

Parakeet: Practical Key Transparency for End-to-End Encrypted Messaging

Harjasleen Malvai^{*†}, Lefteris Kokoris-Kogias^{‡§}, Alberto Sonnino^{‡¶}, Esha Ghosh^{||},
Ercan Oztürk^{**}, Kevin Lewi^{**}, and Sean Lawlor^{**}

^{*}UIUC, [†]IC3, [‡]Mysten Labs, [§]IST Austria, [¶]University College London (UCL), ^{||}Microsoft Research, ^{**}Meta

Abstract—Encryption alone is not enough for secure end-to-end encrypted messaging: a server must also honestly serve public keys to users. Key transparency has been presented as an efficient solution for detecting (and hence deterring) a server that attempts to dishonestly serve keys. Key transparency involves two major components: (1) a username to public key mapping, stored and cryptographically committed to by the server, and, (2) an out-of-band consistency protocol for serving short commitments to users. In the setting of real-world deployments and supporting production scale, new challenges must be considered for both of these components. We enumerate these challenges and provide solutions to address them. In particular, we design and implement a memory-optimized and privacy-preserving verifiable data structure for committing to the username to public key store.

To make this implementation viable for production, we also integrate support for persistent and distributed storage. We also propose a future-facing solution, termed “compaction”, as a mechanism for mitigating practical issues that arise from dealing with infinitely growing server data structures. Finally, we implement a consensusless solution that achieves the minimum requirements for a service that consistently distributes commitments for a transparency application, providing a much more efficient protocol for distributing small and consistent commitments to users. This culminates in our production-grade implementation of a key transparency system (Parakeet) which we have open-sourced, along with a demonstration of feasibility through our benchmarks.

I. INTRODUCTION

The use of end-to-end encrypted (E2EE) messaging ([29], [30]) ensures that conversations between users remain private. It is just as vital, though, to make sure these interactions happen with the right people. To avoid being compromised by a man-in-the-middle attack, users must discover the necessary public keys to identify each other. In reality, this is accomplished either by scanning QR codes in person to validate recipient public keys or by depending on a service provider to generate the relevant public key of a communication partner. The term *key transparency* refers to a system in which a service provider stores the public keys of individuals on a publicly accessible key directory server so that users can query this server for messaging individuals in their contact list.

However, since the server can supply users with outdated or fabricated keys, this strategy makes the server a single point

of failure. To counteract this, three complimentary methods have been discussed in existing literature. Firstly, users must be able to verify public keys served by the identity provider as well as their own key(s) on a regular basis. Secondly, the key changes must be publicly auditable, to ensure that the server is adhering to certain update rules. Finally, a consistency enforcement protocol must be used to prevent the server from serving different versions of the directory to different users.

To address the problem of providing auditable identity bindings, Chase et al. [15] introduced a primitive called *verifiable key directory* (VKD) and defined its required security properties. Their system, called SEEMless, instantiates a VKD and proves the correctness of the underlying scheme. Although the evaluation of SEEMless indicates that it veers toward practicality, SEEMless makes the following simplifying assumptions, omitting issues that must still be resolved in order to be applicable to a large-scale, real-world deployment. SEEMless assumes that there exists a mechanism for clients to obtain consistent views of a small server commitment, at every server epoch. Their implementation relies on local RAM storage to store all of the data ever required for key transparency. Scalable implementations need a separate, optimized storage layer. The data structure that SEEMless uses to keep state, which they call a *history tree*, assumes that it is feasible to perpetually store enough data to reconstruct the state of the server’s cryptographic data structure from any epoch. Even though their data structure is well-optimized to satisfy the requirement of storing this much data, it becomes a bottleneck when scaling to billions of users.

Hence, while a good starting point, these assumptions render SEEMless, as is, infeasible for large-scale deployments. For instance, WhatsApp, one of the most popular end-to-end encrypted messaging apps, has a monthly active user base of 2 billion [2] and has been downloaded on average 0.66 billion times [3] in the last five years. With the rollout of multi-device [1] where each user device has a separate key, the daily new key creations can be estimated to be roughly 5.4 million (based on an average of 3 devices/browsers per user and download numbers). If we assume that existing users update their keys at a similar volume, a rough estimate doubles this to approximately 10 million. This leads SEEMless’s data structure to require approximately 27TB (counting existing and new keys based on number of key updates/creations) of storage within the first year (See full version for details).

