

Hash-Based TPM Signatures for the Quantum World

Megumi Ando¹(✉), Joshua D. Guttman¹, Alberto R. Papaleo¹, and John Scire²

¹ The MITRE Corporation, Bedford, MA, USA
{mando,guttman,apapaleo}@mitre.org

² Stevens Institute of Technology, Hoboken, NJ, USA
jscire@stevens.edu

Abstract. Trusted Platform Modules (TPMs) provide trust and attestation services to the platforms they reside on, using public key encryption and digital signatures among other cryptography operations. However, the current standards mandate primitives that will be insecure in the presence of quantum computers. In this paper, we study how to eliminate these insecure primitives. We replace RSA-based digital signatures with a hash-based scheme. We show that this scheme can be implemented using reasonable amounts of space on the TPM. We also show how to protect the TPM from rollback attacks against these state-sensitive signature operations.

Keywords: Post-quantum · Trusted Platform Module (TPM) · Attestation Identity Key (AIK) · Merkle trees

1 Introduction

A Trusted Platform Module (TPM) is a low-cost cryptographic microprocessor for enabling trusted computing functionalities. TPMs follow a set of global industry standards laid out by the Trusted Computing Group (TCG) also called the TCG standards (versions 1.2 and 2.0). As of this writing, there are a number of vendors which supply hardware TPMs: AMD, Atmel, Broadcom, IBM, Infineon, Intel, Lenovo, National Semi, Nationz, Qualcomm, STMicroelectronics, Samsung, Sinosun, Texas Instruments, Nuvoton Technology, and Fuzhou Rockchip [24]; and most personal computers and a few mobile devices (with the notable exception of Apple computers) are shipped with a TPM [21]. It is assumed that a TPM is implicitly trustworthy and tamper-resistant. This trust is bootstrapped to enable two trusted computing capabilities, secure encrypted storage and remote attestation, explained below.

1.1 Secure Encrypted Storage

Costing a couple of dollars each, a TPM is an inexpensive solution for securing sensitive data even under the threat that the hosting platform may become

corrupted [21]. This is achieved by storing sensitive information outside the TPM only in encrypted form, with the corresponding encrypting TPM key either in the TPM’s hardware-protected area or also encrypted by another TPM key. In this way, a TPM serves as a tamper-resistant hardware Root-of-Trust for Storage (RTS) [2, 14, 21].¹

Because the TPM has a limited amount of non-volatile RAM (NVRAM), only a few special TPM keys are kept in protected space within the TPM. Other keys are stored in encrypted form in unprotected areas on the platform. They exist in unencrypted form only transiently as needed and only within the TPM. The storage keys thus form a key hierarchy (a tree), where each non-root key is encrypted under its parent key.

In TPM version 2.0, the encryption of child nodes can be done either symmetrically or asymmetrically [2, 21]. Thus, we can trivially avoid a quantum-insecure storage mechanism by using only symmetric storage keys. A single caveat to the solution is duplication, or a mechanism for transferring or backing up TPM keys. Currently, duplication is handled by decrypting a specific key (or even an entire subtree) and then re-encrypting it with the public portion of another TPM’s RSA key [2, 14, 21]. Symmetric encryption cannot replace RSA encryption without preplacing shared TPM keys across platforms.

1.2 Attestation

A platform state measurement indicates whether a platform is in an expected state and therefore still trustworthy. Through a separate mechanism, the hash of this measurement can be written into special TPM registers called Platform Configuration Registers (PCRs). This is done using a trusted hardware component: a Root-of-Trust for Measurement (RTM; e.g., Intel TXT for boot measurements) [20].

A TPM can be used for securely reporting the measurement stored in the PCRs [2, 10, 14, 17, 20, 21]; and so it serves also as a hardware Root-of-Trust for Reporting (RTR). This is done using two types of TPM keys: an Attestation Identity Key (AIK) and an Endorsement Key (EK). An AIK is an asymmetric signature key corresponding to a user and or application on the platform and is used for enabling anonymity. There may be many AIKs per platform. To be useful, an AIK must satisfy the property that it is infeasible for an adversary to forge a valid signature (under the key) using any information stored on the platform outside of the TPM. An EK is an asymmetric encryption key with the unforgeability property that it is infeasible for an adversary to create a valid encryption under the key. It is unique to the platform, created randomly, and properly certified by the manufacturer. The EK is kept inside the TPM, never leaves the TPM, and is used for certifying AIKs as genuine.

The AIK is used to sign a quote, which includes the measurements stored in the PCRs as well as a verifier-provided nonce for freshness. For the signed

¹ Note that while a TPM can guarantee confidentiality and detect modification of securely stored data, it cannot retrieve secured data in the event that data is damaged; some other mechanism should be implemented to mitigate data loss.

quote to be meaningful, the AIK itself is certified in order to prove that the AIK belongs to a genuine TPM. The signed quote, along with the AIK’s certificate, is sent to a verifying party.

Each TPM is shipped with a vendor-certified EK (version 1.2) or Primary Seed (version 2.0) [2, 14, 21]. All AIKs are securely derived from this key or seed as needed and certified by a trusted Privacy-CA (PCA). The Direct Anonymous Attestation (DAA) protocol also ensures anonymity [6]. This alternative does not require a trusted third party but incurs a large cost in cryptographic complexity instead.

1.3 Our Contributions

Current TPM solutions were designed to be secure against classical adversaries, as opposed to quantum adversaries. This is problematic since some of the trust arguments rely on the intractability of computational problems with known efficient quantum algorithms. For example, TPM keys were originally (in version 1.2) all RSA keys, which are insecure given Shor’s quantum integer factorization scheme [22], and both PCA and DAA protocols use RSA encryption.

