

# A Cryptographic View of Deep-Attestation, or how to do Provably-Secure Layer-Linking\*

Ghada Arfaoui<sup>1</sup>, Pierre-Alain Fouque<sup>2</sup>, Thibaut Jacques<sup>2</sup>, Pascal Lafourcade<sup>3</sup>, Adina Nedelcu<sup>1,2</sup>, Cristina Onete<sup>4</sup>, and Léo Robert<sup>3</sup>

<sup>1</sup> Orange Labs

<sup>2</sup> IRISA, University of Rennes 1

<sup>3</sup> Université Clermont-Auvergne, CNRS, Mines de Saint-Étienne, LIMOS

<sup>4</sup> XLIM, University of Limoges

**Abstract.** Deep attestation is a particular case of remote attestation, *i.e.*, verifying the integrity of a platform with a remote verification server. We focus on the remote attestation of hypervisors and their hosted virtual machines (VM), for which two solutions are currently supported by ETSI. The first is single-channel attestation, requiring for each VM an attestation of that VM and the underlying hypervisor through the physical TPM. The second, multi-channel attestation, allows to attest VMs via virtual TPMs and separately from the hypervisor – this is faster and requires less overall attestations, but the server cannot verify the *link* between VM and hypervisor attestations, which comes for free for single-channel attestation.

We design a new approach to provide linked remote attestation which achieves the best of both worlds: we benefit from the efficiency of multi-channel attestation while simultaneously allowing attestations to be linked. Moreover, we formalize a security model for deep attestation and *prove* the security of our approach. Our contribution is agnostic of the precise underlying secure component (which could be instantiated as a TPM or something equivalent) and can be of independent interest. Finally, we implement our proposal using TPM 2.0 and vTPM (KVM/QEMU), and show that it is practical and efficient.

**Keywords:** deep attestation, layer linking, TPM, vTPM

## 1 Introduction

*Network Function Virtualization* (NFV) is a technology that promises to provide better versatility and efficiency in large-scale networks.

---

\* This study was partially supported by the French ANR project ANR-18-CE39-0019 (MobiS5), the French government research program “Investissements d’Avenir” through the IDEX-ISITE initiative 16-IDEX-0001 (CAP 20-25), the IMobS3 Laboratory of Excellence (ANR-10-LABX-16-01), the French ANR project DECRYPT (ANR-18-CE39-0007) and SEVERITAS (ANR-20-CE39-0009).

The core idea is to move from architectures in which physical machines are set up to perform various roles in a network, to a design in which that configuration is done virtually. As such, a machine could be configured and re-configured at distance, and, by judicious use of virtual machines, it could perform a variety of roles within the network infrastructure.

Virtualized platforms are set up in layers, including the following basic components: physical resources, the virtualization layer and infrastructures, *virtualized network functions* (VNFs), and the NFV management and orchestration module.

At the bottom of the infrastructure are real, physical components, meant for computations, storage and physical network functions. The virtualization layer (also called *hypervisor*) manages the mapping between those physical components and virtual equivalents. As such, the virtualized network functions – hosted by virtual machines running inside the NFV infrastructure– never have direct access to the physical resources. Instead, the VNFs access the virtual resources. The NFV management and orchestration module runs the combined infrastructure, including: the lifecycle of the instantiated VNFs, resource allocation for VNFs, or overall management in view of particular, given network services.

**Deep Attestation (DA).** Virtualization enables efficient, versatile remote network configuration and administration; however, the fact that multiple virtual processes share resources can introduce hazards to security. One way to ensure that a component runs correctly is by using *attestation*. Attestation is a process complementary to authentication: whereas the latter allows a platform to prove that it is the entity it claims to be, the former ensures that the platform runs a trustworthy code, *i.e.*, it has not been breached. As described in [11], “*Attestation is the process through which a remote challenger can retrieve verifiable information regarding a platform’s integrity state.*” A property can be for instance software integrity, geolocalisation, access control, etc.

Attestation relies on a *root of trust* (RoT), usually instantiated through a *trusted platform module* (TPM) – or an equivalent mechanism. The root of trust is responsible, amongst other things, for protecting sensitive cryptographic materials (such as private keys) and for running cryptographic operations in an isolated way. The virtualization layer (hypervisor) has direct access to the RoT, but the virtual machines it manages do not; instead they will have access to the RoT by means of *virtual Roots of Trust* (vRoTs). Virtual Roots of Trust are a combination of resources, some provided by the physical RoT, and other managed by the hypervisor, which directs and mediates access to the RoT.

In a nutshell, attestation is a process which allows an independent, remote verifier to check that a target platform still behaves in the desired way. This is done by first authenticating the RoT, then by comparing a measurement of the current state of the component to a presumably-correct state, as indicated in a *Root of Trust for Storage* (RTS).

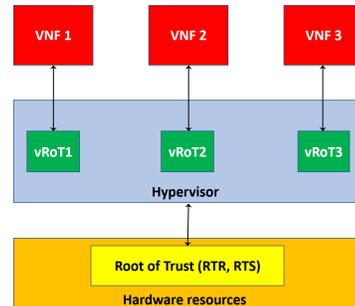


Fig. 1: The setup for DA.

In addition, a guarantee must be given of the correctness of the RTS, which is done by means of a *Root of Trust for Reporting* (RTR). The functionalities of RTS and RTR can effectively be provided by a TPM. A TPM is an example of implementation that could provide RTR and RTS by leveraging the specific tampering detection properties of its Platform Configuration Registers (PCR) and issuing signed reports, or *quotes*, of their content.

We consider the attestation of two types of components: virtual machines (VMs), such as VNFs, and the hypervisor managing them, whose underlying physical component includes a RoT providing an RTR and an RTS. This architecture is depicted in Figure 1.

To verify that the VMs and the hypervisor are running correctly, both these types of components must undergo remote attestation. Not only must they be attested in isolation, but a statement must be attested on the *layer-binding*: notably, we should know which VMs run on which hypervisor. This is known as *deep attestation* (DA). There are two typical ways of achieving deep attestation (as described by ETSI standardization documents [11]): single- and multi-channel VM-Based Deep Attestation.

**Single/Multi-channel Deep Attestation.** In single-channel deep attestation the attestation is run only between the remote verifier and the virtual machines. At each attestation, the VM (by querying its associated virtual TPM, or vTPM) provides not only an attestation for itself, but also the hypervisor it runs on.

Specifically the response forwarded by the VM to the remote verifier includes the (independent) attestation of the hypervisor, and the layer-binding attestation between the VM and its hypervisor. This is depicted in Figure 2, on the left-hand-side. Note that the quotes in this case are both obtained from the (slow) physical TPM.

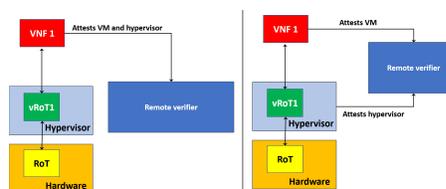


Fig. 2: Single vs multi-channel deep attestation

From the point of view of security, this solution is optimal; however, it scales poorly. Given as few as 1000 VMs running on top of the hypervisor, we would require that the hypervisor be attested 1000 times, once for each VM.

By contrast, in multi-channel deep attestation, the VMs are attested separately and independently from the hypervisor. In this scenario, the VMs attest to the remote verifier, thus proving they were not tampered with. Separately, the hypervisor also attests to the remote verifier. This can be seen on the right hand side of Figure 2. In this case, the efficiency is optimal: for 1000 VMs, we have 1000 VM-attestations and 1 hypervisor attestation. However, there is virtually no layer-binding between the VMs and their hypervisor: there is no guarantee that the VMs are really managed by the hypervisor. An attacker could therefore “convince” a party (such as the owner of the infrastructure) that a VM still exists on a given physical machine when it has, in fact, been removed.

**Our solution.** We take the middle path between single- and multi-channel deep attestation to obtain layer-binding between VMs and hypervisors with reasonable efficiency. Our solution is simple, yet elegant, using standard cryptography to ensure that a hypervisor’s single attestation is linkable to any number of attestations of VMs managed by it. We give three contributions:

**A cryptographic scheme.** Our scheme ensures secure and efficient linked DA. The hypervisor and VMs each attest only once. However, we also embed a list of public keys (associated with the VMs managed by the hypervisor) *within* the hypervisor attestation, which is established by the root of trust. In order to authenticate the list of forwarded keys, we embed them into the attestation nonce, forwarded by the attestation server. If the hypervisor’s attestation verifies, then the attestation server can link that hypervisor with the (subsequently attesting VMs) which use keys in the forwarded list. If the hypervisor’s attestation fails, then the public keys cannot be trusted.

**Provably secure authorized linked attestation.** An important advantage of our approach is that we have a fully-formalized provable-security guarantee. We use a composition-based approach, constructing primitives that are increasingly stronger out of weaker ones. At the basis of our construction is a yea-or-nay *basic attestation scheme*, which we construct so that it is “secure” by assumption. Its functionality is simple: whenever a component is compromised, the basic attestation scheme tells us so by outputting a faulty attestation; whenever the component is honest, the basic attestation outputs a correct attestation. In other words, this basic attestation scheme is a compromise-oracle: when queried it (indirectly) produces a proof of whether a component has been tampered with or not.

Based on this assumption, we build a sequence of cryptographic mechanisms that add security against stronger adversaries. A first step is to build *authenticated attestation*: a scheme which allows us to authenticate the component that provides the attestation, and additionally ensures that this component’s attestations always verify prior to corruption, but fails to verify as soon as a compromise occurs. Then, we consider *linked attestation*: a scheme that introduces the hypervisor-VM relationship described above, and permits not only the verification of individual attestations, but also (publicly) linking attestations.

We construct *authorized linked attestation*: a primitive that permits attestation protocols between a component and authorized attestation server. The latter will also then be able to link certain attestations together. Authorized linked attestation will have three properties: *authorization* (only an authorized server can query an attestation quote); *indistinguishability* (no Person-in-the-Middle adversary can know even a bit of a quote exchanged during a legitimate protocol with probability significantly better than  $\frac{1}{2}$ ); and *linkability* (no one can fool an attestation server into believing that two unlinked components are in fact linked).

**Implementation.** We used a regular laptop equipped with TPM 2.0 (as a root of trust). We set up an architecture with one hypervisor and multiple VMs. The VMs used full virtual TPM as a virtual root of trust. We made over 100

experiments. This showed that our solution is more efficient than single channel approach and adds only insignificant charge (a hash function computation) compared to traditional multi-channel DA.

Our work is, to our best knowledge, the first that attempts to provide a sound cryptographic treatment of deep attestation. In many ways, this is much harder than designing the scheme that we present, because attestation is a generic term comprising an entire class of algorithms that have different goals. As such, we are only scratching the surface here, and believe that –aside from the real, and practical advantages of our presented construction– our cryptographic treatment, primitives, and proofs, may be of independent interest to this line of research.

**Limitations.** A first fundamental limitation is the fact that we assume, in our constructions, the existence of a basic attestation primitive that works infallibly like an oracle, telling us if a component is compromised or not. Attestation is a complex process, allowing for components to attest to many properties *e.g.*; “are you based in the US?”. We explicitly require the strongest attestation possible, which takes into account the state of an entire system, and then has a way of telling whether it has been compromised or not.

We treat the existence of basic attestation as an assumption because we do not see a way of constructing it with cryptographic tools. The cryptography we put on top adds a lot of new properties: authenticity, confidentiality, authorization, linkability, but *not* the simple fact of distinguishing a compromised component from an honest one. Our result should therefore be interpreted as a need for such a scheme to exist, as in fact required by ETSI [11].