A. Our Contributions

In this work, we present Parakeet, a key transparency system which is designed specifically to overcome the scalability challenges that limit prior academic works from being practical for real-world encrypted messaging applications.

We achieve this by constructing a more efficient VKD, extending the VKD to support compaction, and then leveraging a simple consistent broadcast protocol to avoid the reliance on clients having to connect to blockchains (and the scalability issues associated with this dependence) to achieve consistency. Below, we elaborate on these three central contributions:

- 1) **Building an efficient VKD.** A VKD is a cryptographic primitive, defined and implemented by [15], which allows an identity provider to store an evolving set of label-value pairs, commit to this set, and respond (using cryptographic proofs) to client queries about this committed set and its updates. At a scale of billions of users, however, the efficiency requirements of a system implementing a VKD become more stringent than [15]’s implementation is able to meet. We address these limitations by implementing a cryptographic construction called an *ordered zero-knowledge set* (oZKS), which allows us to more efficiently realize a VKD with storage improvements of up to an order of magnitude when compared against existing solutions. In addition to the storage optimizations for the VKD, we also present a modular and flexible data-layer API, called StorageAPI, which can be implemented using any distributed database solution. Our VKD implementation, written in Rust, has been published as an open-source library [5].
- 2) **Supporting compaction.** Many works on verifiable data structures that support updates require append-only data structures ([40], [15]). In large-scale practical contexts, requiring the support of an ever-growing storage and in-memory system that supports append-only data structures can be a barrier for the deployment of key transparency. We introduce the term “compaction” to refer to an operation that allows reducing the data stored on the server by purging ancient and obsolete entries. Secure compaction loosens the requirement for append-only data structures, as introduced in [15], through a minimal additional assumption. By extending the existing oZKS functionality to enable secure deletions, we present our *two-phase compaction* paradigm of a VKD, which can be brought on par to existing works with comparable performance. The novelty of our oZKS construction consists of the mechanism for enabling secure deletions. We include a theoretical discussion of our construction here as we believe it is of independent interest for any applications which may be limited by the storage requirements of append-only verifiable data structures.
- 3) **Serving commitments.** All of the equivocation-based security definitions for transparency require users being able to access a small shared commitment. Some works assume a shared public ledger, such as a blockchain [40], [12]. Others rely on out-of-band gossip (see, e.g. [31]). However, at a scale of billions of users, these mechanisms all have

drawbacks. For example, across platforms and geographic regions, picking good out-of-band mechanisms, which do not result in disconnected partitions of sets of users, is a challenge. As [40] analyzed, using a blockchain to respond a large number of queries could eventually result in flooding the open ports of a large fraction of the nodes of this chain. The alternative described by [40] is to use a *header relay network*, i.e. a small number of nodes serving a specific transparency application, essentially, a centralized service. Since none of these solutions have tackled this problem at a scale of billions of clients, finding the minimal set of requirements for a consistency protocol remains an open question. To this end, we propose a lightweight *consensusless* consistency protocol that provides these guarantees with low performance overhead.

Paper organization. In Section II, we present an overview of the Parakeet system, which consists of the VKD construction, the consistency protocol and an interface between the two. In Section III, we review the original VKD definition and present a more storage-efficient construction based on an oZKS. In Section IV, we extend VKDs to support the compaction operation to address storage limitations in practical systems. In Section V, we describe the consistency protocol used by Parakeet to serve commitments. In Section VI, we provide microbenchmarks and comparisons to prior works, run against our open-source implementation of the system. We discuss related works in Section VII and conclude in Section VIII.

II. OVERVIEW

A central entity called the *identity provider* (IdP) keeps a local database linking users’ identifiers (e.g. phone numbers) to their public keys. It periodically sends a commitment of its latest state to a distributed set of authorities, called *witnesses*, that ensure its correct behavior. The witnesses store the latest commitment and also communicate with clients to make sure the identity provider is not censoring or eclipsing them. The *users* query the identity provider to look up the public key associated with a specific user identifier. Each user can additionally monitor the history of their own public key. The identity provider is trusted for ensuring that only permissioned users can perform particular actions.

A. System & Threat Model

Participants. Parakeet is run by the following participants:

- **Identity provider (IdP):** The entity keeping the identifier-to-key binding for every user and replying to users’ queries.
- **Witnesses:** A distributed set of authorities that cross-check the IdP.
- **Users:** Ask the IdP to store a specific key bound to their identifier, and query it to look up the key bound to specific identifiers. Users can also periodically check the history of their own key, to ensure that the server did not make any unwarranted changes to their key.

