HConfig}$ value input to NewSession. Permissions are listed in the tabular style of $acl_{HConfig}$.

$\pi_i^m.m.bid$. The *master binding identifier* $\pi_i^m.m.bid$ is a list of tuples indexed by parties P_j in the configuration, containing binding identifiers $\pi_j^\ell.bid$ of instances π_j^ℓ involved in that handshake. We include P_i 's values as well; in particular, $\pi_i^m.bid$ is also included.

$\pi_i^m.\mu$. The *master configuration-acceptance bit* $\pi_i^m.\mu$ of a master instance π_i^m , with $P_i \in \mathcal{C} \cup \mathcal{S}$ is a value $\{0, 1, \perp\}$. Initially $\pi_i^m.\mu$ is set to \perp , and may later change to 1 or 0, depending on whether that master instance has validated a given configuration or not.

$\pi_i^m.m.ctx.keys$. The *master context key list* $\pi_i^m.m.ctx.keys$ is a list of $|\Gamma|$ pairs of read and write keys, one for each possible context.

Partnering. We extend 2-party partnering to our proxied case in Def. 1. For an instance π_i^m we require two sets: a set $\pi_i^m.PSet$ of entities partnered to π_i^m (incl. P_i), and a set $\pi_i^m.InstSet$ of instances partnered to π_i^m (incl. π_i^m).

Definition 1. (Partnering in ACCE-AP.) Let π_i^m be an instance of P_i . We define $\pi_i^m.PSet$ and $\pi_i^m.InstSet$ as:

If $P_i \in \mathcal{C} \cup \mathcal{S}$, $\pi_i^m.InstSet$ contains all instances π_j^n such that $\pi_j^n.bid \in \pi_i^m.m.bid$ (incl. π_i^m). The set $\pi_i^m.PSet$ includes all P_j with instances in $\pi_i^m.InstSet$ (incl. P_i).

Let $P_i \in \mathcal{M}\mathcal{W}$. Find the unique $\pi_j^n.m.bid$ of a party $P_j \in \mathcal{C} \cup \mathcal{S}$ such that $\pi_i^m.bid \in \pi_j^n.m.bid$. The set $\pi_i^m.InstSet$ is equal to the set $\pi_j^n.InstSet$ (computed as in the previous bullet point), and $\pi_i^m.PSet := \pi_j^n.PSet$.

Correctness. We cannot limit the correctness of proxied

handshakes to the correctness of partnering and keys. Indeed, we will also need to account for the aspects of session binding, cumulative and master configurations, and proxy permissions.

Definition 2. (Correctness). Consider a handshake executed for a configuration HConfig and a contextual access list $acl_{HConfig}$. Consider the parties $P_i \in HConfig$ and their instances in this handshake. The following conditions must hold simultaneously:

A. Partnering correctness. This is three-fold:

Instances. For each instance π_i^m , there must exist an instance of $\pi_i^m.pid$ whose session identifier sid equals $\pi_i^m.sid$ (the instance's partner is in that handshake).

Parties. All instances of P_i correctly accumulate the partner- and session-information of their siblings (cumulative attributes are consistent w.r.t. each sibling's partnering and session information).

Configuration. Consider a NewSession query outputting instances π_j^n , including master instances π_i^m of endpoints P_i (client or server), and middlebox instances π_k^ℓ of middleboxes.

We require: (i) All instances accept partner authentication and the handshake configuration; (ii) all accepting instances store the same configuration; (iii) the master binding identifier of master instances π_i^m includes exactly twice the binding identifier $\pi_j^n.bid$ of each instance π_j^n output by the same NewSession query (this includes the binders of P_i)⁴.

B. Access-control correctness. We require the consistency of permissions allowed by the endpoints to all middleboxes. Consider a NewSession query yielding a handshake for which the instances complete in an accepting state. The following must hold simultaneously:

Master context. The master permission set $\pi_i^m.m.perm$ of both endpoints coincides to $acl_{HConfig}$.

Middlebox context. The permissions for the handshake stored by each proxy, stored in $\pi_i^m.c.perm$, must be consistent with those given in $acl_{HConfig}$.

C. Key correctness. Partnered instances are also required to compute the same channel keys, and all read/write keys for a given permission must coincide.

Channel keys. Instances with the same session identifier compute the same channel key $\pi_i^m.ck$.

Master keys. All master instances compute the same master context keys.

Middleware keys. Each middlebox stores keys consistent with the permissions and context keys of master instances.

ACCE-AP security. We define security in terms of security games played by an adversary against a challenger. The adversary will have access to a number of oracles including

⁴Each binder is stored twice because our handshakes consist of tuples of 2-party handshakes.

traditional 2-ACCE oracles (described in the Appendix). However, we modify the NewSession oracle to take into account the transition from 2 to k parties:

NewSession(HConfig, acl_{HConfig}). The query takes as input a handshake configuration HConfig and a contextual access list acl_{HConfig}. The output is a list of k sets of tuples $(\pi_i^j, \pi_i^j.\rho, \pi_i^j.\text{pid})$. Each set $\{(\pi_i^j, \pi_i^j.\rho, \pi_i^j.\text{pid})\}_{j=1}^{\ell_i}$ corresponds to a party P_i in the HConfig. Here $\ell_i \geq 1$ is public, but not adversarially-fixed; it depends on the protocol and number of participants⁵. Each set $\{(\pi_i^j, \pi_i^j.\rho, \pi_i^j.\text{pid})\}_{j=1}^{\ell_i}$ is a series of oracle-instances π_i^j (with roles ρ and partner identifiers pid) of the party P_i .

The remaining, ACCE-like oracles we use are:

Send. The adversary uses this oracle to send a message to a given, existent instance, forwarding its reply.

Reveal. This oracle allows the adversary to learn the channel keys of a given party instance.

Corrupt. By this oracle, the adversary corrupts either middleboxes or servers, obtaining their long-term keys.

Encrypt. This is a left-or-right oracle, which allows the adversary to encrypt one of two possible messages of equal length, depending on a hidden bit b .

Decrypt. This oracle allows the adversary to decrypt only ciphertexts not returned by Encrypt.

The security properties. We define the security of a proxied handshake in terms of three properties: instance entity-authentication, instance channel-security, and configuration soundness. The first two are backward compatible with traditional 2-party (S)ACCE models, with a difference in the partnering. In the ACCE-AP model, each instance has more partners, none of which can be corrupted.

Each game will begin with a setup phase, in which all the n_P parties are instantiated, with corresponding long-term keys and certificates, returning the public parameters to the adversary \mathcal{A} . The adversary's winning advantage is quantified over the randomness of all participants.

Definition 3. (Entity Authentication). In the entity authentication game, after setup, \mathcal{A} is given arbitrary access to the new NewSession oracle above, and also to Send, Corrupt, and Reveal. Finally \mathcal{A} halts, and wins if there exists an instance π_i^m such that all the following conditions occur:

- π_i^m has ended in an accepting state: $\pi_i^m.\alpha = 1$, with partner $\pi_i^m.\text{pid} = P_j \in \mathcal{M}\mathcal{W} \cup \mathcal{S}$;
- No party in $\pi_i^m.\text{PSet}$ (partnered to π_i^m) is corrupted;
- There exists no session π_j^n partnered with π_i^m ($\pi_j^n \notin \pi_i^m.\text{InstSet}$).

The adversary's advantage is its success probability.

Definition 4. (Channel Security). For channel security, after setup, \mathcal{A} is given access to the new NewSession oracle and

⁵If $\ell_i = 1$, the output coincides with what would be returned by the standard NewSession query for each participant P_i .

to Send, Corrupt, Reveal, Encrypt, and Decrypt. Whenever a new instance π_i^m is created via NewSession, the attribute storing the secret bit used in Encrypt, denoted $\pi_i^m.\text{b}$, is chosen uniformly at random. Finally, \mathcal{A} halts, outputting a tuple consisting of an instance π_i^m for $P_i \in \mathcal{C}$ and a guess-bit d . We say \mathcal{A} wins if the following conditions occur simultaneously:

- $\pi_i^m.\text{b} = d$;
 - No Corrupt query was made on any party in $\pi_i^m.\text{PSet}$;
 - No Reveal query is made on instances in $\pi_i^m.\text{InstSet}$.
- \mathcal{A} 's advantage is defined as $|p_{\mathcal{A}} - 1/2|$, where $p_{\mathcal{A}}$ is \mathcal{A} 's success probability.

Configuration soundness. We now describe a new security property capturing configuration soundness with respect to the input of NewSession. Configuration soundness has three components, formalized separately: (1) the soundness of the configuration w.r.t. fully-honest partners; (2) the soundness of the access-control-key distribution in the presence of corrupted middleboxes (and honest endpoints); (3) the soundness of the key-distribution even for a corrupted endpoint. We do not consider security against a collusion between endpoints and middleboxes: as soon as such a collusion occurs, we can guarantee practically no security for the middleboxes situated between the two colluding partners. Instead, the security level becomes somewhat equivalent to that of a smaller handshake, in which the colluding middlebox acts as the endpoint.

The three games begin with the setup described above. Then, \mathcal{A} is given access to the new NewSession oracle and also to Send, Corrupt, Reveal, Encrypt, and Decrypt.

Definition 5. (Config. soundness: partnering). The game is played as above. The adversary ends by outputting a master instance π_i^m , with $P_i \in \mathcal{C}$, ending in an accepting state with respect to its configuration, i.e., $\pi_i^m.\mu = 1$. Let $\pi_j^n \in \pi_i^m.\text{InstSet}$ be such that $P_j \in \mathcal{S}$, $\pi_i^m.\text{sid} = \pi_j^n.\text{sid}$ and $\pi_j^n.\mu = 1$. We write HConfig for the configuration input to the NewSession query that outputs π_i^m . The adversary wins if the following conditions occur simultaneously.