This paper presents our initial investigation into architecting a quantum-secure TPM. We have identified the following set of current TPM mechanisms to be those which rely on RSA and, thus, known to be vulnerable to quantum-attacks:

1. Secure (encrypted) storage
2. Duplication
3. TPM signing keys
4. PCA protocol
5. DAA protocol
6. Encrypted transport session

(By encrypted transport session, we mean transporting information, e.g., a TPM command, securely to a TPM.)

Although RSA is insecure in a post-quantum era, there are some notable alternatives without any known attacks. One such alternative is hash-based digital signature schemes, such as Merkle’s tree authentication using a one-time signature (OTS) scheme. We claim that Merkle tree signatures are a practical alternative to RSA authentication, solving issues 3 and 4 above. (**Disclaimer:** The signing-based PCA protocol described in this paper is weaker than the original; we lose repudiability of a claim that two AIKs are linked.)

In this paper:

- We present our hash-based TPM signing key construction: QUAntum Secure Hash (QUASH). Given the space limitation of the TPM, our solution offloads most of the storage to the untrusted platform in a way that preserves security.
- We also show that our solution prevents replay attacks; it prevents a particular OTS key from being used to sign multiple messages.
- Lastly, we provide recommendations for system parameters.

1.4 Related Work

Efficient quantum algorithms [11, 22] break RSA, Digital Signature Algorithm (DSA), and Elliptic Curve Digital Signature Algorithm (ECDSA) in a quantum-world setting [4]. Fortunately, a number of cryptographic techniques are believed to be quantum-secure despite these known algorithms, including: hash-based, code-based, lattice-based, and multivariate-quadratic-equations cryptography [1, 4]. In this paper, we propose a hash-based TPM signing solution, whose security relies solely on the existence of collision-resistant hash functions and pseudorandom number generators and assess its practicality as a system. Our approach for storing data outside the TPM is similar to the technique presented in [19] for creating virtual monotonic counters for TPMs.

Lattice-based cryptography provides a strong alternative candidate for enabling quantum-secure TPM signatures. A recent paper [1] presents an efficient signature scheme Tightly-secure Efficient Signatures from standard LAttices (TESLA) that relies on the intractability of standard lattices as opposed to ideal lattices. The paper presents *fixed* parameters for quantum-security. Our approach uses an efficient, stateful hash-based signature scheme (e.g., [7, 8, 23]) as a primitive. It has the benefit that its security is based solely on the existence of collision-resistant hash functions and pseudorandom number generators, whereas the lattice-based approach relies on the intractability of lattice problems. However, our hash-based approach has a key management problem that the lattice-based approach avoids. A more in-depth comparison between our approach and one that uses this lattice-based signature scheme is outside the scope of this paper.

Road Map: Sect. 2 contains our problem statement (the definitions and system model we adopt for our construction) and the technical background necessary for understanding our post-quantum attestation solution. In Sects. 3 and 4, we provide our main results: our hash-based TPM signing key construction, QUASH, and a practicality assessment of this solution. In Sect. 5, we conclude with a summary of our work.

2 Problem Statement and Preliminaries

Our goal is to architect hash-based TPM signing keys, defined below.

Unlike the RSA signature scheme, the most efficient hash-based signature schemes maintain some key state information to work.² At a minimum, the leaf number is needed in order to avoid reusing an OTS key pair; and some auxiliary information is required for efficient signing. We also introduce the notion of an Endorsement *Signing Key* (ESK), which is a vendor-certified TPM *signing* key for certifying AIKs.

Definition 1. *A tamper-evident hash-based AIK is a TPM signing key with the additional property that it is infeasible for an adversary to change an AIK state value without the change being detected.*

² A recent paper [5] presents a practical stateless hash-based signature scheme.

Definition 2. *The hash-based ESK is an unforgeable TPM signing key derived from a Primary Seed, with the property that it is infeasible for an adversary to change the ESK state value. (A Primary Seed is a large random number which is created by the vendor and protected by the TPM.)*

If the number of AIKs is limited and relatively small, we can simply store all required key state information within the protected NVRAM space of the TPM. Our goal is to extend this solution to the case where an *unbounded* number of AIKs can be created, despite the hard space limit within the NVRAM. We do this by storing the key state information outside of the TPM in balanced binary (e.g., red-black trees or treaps) hash trees and by keeping only digests of the state information within the NVRAM for integrity protection.

Our construction, QUASH, has the following set of desirable properties.

- **Tamper-evidence: prevention of OTS key reuse.** The key state information is integrity-protected by storing hash digests in the TPM NVRAM. Thus, our tamper-evident signing solution prevents an OTS key from being used more than once, a concern if we use hash-based signatures.
- **Availability: localization of data loss.** While a TPM can guarantee confidentiality and data modification detection of securely stored data, it cannot retrieve secured data in the event that encrypted data is lost. Thus, a modification to any of the AIKs’ state information renders *multiple* AIKs unusable. Given “registers” for holding multiple integrity check values within the TPM’s NVRAM, we minimize the number of AIKs that are lost when the state information of a single AIK is corrupted.
- **Efficiency: TPM space requirements.** We show that our construction can work with a space-limited TPM, despite using a hash-based signature scheme requiring key state storage and producing large signatures.
- **Efficiency: integrity checks.** Looking up and updating key state information on the platform is efficient. Our solution uses a balanced binary tree structure that requires only $\mathcal{O}(\log N)$ steps, where N denotes the number of AIKs. More critically, the integrity checks executed by the TPM are efficient, also requiring only $\mathcal{O}(\log N)$ steps.
- **Efficiency: AIK re-generation.** Hash-based signatures require generating fresh OTS keys post-provisioning since only a finite number of OTS keys can be created at set-up time. We amortize the key re-generation time by creating fresh OTS keys for a subsequent Merkle tree during signing.