The IdP and each of the witnesses generate a key pair consisting of a private signature key and the corresponding public verification key, so that their identity is known.

By definition, an *honest* witness always follows the Parakeet protocol, while a *faulty* (or Byzantine) one may deviate arbitrarily. We present the Parakeet protocol for $3f+1$ equally-trusted witnesses, assuming a fixed (but unknown) subset of at most f Byzantine witnesses. In this setting, a quorum is defined as any subset of $2f + 1$ witnesses. (As for many BFT protocols, our proofs only use the classical properties of quorums thus apply to all Byzantine quorum systems [28].) When a protocol message is signed by a quorum of witnesses, it is said to be certified: we call such a jointly signed message a *certificate*. Additionally, we assume the network is fully asynchronous [19]. The adversary may arbitrarily delay and reorder messages; however, messages are eventually delivered.

We discuss the various guarantees for these parties below.

Properties. The consistency protocol of Parakeet satisfies the following properties:

- **Consistency:** For a given epoch t , if a user outputs a commitment com_t and a different user outputs a commitment com'_t , then $\text{com}_t = \text{com}'_t$.
- **Validity:** If an honest witness outputs com_t as valid, then com_t was proposed by the identity provider.
- **Termination:** If the identity provider runs the protocol honestly for epoch t , commitment com_t , eventually the identity provider will produce a certificate cert_t for t .

The security properties of the VKD of Parakeet are more complicated to formally state, so we defer a more formal treatment to appendix B and section III. At a high level, we require that the construction of a VKD satisfy:

- **Completeness:** If an identity provider honestly serves values, then, for any label label , registered with the identity provider, all users should receive (and accept) consistent views of the value associated with label .
- **Soundness:** As stated before, the soundness definition used in this paper is that of non-equivocation, i.e., if, at an epoch t , Bob accepts a value val as Alice’s key, the server cannot convince Alice that her key was val' for the same epoch t with $\text{val}' \neq \text{val}$. Note, this is in the presence of auditors.
- **Privacy:** Privacy for a VKD is defined with respect to all parties who are not the identity provider itself. The responses to all API calls made by parties which are not the identity provider are zero-knowledge with a well-defined, permissible leakage function.

B. Overview of our VKD Solution

Our VKD solution, aimed at real-world, large-scale applications, removes a majority of the assumptions made by previous works in this domain, such as, [15], [32]. We modify the constructions and definitions presented in [15] to use a primitive called ordered zero-knowledge set (oZKS) [6], which, together with a secure commitment scheme (cCS) and a verifiable random function (sVRF), allow us to instantiate a VKD in space linear in the number of updates for labels in the VKD. This is in contrast to [15]’s construction, that additionally requires space linear in the number of server epochs, with a high constant factor (linear in the security

parameter). We further augment these constructions to provide a method to allow compacting the underlying data-structures, i.e., deleting values which are no longer in use, while still requiring very little monitoring from the user to provide security. This combines the space efficiency gains of a VKD construction which relies on the assumption that users monitor their keys constantly (e.g. [32]), with the security gains of a construction based on append-only data structures (e.g.[15]).

C. Overview of our Consistency Protocol

Our first result is to debunk the common belief that such systems require consensus [34], [40], [15]. A first correct but inefficient solution would be to use simple Reliable Broadcast [13] to achieve all required properties in full asynchrony, a setting where deterministic consensus is actually impossible [20]. This, however, is inefficient. Instead, we design a tailor-made solution for Parakeet (Section V). The resulting system has low-latency, and additionally the identity provider can be natively sharded across many machines (unlike consensus-based solutions) to allow unbounded horizontal scalability.

D. Bridging the VKD & the Consistency Protocol

For the most part, we handle the consistency and VKD components of Parakeet separately. However, they do need to communicate with each other. Here, we provide an *application programming interface* (API) to bridge this gap. The IdP must provide small commitments to the witnesses and the users need to communicate with the witnesses to receive these up-to-date commitments. To this end, we define a simple witness API.