- I. No Corrupt query was made on parties in $\pi_i^m.\text{PSet}$;
- II. No Reveal query is made on instances in $\pi_i^m.\text{InstSet}$.
- III. One of the following conditions occurs:
 - a. The master configuration $\pi_i^m.\text{m.config} \neq \pi_j^n.\text{m.config}$, or $\pi_i^m.\text{m.config} = \pi_j^n.\text{m.config} \neq \text{HConfig}$;
 - b. The cumulative configuration of middlebox instances $\pi_k^\ell \in \pi_i^m.\text{InstSet}$ differs from $\pi_i^m.\text{m.config}$ or HConfig;
 - c. There is an instance $\pi_k^\ell \in \pi_i^m.\text{InstSet}$ such that $\pi_k^\ell.\text{bid} \notin \pi_i^m.\text{m.bid}$;
 - d. There is an instance $\pi_k^\ell \in \pi_i^m.\text{InstSet}$ s.t. $\pi_k^\ell.\beta = 0$.

Condition (III.d) sometimes implies (III.c): it does for our protocol. However, this is not always true: the binding of all instances $\pi_k^\ell \in \pi_i^m.\text{InstSet}$ could be correct, but there may exist an instance $\pi_k^\ell \in \pi_i^m.\text{InstSet}$ that rejects that configuration, that is $\pi_j^n.\beta = 0$.

Definition 6. (Config. soundness: malicious MW). The game is played as above. The adversary will output, in a test phase, a tuple consisting of an instance π_i^m with $\pi_i^m.\mu = 1$ and a context ctxt , such that $P_i \in \mathcal{C}$ and there exists an instance $\pi_j^n \in \pi_i^m.\text{InstSet}$ such that $P_j \in \mathcal{S}$ and $\pi_i^m.\text{sid} = \pi_j^n.\text{sid}$. The challenger will output either the true context key K_{ctxt} in $\pi_i^m.\text{m.ctxt.keys}$ or a random key of the same length, according to a hidden bit b . Then \mathcal{A} may continue querying oracles, finally outputting a bit d , winning if the following occurs simultaneously:

- $d = b$;
- All instances $\pi_i \in \pi_i^m.\text{InstSet}$ end in an accepting state with respect to the configuration, i.e., $\pi_i.\beta = 1$;
- No Corrupt query was made on any $P_k \in \pi_i^m.\text{PSet}$ s.t. ctxt is registered as a permission for P_k in $\text{acl}_{\text{HConfig}}$;
- No Reveal query is made on instances in $\pi_i^m.\text{InstSet}$.

Finally, we guarantee that honest middleboxes accepting the configuration cannot distinguish from random the keys they are not entitled to, even if one endpoint is malicious.

Definition 7. (Config. soundness: malicious endpoint). The game is played as above, and \mathcal{A} eventually halts. It wins if there exists a client or server instance π_i^m storing a configuration $\pi_i^m.\text{m.config}$ s.t. the following conditions apply simultaneously.

- All instances π_j for $P_j \in \mathcal{M}\mathcal{W} \cap \pi_i^m.\text{PSet}$ accepted the configuration;
- Parties $P_j \in \mathcal{M}\mathcal{W} \cap \pi_i^m.\text{PSet}$ are uncorrupted;
- No Reveal query is made on instances in $\pi_i^m.\text{InstSet}$.
- One of the following two conditions holds:
 - If Corrupt was queried on a server $P_k \in \{\mathcal{S} \cap \pi_i^m.\text{m.config}\}$ (with $P_k \neq P_i$), then the following holds simultaneously: (a) $P_i \in \mathcal{C}$; (b) there are instances $\pi_i^\ell, \pi_k^s \in \pi_i^m.\text{InstSet}$ with $\pi_i^\ell.\text{sid} = \pi_k^s.\text{sid}$; (c) all instances $\pi_i^t \in \pi_i^m.\text{InstSet}$ ended in an accepting state w.r.t. the configuration, i.e., $\pi_i^t.\mu = 1$; (d) there exists $P_j \in \mathcal{M}\mathcal{W} \cap \pi_i^m.\text{PSet}$ and some instance $\pi_j^z \in \pi_i^m.\text{InstSet}$ with $\pi_i^\ell.\text{m.perm} \neq \pi_j^z.\text{c.perm}$.
 - If no Corrupt query is made on the server $P_k \in \{\mathcal{S} \cap \pi_i^m.\text{m.config}\}$, then the following holds simultaneously: (a) $P_k = P_i$ (the party P_i was the handshake server); (b) all instances $\pi_i^z \in \pi_i^m.\text{InstSet}$ accepted the configuration ($\pi_i^z.\mu = 1$); (c) there exists $P_j \in \mathcal{M}\mathcal{W} \cap \pi_i^m.\text{PSet}$ and an instance $\pi_j^t \in \pi_i^m.\text{InstSet}$ s.t. $\pi_i^m.\text{m.perm} \neq \pi_j^t.\text{c.perm}$.

IV. AN ACCE-AP-SECURE DESIGN

We now propose a protocol that is provably ACCE-AP secure. In fact, we introduce a generic *design*, to be instantiated with one or multiple accountable proxies. For one middlebox, this yields Figure 4; a 2-middlebox case is presented in the full version. In Appendix B, we also depict how to optimally apply our design to TLS 1.3.

A. The ACCE-AP-secure Protocol Construction II

We denote by Π a generic secure-channel establishment protocol with accountable proxies, like the one in Figure 4. We use a modular construction, with a careful composition of independent, unmodified 2-party protocols guaranteeing slightly more than 2-party SACCE and respectively ACCE-security.

Our protocol consists of three phases: 2-ACCE, binding, and access-control.

2-ACCE Phase. This phase consists of running a number of parallel 2-SACCE and ACCE sessions, thus establishing channel and *exported* keys. Although not standardly output by AKE protocols, exported keys feature in many real-world AKE protocols, including TLS 1.3 (which standardizes an exporter secret)⁶.

Each party in the protocol starts the following instances:

Transport instances. These instances are run by each entity with its direct neighbor(s): the endpoints only open one such instance, whereas each middlebox opens two. After successful completion, the transport sessions are used to securely exchange binding and access-control information. They also serve as an outer channel for record-layer message-exchanges.

Master instance. Only the client and the server open one such instance (irrespective of how many proxies take part in the handshake). In this session, the configuration and access control data are agreed upon, and the read/write keys are computed.

Key-transfer instance. Finally, each endpoint must open one instance to each middlebox it is not directly connected to via a transport session. Key-transfer instances are used to provide access-control keys and confirm the configuration with the middlebox so that the latter may accept or reject the handshake.

For the 1-middlebox case in Figure 4, the key transfer is done in the transport sessions: the client, the middlebox, and the server each open precisely two instances, yielding two transport sessions and the master session. For two middleboxes, each party opens three instances: the endpoints create one transport instance (which also serves to transfer keys to the closest proxy), one key-transfer instance (to the farthest middlebox), and the master instance. Each proxy opens two transport sessions (one to the neighboring endpoint, one to the second middlebox) and one key-transfer instance (to the non-neighboring endpoint).

Note that each instance will first need to finish in an accepting state with respect to authentication before the binding phase proceeds.

Binding Phase. We bind sibling instances together by using the exported keys output in each of the parallel sessions.

⁶The TLS 1.2 key exporter construction [11] is insufficient: we will need our exported keys to be unique (for instance, they could depend on the full session hash, not just the nonces).

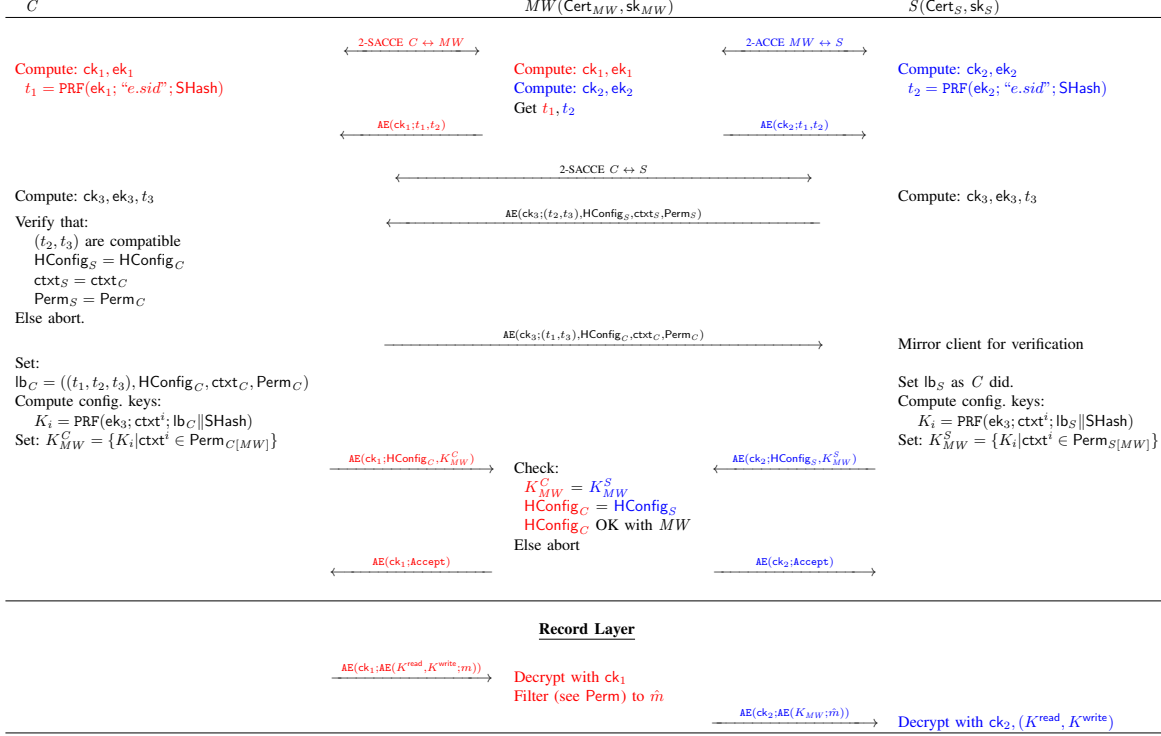


Figure 4: Our generic, ACCE-AP-secure protocol Π , illustrated with a single middlebox

Each session's *binding identifier* is output by a pseudorandom function (like HKDF) keyed with an exported key, on input a label concatenated with that session's hash. It is essential for configuration soundness that the middlebox and server can verify that they are talking to the same client (since the client instances have so far been independent). We force the client to prove its knowledge of its instance binders to the server, and give it the binder for the remaining sessions; this will enable the two endpoints to agree on configuration and access-control details for that handshake. Both endpoints must agree on a configuration and an access control list before proceeding to the next phase.