2.1 System Model

We assume that the TPM is trusted and tamper-resistant, but the hosting platform is untrusted. The TPM’s functionalities are augmented to include hash-based capabilities. Specifically, it can create a Merkle signature scheme key, integrity check a current AIK state, and sign a message using such an integrity-checked current key state. Additionally, its NVRAM is equipped with a small number of special registers, which we call Integrity Registers.

In today’s implementation, the TPM’s NVRAM is used for storing root keys for certificate chains, the EK (an RSA encryption key), the expected measurement of the machine launch state, and decryption keys used before the disk is made available [9]. The minimum NVRAM size required by the TPM (version 1.2) spec is 1280 bytes [18]. We assume that the TPM RAM size is also fairly limited.

2.2 Merkle Tree Authentication

Merkle tree authentication is specified by two algorithms: a one-time signature (OTS) scheme and a Merkle signature scheme. Merkle signature schemes differ from each other in their traversal algorithms for constructing the “next authentication path” efficiently. We describe the Merkle signature scheme generically, intentionally hiding the details for the traversal algorithm for readability, and also because it is well-understood how the traversal time and storage size vary depending on the algorithm used, see [4].³

Moreover, real experimental results reveal that the overall performance of a tree authentication scheme (using a state-of-the-art traversal algorithm) is dominated by the performance of the underlying OTS scheme [16]. While the original Lamport-Diffie scheme (LD-C in this paper) was shown to be optimal in number of hash computations [3], LDWM is the usual go-to OTS scheme, because it decreases both the signature and storage sizes.

Within the context of this paper, $h(\cdot)$ is a cryptographic one-way hash function, and $r(\cdot)$ is a cryptographically secure pseudorandom number generator (PRNG) as defined below.

Definition 3. A cryptographic hash function $h : \{0, 1\}^* \rightarrow \{0, 1\}^n$ maps arbitrary length strings to strings of length n , such that the following properties are satisfied:

- (Easy to compute) Given any $x \in \{0, 1\}^*$, it is easy to compute its hash $y = h(x)$.
- (Pre-image resistance) Given any y in the image, it is computationally infeasible to find any x such that $h(x) = y$.
- (Second pre-image resistance) Given a $x_1 \in \{0, 1\}^*$, it is computationally infeasible to find any $x_2 \neq x_1$ such that $h(x_2) = h(x_1)$.
- (Collision resistance) It is computationally infeasible to find any pair $x_1, x_2 \in \{0, 1\}^*$ such that $h(x_2) = h(x_1)$.

Definition 4. A cryptographically secure pseudorandom number generator

$$r : \{0, 1\}^m \rightarrow \{0, 1\}^n$$

is a function that generates an n -bit output from a truly random m -bit seed and satisfies the next-bit test: Given the first polynomial number of output bits, it is computationally infeasible to predict the next bit of the output with probability non-negligibly larger than $\frac{1}{2}$.

³ Merkle’s original construction requires $\mathcal{O}(H^2)$ space and $\mathcal{O}(H)$ time [15], but recent constructions provide more efficient results [7, 8, 13, 23].

Setup. A Merkle signature scheme (MSS) public key is equivalent to the result of the following computation. First we initialize a tree of constant branching factor D and height H and generate D^H OTS key pairs. The OTS key pairs are derived from the hash of the AIK name id , the re-generation number i , the leaf index ℓ , and the private-value index j : The AIK seed for AIK id is

$$s_{\text{AIK}}(id) = h(id);$$

the seed for the i -th Merkle tree for id is

$$s_{\text{MSS}}(id, i) = s_{\text{AIK}}(id)|i = h(id)|i;$$

and the seed for the ℓ -th leaf of the i -th Merkle tree for id is

$$s_{\text{LEAF}}(id, i, \ell) = s_{\text{MSS}}(id, i)|\ell = h(id)|i|\ell.$$

The concatenation of a leaf seed and a private-value index is inputted into the pseudorandom number generator $r(\cdot)$ to create an OTS private key value.⁴ We store the hash of the ℓ -th OTS public key in the ℓ -th leaf of the tree, and the value of each non-leaf node is computed as the hash of the concatenation of its children's stored values. The MSS public key corresponds to the root value of this tree.

The approach above—computing values for all D^H leaves and subsequently hashing up the Merkle tree—is unnecessarily space consuming in practice. The convention is to keep a much smaller number of hashes in a stack, see [4].

Signing. Given a leaf in the Merkle tree, its authentication path is the sequence of sibling nodes along the path from it to the root. For example, the authenticating path for leaf 12 in Fig. 1(a) is $(\nu(11), \nu(6), \nu(1))$, where $\nu(i)$ denotes the hash value at node i .

An MSS signature is valid iff two conditions hold: the OTS signature σ verifies, and the OTS public key value \mathbf{y} is consistent with the MSS public key Y and an authentication branch B . Thus, an MSS signature Σ is the quadruple:

$$\Sigma = (\ell, \sigma, \mathbf{y}, B), \tag{1}$$

where ℓ denotes the leaf index number. The OTS public key \mathbf{y} and signature σ can be computed from the OTS setup and signing algorithms, respectively. A traversal algorithm is implemented to compute the next authentication path B , and some key state information is saved to do this efficiently.