Another limitation of our scheme lies in the way we modeled our linked-attestation component. We consider classes of components which can be linked. At the registration of each piece of hardware, a number of subcomponents of each type is indicated – and (unique) keys are given to those components. As a result, we cannot account for having two hypervisors that manage the same VM on a given infrastructure. A future work could be to consider *multi-hypervisor VM* as introduced in [12].

**Related Work.** Many attacks have been recently reported on remote attestation mechanism [9] or 5G standards standards [13]. Many tools such as formal methods or cryptography can be used to model and prove the security of such standards. However, this lack of formalization must today be now addressed otherwise we will have more and more attacks. Provable cryptography is a nice solution to solve this problem since it allows to better understand the security model, what is the adversary goal and its means, which oracle can he query. Some cryptographic primitives have already be nicely formalized such as Direct Anonymous Attestation (DAA) which enables remote authentication of a trusted computer (TPM for instance) while preserving the privacy of the platform’s user in [8] by Brickell et al. It is a group signature without the feature that a signature can be opened, *i.e.*, the anonymity is not revocable. Such primitive are well described using cryptography as a variant of signature scheme. However, provable cryptography has also been used successfully to formalize security protocols

as authenticated key exchange [6, 10]. This is precisely our goal here to model the different security components independently and to compose them to prove the security of a new security mechanism. Indeed, the attestation server must authenticate the whole platform, *i.e.*, the hypervisor and the NFV running on top. This problem has been addressed by others in the context of secure boot or for instance in [5], where the authors propose an attestation mechanism for swarms of device softwares in IoT and embedded environment. Software attestation is different from remote attestation, as said in [4] since it cannot rely on cryptographic secrets to authenticate the prover device. The first to have taken into account deep attestation are Lauer and Kuntze in [14] but their solution miss a security proof and a rigorous analysis.

## 2 Towards authorized linked attestation

Our core contribution provides layer-binding in deep attestation. Cryptographically, we view this as a new primitive, which we call *authorized linked attestation*, built in steps from increasingly-stronger primitives. We use a basic-attestation primitive as a black box. This primitive is an abstraction of the algorithm by which a single party (like a component of a virtualized platform) generates an attestation of its state, given a fresh, honestly-generated nonce. To ease notation, we assume that all the registers are attested at each attestation, and that the property we are attesting is that the entire component has not been compromised.

Authenticated attestation builds on basic attestation by associating parties with identities. The attestation must now no longer indicate whether the party is compromised: it must also authenticate the component. One step further, the linked-attestation primitive built from authenticated attestation will allow two different components to (a) attest their own states; (b) provide auxiliary material that will make two separate attestations linkable.

We also add a new party into the system: the attestation server that serves as a verifier. We then *compose* the linked-attestation primitive with a unilaterally-authenticated authenticated key-exchange protocol, which will authenticate the attestation server and permit the attestation itself to remain confidential with respect to a Person-in-the-Middle (PitM) adversary.

### 2.1 Basic attestation

During basic attestation a single honest party is generated. This party can be later compromised. A quote-generation algorithm will output a quote if the party is still honest at that time, or a special symbol if it is not. Finally a (public) verification algorithm will yield 1 (the component is honest) or 0 (otherwise).

Note that a party such as the one we describe could correspond in practice to a combination of two parties: a virtual entity (like a VM or the hypervisor) and an underlying, uncorruptible, secure part (the TPM), which actually generates

the quote. At this stage, we importantly do not associate these entities with keys as authentication will only appear in our next step (Section 2.2).

What we want to capture, formalized by the security of basic attestation, is the minimal assumption that a compromised component will always yield an attestation that will fail the verification. This is why, when basic attestation is run for a compromised component, it will yield the special symbol  $\aleph$ . We also demand correctness: when a non- $\aleph$  quote is generated, the latter will automatically verify. Our basic attestation component thus becomes the minimal non-cryptographic assumption that we need to make to prove our scheme secure.

**Formalization.** We consider an environment parametrized by a security parameter  $\lambda$ , in which we have a single party  $P$ . This party keeps track of a single *attribute*, namely a *compromise bit*  $\gamma$  originally set to 0. Once this bit is flipped to 1, it can never go back to 0. We define a primitive **BasicAtt** as a tuple of algorithms: (**aBSetup**, **aBAttest**, **aBVerify**):

- **aBSetup**( $1^\lambda$ )  $\rightarrow$  **ppar**: on input the security  $1^\lambda$  (in unary), this algorithm outputs some public parameters **ppar**.
- **aBAttest**(**ppar**)  $\rightarrow$  **quote**: on input the public parameters **ppar**, if  $P.\gamma = 0$ , then this algorithm outputs an attestation quote **quote**  $\neq \aleph$  for  $P$ , and if  $P.\gamma = 1$ , then it outputs  $\aleph$ .
- **aBVerify**(**ppar**, (**quote**  $\cup$   $\aleph$ ))  $\rightarrow$   $0 \cup 1$ : on input public parameters **ppar** and a value that is either a quote denoted **quote** or a special symbol  $\aleph$ , this algorithm outputs a bit. By convention, an output of 0 means the attestation fails, while if the output is 1, the attestation succeeds. We require by construction that for all **ppar**: **aBVerify**( $\cdot, \cdot, \aleph$ ) = 0.

This primitive is also depicted in Figure 3. We assume that if  $P.\gamma = 0$  and **quote**  $\rightarrow$  **aBAttest**(**ppar**), then **aBVerify**(**ppar**, **quote**) = 1.

**Security.** The only security we demand from this primitive is that, if a party is compromised, then its attestation will always fail. This will happen by construction (since this is an assumed primitive) and is embedded in the security model. The adversary  $\mathcal{A}$  will play a game against a challenger  $\mathcal{G}$ . Initially, the challenger sets the system up by running **aBSetup** to output **ppar**. This value is given to  $\mathcal{A}$  as well. The unique party is generated, such that its corrupt bit is set to 1 ( $P.\gamma = 0$ ).

Since  $\mathcal{A}$  now has **ppar**, it can now run the **aBAttest** and **aBVerify** algorithms. In addition, it has access to the **OBAttest** oracle: **OBAttest**()  $\rightarrow$  (**quote**  $\cup$   $\aleph$ ). This oracle calls the **aBAttest**() algorithm for the (corrupted) party  $P$  and returns

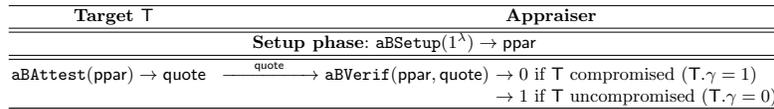


Fig. 3: Basic attestation description with an honestly-generated target. Notice that there is no authentication involved.

the output to the adversary  $\mathcal{A}$ . The challenger stores the result in a database DB. The adversary wins if, and only if, there exists a quote in DB (possibly with  $\text{quote} = \aleph$ ) such that  $\text{aBVerif}(\text{ppar}, \text{quote}) = 1$ . Note that by construction our basic attestation primitive is secure, since once the compromise bit is set, the output is  $\aleph$ , which always yields  $\text{aBVerif}(\text{ppar}, \aleph) = 0$ .

**Basic attestation in reality.** One may wonder at this point what our purpose might be in constructing a security model for a primitive that is by definition correct and secure. We need that security model in our reductions: we will use the attestation primitive to build stronger, linked attestation, and then we will want to make the argument that if an attacker can break the larger primitive, it will also break the smaller primitive. As the smaller primitive is secure by design, this is not possible, and hence, the larger primitive is also secure.

## 2.2 Authenticated Attestation

Basic attestation acts as a foolproof way of telling whether a device is compromised or not. However, the security it provides is very weak. For one thing, it has no authentication guarantees, so potentially one could use a quote that was honestly generated for an honest component to attest a compromised one. Another problem that is more subtle concerns the way components are compromised. Because the basic quotes described in the previous section have no timestamp, nor specific freshness, we cannot take into account adaptive tampering. In the security notion, the party generating the quote is either honest or compromised from the beginning. Yet, ideally we would like a primitive that ensures that a party can start out as honest (and all the quotes generated at that time verify as correct), and later be compromised (and all the quotes generated after that moment will fail). We can do this by deploying cryptographic solutions (two of which are described after the formal definition of the primitive we want to achieve).

A relevant question is why we did not include these security aspects in the basic attestation primitive considered above. To answer this, recall that we have constructed the basic attestation tool to be secure by design. As such, it is an assumption, rather than a solution. To additionally suppose that this primitive also ensures authentication, for instance would be to go against the principle of using minimal assumptions.

**Formalization.** The precise formalization of this primitive is in Appendix A. We consider an environment containing up to  $N$  parties. The parties keep track of the compromise bit  $\gamma$  used also for basic attestation, and a pair of public and secret keys denoted, for each party  $P$ ,  $P.\text{pk}$  (the public key) and  $P.\text{sk}$  (the private key). Intuitively, the security we require for this primitive will be that a valid authenticated quote for a party  $P$  and fresh auxiliary information (used as nonce) is hard to forge by an adversary which knows all the the public information, can register and compromise users, and query an attestation oracle that returns a valid quote or  $\aleph$ . In particular, in a secure scheme, verification should fail if either the authentication or the attestation fails.

**Construction.** We construct an authenticated attestation scheme out of basic authentication, a large set of nonces  $\mathcal{N} := \{0, 1\}^\ell$  (with  $\ell$  chosen as a function of the security parameter  $\lambda$ ), and an EUF-CMA-secure signature scheme  $\text{Sig} = (\text{aSigKGen}, \text{aSigSign}, \text{aSigVerif})$ . We thus instantiate  $\mathcal{AUX} := \mathcal{N}$ , and our  $\text{AuthAtt}$  scheme is as follows:

- $\text{aAuthSetup}(1^\lambda) \rightarrow \text{ppar}$ : this algorithm runs  $\text{aBSetup}(1^\lambda)$  a number  $N$  of times, outputting  $\text{ppar}_1, \text{ppar}_2, \dots, \text{ppar}_N$ . Each time  $\text{ppar}_i$  is created, a party handle  $P_i$  is also created (it will be the party associated with the instance of  $\text{BasicAtt}$  run for those parameters). It sets  $\text{ppar} := (\text{ppar}_1, \text{ppar}_2, \dots, \text{ppar}_N, N)$ , and outputs this value.
- $\text{aAuthKGen}(P_i) \rightarrow (P_i.\text{pk}, P_i.\text{sk})$ : it keeps a counter (starting from 0), which indicates how many times this algorithm has been run. If at the time this algorithm is queried  $\text{counter} < N$ , then  $\text{aAuthKGen}$  runs  $\text{aSigKGen}$  as a black box and outputs the resulting  $(\text{pk}, \text{sk})$  (public and private) keys. It sets  $P_i.\text{pk} := \text{pk}$  and  $P_i.\text{sk} := \text{sk}$ . Party  $P_i$  is then initialized with these keys.
- $\text{aAuthAttest}(\text{ppar}, P.\text{sk}, R) \rightarrow \text{authQuote} \cup \mathbb{N}$ : on input the public parameters  $\text{ppar}$ , a private key  $P.\text{sk}$  of a party  $P$  (which has already been registered), and a value  $R \xleftarrow{\$} \mathcal{N}$ , this algorithm first runs  $\text{quote} \leftarrow \text{aBAttest}(\text{ppar})$ , then the algorithm signs  $\sigma \leftarrow \text{aSigSign}(P.\text{sk}, (\text{quote}, R))$ , that is, it signs a concatenation of the nonce and the obtained quote. The output of this algorithm is  $\text{authQuote} := (\text{quote}, \sigma)$ . If the required party or key does not exist, the value  $\mathbb{N}$  is output by default. If  $\text{quote} = \mathbb{N}$ , then we instantiate  $\text{authQuote} = \mathbb{N}$ .
- $\text{aAuthVerif}(\text{ppar}, P.\text{pk}, R, (\text{authQuote} \cup \mathbb{N})) \rightarrow 0 \cup 1$ : on input public parameters  $\text{ppar}$ , a public key  $P.\text{pk}$  of a party  $P$ , an auxiliary value  $R \in \mathcal{N}$ , this algorithm first checks if the last input is  $\mathbb{N}$ ; if so, the algorithm outputs 0 by default. Else, the algorithm parses  $\text{authQuote} = (\text{quote}, \sigma)$  (with  $\text{quote} \neq \mathbb{N}$  by construction), then runs  $b \leftarrow \text{aSigVerif}(P.\text{pk}, \text{quote}, \sigma)$  and  $d \leftarrow \text{aBVerif}(\text{ppar}, \text{quote})$ . The algorithm outputs  $b \wedge d$  as its response. Notably, 1 is output if, and only if, the signature and the basic attestation verify concomitantly.