In Figure 5, the binding identifiers are the values t_1, t_2, t_3 , computed from the exported keys and the transcript hashes, as the output of a PRF instance, *i.e.*, $t_i = \text{PRF}(ek_i; "e.sid"; \text{SHash})$. Each values t_i ($i \in \{1, 2, 3\}$) is stored individually in an $\pi.\text{bid}$, while $\pi.m.\text{bid}$ stores all three of them.

Access-Control Phase. Finally, the endpoints each compute the access control keys for the handshake by running the same PRF as before, keyed with the exported keys of the master session, on input a label describing the context, and the transcript hash. More precisely, for a context ctxt^i and permissions a , the label lb_i^P is a string $\text{ctxt}^i || a$. The access control keys are computed as $K_i = \text{PRF}(ek_3; \text{ctxt}^i; \text{lb}_i^P || \text{SHash})$. Each middlebox receives the handshake configuration, one set of keys from the client, and one from the server, over a direct,

secure channel (a key-transfer or a transport session). The middlebox verifies that its view is consistent with the received configurations, and that the keys sent by the endpoints coincide. If either verification fails, the middlebox *rejects the configuration*. Otherwise, it accepts the configuration and sends a message to that effect to the endpoints. Only then do the endpoints end in an accepting state, as well.

B. Security analysis

We first discuss the security assumptions we make on our primitives. Then, we state the security of our generic design for accountable proxying, given in Figure 4; due to lack of space, we leave the full proofs of these to the full version.

Security assumptions. The security of our protocol will require a 2-SACCE secure authenticated key-exchange (AKE) protocol Π_1 between the client and the middlebox, a 2-ACCE secure AKE protocol Π_2 between the middlebox and the server, and a 2-SACCE secure AKE protocol Π_3 between the client and the server. Notably, Π_1 and Π_3 offer unilateral authentication, whilst Π_2 is mutually authenticated.

For $i \in \{1, 2, 3\}$, let Ψ_i denote the extension of Π_i to a protocol $\Psi_i(\Pi_i)$ that runs Π_i , outputting ck , but in addition outputs an export key ek . We require that the export keys output $\Psi_i(\Pi_i)$ be indistinguishable from random, and in addition, that no PPT adversary (with access to the usual 2-party oracles NewSession , Send , Reveal , Corrupt) can find two non-partnered sessions for which the output export keys

are identical, *even if* \mathcal{A} can corrupt all servers (for the SACCE protocols) and all parties (for the ACCE protocol). We call the latter property *pseudo-uniqueness of ek*.

Our main reason for this assumption on the export keys is our notion of binding. We need to ensure that each configuration yields a unique set of binders even when a legitimate party in the handshake misbehaves. Such a condition helps us bypass *e.g.* attacks of the triple-handshake kind, like those on mcTLS shown in Section II-C. Two important aspects should be mentioned, regarding the pseudo-uniqueness assumption: (1) it is non-standard in AKE security (though interestingly it holds for TLS 1.3); (2) we do not require this same condition for the channel keys ck .

We do not demand pseudo-uniqueness of channel keys since they play no part in binding. As we do not consider resumption, or other short-cuts needing stronger keys, we only need channel keys be pseudo-random.

The ACCE-AP-security of the 1-middlebox-version of Π . We provide two security theorems for the multicontext handshake proposed in Figure 4: Theorem 1 captures entity authentication and channel security, and Theorem 2 captures configuration soundness (in its three aspects). The proofs are described in the full version.

Theorem 1. *Let Π be protocol in Figure 4, for which Π_1, Π_2, Π_3 and Ψ_1, Ψ_2, Ψ_3 are defined as above. If Π_1, Π_3 are 2-SACCE secure, and Π_2 is 2-ACCE secure, then:*

- Π guarantees ACCE-AP entity authentication and channel security;

Theorem 2. *Let Π be the protocol in Fig. 4, for which Π_1, Π_2, Π_3 and Ψ_1, Ψ_2, Ψ_3 are defined as above. If: (1) Π_1, Π_3 are 2-SACCE secure and Π_2 is 2-ACCE secure; (2) Ψ_1, Ψ_2, Ψ_3 extend Π_1, Π_2, Π_3 to also yield export keys ek which are indistinguishable from random and pseudo-unique; (3) the KDF used to compute t_1, t_2, t_3 is a secure PRF; (4) the hash function used for the recurring session hash is collision-resistant, then:*

- Π is configuration sound w.r.t. partnering, malicious middleboxes, and malicious endpoints.

These theorems informally guarantee 4 properties: a network adversary cannot convince an (honest) party it is part of a different configuration than is actually the case; all access-control keys are correctly distributed (middleboxes get the keys they are entitled to, and no collusion of malicious middleboxes can learn anything about keys it is not entitled to know); middleboxes can detect malicious behavior in one endpoint (with respect to the permissions); no network adversary can learn anything about the messages sent at the record layer of our multicontext handshake.

The security of the handshake, however, makes no guarantee w.r.t. the integrity and authenticity of messages sent at the proxied record layer. The secure-channel statement guarantees the integrity and authenticity of messages, but only

within a two-party session. To capture integrity on the end-to-end record layer, we need to consider the *functionalities* and *permissions* of the middleboxes. We do so next.

Record-layer security. For a middlebox MW and a message fragment m (as given by the division into contexts), let $m' = MW(m)$ be the message m' obtained by the honest modification of m by MW . Note that m' could be the same as m , a correctly formatted message other than m , or it could also be an error message \perp .

For the one-middlebox case, we can prove the following result w.r.t. record-layer security. The case of several middleboxes will be discussed in the full paper, since it presents more problems.

Theorem 3. *For any handshake of the protocol in Figure 4, for all instances ending in an accepting state w.r.t. the configuration, for all contexts, consider a message-fragment m sent from one endpoint (which we will call sender) to the other endpoint (here called the receiver). Assume that the receiver receives a ciphertext c^* , for which it accepts the authentication of the message, decrypting it to a message fragment m^* . Then:*

- If the middlebox is corrupted, but has read-only permissions for the context corresponding to fragments m, m^* , then $m^* = m$.
- If the sender is corrupted, but the receiver and middlebox are honest, then $m^* = MW(m)$.

We can make no integrity and authenticity statement for a corrupted middlebox with write access. However, this is a flaw inherent to using read and write keys. The second guarantee in Theorem 3 is something no SSL inspection method, including mcTLS, has so far provided. Indeed, our double-encryption thwarts the middlebox-bypass attack we presented in Section II-C. This is a useful guarantee, as it ensures that honest proxies always perform their due tasks on messages sent to them, even by a dishonest endpoint.

Enhancing integrity. We could also make this integrity guarantee stronger, and ensure, like mcTLS, that a message that is sent unchanged from the honest sender to the honest receiver can be distinguished from a message that modified by a potentially-malicious middlebox with write access. This is easily done by adding a MAC, under, *e.g.*, a key derived from the export key of the master session under a new label. At the record layer, instead of just sending an authenticated encryption under the read and write key of a message, we could also add a MAC of that same message computed with the write key for that fragment. We choose to avoid this overhead, since the guarantees we can offer for the record layer are anyway limited.

C. Prototype Implementation

We built a proof-of-concept implementation of the ACCE-AP-II protocol instantiated with TLS 1.3, as per

depicted in Fig. 5. Our implementation is based on the miTLS [7] library, an implementation of TLS 1.3 (draft 23) in F* [37]. We prefer to use miTLS over more mainstream implementations, such as OpenSSL, for several reasons:

- miTLS is written in a modular way, with well-contained local state in the various modules (*e.g.*, the key schedule, negotiation module, handshake state-machine); this simplifies the writing of protocol extensions, in particular when multiple handshakes must be managed simultaneously, compared to implementations with large amount of ambient state. miTLS has been previously used for similar purposes within the FlexTLS tool [12];
- miTLS is designed to enable formal verification. The miTLS implementation of TLS 1.3 is under active development, but the security and correctness of many of its components have been formally proved, including the underlying cryptographic library [38] and the record layer [5]. Currently, the miTLS implementation does not prove pseudo-uniqueness of exporter keys (which we rely on in our security proofs), but we believe that the key schedule contains all the necessary elements for such a proof, based on the PRF assumptions on HKDF and on the PRF-ODH assumption already in place for key indistinguishability, and that it is likely to be added in the future by the miTLS authors.

Our main change to miTLS is the addition of a new content type and record sub-protocol, which is used for the (t_1, t_3) and (t_2, t_3) exchanges and for transmitting K_{MW}^C and K_{MW}^S . Using a separate content type means that those messages are invisible to the miTLS handshake and do not interfere with its internal invariant and state machine (in particular, the transcript hashing and post-handshake messages). Similarly, these extra messages do not interfere with the flow of application data, following the multi-stream model of the miTLS record [5], and so do not interfere with the current application guarantees of miTLS.

Our changes are mostly contained to two files within miTLS: the top-level interface (TLS.fst) and the new middlebox sub-protocol (Middlebox.fst). Only minor changes are needed to the handshake to give access to the exporter secret-keyed PRF to the middlebox protocol. Our total changes consist of under 1000 lines of F* code, which is less than 5% of the miTLS protocol code. Our prototype implementation is available on Github as a branch of the miTLS repository.⁷ We note that this prototype is restricted to a single middlebox whose certificate is known in advance to the client and server, and which has full read and write permissions.