Because MSS setup is time-consuming, we recommend amortizing the cost of generating the next Merkle tree during signing. A single call to the signing algorithm should produce a signature, while doing a little bit of computation for incorporating *one* additional leaf in setting up the next tree, see Fig. 1(b). We

⁴ In general, the OTS keys do not have to be derived from the same parent seed. We do so here to save storage space.

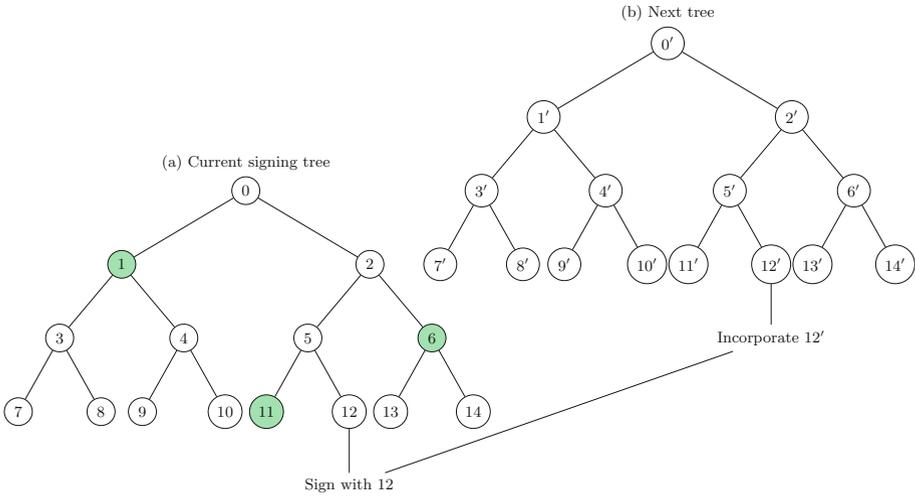


Fig. 1. (a) The nodes for node 12’s authentication path are in green. (b) Computation for the next tree during is amortized during signing. During signing using leaf i , leaf i' for the next tree is created and incorporated into the next tree’s stack.

could increase or decrease the key re-generation rate according to how fast the OTS keys are being used. This would require some other mechanism for determining the key usage rate and slight modifications to the `key_state` structure below.

The `key_state` object has the following fields:

- `ID` \triangleq TPM signing key id
- `parent` \triangleq parent TPM signing key id
- `iter` \triangleq tree iteration or re-generation number
- `height` \triangleq height of Merkle tree
- `bfactor` \triangleq branching factor of Merkle tree
- `ots` \triangleq OTS scheme used
- `lparam` \triangleq loop parameter for the OTS scheme, see appendix
- `leaf` \triangleq leaf number (for signing)
- `stack` \triangleq stack for next tree
- `state` \triangleq initial state for next tree
- `cur_state` \triangleq saved state for updating the current tree’s auth. path

`ID` stores the MSS key’s unique id. `parent` stores the id of the key’s parent. `iter` stores the re-generation number. `height` and `bfactor` store the tree’s height and constant branching factor, respectively. `ots` and `lparam` store the OTS scheme and its corresponding loop parameter. `leaf` stores the leaf index. `stack` stores the stack for setting up the next tree. `state` stores the initial saved state for the next tree, which includes the first authentication path. `cur_state`

stores auxiliary state structures for tree traversal algorithms that are more optimal and may include the authentication path for the next signature. An instance of a `key_state` object is initialized with `iter` and `leaf` set to 1; `lparam` set to \perp ; and with `state` and `cur_state` empty.

3 QUANTUM SECURE HASH (QUASH)

The MSS setup algorithm must be executed within the TPM. Otherwise malicious software may swap out the real public key with any value, including a MSS public key created from its choice for the seed. The MSS signing algorithm must also be executed within the TPM, since the purpose of having a TPM signing key is to prove (to some verifier) that some function (e.g., construction of a quote) was carried out by a trusted component.

However, (stateful) hash-based digital signature schemes are space-intensive, requiring key state and producing large signatures. Thus, the design challenge is to offload the key state storage to the untrusted platform in a way that preserves security. We outline the main design challenges here:

- **Key state must be stored outside TPM:** A single `key_state` object for an AIK of branching factor 2 and height 20 takes 4KB of space, see Table 1 in Sect. 4. To support even twenty such AIKs, either the TPM’s NVRAM size must increase by an order of magnitude, or we must store the required key state information outside of the TPM in such a way that the overall solution remains secure.
- **Preventing rollbacks:** Rollbacks must be carefully managed, since the AIK key states are stored on the untrusted platform, and an OTS key is secure only if it is used once.⁵
- **Efficient integrity checks:** To prevent rollback attacks and issues from unintended data corruptions, we keep hash chains of the AIKs’ key states in the protected space. These hashes are kept in special registers in the TPM’s NVRAM and are used for integrity checks prior to key creation and signing. Since the integrity checks must be computed by the resource-limited TPM, the naive approach of computing the hash chain directly from the AIKs’ key states is impractical; a more efficient method is needed.
- **Resiliency from data corruption:** Each TPM Integrity Register holds a compressed representation of a group of AIKs’ key state information. Thus, damage to a single key state in the group could render all AIKs in the group useless. Our solution allows for recovery of AIKs and TPM registers where possible and minimizes the number of AIKs that are unrecoverable.

⁵ Encrypting the AIKs’ state information does not prevent rollbacks; an adversary could restore an AIK to a previous state by writing over its current encrypted state with a saved previous encrypted state.