**Theorem 21 (Secure Authenticated Attestation)** *The  $\text{AuthAtt}$  scheme is secure assuming that (1)  $\text{BasicAtt}$  scheme is secure (2) the size of  $\mathcal{N}$  is large and (3) the  $\text{Sig}$  signature scheme is EUF-CMA secure.*

The proof is given in Appendix A

### 2.3 Linked Attestation

Authenticated attestation allows the attestation of one (out of many) components, based on that component's unique secret key. If we define now parties as being either VMs or hypervisors, the notion of authenticated attestation suffices to capture the basic guarantees of multi-channel deep-attestation. However, in this paper our goal is to allow parties to *link* their attestations (a hypervisor's attestation should, *e.g.*, be linkable to that of a number of VMs also hosted on that platform).

| Target $T$  | Appraiser   |
|---|---|
| Setup phase: $\text{aAuthSetup}(1^\lambda) \rightarrow \text{ppar}$                           |   |
| $\text{aAuthKGen} \rightarrow (T.\text{pk}, T.\text{sk})$                                     |   |
| $\text{aAuthAttest}(\text{ppar}, T.\text{sk}, \text{aux}) \rightarrow (\text{quote}, \sigma)$ | $\xrightarrow{\text{authQuote}=(\text{quote}, \sigma)} \text{aAuthVerif}(\text{ppar}, T.\text{pk}, \text{aux}, (\text{quote}, \sigma))$<br>$\rightarrow 0$ if $T$ compromised ( $\text{authQuote} = \aleph$ or $\sigma$ invalid)<br>$\rightarrow 1$ if $T$ uncompromised ( $\text{authQuote} \neq \aleph$ and $\sigma$ valid) |

Fig. 4: Authenticated attestation built upon basic attestation (Figure 3).

In this section we describe our next primitive: linked attestation. The latter takes place in an environment where several parties are registered in a linked way – this corresponds to a single platform. A first step is platform registration, by which several parties are linked on the same underlying hardware. Each entity later generates a linkable attestation – verifiable on its own, and linkable with other linkable attestations.

Although our application scenario is that of linking VM and hypervisor attestations, we make our framework more generic than that. Instead of just two types of components, we consider linkable sets  $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_L$ , which resemble equivalence classes. These sets are defined such that any party in one set (say  $P_{\mathcal{S}_1}$ ) can produce an attestation that is linked to attestations produced by parties in sets  $\mathcal{S}_2, \dots, \mathcal{S}_L$ . We write  $P \diamond Q$  to say that two parties are linked. The relation is reflexive ( $P \diamond P$ ), symmetric (if  $P \diamond Q$ , then  $Q \diamond P$ ), and transitive (if  $P \diamond Q$  and  $Q \diamond R$ , then  $P \diamond R$ ).

We formalize a linked-attestation scheme **LinkedAtt** as a tuple of algorithms  $\text{LinkedAtt} = (\text{aLSetup}, \text{aLReg}, \text{aLAttest}, \text{aLVerif}, \text{aLLink})$ , defined for some auxiliary set  $\mathcal{AUX}$ . The detailed formalization is given in Appendix B.

The setup algorithm outputs public parameters  $\text{ppar}$ , including the maximal number  $L$  of sets considered for linking. One can register platforms including subsets of components of each type: this algorithm generates keys for each party. A linked attestation algorithm produces a linked quote  $\text{linkedQuote}$  and an auxiliary linking value  $\text{lkaux}$ . Finally, the verification algorithm checks the attestation in each individual  $\text{linkedQuote}$  and the linking algorithm outputs 1 if several linked attestations seem to belong to the same registered platform, and 0 otherwise. This syntax is also depicted in Figure 5.

The security of linked attestation informally states that an adversary, which has Person-in-the-Middle capabilities and can compromise devices at will, cannot make it appear that two devices are linked when they are not, in fact, so.

A significant limitation on the adversary’s capabilities is that compromising a device will not leak its private keys (which are assumed to be held by a TPM). However, the adversary will gain a limited oracle access to those keys upon compromising the device. The limitations to those queries follow rules of access to an actual TPM.

More formally, we define the security of *linked attestation* as a game  $\text{LinkSec}_{\lambda, F}$  parametrized by a security parameter  $\lambda$  and a set of functions  $F$ , which we call the *permitted key-access functions*. The adversary wins if it is able to make attes-

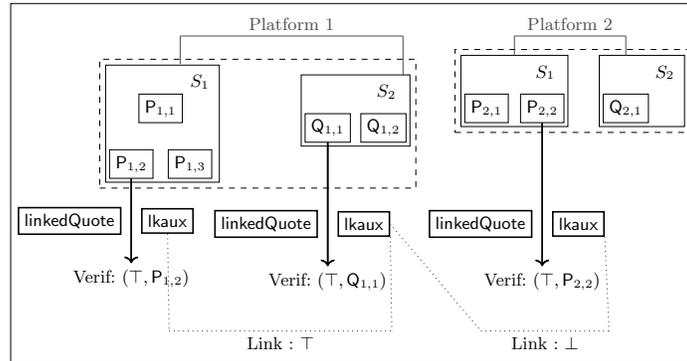


Fig. 5: Linked attestation primitive. The dashed line indicates a platform under the same registration. In this example, both platforms are composed of two subsets (namely  $S_1$  and  $S_2$ ). There are a total of three quote verifications ( $P_{1,2}, Q_{1,1}, P_{2,2}$ ). The link verification outputs true when the devices are registered under the same platform and false otherwise.

tations stored in  $\mathcal{L}_{\text{Att}}$  for parties registered on different platforms ( $P$  and  $Q$ ) link. However, at this point the adversary is constrained to a change-one-change-all kind of game: it cannot, for instance, append an  $\text{lkaux}$  component of its choice to an honestly-generated  $\text{linkedQuote}$ , nor vice-versa.

In the security game, the adversary registers platforms and can compromise some of their components. When a component is compromised, the adversary gets oracle access to a set of permitted functions of the component’s private key. As a result, the strength of the security proof depends on the function space  $F$ . The more functions the adversary is able to query once it compromises a component, the more security our primitive is able to provide. However, note that we cannot give the adversary access to *some* functions, such as the identity function on the component’s private key.

**Construction.** We provide a construction for platforms that have two types of components: virtual machines (VMs) and their managing hypervisor. Thus, in our instantiation,  $L = 2$ . We use an authenticated attestation scheme ( $\text{aAuthSetup}, \text{aAuthKGen}, \text{aAuthAttest}, \text{aAuthVerif}$ ) as a black box. The basic construction is depicted in Figure 6. During setup, our linked-attestation scheme first runs  $\text{aAuthSetup}$  and outputs  $\text{ppar}$  and  $L = 2$ . Note that by construction  $\text{aAuthSetup}$  must output a number  $N$ , denoting the maximal number of parties that can be set up. This counter will represent a global maximum to parties of all types that will exist in our ecosystem. Following setup, one can register a subset of VMs together with a hypervisor. The algorithm runs the key-generation algorithm  $\text{aAuthKGen}$  of the underlying authenticated attestation scheme for each party,

|  |  |   |
|--|--|---|
| <pre> aLSetup(<math>\lambda</math>): ppar' ← aAuthSetup(<math>\lambda</math>) Return ppar ← ppar'  // Attest hypervisor P on platform (s<sub>1</sub>, s<sub>2</sub>) with nonce aux aLAttest(ppar, PK, P.sk, (s<sub>1</sub>, s<sub>2</sub>), aux): Parse PK as PK[1], PK[2] // PK[1] is the set of all VM pks Parse PK[1] as PK<sup>k</sup>, PK<sup>h</sup>, PK<sup>S</sup> // PK<sup>k</sup> contains the keys of all VMs on platform i Set lkaux ← PK<sup>k</sup> with k the index of s<sub>1</sub> in S<sub>1</sub> // lkaux is now the list of all VM keys on that platform aux' ← H(aux  lkaux) // Embed lkaux into a new attestation nonce authQuote ← aAuthAttest(ppar', P.sk, aux') linkedQuote ← authQuote Return (linkedQuote, lkaux) </pre> | <pre> aLReg(s<sub>1</sub>, s<sub>2</sub>): // Registers a platform with a set s<sub>1</sub> of VMs and the hypervisor in s<sub>2</sub> For each i ∈ {1, 2}:   For each j ∈ s<sub>i</sub>:     (P<sub>j</sub>.pk, P<sub>j</sub>.sk) ← aAuthGen(P<sub>j</sub>)   Group all P<sub>j</sub>.pk into PK<sub>i</sub> and all P<sub>j</sub>.sk into SK<sub>i</sub>   Return({PK<sub>1</sub>, SK<sub>1</sub>}, {PK<sub>2</sub>, SK<sub>2</sub>})  // Attesting VM P on platform (s<sub>1</sub>, s<sub>2</sub>) for nonce aux aLAttest(ppar, PK, P.sk, (s<sub>1</sub>, s<sub>2</sub>), aux): Get P.pk matching P.sk from PK lkaux ← P.pk // The linking information is P's public key aux' ← H(aux  lkaux) // Embed lkaux into attestation nonce authQuote ← aAuthAttest(ppar', P.sk, aux') linkedQuote ← authQuote Return (linkedQuote, lkaux) </pre> | <pre> aLVerif(ppar, P.pk, linkedQuote, aux, lkaux): // Verify attestation quote of party P aux' ← H(aux  lkaux); authQuote ← linkedQuote Return aAuthVerif(ppar', P.pk, authQuote, aux')  aLLink(ppar, PK, <math>\Pi_1</math>, <math>\Pi_2</math>): // Link VM quotes from <math>\Pi_1</math> and the hypervisor quote from <math>\Pi_2</math> Initialize AUX<sub>vm</sub> ← <math>\emptyset</math> For each (P<sub>j</sub>.pk, aux, linkedQuote, lkaux) ∈ <math>\Pi_1</math>:   Return 0 if aLVerif(ppar', P<sub>j</sub>.pk, linkedQuote, aux, lkaux) returns 0   Return 0 if lkaux ≠ P<sub>j</sub>.pk   // Linking fails if quotes fail to verify or authenticate each VM   Add lkaux to AUX<sub>vm</sub> // Each lkaux here is a VM public key. Parse <math>\Pi_2</math> as (P<sub>h</sub>.pk, aux, linkedQuote, lkaux) Return 0 if aLVerif(ppar', P<sub>h</sub>.pk, linkedQuote, lkaux) returns 0 AUX<sub>vm</sub> ← lkaux // This lkaux is a list of VM public keys. Return 0 if AUX<sub>vm</sub> is not a subset of AUX<sub>vm</sub> // Linking fails if the hypervisor's list of PKs does not include all VM keys. Return 1 </pre> |
|--|--|---|

Fig. 6: Our linked attestation scheme for platforms with 2 types of components: VMs (stored in  $\mathcal{S}_1$ ) and hypervisors (stored in  $\mathcal{S}_2$ ). Each type of component attests via a different `aLAttest` algorithm, the main difference between them being that the hypervisor embeds a list of public keys in its nonce.

independently (note that this also ensures that the total number of parties remains at most  $N$ ). Finally, keys are grouped by types of parties: keys of VMs are output in a set of public keys  $\text{PK}_1$  and the key of the hypervisor is output as  $\text{PK}_2$ .