To test our implementation, we use the simple HTTP client/server tool provided by miTLS. For the middlebox, we created a new application that exercises the new options added to the top-level interface TLS.fst. Our middlebox

is configured to read and display the traffic, but otherwise performs no function.

We measured the performance of our code when using AES128-GCM record algorithm, HKDF-SHA256 for the key schedule, Curve25519 for the Diffie-Hellman key exchange, and ECDSA on the NIST P-384 curve with SHA-384 for the server and middlebox signatures. We measure a marginal increase in connection time (35ms \rightarrow 38ms), but a significant throughput loss of about 40% (516MB/s \rightarrow 299MB/s) in our experiment, compared to a direct connection between the client and server.

The latency increase is caused by the additional Diffie-Hellman computation at the middlebox, and disappears when PSK or 0-RTT is used. The decrease in throughput is primarily due to the re-encryption overhead, which is done sequentially in our prototype. We believe this overhead can be mitigated by parallelizing encryption and decryption on the middlebox.

V. CONCLUSIONS

In this paper, we proposed a new security definition, as well as a modular, provably-secure construction of an accountably-proxied secure channel. We hope that our design will provide a better example of how accountability can be achieved without harming the authenticity and integrity of the underlying channel. However, our formal model/definitions are complex, and –even so– we achieve limited record-layer guarantees in multi-middlebox setting. This illustrates the intuitive principle that the more middleboxes are interspersed between the client and the server, the weaker the security of incoming and outgoing messages becomes. If many (potentially-malicious) middleboxes are present, the endpoints may no longer rely on each of the proxies to inspect the traffic. As our analysis suggests, the most important middlebox should then be placed closest to the server to guarantee its functionality is observed.

Finally, our proposed protocol design is not meant to encourage widespread active proxying. Instead, we hope to have shown that, beyond the loss of end-to-end security inherent to proxying, it is difficult to construct a proxied handshake in a sound way. The integrity and authenticity properties of the record layer degrade quickly with the number and position of the corrupted middleboxes, and schemes with more than two proxies are not very efficient. As such, our results indicate that care should be taken before introducing a middlebox, and for sensitive end-to-end communications, they should preferably be eliminated.

ACKNOWLEDGMENTS

P.-A. Fouque, K. Bhargavan, and C. Onete are grateful for the support of the ANR through grant 16 CE39 0012 (SafeTLS). K. Bhargavan was also funded by the ERC grant CIRCUS. I. Boureanu was in part funded by the Marie Skłodowska-Curie grant No 661362 under EU’s Horizon 2020 programme.

⁷<https://github.com/mitls/mitls-fstar/tree/middlebox>

REFERENCES

- [1] S. Alt, P. Fouque, G. Macario-Rat, C. Onete, and B. Richard. A cryptographic analysis of UMTS/LTE AKA. In *Proceedings of ACNS*, volume 9696 of *LNCS*, pages 18–35. Springer, 2016.
- [2] M. Bellare and P. Rogaway. Entity authentication and key distribution. In *Advances in Cryptology — CRYPTO*, pages 232–249, 1993.
- [3] K. Bhargavan, B. Blanchet, and N. Kobeissi. Verified models and reference implementations for the TLS 1.3 standard candidate. In *Proceedings of S&P*, pages 483–502. IEEE, 2017.
- [4] K. Bhargavan, I. Boureanu, P. Fouque, C. Onete, and B. Richard. Content delivery over TLS: a cryptographic analysis of Keyless SSL. In *Proceedings of Euro S&P*, 2017.
- [5] K. Bhargavan, A. Delignat-Lavaud, C. Fournet, M. Kohlweiss, J. Pan, J. Protzenko, A. Rastogi, N. Swamy, S. Zanella-Béguélin, and J. K. Zinzindohoué. Implementing and proving the tls 1.3 record layer. Appears in *IEEE S&P (Oakland)*, 2017. Cryptology ePrint Archive, Report 2016/1178, 2016. <http://eprint.iacr.org/2016/1178>.
- [6] K. Bhargavan, A. Delignat-Lavaud, C. Fournet, A. Pironti, and P.-Y. Strub. Triple handshakes and cookie cutters: Breaking and fixing authentication over TLS. In *Proceedings of IEEE S&P 2014*, pages 98–113. IEEE, 2014.
- [7] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, and P. Strub. Implementing TLS with verified cryptographic security. In *Proceedings of IEEE S&P 2013*, pages 445–469, 2013.
- [8] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, P. Strub, and S. Z. Béguélin. Proving the TLS handshake secure (as it is). In *CRYPTO*, pages 235–255, 2014.
- [9] C. Brzuska, M. Fischlin, N. Smart, B. Warinschi, and S. Williams. Less is more: Relaxed yet composable security notions for key exchange. *International Journal of Information Security*, 12(4):267–297, 2013.
- [10] C. Brzuska and H. kon Jacobsen. A modular security analysis of EAP and IEEE 802.11. In *Proceedings of PKC*, volume 10175 of *LNCS*, 2017.
- [11] C. Brzuska, H. kon Jacobsen, and D. Stebila. Safely exporting keys from secure channels: on the security of EAP-TLS and TLS key exporters. In *EuroCrypt*, 2016.
- [12] B. B. A. Delignat-Lavaud, N. Kobeissi, A. Pironti, and Karthikeyan Bhargavan. FLEXTLS: A tool for testing TLS implementations. In *Proceedings of USENIX WOOT 2015, best paper award*, 2015.
- [13] B. Dowling, M. Fischlin, F. Günther, and D. Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In *ACM CCS*, pages 1197–1210, 2015.
- [14] B. Dowling, M. Fischlin, F. Günther, and D. Stebila. A cryptographic analysis of the TLS 1.3 draft-10 full and pre-shared key handshake protocol. IACR ePrint archive, <https://eprint.iacr.org/2016/081>, 2016.
- [15] Z. Durumeric, Z. Ma, D. Springall, R. Barnes, N. Sullivan, E. Bursztein, M. Bailey, J. Halderman, and V. Paxson. The security impact of HTTPS interception. In *Proceedings of NDSS*, 2017.
- [16] M. Fischlin and F. Günther. Replay attacks on zero round-trip time: the case of the TLS 1.3 handshake candidates. In *Proceedings of Euro S&P*, pages 60–75. IEEE, 2017.
- [17] P. Fouque, G. Macario-Rat, C. Onete, and B. Richard. Achieving better privacy for the 3gpp aka protocol. In *Proceedings of PETS (PoPETS)*. De Gruyter, 2016.
- [18] R. Graham. Extracting the SuperFish certificate. <http://blog.erratasec.com/2015/02/extracting-superfish-certificate.html>, 2015.
- [19] F. Günther, B. Hale, T. Jager, and S. Lauer. 0-RTT key exchange with full forward secrecy. In *Advances in cryptology — EUROCRYPT*, volume 10212 of *LNCS*, pages 519–548. Springer, 2017.
- [20] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk. On the security of TLS-DHE in the standard model. In *Advances in Cryptology — CRYPTO*, volume 7417 of *LNCS*, pages 273–293. Springer, 2012.
- [21] T. Jager, J. Schwenk, and J. Somorovsky. On the security of TLS 1.3 and QUIC against weaknesses in PKCS#1 v1.5 encryption. In *Proceedings of ACM CCS 2015*, 2015.
- [22] M. Kohlweiss, U. Maurer, C. Onete, B. Tackmann, and D. Venturi. (de-)constructing TLS. IACR ePrint archive, report 020/2014, 2014.
- [23] M. Kohlweiss, U. Maurer, C. Onete, B. Tackmann, and D. Venturi. (De-)constructing TLS 1.3. In *Proceedings of Indocrypt*, volume 9462 of *LNCS*, pages 85–102. Springer, 2015.
- [24] H. kon Jacobsen. A modular security analysis of eap and ieee 802.11. Ph.D. thesis, 2017.
- [25] H. Krawczyk, K. Paterson, and H. Wee. On the security of the TLS protocol: A systematic analysis. In *Proceedings of CRYPTO 2013*, volume 8042 of *LNCS*, pages 429–448, 2013.
- [26] H. Krawczyk and H. Wee. The OPTLS protocol and tls 1.3. In *Proceedings of Euro S&P*. IEEE, 2016.
- [27] O. Levillain, B. Gourdin, and H. Debar. TLS record protocol: Security analysis and defense-in-depth countermeasures. In *Proceedings of ACM ASIACCS 2015*, pages 225–236. ACM, 2015.
- [28] X. Li, J. Xu, Z. Zhang, D. Feng, and H. Hu. Multiple handshakes security of TLS 1.3 candidates. In *Proceedings of S&P*, pages 485–505. IEEE, 2016.
- [29] A. T. Ltd. Comparing approaches for web and DNS infrastructure security. goo.gl/PBD2N6, 2016.
- [30] D. Migault, K. Ma, Ericsson, R. Rich, Akamai, S. Mishra, V. Communications, O. G. de Dios, and Telefonica. Lurk tls dtls use cases. <https://tools.ietf.org/html/draft-mglt-lurk-tls-use-cases-02>, 2016.
- [31] P. Morrissey, N. Smart, and B. Warinschi. A modular security analysis of the TLS handshake protocol. In *Proceedings of ASIACRYPT 2008*, volume 5350 of *LNCS*, pages 55–73, 2008.
- [32] D. Naylor, R. Li, C. Gkantsidis, T. Karagiannis, and P. Steenkiste. And then there were more: Secure communication for more than two parties. In *the 13th International Conference on emerging Networking EXperiments and Technologies*, pages 88–100, 2017.
- [33] D. Naylor, K. Schomp, M. Varvello, I. Leontiadis, J. Blackburn, Diego R. López, K. Papagiannaki, Pablo Rodriguez Rodriguez, and P. Steenkiste. Multi-Context TLS (mcTLS): Enabling Secure In-Network Functionality in TLS. In *Proceedings of SIGCOMM 2015*, pages 199–212. ACM, 2015.
- [34] M. O’Neill, S. Ruoti, K. Seamons, and D. Zappala. TLS proxies: Friend or foe. In *Proceedings of IMC*, pages 551–557, 2016.
- [35] T. Ormandy. Cloudflare reverse proxies are dumping uninitialized memory. <https://bugs.chromium.org/p/project-zero/issues/detail?id=1139>, 2015.
- [36] K. Paterson, T. Ristenpart, and T. Shrimpton. Tag size does matter: Attacks and proofs for the TLS record protocol. In *Advances in Cryptology — ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 372–389. Springer-Verlag, 2011.