3.1 Data Structures

State Storage Trees. The AIKs’ state information is organized in balanced binary trees outside of the TPM. Each node in a State Storage Tree corresponds to an AIK and contains the information necessary for signing under this key. Specifically, each node stores a pair: a `key_state` object k (introduced in Sect. 2.2) and a subtree hash. The subtree hash is the hash of the concatenation of k and the subtree hashes of the node’s children.⁶

Integrity Registers. Digests of last good key states are stored in the TPM’s NVRAM in registers, which we call Integrity Registers.

Group Membership. The key states of a group of AIKs are hash-chained together to create a single hash digest. Thus, damage to any key state in the group can render all of the keys in this group useless. To mitigate this kind of data loss, an AIK belongs to multiple groups, where the hash chain of each group is stored in a separate Integrity Register. Given $(d \cdot I)$ Integrity Registers, each key state object belongs to I groups. Group membership is decided by the hash value of the key id. A key id is a member of group G_i^b if the i -th d -ary symbol of its hash is b :

$$G_i^b = \{id : i\text{-th } d\text{-ary symbol of } h(id) \text{ is } b\}. \quad (2)$$

For every group G_i^b , there is a corresponding State Storage Tree T_i^b for maintaining the state information of the AIKs in G_i^b and a corresponding Integrity Register R_i^b for storing the tree hash of T_i^b . We denote by `trees`(id) a list of State Storage Trees to which id belongs, ordered lexicographically:

$$\text{trees}(id) = \{T_i^b : id \in G_i^b\}. \quad (3)$$

For example, if $d = 2$, $i = 2$, and the hash of id ends in 01, id belongs to groups G_0^1 and G_1^0 . The State Storage Trees `trees`(id) = (T_0^1, T_1^0) and corresponding Integrity Registers R_0^1 and R_1^0 are shaded in Fig. 2 below.

3.2 AIK Methods

We describe our methods for creating and certifying AIKs and signing under AIKs using the data structures described in the previous subsection. In Sect. 3.3, we show that our solution is tamper-evident and thwarts forgeries. Our practicality assessment of our solution is given in Sect. 4.

⁶ The AIK seeds are never stored anywhere. The Primary Seed s_0 , i.e., the “unsalted ESK seed,” is the only seed stored in the TPM’s NVRAM. A TPM signing key seed is generated from s_0 as needed and only in the protected space in the TPM.) .

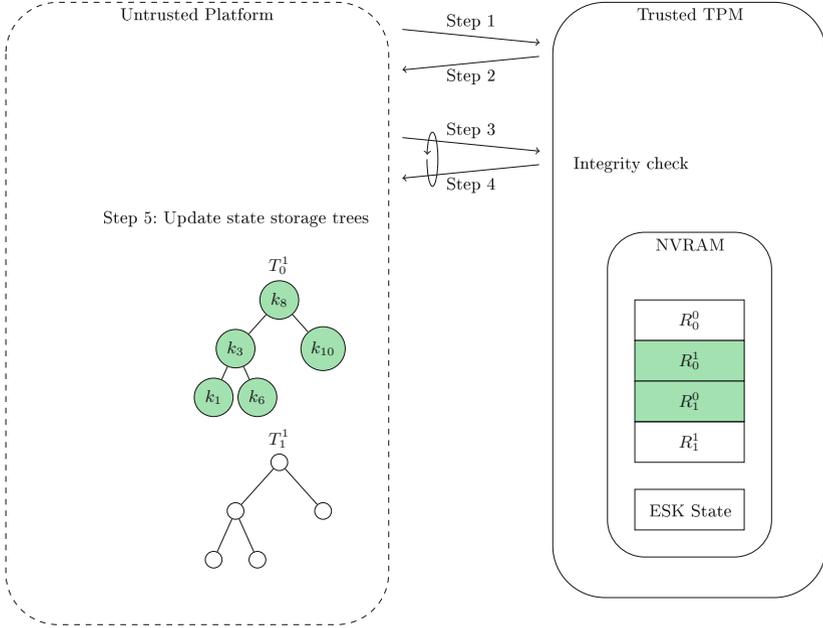


Fig. 2. State Storage Trees track AIK key state information. Integrity Registers in the TPM’s NVRAM store integrity check values.

AIK Creation. To create a new identity key id ,

1. The platform determines the list $\mathbf{trees}(id)$ and sends a request to create an AIK to the TPM via a TPM interface, such as the TCG Software Stack (TSS). The request contains id , the parent id, the OTS scheme, and the Merkle tree height and branching factor.
2. Upon receiving the request, the TPM computes the list $\mathbf{trees}(id)$ and sends an acknowledgement that the request was received to the platform.

Steps 3 and 4 are repeated for every $T_i^b \in \mathbf{trees}(id)$:

3. The platform computes what the authentication path for node $h(id) \in T_i^b$ will be after the node is inserted into the tree. It sends this authentication path to the TPM.
4. To check the integrity of the authentication path, the TPM ensures that (i) the hashes along the authentication path are consistent with each other, and (ii) the hash at the root is equal to the value stored in the Integrity Register R_i^b corresponding to T_i^b .
 - a. For the first authentication path that passes both integrity checks, the TPM instantiates a new `key_state` object k using the MSS setup method. The AIK seed s is determined from the parent signing key seed s_P and id , e.g., $s = s_P || id$.⁷ The TPM computes the hash of the key state k and returns k and the public key Y to the the platform.

⁷ If the parent key is not the ESK, this is determined recursively.

- b. If both integrity checks pass, the TPM updates the Integrity Register R_i^b using the hash of k ; otherwise, if either integrity checks fail, the Integrity Register is marked invalid.
 - c. If T_i^b is the final tree in $\mathbf{trees}(id)$: if a new key was successfully created, the public portion Y of the MSS key is signed by the ESK within the TPM. The TPM returns the ESK endorsement to the platform; otherwise, it returns an error.
 - d. If nothing has been sent yet this round, the TPM sends an acknowledgment to the platform that the round has ended.
5. The platform updates all the trees in $\mathbf{trees}(id)$ with the returned key state k .