The VMs and hypervisor generate linked attestations differently. The hypervisor first fetches the public keys of all the components registered with it on the same platform. It computes a new nonce as the hash of two concatenated values: the original auxiliary value `aux` and the list of the public keys. The component then runs `aAuthAttest` on the public parameters, this new nonce, and its private key, outputting the authenticated quote. By contrast, when a VM attests, it computes a new nonce from the original auxiliary value `aux` and (only) its own public key. The authenticated quote obtained as a result is provided as the VM's linked quote.

A VM (or a set of VMs) are considered to be linked to a hypervisor if, and only if, the following conditions hold simultaneously: (1) the attestations of all the purportedly-linked parties verify individually (if we run `aAuthVerif` it returns 1 for each individual attestation); (2) the public key that was successfully used to verify each of the VMs' attestation is part of the auxiliary value `lkaux` forwarded by the hypervisor.

**Security.** We prove (in AppendixB) the security of our scheme with respect to a single permitted function,  $F_{\text{Sign}}$  that takes in input a message  $M$  from a message space  $\mathcal{M}$  and outputs, when queried for a compromised party  $P$ , a signature on the message  $M$  with the private key  $P.sk$ . We demand that the message space  $\mathcal{M}$  be disjoint from the range of *any* basic attestation scheme.

**Theorem 22 (Secure Linked Attestation)** *The LinkedAtt scheme is secure assuming that AuthAtt scheme is secure and that the hash function  $H$  is collision resistant.*

## 2.4 Authorized Linked Attestation

So far, attestation has been viewed as a primitive, run by a single party (which can be of various types) and outputting an attestation. However, one of the most important requirements of attestation is that the actual quote only be given to authorized parties – which we call *attestation servers* [14].

We will define an *authorized linked attestation* protocol, which allows an attestation server to act as a verification party in the attestation procedures. The same server will also be the one to generate the auxiliary values required for the attestation (this provides freshness to the protocol). The server will also be responsible for linking multiple attestations.

**Intuition.** We provide a full formalization of authorized linked attestation below. However, we also believe it is useful to first give an intuitive understanding of what this primitive *is* and the security it wants to achieve.

In authorized linked attestation we consider a (single) attestation server  $S$  and platforms consisting of several types of components (as shown for linked attestation). The server will keep track of an evolving state, which is initially empty. However, as the server starts to attest various components, at every execution of the authorized attestation protocol, the server will output a verdict (indicating whether the component’s individual attestation has failed or succeeded) and may – or may not – update its internal state. Intuitively, the state is meant to contain the *linking information* provided by each of the attesting components. After a number of attestations have taken place, the server might have enough information in its state to decide whether some of the components are linked or not.

The security notion we require for authorized linked attestation is threefold: (1) we require that parties only provide attestation guarantees to the actual attestation server; (2) we require that the contents of the attestation be actually indistinguishable from random for all unauthorized parties; (3) we require a similar kind of linking security as demanded in linked attestation see Section 2.3. However, as opposed to linked attestation, the adversary in this case can also play a Person-in-the-Middle role between honest components and the honest server, or it may attempt to replay messages or impersonate one or both parties. Finally, the adversary will be able to have oracle access to the secret key of any compromised component (this oracle access is parametrized in terms of a function space  $F$  of allowed functions).

**Formalization.** The complete formalization of authorized linked attestation is given in Appendix C. Components on platforms are either VMs or hypervisor. In addition, we consider a  $S$ , which stores a tuple consisting of a public and a private key  $S.pk = pk$  and  $S.sk = sk$  respectively, and a state  $S.st$ . Parties interact with each other in sessions, which are run by an *instance* of the server and an *instance* of a given component. Instances of each party use that party’s long-term public and private keys, as well as potential local randomness, such as instance-specific nonces. An instance of a component and an instance of the server are *partnered* if they essentially run the same session (formally, if they

share a session identifier, which consists of the concatenation of a number of session-specific values).

*Authorized linking attestation* is defined as the tuple  $\text{ALA} = (\text{ASetup}, \text{AReg}, \text{AAttest}, \text{aALink})$ . The first, second, and last of these are algorithms, while  $\text{AAttest}$  is a protocol. The setup algorithm generates parameters (keys and public system values) for all the involved parties. The registration algorithm allows the VMs and hypervisor on a single platform (defined as sets  $\mathbf{s}_1$  for the VMs on the platform and  $\mathbf{s}_2$  for the hypervisor) to be associated with each other. For administration purposes, the public keys of all VMs on a platform (*i.e.*, all VMs in some  $\mathbf{s}_1$ ) and respectively the public key of the platform's hypervisor (the hypervisor in the corresponding  $\mathbf{s}_2$ ) are stored respectively in subsets  $\text{PK}_1, \text{PK}_2$  ( $i = 1$  for VMs and  $i = 2$  for the hypervisor). Together all the subsets  $\text{PK}_i$  for all the components form a set  $\mathcal{PK}[i]$  (for  $i = 1, 2$ ).

The authorized attestation protocol is run by an instance of a component and an instant of the server, yielding, for the component, an acceptance bit (corresponding to the authentication of its partner as the authorized server) and for the server, a tuple  $\text{verdict}, \text{S.st}$ : the verdict  $\text{verdict}$  is 1 or 0 depending on whether the component attested successfully or not, and the state is an update of the server's current internal state. Finally, the server state can be used on a subset of components in the  $\text{aALink}$  algorithm, yielding either 1 (the components are linked) or 0 otherwise.

**Construction.** Our construction of the  $\text{ALA}$  primitive can be seen in the Figure 7. We consider the existence on an underlying  $\text{LinkedAtt}$  scheme that we use for the  $\text{aLSetup}$ ,  $\text{aLReg}$  and  $\text{aLLink}$  in a straightforward manner. However, the  $\text{aLAttest}$  algorithm is no longer a primitive, but a protocol between two instances of two parties,  $\text{P}$  and  $\text{Q}$ . For simplicity of exposition, we assume that the instance of  $\text{Q}$  is the server attesting the component identified by  $\text{P}$ .

The protocol proceeds as follows. First,  $\text{P}$  and  $\text{Q}$  execute the TLS protocol, with  $\text{P}$  playing the role of the client and  $\text{Q}$  playing the role of the server. The role of the TLS protocol is two-fold: first,  $\text{P}$  authenticates the server, so that they can determine whether this party is allowed to obtain attestation data. Second, it leads to the establishment of a secure channel, such that the following messages can be passed on in a secure manner. Once the traffic key(s) established, the protocol continues as follows. First, the server uniformly randomly

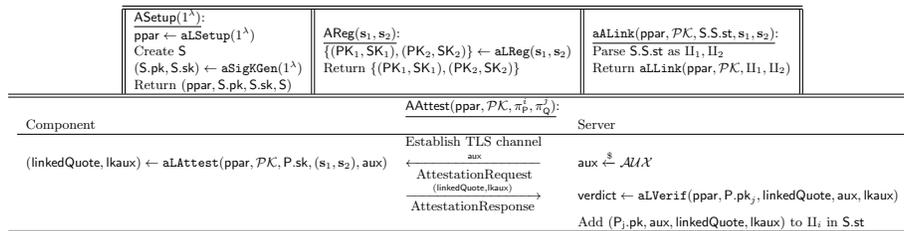


Fig. 7: Our authorized linked attestation scheme for 2 types of components.

samples a nonce `aux`, which is embedded in the first message of the protocol, `AttestationRequest`. In response, the party `P` executes the `aLAttest` algorithm and the output, consisting of a `linkedQuote` and the linkage information `lkaux`, is then sent to the server. The server will subsequently update his state.

In order for two components to be linked by the server successfully, the following conditions have to be met. First, the two components' attestation must be valid (their associated verdicts equals 1). Second, the two `lkaux` must be subsets of each other; essentially, the key that the VM used as part of its attestation must be found in the `lkaux` provided by the hypervisor.

We note that if the server has at some point accepted the attestation of a component (thus updating its state to add the linking information), and if later a failed attestation occurs with respect to that component, the server updates state as follows: it ignores the linking information provided in the second attestation; and it removes prior linking information provided by that component.

**Security.** There are three fundamental properties we want ALA schemes to have: an authenticity guarantee for the attestation server (authorization); a confidentiality guarantee for the contents of the attestation (indistinguishability); and a linkability guarantee for honestly-behaving components (linking-security). The first notion, authorization, captures the fact that before reaching an accepting state, a (non-server) party must be sure that it is speaking to the legitimate server (game `AuthSecλ,F`). The second notion, indistinguishability, essentially covers Person-in-the-Middle confidentiality for the attestation protocol (game `AuthIndλ,FSign`). The last property, linking-security, refers to the fact that no PitM adversary with the ability to compromise components can convince an attestation server that a component is linked to another if that is not the case in reality (game `AuthLinkλ,F`). Although this last property might seem similar to the security notion for our linked attestation primitive, there is one important difference between the two: in linked attestation the adversary has access to essentially two ways to generate an attestation (depending on whether the component is honest or compromised), whereas in *authorized linked attestation* the adversary will have more leeway in combining attestation material across sessions. The stronger adversary in this section will thus make for a stronger primitive in the end. The three security games are defined in Appendix D.

**Theorem 23** *Our construction is `AuthSecλ,FSign` secure if the TLS protocol provides server authentication:  $\Pr[\mathcal{A} \text{ wins } \text{AuthSec}_{\lambda,F}] \leq \varepsilon_{\text{TLS-auth}}$ .*

*Proof.* (sketch) Note that in order to win this game, the adversary must make a party accept a session with the server, such that no matching server instance exists. This is against the server authentication property we assume of the TLS protocol.

**Theorem 24** *Our construction is `AuthLinkλ,FSign` secure if the underlying primitive `LinkedAtt` is `LinkSecλ,FSign` and TLS is at least (s)ACCE secure.*

*Proof.* (sketch) A key observation for this game is that the adversary cannot impersonate a server or determine it to provide bad randomness. Instead, the

adversary can compromise components and run TLS sessions on their behalf with the server, or try to obtain input from honest components instead. We distinguish between two types of adversary behaviours.

Say  $\mathcal{A}$  has never queried `OCompromise` for some party  $P$ . If the adversary prompts  $P$  to run a session, then  $\mathcal{A}$  will not actually know anything about the messages (so it cannot misbehave on the quote, the nonce, or anything else). If  $\mathcal{A}$  runs the TLS session instead of  $P$ , it will learn the channel key, but will not be able to prompt  $P$  for the quote (since  $P$  wants to run TLS and not the attestation protocol, and since  $\mathcal{A}$  cannot impersonate the server).

Say  $\mathcal{A}$  queries `OCompromise` for some party  $P$ . Then the adversary can run TLS sessions on behalf of that party and query `OUseKey` in an attempt to get information on the quotes. However, in that case, the attestation of that component fails, except again if we break linked authentication security.

This essentially means that the adversary has no way to maul honestly-generated input to suit its purposes.

**Theorem 25** *Our construction is  $\text{AuthInd}_{\lambda, \mathcal{F}_{\text{Sign}}}$  secure if the TLS channel provides (minimally) (s)ACCE security. Let  $q_{\text{sessions}}$  be the number of sessions.*

$$\Pr[\mathcal{A} \text{ wins } \text{AuthInd}_{\lambda, \mathcal{F}_{\text{Sign}}}] \leq \frac{1}{q_{\text{sessions}}} \varepsilon^{\text{TLS-sACCE}}.$$

*Proof.* Like in the proof of authorization, the reduction here is immediate. The property of sACCE (which is already provided by TLS 1.2, whereas TLS 1.3 gives even stronger guarantees) implies that messages exchanged across the TLS channel are secure.