- [37] N. Swamy, C. Hritcu, C. Keller, A. Rastogi, A. Delignat-Lavaud, S. Forest, K. Bhargavan, C. Fournet, P.-Y. Strub, M. Kohlweiss, J.-K. Zinzindohoue, and S. Zanella-Béguelin. Dependent types and multi-monadic effects in F^* . In *43rd ACM Symposium on Principles of Programming Languages, POPL 2016*, pages 256–270, 2016.
- [38] J. K. Zinzindohoué, K. Bhargavan, J. Protzenko, and B. Beurdouche. Hacl*: A verified modern cryptographic library. Appears in ACM CCS 2017. Cryptology ePrint Archive, Report 2017/536, 2017. <http://eprint.iacr.org/2017/536>.

APPENDIX A. RELATED WORK

Two-party authenticated key-exchange protocols are an old and well-studied primitive [2] and various notions of channel security have been formalized and applied to TLS 1.2 [20], [25], [8], [22], [31], [36], [27] and 1.3 [13], [14], [19], [16], [3], [28], [23], [26], [21]. However, as Bhargavan et al. [4] showed in the context of TLS, and as Alt et al. [1] and Fouque et al. [17] – in the context of mobile networks, expanding two-party handshakes to include even a single, dedicated middlebox can expose the obtained channel to serious attacks. Our work reuses some parts of [4]’s security model, specifically the elegant way in which multiparty protocols can be viewed as compositions of 2-party AKE schemes; however, the scope of our paper is fundamentally different. The work of Bhargavan et al. centers around a particular type of caching middlebox, with all-or-nothing access rights to server contents. In this paper, our goal is to provide finer-grained access-rights (as defined through contexts) to both client- and server-side middleboxes.

Other approaches to delegated or proxied AKE protocols exist. Composition-centric approaches, such as those of Brzuska et al. [11], [10] and Jacobsen [24] capture the three-party handshake usually deployed in WLANs (also called the 4-way-handshake protocol). The main goals of that work is to prove that a handshake consisting of three separate executions of different AKE protocols, relying on different credentials, can still yield a secure channel. This is different from our case, for several reasons. The handshakes involved in the EAP protocol do not interfere with each other and can be treated independently, whereas in mcTLS and our protocol, keys are simultaneously established and distributed to multiple middleboxes. Furthermore, since our middleboxes have finer-grained access rights, record layer security in our setting is more subtle.

Our construction builds on the existing mcTLS protocol [33], which is a novel and a fundamentally different way of thinking about proxying TLS connections. In parallel to our work, a different variant of mcTLS, called middlebox TLS (mbTLS) [32], has been recently proposed. Like our protocol, mbTLS tries to keep more of the TLS handshake unchanged. However, mbTLS appears more focused on the network-architecture than on (formal) security guarantees. It also dedicates a significant part of its thrust to non-security properties (e.g., deployability, backwards compatibility). Different to our

design, mbTLS looks at achieving certain security guarantees between two hops via the use of attestation, and at integrating the latter in the TLS handshake. But, unlike both mcTLS and mbTLS, we show how our design can be concretely instantiated with TLS 1.3, and we prove our proposal secure formally.

APPENDIX B. INSTANTIATING OUR DESIGN WITH TLS 1.3

Fig. 5 shows how to optimally instantiate our ACCE-AP-secure design called Π , with TLS 1.3⁸.

We presented our modular, ACCE-AP-secure design called Π for 1 middlebox, in an abstract manner, using generic authenticated key-exchange protocols. We will now show how this translates into a more concrete design, when the underlying authenticated key-exchange protocol used as the building block is TLS 1.3; to illustrate the computations therein, we will focus more on the server party S . I.e., one server-instance will “liaise” directly with the middlebox (in Figure 5, computations by this server-instance are written in blue). Another server-instance will “liaise” with the client (in Figure 5, computations by this server-instance are written in black). Similarly, one client-instance will “liaise” directly with the middlebox (written in red, on Figure 5), and one will “liaise” with the server (written in black, on Figure 5). For *each* instance, the normal computations for TLS1.3 take place: a handshake transport key htk is derived (in two stages: first, the TLS1.3 handshake secret hs is derived using $HKDF.extract(\cdot)$ over the TLS1.3 early secret es and using as secret input e.g. the value $g^{yz'}$ produced by S ’s blue instance; second, the $HKDF.expand(\cdot)$ is used over hs , a finished key fk is derived (fk is derived using $HKDF.expand(\cdot)$ of the base-key), and finally a transport key tk and export keys ek are derived (as explained in the TLS 1.3 draft). As per TLS1.3, after each handshake transport key htk is derived, the TLS 1.3 handshake messages thereafter are encrypted using htk , and the finished messages are encrypted with fk . In Figure 5, we use the notation $AE(htk; \cdot)$ to denote the authentication encryption of a message under the appropriate handshake traffic key. We can see that the ACCE part of our design, instantiated herein via TLS1.3 handshake, ends in each session –as expected– with the Finished messages. Clearly, the transport keys tk on Fig. 5 correspond to what in our generic ACCE-AP-secure construction Π was denoted as the ck keys.

Then, the *binding phase* can start, whereby different t_i (with $i \in \{1, 2, 3\}$) are produced by particular nodes of the communication and are sent across in a particular order, encrypted with the transport key. To this end, we write $AE(tk; \cdot)$ for the authenticated encryption under a relevant traffic key tk . Like in our generic ACCE-AP-secure

⁸The newest draft of the TLS 1.3 handshake can be found at <https://tools.ietf.org/html/draft-ietf-tls-tls13-21>.

MW can store the config and permissions as sent by C and retrieve the binders (notably t_1) from K_{MW}^C . To this end, recall from the main manuscript that K_{MW}^C is composed of a series of keys $K_i = \text{PRF}(\text{ek}_3; \text{ctx}^i; \text{mcH}_C \parallel \text{SHash})$, one for each context i . Importantly, having retrieved t_1 , the middleware can send it encrypted with the transport key it shares with the server to the server, *i.e.* $\text{AE}(\text{tk}; t_1)$ (recall that the color blue simply denotes the session $S - MW$).

From $\text{AE}(\text{tk}; \text{mcH}_C)$ and $\text{AE}(\text{tk}; t_1)$, the server can store t_1 but also check that t_1 and t_3 as seen by C are the same as produced by himself (in the case of t_3) and received by the middlebox (in the case of t_1). From here on, if no fault was found by S , then he will compute the access keys at his end and send them to the middlebox, as done previously by C , but –of course– encrypted with the key tk for the $S - MW$ session, *i.e.* $\text{AE}(\text{tk}; \text{HConfig}, K_{MW}^S)$ (recall that the colour blue simply denotes the session $S - MW$). So, at this point the middleware can confront the configs, permissions and keys as received from C and from S , to check that they are the same. From here, upon success, the configuration acceptance messages can be sent and the application layer exchanges can begin thereafter.

APPENDIX C. AKE & ACCE SECURITY

The AKE and ACCE models. We now briefly review the AKE and ACCE security notions, using the notations of Brzuska et al. [11].

Parties and instances. Both models are defined in the context of a set \mathcal{P} of parties, which is formed of two disjoint subset, the client set \mathcal{C} , and the server set \mathcal{S} . Parties are associated with private keys sk and associated public keys pk , which are certified by certifying authorities. Each protocol session in which a party is involved describes an *instance* of that party; π_i^m denotes the m -th instance of party P_i . Each instance is associated with a set of attributes, the session-specific ones intrinsically pertaining to the instances, whereas the long-term ones being “inherited” by instances from their “owning” parties.

For the *ACCE* model, the attributes are as follows: 1) The **private key** $\pi_i^m.\text{sk} := \text{sk}_i$ and **public key** $\pi_i^m.\text{pk} := \text{pk}_i$ of an instance π_i^m of party P_i ; 2) The **role** $\pi_i^m.\rho \in \{\text{init}, \text{resp}\}$ of P_i , *i.e.*, that of an *initiator* (the party that initiates the execution of the protocol) or of a *responder* (the party responding in that execution); 3) The **session identifier** $\pi_i^m.\text{sid}$ of an instance; 4) The **partner identifier**, $\pi_i^m.\text{pid}$ is either a party identifier P_j or, in unilateral authentication, it can be a label “Client” denoting the client with whom the party P_i believes to be communicating. 5) The **acceptance-flag** $\pi_i^m.\alpha$ is 1 or 0 denoting that party respectively accepts or rejects the partner’s authentication; to begin with and prior to the end of the handshake, this attribute is set to \perp ; 6) The **revealed bit** $\pi_i^m.\gamma$ is 1 or 0 denoting that party has had its

session keys revealed or not; to begin with, 0 is assigned to this attribute; 7) The **channel-key**, $\pi_i^m.\text{ck}$ is set to a non-null bitstring when π_i^m ends in an accepting state; prior to the successful end of the session, this attribute is set to \perp ; 8) The **left-or-right** bit $\pi_i^m.\text{b}$, sampled at random when the instance is generated and used in the key-indistinguishability and channel-security games; 9) The **transcript** $\pi_i^m.\tau$ of the instance, containing the suite of messages received and sent by this instance, as well as all public information.