Signing. The protocol for signing a message M using an identity key id is similar:

1. The platform determines the list $\mathbf{trees}(id)$ and sends a sign-message request to the TPM. The request contains id and the message M .
2. Upon receiving the request, the TPM computes the list $\mathbf{trees}(id)$ and sends an acknowledgement that the request was received to the platform.

Steps 3 and 4 are repeated for every $T_i^b \in \mathbf{trees}(id)$:

3. The platform computes the authentication path for node $h(id) \in T_i^b$. It sends this authentication path and the key state k stored at node $h(id) \in T_i^b$ to the TPM.
 4. To check the integrity of the authentication path, the TPM ensures that (i) the hash of k and the hashes along the authentication path are consistent with each other, and (ii) the hash at the root is equal to the value stored in the Integrity Register R_i^b corresponding to T_i^b .
 - a. For the first authentication path that passes both integrity checks, the TPM signs the message M using the MSS signing method and k .
 - b. If both integrity checks pass, the TPM updates the Integrity Register R_i^b using the updated MSS object k ; otherwise, if either integrity checks fail, the Integrity Register is marked invalid.
 - c. If T_i^b is the final tree in $\mathbf{trees}(id)$: if M was successfully signed, the TPM returns the signature and the updated key state k' to the platform; otherwise, it returns an error.
 - d. If the signature used the “last leaf” in the current Merkle tree, the public portion Y of the next MSS key is signed by the ESK and this newly created ESK endorsement is also sent to the platform.
 - e. If nothing has been sent yet this round, the TPM sends an acknowledgment to the platform that the round has ended.
5. The platform updates all the trees in $\mathbf{trees}(id)$ with the returned key state k .

Signing a Key Handle. Each AIK must be certified as having been created by a legitimate TPM. This is done by signing each AIK public value using the ESK. The ESK itself is certified by the vendor, perhaps using a message authentication code (MAC) derived from the Primary Seed s_0 .⁸

The ESK is a TPM hash-based signing key whose seed is derived from s_0 and some user-inputted salt. Both s_0 and the ESK state information is kept in the TPM's NVRAM and never leaves the protected space.

Key and Register Recovery. The TPM itself does not provide a mechanism for recovering data lost during storage, transmission, or processing. Even a low bit error rate can render all AIKs useless without some means of recovering AIKs (when possible) and reverting Integrity Registers to useable states. Here, we describe how this can be accomplished.

First, for each State Storage Tree, the platform computes the hash chain value from scratch (i.e., from the key states, as opposed to from the subtree hashes) and sends an AIK recovery request to the TPM along with the tree hashes. The TPM responds with the set B of State Storage Trees which did not pass the integrity checks and resets the corresponding Integrity Registers. (By resetting, we mean that the registers reflect the valid integrity value corresponding to an empty State Storage Tree.) For every tree $t \in B$, for every node $v \in t$, the platform checks if there exists a State Storage Tree $t' \notin B$, and $t' \in \mathbf{trees}(id)$, where id is the id of v . Note that this can be computed efficiently from the hash of id . If t' exists, the platform marks the AIK as recoverable.

For every tree $t \in B$ and recoverable node $v \in t$, the platform sends to the TPM a request to add node v to t , along with an authenticating path for v in t and a proof of correctness for v (i.e., a tree $t' \notin B$, and $t' \in \mathbf{trees}(id)$ and an authenticating path for v on t'). If the proof verifies, the TPM updates the corresponding Integrity Register accordingly; otherwise, the register is marked as invalid.

From above, we see that AIKs in damaged State Storage Trees can be recovered from undamaged State Storage Trees. An AIK id is recoverable iff there exists a State Storage Tree in $\mathbf{trees}(id)$ that passes all integrity checks. In other words, an AIK is *unrecoverable* iff

$$t \in \mathbf{trees}(id) \implies t \in B.$$

For a fixed number N of AIKs, increasing d and I decreases the chance of an AIK being unrecoverable. However, this obviously increases the number of Integrity Registers in the TPM's NVRAM.

3.3 Correctness and Security Proofs

We provide sketches of correctness and security proofs for our construction here.

⁸ Note that ESK may periodically require a fresh certificate.

Lemma 1. *If an authenticating path passes both integrity checks (step 4 for both AIK setup and signing), then the path does not contain any unauthorized modifications.*

Proof. If the hash value in the Integrity Register holds the intended value, then in order for an adversary to modify the path, it must do so in such a way that the final tree hash at the root retains this intended value. This is infeasible by the (second pre-image resistance) property of the hash function, see Def. 3.

Moreover, if an Integrity Register holds a valid value, it must be the intended value; since the Integrity Register is protected by the TPM, and the TPM updates the register with a valid value iff all integrity checks on the corresponding authenticating branch pass. (**Remark:** Thus when an integrity check fails for a State Storage Tree T_i^b , the corresponding Integrity Register R_i^b must be marked invalid in order for Lemma 1 and Corollaries 1 and 2 to hold, see step 4b in AIK creation and signing.) \square

Corollary 1 (Tamper-resistance). *Any unauthorized modification of an AIK’s key state information (stored outside the TPM) is detectable.*

Corollary 2 (Unforgeability). *It is infeasible for an adversary to forge a signature under a properly certified AIK.*

Theorem 1 (Correctness). *Let id be any hash-based AIK with at least one authentication path which passes both integrity checks. Given any message M of any arbitrarily length, a signature on M under AIK id verifies.*

Proof. Let p be the first authenticating path that passes both integrity checks. From Lemma 1, p does not contain any unauthorized modifications. Thus, the AIK’s key state is the intended key state, and so any signature under it verifies. \square

4 Practicality Assessment

4.1 Space Analysis

We first determine how much space is needed for storing a key state object. The hash output size is assumed to be (256 bits or) 32 bytes. Eight of the `key_state` object fields—`ID`, `parent`, `iter`, `height`, `bfactor`, `ots`, `lparam`, and `leaf`—take up very little space. For concreteness, we have allotted 4 bytes for each of these fields; so $8 \times 4 = 32$ bytes are needed for storing all of these fields.