### 3 Implementation

We provide a proof of concept implementation of our authorized linked attestation scheme. The implementation consists of three parts, a client for the hypervisor, a client for the Virtual Machines, and an attestation server written in Python 3. We do not consider the underlying NFV or cloud infrastructure, since our scheme abstracts those environments and can be used in any kind deep-attestation scenario. Therefore, any computer equipped with a TPM 2.0 (which can also be emulated) and which has virtualization capacities suffices for the purposes of our implementation. We provide our code as well as a detailed tutorial on how to install and configure both the infrastructure [3].

**The infrastructure.** We summarize our testing architecture in Figure 8. Although only two VMs are represented in this diagram, some of our tests will use 3 or 4.

Our *hypervisor* is a laptop running Ubuntu 20.04.1 (kernel version 5.4.0-58) with an Intel i5-10210U CPU, 8GB RAM and a Nuvoton TPM NPCT75X. We used KVM to turn this laptop into a hypervisor.

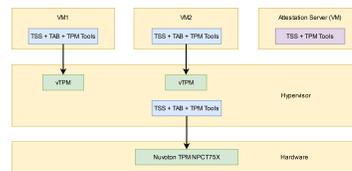


Fig. 8: Architecture for tests.

Table 1: Minimum, median, mean and maximum time in second for attestation of a hypervisor and a virtual machine for 100 trials.

|            | min  | median | mean | max   |
|------------|------|--------|------|-------|
| Hypervisor | 3.22 | 5.30   | 5.68 | 11.55 |
| VM         | 0.66 | 0.97   | 1.03 | 1.41  |

In order to achieve a high attestation performance, we used a full *virtual TPM implementation*, using QEMU [1] with libtpms [7] version 0.7 and swtpm [15] version 0.5.

All *virtual machines* are QEMU virtual machines (version 4.2.1) with 2 cores and 4G RAM running Ubuntu 20.04.1.

The hypervisor, server, and VMs communicate through a virtual network. A virtual bridge on the hypervisor redirects packets on the appropriate virtual interface. Thus, connection time is not considered in our tests.

To communicate with the TPM we used tpm2-tss, tpm2-abrmd and tpm2-tools from the tpm2-software [2]. Note that the tpm2-tss project implements the TPM software stack (TSS), which is an API specified by the Trusted Computing Group to interact with a TPM. The tpm2-abrmd implements the access broker and resources to manage concurrent access to the TPM and manage memory of the TPM by swapping in and out of the memory as needed (hardware TPM have limited memory). tpm2-tools are a set of command-line tools based on TSS, which are used to send commands to the TPM. We used Python to wrap tpm2-tools commands.

The *attestation server* is also a virtual machine, with the same characteristics as those above. This allows us to test our implementation on a single machine. We establish a secure connection between the client and the server by using Python’s SSL library and then sending protocol messages as encoded json strings directly into TLS socket.

**Tests.** We perform three types of experiments. The first is a comparison of hypervisor attestation time and VM attestation time. Although both those processes have some (very small) amount of noise, our values faithfully show the difference between attesting a component through the physical TPM – hypervisor attestation – and attesting it by using a virtual TPM – VM attestation.

We ran 100 attestations for the hypervisor and 100 attestations for a virtual machine. The results have high variance so Table 1 presents the minimum, the maximum, mean and the median value of those 100 trials. As expected, time for an attestation using a hardware TPM is much higher than using a Virtual TPM.

As our second and third experiments we wanted to see how the overall runtime of our scheme evolves with the number of virtual machines that need to be attested, when the attestation is sequential or parallelized for the VM attestations. In both cases, each experiment run first executed the attestation of the

hypervisor, and then (sequentially or in parallel) the attestations of a varying number of VMs (between 1 and 4). For both experiments we did 100 runs of the experiment. The median runtime was calculated in each experiment (sequential or parallel) for each number of VMs (1 through 4), and the results are plotted in Figure 9

We note that the runtime is not entirely linear. This is because in experiments 2 and 3 the initial attestation of the hypervisor (which only occurs once) takes a much larger time than the subsequent VM runtimes.

**Comparison to single-channel attestation.** We did not implement single-channel attestation. However, since we have implemented hypervisor and VM attestations, we can theoretically estimate the runtime of single-channel attestation for a varying number of VMs – which we plot in Figure 9. Indeed, a single-channel attestation process for a single VM includes a VM attestation *and* a hypervisor attestation. If we want to run it for 2 VMs, then we need to perform 2 hypervisor attestations and 2 VM attestations. This cannot be easily parallelized either, because the same TPM has to run the attestations. This yields a much higher runtime, as depicted in Figure 9.

**Comparison to multi-channel attestation.** Although our method follows basic multi-channel attestation approaches, we do add an extra computation (a hash function computation) compared to traditional multi-channel attestation. In addition, we require a little extra memory overhead for both the attestation server and for each platform, so that the additional attestation keys are stored for each VM. There is also a slight transmission overhead, since those keys are also sent upon attestation. However, the transmission overhead is negligible since it only appears for the hypervisor attestation (which occurs only once).

## 4 Conclusions and Future Work

We proposed a layer-binding in deep-attestation without running into the complexity of single-channel attestation. Our construction achieves the best of both worlds, with a complexity similar to that of multi-channel attestation, but with the strong linkage properties provided in single-channel attestation.

We accompany our construction by a proof-of-concept implementation that clearly shows the viability and scalability of our solution, especially if VM attestations are run in parallel.

In addition, we are the first to present a full, formal treatment of our new protocol, which we call *authorized linked attestation*. Our construction of authorized linked attestation is modular, built on primitives which have increasingly

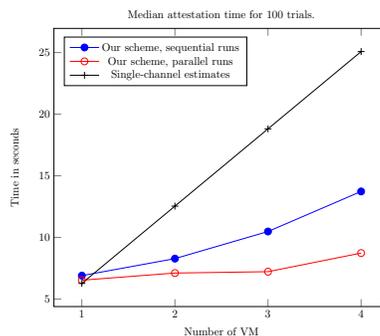


Fig. 9: Attestation time for 1, 2, 3 and 4 VMs.

stronger properties. Our underlying assumption is a primitive called basic attestation. We show that in order to be able to prove security, we need that attestations be able to reflect compromise of the component. In addition, we rely on a collision-resistant hash function, an EUF-CMA-secure signature scheme, and the sACCE security of a TLS protocol (having AKE properties would be even better).

However, our model (and scheme) does not immediately account for other features of virtual infrastructures, such as privacy CAs, migrating VMs, multiple hypervisors managing the same VM, or even replacing TPMs. These aspects are left as future work.

## References

1. The Qemu machine emulator. <https://www.qemu.org/>
2. TPM2\_tools. <https://github.com/tpm2-software>
3. Anonymized: Implementation. <https://github.com/AnonymousDeepAttestation/deep-attestation>
4. Armknecht, F., Sadeghi, A., Schulz, S., Wachsmann, C.: A security framework for the analysis and design of software attestation. In: 2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013. pp. 1–12. ACM (2013)
5. Asokan, N., Brasser, F.F., Ibrahim, A., Sadeghi, A., Schunter, M., Tsudik, G., Wachsmann, C.: SEDA: scalable embedded device attestation. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015. pp. 964–975. ACM (2015)
6. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) *Advances in Cryptology - CRYPTO '93*, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings. Lecture Notes in Computer Science, vol. 773, pp. 232–249. Springer (1993)
7. Berger, S.: Library for TPM tools. <https://github.com/stefanberger/libtpms>
8. Brickell, E.F., Camenisch, J., Chen, L.: Direct anonymous attestation. In: Atluri, V., Pfitzmann, B., McDaniel, P.D. (eds.) *Proceedings of the 11th ACM Conference on Computer and Communications Security, CCS 2004*, Washington, DC, USA, October 25-29, 2004. pp. 132–145. ACM (2004)
9. Buhren, R., Werling, C., Seifert, J.: Insecure until proven updated: Analyzing AMD sev's remote attestation. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019*, London, UK, November 11-15, 2019. pp. 1087–1099. ACM (2019)
10. Canetti, R., Krawczyk, H.: Universally composable notions of key exchange and secure channels. In: Knudsen, L.R. (ed.) *Advances in Cryptology - EUROCRYPT 2002*, International Conference on the Theory and Applications of Cryptographic Techniques, Amsterdam, The Netherlands, April 28 - May 2, 2002, Proceedings. Lecture Notes in Computer Science, vol. 2332, pp. 337–351. Springer (2002)
11. ETSI: Network functions virtualisation(nfv); trust; report on attestation technologies and practices for secure deployments (10 2017)

12. Gopalan, K., Kugve, R., Bagdi, H., Hu, Y., Williams, D., Bila, N.: Multi-hypervisor virtual machines: Enabling an ecosystem of hypervisor-level services. Proceedings of the 2017 USENIX Annual Technical Conference, USENIX ATC 2017 pp. 235–249 (2019)
13. Hussain, S.R., Echeverria, M., Karim, I., Chowdhury, O., Bertino, E.: 5greasoner: A property-directed security and privacy analysis framework for 5g cellular network protocol. In: Cavallaro, L., Kinder, J., Wang, X., Katz, J. (eds.) Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019, London, UK, November 11-15, 2019. pp. 669–684. ACM (2019)
14. Lauer, H., Kuntze, N.: Hypervisor-based attestation of virtual environments. In: 2016 IEEE UIC/ATC/ScalCom/CBDCCom/IoP/SmartWorld. pp. 333–340 (2016)
15. Safford, D., Berger, S.: Software TPM emulator – swtspm. <https://github.com/stefanberger/swtspm>
16. Shoup, V.: Sequences of games: a tool for taming complexity in security proofs. Cryptology ePrint Archive, Report 2004/332 pp. 1–33 (2006)

## A Authenticated Attestation

We consider an environment parametrized by a security parameter  $\lambda$ , which will contain up to  $N$  parties. Parties keep track of two attributes: the compromise bit  $\gamma$  used also for basic attestation, and a set of public and secret keys  $(\mathbf{pk}, \mathbf{sk})$ : a public key  $\mathbf{pk}$  (assumed to be unique per party, and known to all other parties including the adversary) and a private key  $\mathbf{sk}$  known only to that party. We use  $P.\mathbf{pk}/P.\mathbf{sk}$  to indicate the public/secret key of party  $P$ . We require that this primitive be *correct* in two ways:  $\mathbf{aAuthVerif}(\cdot, \cdot, \cdot, \aleph) = 0$  for all possible input values in the first three parameters, and: for all  $\mathbf{ppar} \leftarrow \mathbf{aAuthSetup}(1^\lambda)$ ,  $(\mathbf{pk}, \mathbf{sk}) \leftarrow \mathbf{aAuthKGen}(P)$ ,  $\mathbf{aux} \in \mathcal{AU}\mathcal{X}$ , if  $\mathbf{authQuote} \rightarrow \mathbf{aAuthAttest}(\mathbf{ppar}, P.\mathbf{sk}, \mathbf{aux})$  and  $\mathbf{authQuote} \neq \aleph$ ,  $\mathbf{aAuthVerif}(\mathbf{ppar}, P.\mathbf{pk}, \mathbf{aux}, \mathbf{authQuote}) = 1$ .