A crux notion of the ACCE models is that of *partnering*: two instances π_i^m and π_j^n are *partnered* if their identifiers $\pi_i^m.\text{sid}$ and $\pi_j^n.\text{sid}$ are equal and non- \perp .

ACCE/AKE games and adversarial queries. In the ACCE model, a MiM adversary can interact with parties in concurrent or sequential session. His interaction is formally captured via a series of oracles. *I.e.*, the $\text{NewSession}(P_i, \rho, \text{pid})$ oracle-query produces a new instance of P_i . Sending a message mes to such an instance π_i^m is modelled via the MiM issuing a $\text{Send}(\pi_i^m, \text{mes})$ query. The $\text{Corrupt}(P_i)$ queries encapsulate the MiM learning party P_i ’s secret key. Querying $\text{Reveal}(\pi_i^m)$ denotes the MiM finding out the channel keys for an accepting instance π_i^m .

In the AKE models, the MiM has access to all the above ACCE queries, as well as to the $\text{Test}(\pi_i^m)$ oracle. On such a query, the output is either the real channel keys $\pi_i^m.\text{ck}$ computed by the accepting instance π_i^m or some random keys of the same length.

To define ACCE/AKE security on top of a MiM adversary \mathcal{A} having access to all the oracles above, we also need to employ the notion of session freshness.

Session freshness. A session π_i^m is *fresh* with intended partner P_j , if the following holds: 1) upon the last query of the adversary \mathcal{A} , the uncorrupted instance π_i^m has finished its session in an accepting state, with $\pi_i^m.\text{pid} = P_j$; 2) P_j is uncorrupted; 3) no Reveal query was made for the π_i^m, π_j^n .

AKE Entity Authentication (EA). In the entity authentication game, the adversary queries the $\text{NewSession}(P_i, \rho, \text{pid})$, $\text{Send}(\pi_i^m, \text{mes})$, $\text{Reveal}(\pi_i^m)$, and $\text{Corrupt}(P_i)$ oracles and its goal is to make one instance, π_i^m , end in an accepting state, with partner ID P_j , which must be uncorrupted, such that no other unique instance of P_j partnering π_i^m exists. The adversary’s advantage in this game is its winning probability.

AKE Key-indistinguishability (KI). In the KI game, \mathcal{A} makes queries to the $\text{NewSession}(P_i, \rho, \text{pid})$, $\text{Send}(\pi_i^m, \text{mes})$, $\text{Reveal}(\pi_i^m)$, and $\text{Corrupt}(P_i)$ oracles and makes a single Test query for a *fresh* instance π_i^m . It wins if it can guess this party’s correctly randomly-chosen bit $\pi_i^m.\text{b}$. The adversary’s advantage is the absolute difference between its winning probability and $\frac{1}{2}$. The AKE security of a protocol is given by the sum of the advantages an adversary has to break either of the two properties, EA or KI.

ACCE security. In the ACCE model, the adversary is

slightly different to the AKE adversary: 1) it does not have access to the Test oracle; 2) it has access to the oracles $\text{Encrypt}(\pi_i^m, l, M_0, M_1, H)$ and $\text{Decrypt}(\pi_i^m, C, H)$, which simulate the access to the secure channel established by party instances. Intuitively, if the Decrypt oracle is queried on an adversarially created ciphertext (without using the Encrypt oracle), it will implicitly reveal the value of the bit b .

ACCE security is defined in terms of two games: the EA game presented as in the ACCE model, and the following game expressing the requirement of security for the established channel.

ACCE Security of the Channel (SC). In this game, the adversary \mathcal{A} can use the $\text{NewSession}(P_i, \rho, pid)$, $\text{Send}(\pi_i^m, mes)$, $\text{Reveal}(\pi_i^m)$, and $\text{Corrupt}(P_i)$ oracles. For a fresh instance π_i^m , \mathcal{A} must output, the bit $\pi_i^m.b$ of that instance. The adversary's advantage is the absolute difference between its winning probability and $\frac{1}{2}$.

ACCE-security for TLS. It is known that TLS 1.2 in DHE mode is (S)ACCE secure with unilateral authentication and ACCE secure if the mutual authentication mode is used [20], [25]. A recent compositional result by Brzuska et al. [11] describes the construction of AKE-secure export keys from an ACCE-secure key-exchange protocol such as TLS 1.2. We note, however, that these results only prove the security of fresh sessions, in which both partners are honest and legitimate. In particular, the fact that the sessions keys are computed only on the nonces (and not, e.g., over the entire session transcript incl. the exchanged certificates) implies a lack of uniqueness. In particular, a malicious, but legitimate server can force two different sessions (one with a client and another with an honest server) to share keys [6]. In addition, moreover, the client may also be fooled, by a malicious server, into thinking it is negotiating a handshake with a different, honest server. In our case, such attacks are particularly relevant, since we wish to compose handshakes, in a scenario in which some of the participants *may* be malicious. This weakness extends to the export-key construction of Brzuska et al. [11].

For TLS 1.3 (up to and including draft 10), Dowling et al. [14] have proved the (multi-stage) AKE security of the keys output in TLS 1.3. Although the protocol has changed since, it can still be inferred that the full mode of TLS 1.3 provides the same level of security for the obtained keys. Note that by construction, TLS 1.3 offers the possibility of computing export keys, and the proof of Dowling et al. also postulates their indistinguishability from random. As opposed to TLS 1.2, the TLS 1.3 handshake does provide the uniqueness of session keys, even in the presence of malicious, legitimate servers. This is because the TLS 1.3 key material is always computed on the entire session hash, including the authenticating information provided by the client and server.

On a different note, whereas in TLS 1.2 the server can simply choose DHE parameters as it wishes, in TLS 1.3 it is

the client that chooses the DHE groups from amongst a list of possible groups that are considered secure. Although this does not impact our result, it does mean that our multicontext protocol is secure when instantiated with *any* of those groups. Should we wish to instantiate it with TLS 1.2, the security of the multicontext handshake would depend on which DHE parameters the server chooses.

APPENDIX D.

FULL THEOREMS AND PROOFS

Full Theorem 1. Let P be a protocol following the mc - Π scheme. Let Π_1 denote a unilaterally-authenticated AKE protocol run between clients and middleware. Let Π_2 denote a mutually-authenticated AKE protocol run between middleware and servers. Finally, let Π_3 be a unilaterally authenticated AKE protocol run between clients and servers.

I. If the protocols Π_1, Π_3 are 2-SACCE secure and Π_2 is 2-ACCE secure, then P guarantees mc – ACCE entity-authentication.

II. If the protocols Π_1, Π_3 are 2-SACCE secure and P_2 is 2-ACCE secure, then P guarantees mc – ACCE channel-security.

Proof: I. We start by proving the entity-authentication first; this proof almost follows directly from the security of each underlying protocol.

\mathbb{G}_0 : This is the original entity authentication game;

\mathbb{G}_1 : In this game, the adversary guesses the party P_i for which the adversary will win, as well as its partner P_j , either a middlebox or a server. Denoting by N the total number of parties, this yields a loss of a factor $\frac{1}{N^2}$, notably:

$$\mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_0] \leq \frac{1}{N^2} \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_1].$$

Winning \mathbb{G}_1 : In the following we show that if there exists an adversary \mathcal{A} winning \mathbb{G}_1 with some non-negligible probability, then we can construct one of the following adversaries, winning with non-negligible probability as well: an adversary $\mathcal{B}^{2\text{-SACCE}}$ against the 2-SACCE property of Π_1 , an adversary $\mathcal{B}^{2\text{-ACCE}}$ against the 2-ACCE property of Π_2 , or an adversary $\mathcal{C}^{2\text{-SACCE}}$ against the 2-SACCE property of Π_3 .

We construct one of those adversaries depending on the identities of the guessed P_i, P_j . There are three possibilities: (1) $P_i \in \mathcal{C}, P_j \in \mathcal{MW}$ (and we construct $\mathcal{B}^{2\text{-SACCE}}$); (2) $P_i \in \mathcal{C}, P_j \in \mathcal{S}$ (and we construct $\mathcal{C}^{2\text{-SACCE}}$); and (3) $P_i \in \mathcal{MW}, P_j \in \mathcal{S}$ or $P_i \in \mathcal{S}, P_j \in \mathcal{MW}$ (constructing $\mathcal{B}^{2\text{-ACCE}}$). We describe only the first of these reductions, since the others follow the same structure.

Assume now that $P_i \in \mathcal{C}, P_j \in \mathcal{MW}$, and let \mathcal{A} be an adversary against the mc -EA property of the protocol Π in Figure 4. The adversary $\mathcal{B}^{2\text{-SACCE}}$ must simulate \mathcal{A} 's environment.

Note that the authentication-acceptance bit in our mc-ACCE protocol Π is set before the configuration-acceptance bit. In other words, an instance might accept the authentication of its partner, but not the configuration that was proposed.

The adversary $\mathcal{B}^{2\text{-SACCE}}$ begins its simulation by corrupting all entities except P_i, P_j . Every time \mathcal{A} queries its NewSession oracle for a given configuration, the adversary $\mathcal{B}^{2\text{-SACCE}}$ proceeds as follows:

- If neither P_i , nor P_j are part of that configuration, $\mathcal{B}^{2\text{-SACCE}}$ simulates creating instances of those parties. Throughout the game, for each Corrupt, Reveal, and Send query, $\mathcal{B}^{2\text{-SACCE}}$ will respond using the parameters it learned by corruption.
- If at least one of P_i, P_j are part of the configuration, the adversary $\mathcal{B}^{2\text{-SACCE}}$ queries its 2-party NewSession oracle twice per each of party. Since P_i is a client, if the configuration includes that party, $\mathcal{B}^{2\text{-SACCE}}$ must query it once with its middlebox-partner in the configuration (whether P_j or some other middleware), and once with the server in the configuration. For the middleware P_j , if the configuration includes it, $\mathcal{B}^{2\text{-SACCE}}$ must query NewSession twice: once, with the partner identifier for this currently created session being the client which is P_j 's partner in the configuration (either P_i or another client), and once with the partner identifier for this currently created session being the server in the configuration.