The maximum total storage size is dominated by the storage requirements for the remaining fields—`stack`, `state`, and `cur_state`—which depend on the traversal algorithm being used and the dimensions of the Merkle tree. We use the state-of-the-art traversal scheme from [8]; and our Merkle tree is binary with height $H = \text{height}$.

During Merkle tree setup, an initial `state` consisting of H “authentication” nodes, $H - 2$ “treehash” nodes, and a single “retain” node are stored [8],

hence, $2H - 1$ nodes. The `stack` for this computation requires a maximum of H nodes.

Additionally, the traversal algorithm requires storing at most $3.5H - 4$ nodes in `cur.state` at any given step [8]. Thus, the total maximum space (in bytes) required for storing a key state object is bounded by $\Gamma(H) =$

$$32 + (32 \times (6.5H - 5)) = 208H - 128. \quad (4)$$

Table 1 shows the key state storage size for binary Merkle trees of various sizes.

Table 1. Key state storage size in KB.

		Targeted Range		
		$H = 15$	$H = 20$	$H = 25$
# OTS keys	2^H	$3 \cdot 10^4$	$1.05 \cdot 10^6$	$33.55 \cdot 10^6$
Storage size $\Gamma(H)$	$208H - 128$	3.0 KB	4.0 KB	5.1 KB

TPM Space for AIK Creation. Here, we provide an estimate for the TPM space needed to create an AIK. This estimation is meant to approximate the maximum space used by a space-efficient implementation.

In the analysis below, a unit is 32 bytes. For simplicity, we assume that the Merkle trees are binary of height H ; we use the traversal algorithm described in [8] and the space-efficient [12] SHA-256 for the hash function; the space required for running the PRNG likewise is minimal; and the length of each OTS private x -value is 1 unit.

In creating a TPM signing key, the maximum resident TPM space is needed when the public key is being signed. We estimate the space required as follows:

- 1 unit for storing the public key (input);
- $\Gamma(H) = 6.5H - 4$ units for storing the ESK key state (input), see (4);
- 1 unit for storing the leaf index ℓ (output);
- $N(ots)$ units for the maximum OTS signature σ size (output);
- $N(ots)$ units for the maximum OTS public-key portion \mathbf{y} size (output); and
- H units for the maximum authentication path size (output);

where the loop parameter $N(ots)$ depends on the OTS scheme. At this point, we need not retain the input authentication path, the newly created key state, its hash, nor the updated register value. (**Remark:** In order to save TPM space, the newly created AIK state is returned to the platform before signing the retained public key value using the ESK state, see **AIK Creation** in Sect. 3.2.)

We need only negligible scratch space for running the MSS signing method: Only negligible additional space (beyond the space allotted for storing the output) is required for computing the OTS signature and public-key. Likewise, only negligible additional space (beyond space for storing the inputted ESK state) is required for updating `cur.state`, `state`, and `stack`.

Table 2 shows the maximum TPM space needed for creating an AIK of various sizes. Note that the total space required is not impacted by the number N of AIKs; (unless N is impractically large).

Table 2. TPM space required for AIK creation in KB: In LD-C, the complement of the message is also encoded to prevent man-in-the-middle attacks. LD-Z [15] encodes the zero count of the message digest to prevent forgeries.

	LD-C	LD-Z	LDWM $w = 2$	LDWM $w = 4$	LDWM $w = 8$
$H = 15$	36.4	20.5	12.1	7.9	5.8
$H = 20$	37.6	21.7	13.3	9.1	7.0
$H = 25$	38.8	22.9	14.5	10.3	8.2

TPM Space for AIK Signing. The maximum space in the TPM needed for signing a message is identical to that of setup.

4.2 Time Analysis

Since the State Storage Trees are balanced binary hash trees,

- Determining (the key state and) the authentication path takes $I \cdot \log(N)$ time, where N is the number of AIKs, and an id belongs to I groups.
- Determining the root value from (the key state and) the authentication path also takes $I \cdot \log(N)$ time.
- To update a State Storage Tree, only the nodes along a path to the root need to be reevaluated. Moreover, reevaluating a subtree hash requires only one hash operation. So, updating the State Storage Trees takes $I \cdot \log(N)$ time.

There are traversal algorithms that require $\mathcal{O}(H)$ space and $\mathcal{O}(H)$ time with low constant factors [8, 23]. Therefore, by using an optimal traversal algorithm, the cost of MSS signing is dominated by the cost of executing the underlying OTS scheme once. Amortizing the re-generation cost only increases the signing time by roughly a factor of two.

The only time-inefficient step is MSS setup, which necessarily requires time proportional to the number of OTS keys created at setup time. For the ESK, this can occur at provisioning time. For AIKs, we would like to set up on-the-fly, as needed. To mitigate this setup cost, we can setup a new AIK with a small Merkle tree and ramp up its size for subsequent tree(s).