**Security.** Formally,  $\mathcal{A}$  will play the  $\mathbf{AuthSec}$  game against a challenger  $\mathcal{G}$ , which begins the game by running  $\mathbf{aAuthSetup}$  and outputting  $\mathbf{ppar}$  to  $\mathcal{A}$ . The challenger also initializes  $N$  to 1. The adversary then has access to the following oracles:

- $\mathbf{OAuthReg}() \rightarrow (P_i, P_i.\mathbf{pk})$  : if  $i \leq N$ , it runs  $\mathbf{aAuthKGen}(P_i) \rightarrow (P_i.\mathbf{pk}, P_i.\mathbf{sk})$ . It outputs  $P_i.\mathbf{pk}$  to all parties and keeps  $P_i.\mathbf{sk}$  private, stored in the keys attribute of party  $P_i$ . A handle (in practice the index) of this party is also returned to  $\mathcal{A}$ .
- $\mathbf{OAuthAttest}(P_i, \mathbf{aux}) \rightarrow \mathbf{authQuote} \cup \aleph$  : this oracle runs the  $\mathbf{aAuthAttest}$  algorithm on  $\mathbf{ppar}$ ,  $P_i.\mathbf{sk}$  and input  $\mathbf{aux}$ , and returns the output. On adversarially chosen input  $P_i$  and  $\mathbf{aux}$ , the oracle updates a list  $\mathcal{L}_{\text{Att}} \leftarrow \mathcal{L}_{\text{Att}} \cup (P_i, \mathbf{aux}, \mathbf{authQuote})$ .
- $\mathbf{OCompromise}(P_i) \rightarrow \text{OK}$  : this oracle allows an adversary to compromise party  $P_i$ , thus changing  $P_i.\gamma$  to 1.
- $\mathbf{OAuth}(P_i, M) \rightarrow \sigma_M$  : this oracle can only be queried for a party whose compromise bit is 1, and it outputs an EUF-CMA-secure signature keyed with  $P_i.\mathbf{sk}$  on a message  $M$ . We require that  $M$  be outside the range of any basic

attestation scheme. This last oracle reflects the fact that compromised parties can access a signing function within the TPM.

Finally the adversary outputs a tuple  $(P, \text{aux}, \text{authQuote})$ . It is said to *win* if and only if the following condition holds:  $\text{aAuthVerif}(\text{ppar}, P, \text{aux}, \text{authQuote}) = 1$  and there exist no tuples  $(P, \text{aux}, \text{authQuote})$  such that  $(\text{authQuote} \neq \aleph) \rightarrow \text{OAuthAttest}(\text{ppar}, P, \text{aux})$  for the current public parameters  $\text{ppar}$  (output by the challenger  $\mathcal{G}$ ).

**Proof of Theorem 21.**

*Proof.* Let  $\mathcal{A}$  be a probabilistic polynomial-time algorithm. The goal of  $\mathcal{A}$  is to provide a signed quote with correct auxiliary value (a nonce) such that the quote and the signature are valid for a fresh nonce.

We propose a proof using a sequence of game hops as introduced in [16]. The initial game corresponds to the security game **AuthSec**. The successive games are slight modification to its previous one to end up with a game corresponding to generic primitive game (such as EUF-CMA).

**Game 0.** This is the original security game

**Game 1 (transition based on indistinguishability).** This game is defined

as the previous one except that the challenger aborts the game if a compromised component is able to generate a valid attestation quote. Suppose that  $\mathcal{A}$  has a non-negligible advantage  $\epsilon_B$  of winning the basic attestation game. This means that there exists a party  $P$  such that  $P.\gamma = 1$  (*i.e.*,  $P$  is compromised) but  $\text{quote} \in \text{DB}$  with  $\text{aBVerif}(\text{ppar}, \text{quote}) = 1$  (note that  $\text{quote} = \aleph$  potentially). By difference lemma we have:  $|\Pr[\mathcal{A} \text{ wins } G_0] - \Pr[\mathcal{A} \text{ wins } G_1]| \leq \epsilon_B$

**Game 2.** This game is defined as the previous one except that the game aborts if

$\mathcal{A}$  can generate a valid signature. We show that:  $|\Pr[\mathcal{A} \text{ wins } G_1] - \Pr[\mathcal{A} \text{ wins } G_2]| = \frac{\epsilon_{\text{EUF-CMA}}}{N}$  where  $\epsilon_{\text{EUF-CMA}}$  is the advantage of EUF-CMA security game. The proof is done by reduction. Assume that  $\mathcal{A}$  can generate a valid  $\text{authQuote}^*$ , *i.e.*,  $\text{aAuthVerif}(\text{ppar}, P^*, \text{aux}^*, \text{authQuote}^*) = 1$  with  $(P^*, \text{aux}^*, \text{authQuote}^*) \notin \mathcal{L}_{\text{Att}}$ . We then show that there exists adversary  $\mathcal{B}$  using  $\mathcal{A}$  as a sub-routine with non-negligible advantage of winning the EUF-CMA security game.

Adversary  $\mathcal{B}$  simulates the game of  $\mathcal{A}$  thus acting as the challenger in the **AuthAtt** game. The behavior of  $\mathcal{B}$  is defined as follows:

- receives  $pk$  from its own challenger of the EUF-CMA game.
- runs the **aAuthSetup** algorithm to get  $\text{ppar}$  (and also  $N$ ).
- randomly selects  $i^* \xleftarrow{\$} \{1, \dots, N\}$ . Two cases need to be studied depending on the  $i^{\text{th}}$  query of  $\mathcal{A}$ :
  - case 1:**  $i \neq i^*$ . When  $\mathcal{A}$  calls oracle **OAuthReg**() then  $\mathcal{B}$  runs **aAuthKGen**( $P_i$ ) to retrieve  $(P_i.\text{pk}, P_i.\text{sk})$ .  $\mathcal{B}$  sends back  $P_i.\text{pk}$  to  $\mathcal{A}$ . When  $\mathcal{A}$  calls oracle **OAuthAttest**() then  $\mathcal{B}$  runs algorithm **aAuthAttest** (which is possible since  $\mathcal{B}$  has the corresponding secret key).  $\mathcal{B}$  sends back to  $\mathcal{A}$  the output of the algorithm. Note that in this case, the simulation is the same as the original game since  $\mathcal{B}$  uses the same algorithm of the oracles.

**case 2:**  $i = i^*$ . In this case,  $\mathcal{B}$  will inject its own material to use it in its EUF-CMA game. When  $\mathcal{A}$  calls oracle  $\text{OAuthReg}(P_{i^*})$  then  $\mathcal{B}$  simply returns  $pk$ . Note that in this case,  $\mathcal{B}$  does not have access to  $sk$ . Thus when  $\mathcal{A}$  calls oracle  $\text{OAuthAttest}()$ ,  $\mathcal{B}$  cannot sign the quote. Instead,  $\mathcal{B}$  runs  $\text{aBAttest}(\text{ppar}_{i^*})$  and sends to its challenger  $\text{quote}||\text{aux}$ . In response, its challenger will send a signed value of it using  $sk$ .  $\mathcal{B}$  then forward this to  $\mathcal{A}$ . The view of  $\mathcal{A}$  that is different from the original game in this case is the output of the  $\text{OAuthAttest}()$  oracle. The latter runs algorithm  $\text{aAuthAttest}$  which generate a quote  $\text{quote} \leftarrow \text{aBAttest}(\text{ppar}_i)$  and a signature  $\sigma \leftarrow \text{aSigSign}(P_i.sk, (\text{quote}, \text{aux}))$ . In the simulation,  $\mathcal{B}$  has access to algorithm  $\text{aBAttest}()$  but the signature scheme is different. Yet, both signature schemes are EUF-CMA thus their outputs are indistinguishable (meaning that  $\mathcal{A}$  cannot decide from which schemes the output comes from with non-negligible probability) since the keys have the same probability distribution. Hence, the simulation of the game by  $\mathcal{B}$  and the real game are indistinguishable.

- When  $\mathcal{A}$  returns its forgery on query  $i$ ,  $\mathcal{B}$  parses  $\text{authQuote}_i$  as  $\text{authQuote}_i := (m^*||\text{aux}^*, \sigma^*)$ .
- Finally  $\mathcal{B}$  returns  $(m^*||\text{aux}^*, \sigma^*)$  to its challenger and wins if  $\mathcal{A}$  forges the  $i^*$  query (meaning that  $i = i^*$ ).

By combining the results, we have:  $\Pr[\mathcal{A} \text{ wins } G_0] \leq \varepsilon_B + \frac{\varepsilon_{\text{EUF-CMA}}}{N}$  which is negligible.

## B Linked attestation

We formalize a linked-attestation scheme  $\text{LinkedAtt}$  as a tuple of algorithms  $\text{LinkedAtt} = (\text{aLSetup}, \text{aLReg}, \text{aLAttest}, \text{aLVerif}, \text{aLLink})$ , defined for some auxiliary set  $\mathcal{AUX}$ .

- $\text{aLSetup}(1^\lambda) \rightarrow \text{ppar}$ : on input the security parameter  $1^\lambda$  (in unary), this algorithm outputs public parameters  $\text{ppar}$ . This security parameter includes the maximal number of allowed disjoint linkable sets, which we denote as  $L$ .
- $\text{aLReg}(\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_L) \rightarrow \{(\text{PK}_1, \text{SK}_1), \dots, (\text{PK}_L, \text{SK}_L)\}$ : this algorithm keeps as state a number  $L$  of sets  $\mathcal{S}_i$  originally set to  $\emptyset$ , and a vector of sets of public keys  $\mathcal{PK}$  (also initialized to  $\emptyset$ ). On input a number of subsets  $\mathbf{s}_i$  ( $i = 1, 2, \dots, L$ ), this algorithm first checks that  $\forall i, j, \mathbf{s}_i \cap \mathbf{s}_j = \emptyset$  (else the algorithm outputs  $\perp$ ). If the relation is true, then the algorithm generates for each party  $P_j \in \mathbf{s}_i$ , (for all  $j, i$ ) a tuple of public and private keys  $P_j.pk, P_j.sk$ , initializing  $P_j$  with those keys. We require the unicity of all the generated public keys. The subsets  $\mathbf{s}_i$  are each added to greater sets  $\mathcal{S}_i$ . The algorithm groups the keys of all parties  $P_j \in \mathbf{s}_i$  in a pair of private/public key subsets:  $(\text{PK}_i, \text{SK}_i)$ , updating the  $i$ -th component  $\mathcal{PK}[i]$  of  $\mathcal{PK}$  as  $\mathcal{PK}[i] \cup \text{PK}_i$ . All parties are given access to the public-key subsets (and more generally, to  $\mathcal{PK}$ ).
- $\text{aLAttest}(\text{ppar}, \mathcal{PK}, P.sk, (\mathbf{s}_1, \dots, \mathbf{s}_L), \text{aux}) \rightarrow (\text{linkedQuote} \cup \mathfrak{N}, \text{lkaux})$ : on input public parameters  $\text{ppar}$ , the current set of public keys  $\mathcal{PK}$ , the private key  $P.sk$  of some party  $P$ , subsets  $\mathbf{s}_i \in \mathcal{S}_i$ , and an auxiliary value  $\text{aux} \in \mathcal{AUX}$ , this algorithm outputs either a linked quote  $\text{linkedQuote}$  or a special failure

symbol  $\aleph$ , and a different value  $\text{lkaux}$  (this last entry could be used to store linkage-related information).

- $\text{aLVerif}(\text{ppar}, \text{P.pk}, (\text{linkedQuote} \cup \aleph), \text{aux}) \rightarrow 0 \cup 1$ : On input the public parameters  $\text{ppar}$ , a public key  $\text{P.pk}$ , a linked quote (or a failure symbol  $\aleph$ ), and an auxiliary value  $\text{aux}$ , this algorithm outputs a verification bit. By convention, 0 means failure and 1 means acceptance of the attestation.
- $\text{aLLink}(\text{ppar}, \mathcal{PK}, \Pi_1, \dots, \Pi_L) \rightarrow 0 \cup 1$ : on input the public parameters  $\text{ppar}$ , the set of public keys  $\mathcal{PK}$ , and subsets  $\Pi_i$  containing elements of the form  $(\text{P}_j.\text{pk}, \text{aux}, (\text{linkedQuote} \cup \aleph), \text{lkaux})$ , this algorithm outputs 1 if the quotes in all the indicated subsets can all be linked (thus also indicating the parties are linked) or 0 otherwise.