For the simulation in the second bullet point, it is important to recall that $\mathcal{B}^{2\text{-SACCE}}$ knows all the long-term secrets of all the parties apart from P_i, P_j , and hence can perfectly simulate their messages (up to and including the key-computation). Thus, whenever \mathcal{A} queries Send to instances of either P_i or P_j , $\mathcal{B}^{2\text{-SACCE}}$ queries its own Send oracle with the same input as \mathcal{A} . By contrast, whenever \mathcal{A} queries that oracle for instances of any other party, $\mathcal{B}^{2\text{-SACCE}}$ simulates the correct response. Note that \mathcal{A} may not query Corrupt on either P_i , nor P_j , but it may query Reveal arbitrarily. Whenever \mathcal{A} queries Corrupt, $\mathcal{B}^{2\text{-SACCE}}$ has the correct response (since it has corrupted all parties that are not P_i, P_j). Similarly, $\mathcal{B}^{2\text{-SACCE}}$ can use its Reveal oracle on any instance of P_i, P_j and otherwise can reply with the honestly generated values.

We note that this simulation is perfect, and thus, in the case that $P_i \in \mathcal{C}$ and $P_j \in \mathcal{M}\mathcal{W}$, $\mathcal{B}^{2\text{-SACCE}}$ wins with equal probability to \mathcal{A} .

In the same way, we can construct adversaries $\mathcal{B}^{2\text{-ACCE}}$ or $\mathcal{C}^{2\text{-SACCE}}$. In particular, note that we can construct a single adversary $\mathcal{B}^{2\text{-ACCE}}$ for both the case that $P_i \in \mathcal{M}\mathcal{W}, P_j \in \mathcal{S}$ and for the inverse case, since the simulation is symmetric in the 2-ACCE case.

As a result, we can write:

$$\begin{aligned} \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_0] &\leq \frac{1}{N^2} \cdot \left(\mathbb{P}[\mathcal{B}^{2\text{-SACCE}} \text{ wins 2-SACCE of } P_1] \right. \\ &\quad \left. + \mathbb{P}[\mathcal{B}^{2\text{-ACCE}} \text{ wins 2-ACCE of } P_2] \right. \\ &\quad \left. + \mathbb{P}[\mathcal{B}^{2\text{-SACCE}} \text{ wins 2-SACCE of } P_3] \right). \end{aligned}$$

This can also be written as:

$$\text{Adv}_{\Pi}^{mc\text{-EA}}(\mathcal{A}) \leq \frac{1}{N^2} \left(\text{Adv}_{P_1}^{2\text{-SACCE}}(\mathcal{B}_1) + \text{Adv}_{P_2}^{2\text{-ACCE}}(\mathcal{B}_2) + \text{Adv}_{P_3}^{2\text{-SACCE}}(\mathcal{B}_3) \right),$$

denoting by $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ the adversaries $\mathcal{B}^{2\text{-SACCE}}, \mathcal{B}^{2\text{-ACCE}}, \mathcal{C}^{2\text{-SACCE}}$ respectively. This ends the proof for the mc-EA property, I., in the text of the theorem.

II. The proof for the the security of the channel is analogous to that of the mc-EA theorem above. \blacksquare

We note that the mc-CS property only indicates the security of the channel in the presence of honest participants to the handshake. No guarantees are given w.r.t. the confidentiality, integrity, and authenticity of messages in the presence of malicious parties.

Full-text Theorem 2, configuration soundness — partnering. Let P be a protocol following the $mc\text{-}\Pi$ scheme. Let Π_1 denote the unilaterally-authenticated TLS 1.2 protocol run between clients and middleware. Let Π_2 denote the mutually-authenticated TLS 1.2 protocol run between middleware and servers. Let Ψ be the transformation of Π_1 and Π_2 into AKE protocols by respective computation of the export keys denoted ek , indexed per session. If the following conditions hold:

- Π_1, Π_3 are 2-SACCE secure and P_2 is 2-ACCE secure;
- If $\Psi(P_1), \Psi(P_2), \Psi(P_3)$ yield pseudorandom keys;
- If the KDF used to compute the binders t_1, t_2, t_3 is a secure pseudorandom function;
- The hash function used for the computation of the keys is collision-resistant.

Then, P guarantees $mc\text{-ACCE}$ configuration soundness with respect to partnering.

Proof: Our proof has the following game hops.

\mathbb{G}_0 : This is the original game for configuration soundness-partnering. The adversary \mathcal{A} has an advantage denoted by $\text{Adv}_{\Pi}^{\text{HConfig}}(\mathcal{A})$.

\mathbb{G}_1 : We modify \mathbb{G}_0 by making \mathcal{A} now guess the participants in the challenge configuration for which the game will be won. This yields a loss of a factor equal to $\frac{1}{N^3}$

$$\mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_0] \leq \frac{1}{N^3} \cdot \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_1].$$

\mathbb{G}_2 : In this game, we exclude impersonations of either the middlebox or the server in the challenge configuration (in particular note that the session identifiers of two independent sessions are quasi-unique). Note that this automatically rules out an adversary winning by making the endpoints accept for a configuration different than HConfig. Since we have already guessed the challenge entities (see the mc-EA proof), the security loss is translated to the following:

$$\begin{aligned} \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_1] &\leq \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_2] \\ &\quad + \mathbb{P}[\mathcal{B}_1 \text{ wins 2-SACCE of } \Pi_1] \\ &\quad + \mathbb{P}[\mathcal{B}_2 \text{ wins 2-ACCE of } \Pi_2] \\ &\quad + \mathbb{P}[\mathcal{B}_3 \text{ wins 2-SACCE of } \Pi_3]. \end{aligned}$$

\mathbb{G}_3 : We modify \mathbb{G}_2 to replace the export keys ek output by $\Psi(P_1), \Psi(P_2), \Psi(P_3)$ for the challenge configuration (for as many instances of it as we run) by truly random, but consistent values (*i.e.*, two entities sharing a given ek will still share the random value yielded instead). Furthermore, we assume these values are unique per session (since we have already ruled out collisions of sessions). Thus:

$$\begin{aligned} \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_2] &\leq \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_3] + \text{Adv}_{\Psi(\Pi_1)}^{\text{S-AKE}}(\mathcal{B}_1) \\ &\quad + \text{Adv}_{\Psi(\Pi_2)}^{\text{AKE}}(\mathcal{B}_2) + \text{Adv}_{\Psi(\Pi_3)}^{\text{S-AKE}}(\mathcal{B}_3) \\ &\quad + 2 \times \binom{q_{\text{NewSession}}}{2} \times 2^{-|ek|}. \end{aligned}$$

\mathbb{G}_4 : We modify \mathbb{G}_3 to replace all binding identifiers for the challenge configuration by truly random values.

$$\mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_3] \leq \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_4] + \text{Adv}_{\text{PRF}}^{\text{PRF}}(\mathcal{B}_5),$$

where \mathcal{B}_5 is an attacker against the PRF-assumption on the KDF producing the binders t_1, t_2, t_3 .

Winning \mathbb{G}_4 : We now note that instances running in the challenge configuration HConfig have truly random binding identifiers, and we have excluded entity authentication. We now rule out adversaries that win by ensuring that either the client, or the server, have a wrong binding identifier in their master binding identifier attribute.

We first rule out the possibility that \mathcal{A} breaks the channel security of any of the three sessions, thus losing a term

$$\begin{aligned} &\mathbb{P}[\mathcal{B}_6 \text{ wins 2-SACCE of } \Pi_1] \quad + \\ &\mathbb{P}[\mathcal{B}_7 \text{ wins 2-ACCE of } \Pi_2] \quad + \\ &\mathbb{P}[\mathcal{B}_8 \text{ wins 2-SACCE of } \Pi_3], \text{ i.e. we have that:} \end{aligned}$$

$$\begin{aligned} \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_4] &\leq \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_3] \\ + \mathbb{P}[\mathcal{B}_6 \text{ wins 2-SACCE of } \Pi_1] &+ \end{aligned}$$

$$\begin{aligned} &\mathbb{P}[\mathcal{B}_7 \text{ wins 2-ACCE of } \Pi_2] \quad + \\ &\mathbb{P}[\mathcal{B}_8 \text{ wins 2-SACCE of } \Pi_3], \text{ where} \end{aligned}$$

\mathcal{B}_6 is the channel-security attacker (in the 2-SACCE game) against Π_1 , \mathcal{B}_7 is the channel-security attacker (in the 2-ACCE game) against Π_2 , and \mathcal{B}_8 is the channel-security attacker in the 2-SACCE game against Π_3 .

Thus far, we have ruled out that the adversary can insert wrong binders into the transmissions between the parties in the challenge configuration, and the advantage of the adversary $\mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_4]$ is so far bounded as per the above. However, at this point in our proof, one could conceive that a network adversary could still try to mismatch instances from different handshakes. We argue that this is not possible, since that would be equivalent to a breach in channel security or an impersonation attack, *i.e.*,

$$\begin{aligned} \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_4] &\leq \mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_3] + \\ &\mathbb{P}[\mathcal{B}_6 \text{ wins 2-SACCE of } \Pi_1] \quad + \\ &\mathbb{P}[\mathcal{B}_7 \text{ wins 2-ACCE of } \Pi_2] \quad + \\ &\mathbb{P}[\mathcal{B}_8 \text{ wins 2-SACCE of } \Pi_3] \\ &+ \text{Adv}_{\Pi}^{\text{mc-EA}}(\mathcal{A}). \end{aligned}$$

The adversary now can still win by ensuring that the middlebox rejects this configuration, whereas the endpoints accept it. However, this is not possible in our design (by our construction of the configuration-acceptance process, *i.e.*, see the last 2 exchanges C – MW, MW – S in our design, before the record-layer transmissions start). In other words, $\mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_4]$ is 0, and – by the above, $\mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_0]$ is no bigger than $\mathbb{P}[\mathcal{A} \text{ wins } \mathbb{G}_4]$ but for some negligible amount.