5 Conclusion

We conclude that our construction for hash-based TPM signing keys is practical, if a state-of-the-art traversal algorithm, such as [8], is implemented.

QUASH can be implemented with a fairly small NVRAM: If an ESK Merkle tree includes roughly 6.5 million OTS key pairs, it can certify 100 AIKs, each requiring at most ($2^{16} =$) 65 thousand re-generations without needing to be re-certified by the vendor. This can be accomplished by setting a binary Merkle tree height to 23. Given these Merkle tree parameters, the maximum storage space required for storing the `key_state` object is 4.7 KB. Supporting 6 Integrity Registers would require an additional 192 bytes.

Acknowledgment. The authors would like to thank Anna Lysyanskaya for her suggestion on how to derive the ESK from a Primary Seed and Chris Eliopoulos, Alicea, Joseph J. Ferraro, John D. Ramsdell, and the anonymous reviewers for helpful comments.

References

1. Alkim, E., Bindel, N., Buchmann, J., Dagdelen, Ö.: TESLA: Tightly-Secure Efficient Signatures from Standard Lattices. Cryptology ePrint Archive, Report 2015/755 (2015)
2. Arthur, W., Challener, D., Goldman, K.: A Practical Guide to TPM 2.0: Using the Trusted Platform Module in the New Age of Security. Apress, Berkeley (2015)
3. Barak, B., Mahmoody-Ghidary, M.: Lower bounds on signatures from symmetric primitives. In: Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science, pp. 680–688, October 2007
4. Bernstein, D.J., Buchmann, J., Dahmen, E. (eds.): Post-Quantum Cryptography. Springer Science & Business Media, Heidelberg (2009)
5. Bernstein, D.J., et al.: SPHINCS: practical stateless hash-based signatures. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 368–397. Springer, Heidelberg (2015)
6. Brickell, E., Camenisch, J., Chen, L.: Direct anonymous attestation. In: Proceedings of the 11th ACM Conference on Computer and Communications Security, CCS 2004, pp. 132–145. ACM, New York, NY, USA (2004)
7. Buchmann, J., Dahmen, E., Hülsing, A.: XMSS—a practical forward secure signature scheme based on minimal security assumptions. In: Yang, B.-Y. (ed.) PQCrypto 2011. LNCS, vol. 7071, pp. 117–129. Springer, Heidelberg (2011)
8. Buchmann, J., Dahmen, E., Schneider, M.: Merkle tree traversal revisited. In: Buchmann, J., Ding, J. (eds.) PQCrypto 2008. LNCS, vol. 5299, pp. 63–78. Springer, Heidelberg (2008)
9. Challener, D., Yoder, K., Catherman, R., Safford, D., Van Doorn, L.: A Practical Guide to Trusted Computing. Pearson Education, Upper Saddle River (2007)
10. Coker, G., Guttman, J., Loscocco, P., Herzog, A., Millen, J., O’Hanlon, B., Ramsdell, J., Segall, A., Sheehy, J., Sniffen, B.: Principles of remote attestation. *Int. J. Inf. Secur.* **10**(2), 63–81 (2011)
11. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing, STOC 1996, pp. 212–219. ACM, New York, NY, USA (1996)
12. Ideguchi, K., Owada, T., Yoshida, H.: A study on RAM requirements of various SHA-3 Candidates on Low-cost 8-bit CPUs. IACR Cryptology ePrint Archive (2009)
13. Jakobsson, M., Leighton, T., Micali, S., Szydlo, M.: Fractal merkle tree representation and traversal. In: Joye, M. (ed.) CT-RSA 2003. LNCS, vol. 2612, pp. 314–326. Springer, Heidelberg (2003)
14. Kinney, S.: Trusted Platform Module Basics: Using TPM in Embedded Systems. Elsevier Inc., Burlington (2006)
15. Merkle, R.C.: Advances in Cryptology—CRYPTO 1989 Proceedings, chapter A Certified Digital Signature, pp. 218–238 (1990)
16. Naor, D., Shenhav, A., Wool, A.: One-Time Signatures Revisited: Have They Become Practical? IACR Cryptology ePrint Archive (2005)

17. Parno, B., McCune, J.M., Perrig, A.: Bootstrapping trust in commodity computers. In: 2010 IEEE Symposium on Security and Privacy (SP), pp. 414–429. IEEE, May 2010
18. Parno, B., McCune, J.M., Perrig, A.: *Bootstrapping Trust in Modern Computers*. Springer Science & Business Media, New York (2011)
19. Sarmanta, L.F., van Dijk, M., O'Donnell, C.W., Rhodes, J., Devadas, S.: Virtual monotonic counters and count-limited objects using a TPM without a trusted OS. In: *Proceedings of the First ACM Workshop on Scalable Trusted Computing, STC 2006*, pp. 27–42. ACM, New York, NY, USA (2006)
20. Scarlata, V., Rozas, C., Wiseman, M., Grawrock, D., Vishik, C.: *Trusted Computing: Ein Weg zu neuen IT-Sicherheitsarchitekturen*, chapter TPM Virtualization: Building a General Framework, pp. 43–56. Vieweg+Teubner (2008)
21. Segall, A.: *Trusted platform modules: When, Why, and How to Use Them*. Version: 21 June 2015
22. Peter, W.: Shor.: polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Rev.* **41**(2), 303–332 (1999)
23. Szydło, M.: Merkle tree traversal in log space and time. In: Cachin, C., Camenisch, J.L. (eds.) *EUROCRYPT 2004*. LNCS, vol. 3027, pp. 541–554. Springer, Heidelberg (2004)
24. TCG: TCG Vendor ID Registry, September 2015. <http://www.trustedcomputinggroup.org>