By convention, we allow the use of  $\emptyset$  to indicate that any of the input or output (sub)sets to also be empty.

**Formalization.** The security of linked attestation informally states that an adversary, which has Person-in-the-Middle capabilities and can compromise devices at will, cannot make it appear that two devices are linked when they are not, in fact, so.

A significant limitation on the adversary’s capabilities is that compromising a device will not leak its private keys (which are assumed to be held by a TPM). However, the adversary will gain a limited oracle access to those keys upon compromising the device. The limitations to those queries follow rules of access to an actual TPM.

More formally, we define the security of Linked Attestation as a game  $\text{LinkSec}_{\lambda, F}$  played by an adversary  $\mathcal{A}$  against its challenger  $\mathcal{G}$ . The game is parametrized by a security parameter  $\lambda$  and a set of functions  $F$ , which we call the *permitted key-access functions*. The challenger begins by running  $\text{aLSetup}(1^\lambda)$ , returning  $\text{ppar}$  to the adversary, and then it instantiates two lists: a list of parties  $\mathcal{L}_{\text{Reg}} = \emptyset$  and a list of linkable attestations  $\mathcal{L}_{\text{Att}} = \emptyset$ . The adversary then plays its game by using the following oracles adaptively:

- $\text{OLReg}(n_1, \dots, n_L) \rightarrow (\text{PK}_1, \dots, \text{PK}_L)$ : the linked user-registration oracle creates a linked platform consisting of  $n_i$  components of the type indicated by  $\mathcal{S}_i$ . The challenger first instantiates a counter  $N_i = 0$  for all  $i$ ; it also instantiates subsets  $\mathbf{s}_i$  as a tuple of  $n_i$  handles  $\text{P}_{i,j}$ , with  $N_i + 1 \leq j \leq N_i + n_i$  and then runs the algorithm  $\text{aLReg}(\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_L)$ , instantiating the parties with their keys and outputting the public keysets to the adversary. The subset consisting of the list of subsets is added to  $\mathcal{L}_{\text{Reg}}$ . We note that this way of registering parties ensures by construction that no party finds itself in multiple sets, nor on multiple platforms.
- $\text{OLHAttest}(\text{P}, (\mathbf{s}_1, \dots, \mathbf{s}_L), \text{aux}) \rightarrow (\text{linkedQuote} \cup \aleph, \text{aux}^*)$ : this oracle first verifies that  $\text{P}.\gamma = 1$ . If the condition is false ( $\text{P}$  is compromised), then this oracle outputs an error symbol  $\perp$  (compromised parties must use the oracle  $\text{OLCAttest}$  described below). If the condition is true, then this algorithm runs  $\text{aLAttest}(\text{ppar}, \mathcal{PK}, \text{P.sk}, (\mathbf{s}_1, \dots, \mathbf{s}_L), \text{aux})$ , and returns the output to the adversary. The tuple  $(\text{P}, (\mathbf{s}_1, \dots, \mathbf{s}_L), \text{aux}, \text{linkedQuote}, \text{aux}^*)$  is stored in  $\mathcal{L}_{\text{Att}}$ .

- $\text{OCompromise}(\mathsf{P}) \rightarrow \text{OK}$ : this oracle allows an adversary to compromise party  $\mathsf{P}$ , thus changing  $\mathsf{P}.\gamma$  to 1.
- $\text{OLCAttest}(\mathsf{P}, (\mathbf{s}_1, \dots, \mathbf{s}_L), \text{aux}, f) \rightarrow (\aleph, \text{aux}^*)$ : this oracle first checks that  $\mathsf{P}.\gamma = 1$  (else,  $\perp$  is returned as an output). If the condition holds, then this oracle first checks that  $f \in F$  and if so, it runs  $f$  on  $\mathsf{P}.\text{sk}$  and input  $\text{aux}$  to output  $\text{aux}^*$ . Then it runs  $\text{OLHAttest}(\mathsf{P}, (\mathbf{s}_1, \dots, \mathbf{s}_L), \text{aux})$  to obtain  $\text{linkedQuote}$  (the second output is discarded). Note that by the security of the linked attestation primitive, we will have that  $\text{linkedQuote} = \aleph$ . The tuple  $(\mathsf{P}, (\mathbf{s}_1, \dots, \mathbf{s}_L), \text{aux}, \text{linkedQuote}, \text{aux}^*)$  is added to  $\mathcal{L}_{\text{Att}}$  and  $(\text{linkedQuote}, \text{aux}^*)$  is returned to  $\mathcal{A}$ .

At the end of its interaction,  $\mathcal{A}$  outputs a party  $\mathsf{P}$  and a tuple of subsets  $(\tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_L)$  with an index  $i^*$  such that  $\forall i \neq i^*, \mathbf{s}_i := \tilde{\mathbf{s}}_i$  and  $\mathbf{s}_{i^*} := \tilde{\mathbf{s}}_{i^*} \cup \{\mathsf{P}\}$ . In addition the adversary outputs for every party  $\mathsf{P} \in \mathbf{s}_1 \cup \dots \cup \mathbf{s}_L$  (parties being indexed as  $\mathsf{P}_{i,j}$ ) a tuple  $(\text{aux}, \text{linkedQuote}, \text{aux}^*)$  such that  $(\cdot, \cdot, \text{aux}, \text{linkedQuote}, \text{aux}^*) \in \mathcal{L}_{\text{Att}}$ .

We say the adversary *wins* if all the following conditions hold simultaneously:

- For each  $\mathbf{s}_i$  the parties inside this set are all registered *i.e.*, they were output by  $\text{OLReg}$ . In addition  $\mathsf{P}$  is registered;
- There exists at least one party  $\mathsf{Q} \in \mathbf{s}_j$  such that  $\mathsf{P}$  and  $\mathsf{Q}$  were issued from different  $\text{OLReg}$  queries;
- By setting  $\Pi_i := (\mathsf{P}.\text{pk}, \text{aux}, (\text{linkedQuote} \cup \aleph), \text{aux}^*)$  and, for  $k \neq i$ , for all  $\mathsf{P}_{k,j} \in \mathbf{s}_j$ ,  $\Pi_k := (\mathsf{P}_{k,j}.\text{pk}, \text{aux}, (\text{linkedQuote}_{k,j} \cup \aleph), \text{aux}_{k,j}^*)$ , it holds that  $\text{aLLink}(\text{ppar}, \mathcal{PK}, \Pi_1, \dots, \Pi_L) = 1$ .

In other words, the adversary wins if it is able to make attestations stored in  $\mathcal{L}_{\text{Att}}$  for parties registered on different platforms ( $\mathsf{P}$  and  $\mathsf{Q}$ ) link. Note that there are two ways that an attestation can end up in  $\mathcal{L}_{\text{Att}}$ : either it is issued for an honest component (and then it should hold that  $\text{linkedQuote} \neq \aleph$ ), or it is issued for a compromised party, for an adversarially-chosen evaluation of a permitted function  $f$  on a secret key (in which case  $\text{linkedQuote} = \aleph$ ). In other words, at this point a compromised component cannot just bind the output  $\text{aux}^*$  from the function evaluation oracle with a different quote. The adversary will gain this ability at the next step (when the freshness  $\text{aux}$  will no longer be chosen by the adversary).

**Proof of Theorem 2.** We will now prove our construction is secure with respect to the  $\text{LinkSec}_{\lambda, \text{F}_{\text{Sign}}}$  experiment.

*Proof.* **Game 0.** The original security game  $\text{LinkSec}_{\lambda, \text{F}_{\text{Sign}}}$ .

**Game 1.** We guess parties  $\mathsf{P}, \mathsf{Q}$  output by the adversary in the last part of its game. In other words, the challenger must draw at random two values between 1 and  $N$ , such that those values correspond to those chosen by the adversaries. We lose a factor  $\frac{1}{N^2}$ .

**Game 2.** We now rule out that  $H(\cdot, \text{lkaux}) = (\cdot, \text{lkaux}')$  for any  $\text{lkaux} \neq \text{lkaux}'$ .

Trivially, if the converse were true, we could break the collision resistance of  $H$  with equal probability.

Note that now, since parties P and Q are registered on different platforms, since  $\text{lk}_{\text{aux}}$  keys are unique, and since we have ruled out collisions, any honestly-generated attestations for P and Q will not link. The adversary’s only hope is to forge an attestation for either one of those parties.

**Game 3.** At this point we rule out the fact that P’s tuple  $(\text{aux}, \text{linkedQuote}, \text{aux}^*)$  was in fact part of a tuple  $(P', \cdot, \text{aux}, \text{linkedQuote}, \text{aux}^*) \in \mathcal{L}_{\text{Att}}$  (with  $P \neq P'$ ). If that were so, we could construct an adversary against the authenticated attestation scheme (since P purports to be  $P'$ ). In so doing an important oracle will be the signature oracle  $\text{OAuth}$  added artificially in the authenticated attestation primitive; the latter will allow us to simulate  $\text{OLCAttest}$  queries. We lose  $\varepsilon_{\text{Auth-attest}}$ .

**Game 4.** We repeat the previous game hop for party Q, and lose  $\varepsilon_{\text{Auth-attest}}$ .

At this point, the adversary can no longer win the game.

## C Formalization of Authorized Attestation

We will consider parties of multiple categories as for the linked-attestation primitive. We also define a special server entity, denoted S, which stores the following attributes:

- $(\text{pk}, \text{sk})$ : a tuple consisting of a public key  $\text{pk}$  (assumed to be unique and known to all other parties including the adversary) and a private key  $\text{sk}$  known only to the server. We use  $\text{S.pk}$  to indicate the public key of party S, and  $\text{S.sk}$  to indicate its private key.
- $\text{S.st}$ : a value called *state*, which stores tuples of linked attestations which are susceptible to be linkable to each other.

We will consider an environment in which parties interact with each other in *sessions*. The session is run by two party *instances*, one of the attesting party and the other, of the attestation server. For a party P we denote by  $\pi_P^i$  the  $i$ -th instance of party P.

Just as in the case of linked attestation, parties store a set of keys  $(\text{pk}, \text{sk})$ , as well as a compromise bit  $\gamma$ .

Party instances use the same keys and have the same compromise bit as the party itself, but in addition keep track of the following session-specific attributes:

- $\text{sid}$ : a session identifier, which will be useful in understanding which two party instances converse together.
- $\text{pidpk}$ : the public key belonging to this instance’s intended communication partner.
- $\text{T}$ : a transcript of messages exchanged throughout a protocol session, in plaintext. Even if encryption is used at some point, parties append messages to their transcripts only after decrypting.
- $\alpha$ : this bit is originally set to 0, but can be changed to 1 if this party instance has accepted its partner as a legitimate entity to run the authorized linked attestation with.

- **lst**: this local state variable stores instance- (and protocol-) specific values, such as encryption keys, randomness, etc.

In addition server instances  $\pi_S^j$  keep track of the following attribute, which is the output of the immediate attestation process taking place:

- **verdict**: this attribute stores a bit, initially set to 0, which is flipped to 1 if the attestation server's instance has accepted the attestation received during that session.

We call two instances  $\pi_P^i$  and  $\pi_Q^j$  *partnered* if, and only if, the following conditions hold simultaneously: exactly one of P and Q is in fact the attestation server S;  $\pi_P^i.\text{pidpk} = Q.\text{pk}$  and  $\pi_Q^j.\text{pidpk} = P.\text{pk}$ ; and  $\pi_P^i.\text{sid} = \pi_Q^j.\text{sid}$ .