This completes our proof. \blacksquare

Full-text Theorem 2, configuration soundness — the malicious-MW case.

Let P be a protocol following the mc - Π scheme. Let P_1 denote the unilaterally-authenticated TLS 1.2 protocol run between clients and middleware. Let P_2 denote the mutually-authenticated TLS 1.2 protocol run between middleware and servers. Let Ψ be the transformation of P_1 and P_2 into AKE protocols by respective computation of the export keys denoted ek , indexed per session. If the following conditions hold:

- P_1, P_3 are 2-SACCE secure and P_2 is 2-ACCE secure;
- If $\Psi(P_1), \Psi(P_2), \Psi(P_3)$ yield pseudorandom keys;
- If the KDF used to compute the binders t_1, t_2, t_3 is a secure pseudorandom function;
- The hash function used for the computation of the keys is collision-resistant.

Then, P guarantees mc – ACCE configuration soundness with respect to malicious middleware.

Proof: We use essentially the same strategy as above to prove the indistinguishability from random of the context

keys not revealed to the malicious middlebox. The strategy goes as follows:

- \mathbb{G}_0 : The original game of configuration soundness in the presence of a malicious middlebox.
 - \mathbb{G}_1 : We again guess the configuration for which the adversary will win, losing a factor $\frac{1}{N^3}$.
 - \mathbb{G}_2 : We assume unique session identifiers.
 - \mathbb{G}_3 : We replace ek , and then the binders, by truly random, consistent, and unique values.
 - \mathbb{G}_4 : We assume the security of the master-session channel (note that neither endpoint may be corrupted).
 - \mathbb{G}_5 : We replace the context keys computed by the endpoints by truly random, independent, though consistent values. Since the export keys are unique and truly random, at this point we lose the advantage of the key-derivation algorithm PRF.
- At this point, the adversary can only win by guessing. ■

Full-text Theorem 2, configuration soundness — the malicious-endpoint case. Let P be a protocol following the mc -II scheme. Let P_1 denote the unilaterally-authenticated TLS 1.2 protocol run between clients and middleware. Let P_2 denote the mutually-authenticated TLS 1.2 protocol run between middleware and servers. Let Ψ be the transformation of P_1 and P_2 into AKE protocols by respective computation of the export keys denoted ek , indexed per session. If the following conditions hold:

- P_1, P_3 are 2-SACCE secure and P_2 is 2-ACCE secure;
- If $\Psi(P_1), \Psi(P_2), \Psi(P_3)$ yield pseudorandom keys;
- If the KDF used to compute the binders t_1, t_2, t_3 is a secure pseudorandom function;
- The hash function used for the computation of the keys is collision-resistant.

Then, P guarantees mc – ACCE configuration soundness with respect to a malicious endpoint.

Proof: The proof here relies on the fact that an honest middlebox will reject configurations in which the client and server supply it with a different subset of keys. Note, however, that we only assume one corrupted entity (either the server, or the client) for the duration of the game. We will need to guess which entity will be malicious (the client, the server, or none). If no party is corrupted, the security statement follows essentially from the security of the partnering. If exactly one party is corrupted, then the proof is much easier, since we rely on the definition of the configuration-acceptance bit (in this latter case, we can no longer assume the security of any keys which the corrupted party knows). ■

Theorem 3, record-layer security — the 1-MW case. For any handshake of the protocol mc II in Figure 4, for all instances ending in an accepting state w.r.t. the configuration, consider a message-fragment m (with respect to the established contexts) that is sent from one endpoint (which we will call *sender*) to the other endpoint (here called the

receiver). Assume that in one handshake, the receiver receives a ciphertext c^* , for which it accepts the authentication of the message, decrypting it to a message fragment m^* . Then: 1) If for that handshake the middlebox is corrupted, but has read-only permissions for the context corresponding to the fragments m, m' , then $m^* = m$. 2) If the sender is corrupted, and the middlebox as well as the receiver are honest, then $m^* = MW(m)$.

Proof: This proof is straightforward. We start by proving the first point. The configuration was accepted, the mc II is configuration-sound (w.r.t. malicious middleware) and authentication within c^* succeeded upon receipt by the receiver, so c^* did get to the receiver from MW. The configuration was accepted and sound (w.r.t. malicious middleware) so MW did get the message m (encrypted/MAC-ed) from the honest sender. On top, MW only has read rights, so MW cannot modify m . So, c^* must decrypt to m , i.e. $m^* = m$.

We now prove the second point. The configuration was accepted, the mc II is configuration-sound (w.r.t. malicious endpoint) and the permission in question is w.r.t. m 's context. Given these, when the message arrives from the sender to MW, if MW does not abort, then the honest MW applies his transformation to the context in course which is w.r.t. m , i.e. MW produces $MW(m)$ and sends it on to the receiver. Moreover, if the receiver gets c^* and does not abort, then it must have come from MW (i.e., the malicious sender cannot bypass the middleware due to the interleaving of messages and the double encryption on the record layer). So, the honest receiver, who has accepted the sound configuration, will receive and authenticate the receipt of the record-layer, access-controlled encryption of $MW(m)$ from MW, i.e. $m^* = MW(m)$. ■

Theorem 3, record-layer security — the multi-MW case. Consider an mc ACCE protocol P featuring arbitrarily many middleboxes in each session. For any handshake, for all instances which end in an accepting state with respect to the configuration, consider a message fragment m (with respect to the established contexts) that is sent from one endpoint (which we will call *sender*) to the other endpoint (here called the *receiver*). Assume that in one handshake, the receiver accepts a ciphertext c^* , for which it accepts the authentication of the message, decrypting it to a message fragment m^* . Let that handshake have $k \geq 2$ middleboxes, indexed MW_1, \dots, MW_k , in ascending order from the sender to the receiver. We can guarantee the following properties:

- A. If for that handshake, the configuration is sound and all corrupted middleware have at most read-only permissions, then $m^* = m$. More precisely, 1) If the sender is corrupted and all middleware have (at most) read-only rights, then all middleboxes have performed, in chronological order, their functionalities on m , and $m^* = m$. 2) All middleboxes between the sender and the first corrupted middleware have performed,

in increasing order of the indeces from sender-to-receiver, their functionalities on m , and $m^* = m$. 3) All middleboxes between the last corrupted middleware and the receiver have performed, in increasing order of the indeces from left-to-receiver,, their functionalities on m , and $m^* = m$.

- B. For that handshake, let some corrupted parties have write access for the fragment m . Orienting an order from the sender to the receiver, denote by MW_ℓ the last middleware to be corrupted and by m' the message that MW_ℓ has encrypted and sent to its neighbour $MW_{\ell+1}$. Then, it holds that
- $$m^* = MW_k(\dots(MW_{\ell+2}(MW_{\ell+1}(m')))).$$

Proof: This proof is straightforward. We start by proving the first sub-point, A.1. The configuration was accepted, the $mc\Pi$ is configuration-sound (in this case, w.r.t. malicious endpoint) and the permission in question is w.r.t. m 's context. Given these, when the message arrives from the sender to the first MW, if this MW does not abort, then MW applies his transformation to the context in course which is w.r.t m , i.e. MW produces $MW(m)$ which leaves m unchanged since this MW only has (at most) read rights. Then, this MW sends m (encrypted) onto the next MW or, if there is none, to the receiver. If there is an inner MW, the same arguments that held for the previous MW hold for this one. So, over authenticated exchanges and accepted, sound configuration, the receiver will eventually decrypt his received message to m , i.e. $m^* = m$.

Sub-points A2 and A.3 just follow as per A.1 above from the configuration soundness (in this case, w.r.t. malicious MW) and the fact that in middleboxes cannot be bypassed due to the double encryption we use at the record-layer.

We now prove the second point, B. Let $i \geq 0$ such that there are middleware $MW_\ell, \dots, MW_{\ell+i}, \dots, MW_k$ in the configuration "sender" – MW_ℓ – $MW_{\ell+i}$ – \dots MW_k at hand, i.e., there are at least 2 MW in the configuration. We do this part of the proof by induction on the length of the configuration.

Let $i=0$ for the base step. The configuration was accepted, the $mc\Pi$ is configuration-sound (in this case, w.r.t. malicious MW). So, over authenticated exchanges and accepted, sound configuration, MW_k receives an inner-message which is m' from MW_ℓ , m' corresponds to the right context under discussion (e.g. the HTTP header to read/write over). If MW_k does not abort, since MW_k is honest (as MW_ℓ was the last corrupted MW in the sender-to-receiver ordered configuration-list), then MW_k applies his transformation to the m' yielding $MW_k(m')$ and passing this on to the sender. So, $m^* = MW_k(m')$.

Let $i \geq 1$ for the inductive step. Let $MW_{\ell+i}$ receive a message, with MW_ℓ being the last corrupted MW in the ordered list $MW_\ell, \dots, MW_{\ell+i}, \dots, MW_k$. The configuration was accepted, the $mc\Pi$ is configuration-sound (in this

case, w.r.t. malicious MW). So, over authenticated exchanges and accepted, sound configuration, $MW_{\ell+i}$ receives an inner-message which is $MW_{\ell+(i-1)}(m')$, and this message corresponds to the right context under discussion (e.g. the HTTP header to read/write over). If $MW_{\ell+i}$ does not abort, since $MW_{\ell+i}$ is honest (since it is closer than MW_ℓ to the receiver), then $MW_{\ell+i}$ applies his transformation to the $MW_{\ell+(i-1)}(m')$, yielding $MW_{\ell+i}(MW_{\ell+(i-1)}(m'))$ and passing $MW_{\ell+i}(MW_{\ell+(i-1)}(m'))$ on to $MW_{\ell+(i+1)}$ over the configuration.

By induction, this proves our case B of the theorem. ■