*Authorized linking attestation* is defined as the tuple of algorithms and protocols  $\text{ALA} = (\text{ASetup}, \text{AReg}, \text{AAttest}, \text{aALink})$  described as follows:

- $\text{ASetup}(1^\lambda) \rightarrow (\text{ppar}, S.\text{pk}, S.\text{sk}, S)$  : on input the security parameter  $1^\lambda$  (in unary), this algorithm outputs public parameters **ppar**, as well as the server handle S (such that S is equipped with newly generated keys **pk**, **sk**). The value **ppar** includes the maximal number of allowed disjoint linkable sets of parties, which we denote as  $L$ . The values **ppar**,  $S.\text{pk}$ , and S are public,  $S.\text{sk}$  remains private. The value  $S.\text{pk}$  is added as the first value in the set  $\mathcal{PK}$ .
- $\text{AReg}(\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_L) \rightarrow \{(\text{PK}_1, \text{SK}_1), \dots, (\text{PK}_L, \text{SK}_L)\}$ : this algorithm keeps as state a number  $L$  of sets  $\mathcal{S}_i$  originally set to  $\emptyset$ , and a vector of sets of public keys  $\mathcal{PK}$  (also initialized to  $\emptyset$ ). On input a number of subsets  $\mathbf{s}_i$  ( $i = 1, 2, \dots, L$ ), this algorithm first checks that  $\forall i, j, \mathbf{s}_i \cap \mathcal{S}_j = \emptyset$  (else the algorithm outputs  $\perp$ ). If the relation is true, then the algorithm generates for each party  $P_j \in \mathbf{s}_i$ , (for all  $j, i$ ) a tuple of public and private keys  $P_j.\text{pk}, P_j.\text{sk}$ , initializing  $P_j$  with those keys. We require the unicity of all the generated public keys. The subsets  $\mathbf{s}_i$  are each added to greater sets  $\mathcal{S}_i$ . The algorithm groups the keys of all parties  $P_j \in \mathbf{s}_i$  in a pair of private/public key subsets:  $(\text{PK}_i, \text{SK}_i)$ , updating the  $i$ -th component  $\mathcal{PK}[i]$  of  $\mathcal{PK}$  as  $\mathcal{PK}[i] \cup \text{PK}_i$ . All parties are given access to the public-key subsets (and more generally, to  $\mathcal{PK}$ ).
- $\text{AAttest}(\text{ppar}, \mathcal{PK}, \pi_P^i, \pi_Q^j) \rightarrow (\text{verdict}, S.\text{st})$ : this protocol is an interaction between two party oracles, such that exactly one of P, Q is S. The protocol yields a tuple of values **verdict** and  $S.\text{st}$  to the server (and no output for the other party). Both party oracles are assumed to update their attributes accordingly as the protocol unfolds.
- $\text{aALink}(\text{ppar}, \mathcal{PK}, S.S.\text{st}, \mathbf{s}_1, \dots, \mathbf{s}_L) : 0 \cup 1$  : given the public parameters and public-key set, the server's current state, and a number of subsets of (purportedly-linked) parties, this algorithm outputs either 0 (the parties are not linked) or 1 (the parties are linked).

We require two types of *correctness* properties. First, we require that running the protocol between two honest parties yields a verdict of 1 (accept) on the side of the attestation server. Secondly, we require that components that are linked at registration will be viewed as linked by the **aALink** algorithm. More formally, we require that schemes  $\text{ALA} = (\text{ASetup}, \text{AReg}, \text{AAttest}, \text{aALink})$  be such that:

- For all  $(\text{ppar}, S.\text{pk}, S.\text{sk}, S) \leftarrow \text{ASetup}(1^\lambda)$  and for all parties  $P \in \mathcal{S}_i$  for some  $1 \leq i \leq L$ , it holds that  $(\text{verdict}, \cdot) = \text{AAttest}(\text{ppar}, \mathcal{PK}, \pi_P^i, \pi_S^j)$  (any legitimate party will successfully attest to the legitimate server);

- For all  $(\text{ppar}, \cdot, \cdot, S) \leftarrow \text{ASetup}(1^\lambda)$ , for all subsets  $\tilde{\mathbf{s}}_1, \tilde{\mathbf{s}}_2, \dots, \tilde{\mathbf{s}}_L$  such that there exist sets  $\mathbf{s}_i$  for  $i = 1, 2, \dots, L$  such that  $\tilde{\mathbf{s}}_i \subset \mathbf{s}_i$  and  $\text{AReg}(\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_L)$  was called and did not result in  $\perp$ , it holds that:  
 $\text{aALink}(\text{ppar}, \mathcal{PK}, \text{S.S.st}, \tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_L) = 1$  (parties that are registered together can be linked through the server’s state).

## D Security games for AuthAtt

The three security games we define are parametrized by a function space  $F$  and a security parameter  $\lambda$ . They start with the challenger running the setup algorithm and outputting  $\text{ppar}$  as well as the handle  $S$  and its public key  $S.\text{pk}$  to the adversary. Note that this will not give the adversary black-box access to  $S$ ’s attributes: it simply allows the adversary to later instantiate new attestation-protocol sessions for that server.

The adversary will then have access to some, or all of the following oracles:

- $\text{OALAReg}(n_1, \dots, n_L) \rightarrow (\text{PK}_1, \dots, \text{PK}_L)$  : the authorized linked user-registration oracle creates a linked platform consisting of  $n_i$  components of the type indicated by  $\mathcal{S}_i$ . The challenger first instantiates subsets  $\mathbf{s}_i$  as a tuple of  $n_i$  handles  $P_{i,j}$ , with  $1 \leq j \leq n_i$  and then runs the algorithm  $\text{AReg}(\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_L)$ , instantiating the parties with their keys and outputting the public keysets to the adversary. The subset consisting of the list of subsets is added to  $\mathcal{L}_{\text{Reg}}$ .
- $\text{ONewSession}(P, Q) \rightarrow \pi_P^i$  : on input the identity of a target party  $P$  and a partnering party  $Q$ , if both entities are correctly registered and exactly one of them is the server, then this oracle instantiates an instance of  $P$  whose partner will be instantiated as  $\text{pidpk} = Q.\text{pk}$ . Note that in order to observe an honest session between two parties, an adversary would have to create two partnered instances, one of  $P$  and the other of  $Q$ .
- $\text{OSend}(M, \pi_P^i) \rightarrow M'$  : this oracle simulates sending a message  $M$  to an instance  $\pi_P^i$ , and outputs the response  $M'$  of the party instance. If the input message takes a special value  $M = \text{prompt}$  and  $P$  is the initiator of the protocol (*i.e.*, the first party to have to send a message), this will trigger  $\pi_P^i$  to output the first message in the protocol. We note that some messages, when sent, might trigger errors, leading to an output  $M' = \perp$ . Other messages might trigger the attributes of the party (or instance) to be modified.
- $\text{ORevealState}(\pi_P^i) \rightarrow \text{lst}$  : on input a valid party instance, this oracle returns the value stored by the attribute  $\text{lst}$  of that instance.
- $\text{OUseKey}(P, f, \text{aux}) \rightarrow \text{aux}'$  : on input a compromised party  $P$  (not the server), a function  $f \in F$ , and an auxiliary input value  $\text{aux}$ , this oracle evaluates  $f$  on  $P.\text{sk}$  and  $\text{aux}$ .
- $\text{OCompromise}(P) \rightarrow \text{OK} \cup \perp$  : on input a registered party  $P \neq S$ , this oracle turns the party’s compromise bit to 1 and returns  $\text{OK}$ . If  $P = S$  or the party has not been registered, the output is an error symbol  $\perp$ .

We now proceed to describe each of the three security experiments we consider for our authorized linked attestation primitive.

In the *authorization* game  $\text{AuthSec}_{\lambda, \mathcal{F}}$ , after the challenger runs the setup algorithm, the adversary  $\mathcal{A}$  gets access to all the oracles described above. It ultimately stops with a `stop` message. We say  $\mathcal{A}$  wins if, and only if, there exists an instance  $\pi_{\mathcal{P}}^i$  such that  $\mathcal{P} \neq \mathcal{S}$ , for which the following conditions hold simultaneously:

- $\pi_{\mathcal{P}}^i$  ends in an accepting state, *i.e.*,  $\pi_{\mathcal{P}}^i.\alpha = 1$ ;
  - There exists no server instance  $\pi_{\mathcal{S}}^j$  such that  $\pi_{\mathcal{S}}^j$  is partnered with  $\pi_{\mathcal{P}}^i$ .
- In other words, the adversary wins if it can make a registered party believe it has talked to the server when this is not the case.

In the *linking* game  $\text{AuthLink}_{\lambda, \mathcal{F}}$ , after the challenger runs the setup algorithm, the adversary  $\mathcal{A}$  gets access to all the oracles above. It ends by outputting a tuple  $(\mathcal{P}, \mathbf{s}_1, \dots, \mathbf{s}_L)$  : such that for all  $1 \leq i \leq L$ ,  $\mathbf{s}_i \subset \mathcal{S}_i$  and there exists a unique  $i^*$  such that  $\mathcal{P} \in \mathcal{S}_{i^*}$  and  $\mathbf{s}_{i^*} = \emptyset$ . The challenger sets  $\tilde{\mathbf{s}}_i := \mathbf{s}_i$  for all  $i \neq i^*$ , and  $\tilde{\mathbf{s}}_{i^*} := \mathcal{P}$ . Then the challenger evaluates:  $b \leftarrow \text{aALink}(\text{ppar}, \mathcal{PK}, \mathcal{S.S.st}, \tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_L)$ . The adversary is said to *win* if, and only if the following conditions hold simultaneously:

- $b = 1$ ;
- There exists a party  $\mathcal{Q}$  and an index  $j^* \neq i^*$  such that  $\mathcal{Q} \in \tilde{\mathbf{s}}_{j^*}$  and  $\mathcal{P}$  and  $\mathcal{Q}$  were not output by the same `OALAReg` query.

In other words, for this second game, the adversary has to run several sessions between (potentially compromised) parties and the (honest) server, thus bringing the server's state to a point where linkage can be verified based on that state.

Finally, for the *indistinguishability* game  $\text{AuthInd}_{\lambda, \mathcal{F}}$ , once the challenger has finished the setup, it also draws a bit  $b$  at random. The adversary gets once more access to the oracles described above. It finally outputs a tuple  $(\pi_{\mathcal{P}}^i, m_0, m_1)$ , consisting of a party instance and two messages, such that:  $|m_0| = |m_1|$  and  $\pi_{\mathcal{P}}^i.\alpha = 1$ . The challenger uses its knowledge of  $\pi_{\mathcal{P}}^i$ 's state on input  $m_b$  (which is  $m_0$  or  $m_1$  depending on the challenger's hidden bit) to simulate outputting a message  $M_b$  which corresponds to the next protocol message of  $\pi_{\mathcal{P}}^i$  as that party would have sent it. Clearly if the protocol requires messages be sent in plaintext,  $M_b = m_b$ . The instance  $\pi_{\mathcal{P}}^i$ , as well as any of its partnering instances, are closed and may no longer be used in any oracle. The adversary may subsequently continue to use oracles at will (except on the instances closed above) and eventually outputs a guess  $d \in \{0, 1\}$ . We say the adversary *wins* if, and only if, the following conditions hold simultaneously:

- $d = b$ ;
- No `ORevealState` query was made for either  $\pi_{\mathcal{P}}^i$ , nor for any instance  $\pi_{\mathcal{S}}^j$  of the server such that  $\pi_{\mathcal{P}}^i$  and  $\pi_{\mathcal{S}}^j$  are partnered.