As stated above, the witnesses run the consistency protocol to certify a commitment to the internal state of the VKD. So, our witness API includes an algorithm for the IdP to propose a commitment at each server epoch and prove that it correctly updated its VKD (WitnessAPI.ProposeNewEp). We also allow any party to query the witnesses to retrieve the commitment and certificate for an epoch (WitnessAPI.GetCom). Finally, the API call WitnessAPI.VerifyCert verifies the certificate for a commitment. The details of this API are in the full version.

III. VERIFIABLE KEY DIRECTORY

In this section, we will discuss the primitive *Verifiable Key Directory* (VKD), defined by [15], and its properties. Recall that a VKD consists of three types of parties: (1) an identity provider (or, server), (2) users, and (3) independent auditors. In a VKD, each user has an associated label, denoting their username. The server stores a directory Dir of label-value pairs. Each value corresponds to a public key. The clients can request updates to their own public keys – equivalent to requesting a change to the state of Dir . For efficiency, many such requests are batched together, with updates going into effect at discrete timesteps (*epochs*). So, Dir is stateful, has of an ordered sequence of states Dir_t , one state per epoch t .

To support verifiability in a VKD, each state of the directory needs a corresponding commitment com_t . The commitment com_t is made public using the WitnessAPI (see section II-D). This model assumes that any changes to the directory go into

effect when the corresponding epoch t goes into effect and the commitment, com_t , for this epoch is published.

The security of a VKD crucially relies on at least one honest auditor per epoch checking the latest update for validity. This assumption is common across a majority of the work in this area (see, e.g. [42], [15]), in order to maintain client-side efficiency. In this work, we realize this assumption using the witnesses, each of which runs the auditing operation as part of executing `WitnessAPI.ProposeNewEp`. Note that the `WitnessAPI` solution can be used *in addition to* any other solution, such as gossip for distributing commitments [31] or client auditing [42], also discussed in Section VII. When serving billions of users across different geographical regions and platforms, it is not always reasonable to expect clients to run audit operations or to participate in a connected gossip network, so using a set of witnesses only increases security.

A. Outline of this Section

The rest of this section is devoted to revisiting the VKD primitive, with improvements targeted for production. In section III-B, we recall the various algorithms for this primitive and its properties. For our storage-optimized VKD, we need the `oZKS` primitive, whose motivation and properties we describe in section III-C. Then, in section III-D, we describe our VKD construction using the `oZKS`, which achieves properties identical to [15]. Even with improved efficiency, several practical problems remain, including: (1) allowing users to have multiple devices linked to their accounts, and multiple updates for the same user in an epoch, (discussed in section III-E), and (2) introducing a separate, efficient storage layer for a scalable VKD implementation (discussed in section III-F).

B. VKD Definition

The primitive verifiable key directory (VKD) was first defined by Chase, et. al. [15]. As stated before, the VKD server holds a directory mapping labels to values with one state per epoch. The commitment to this state (and the proof of correct update) is published by the server, using `VKD.Publish`.

When users want to lookup the public key for a particular label, they request this by calling `VKD.Query`. The response to a lookup comes with a proof of correctness, which the requesting user can verify (with `VKD.VerifyQuery`). Each user can also check the history of her own key, implicitly, for her label, getting a mapping $t \rightarrow \text{val}_t$ of the state of the associated value at every epoch. The user can get this history and its proof (`VKD.KeyHistory`), as well as verify it (`VKD.VerifyHistory`). Any party can verify the output of a `publish` (`VKD.VerifyUpd`) or a sequence of `publish` operations (`VKD.Audit`).

Soundness and privacy of a VKD. The soundness definition for the VKD primitive is exactly the same as that of [15]. This definition captures the non-equivocation property stated in section II-A, assuming that the server never deletes any existing records. We also define the privacy of a VKD the same way as [15]: all operations are zero knowledge with a well defined leakage function.

C. Ordered (Append-Only) Zero-Knowledge Set with Deletion

As a step towards a more space-efficient implementation of a VKD than that of SEEMless, we replace their `aZKS` building block with a primitive which we call an *ordered append-only zero-knowledge set* (`oZKS`). Here, we discuss the properties of `oZKS` and briefly describe our implementation.

An `oZKS` is actually a further generalization of the append-only zero-knowledge set (`aZKS`) primitive, which was first introduced in [15]. [6] presents an implementation (and corresponding informal definition) of an `oZKS` construction. The `oZKS` primitive is used in a setting where a party (often called a *server*) holds a data-store of label-value pairs (with unique labels), and is trusted for privacy, but not to serve consistent views of label-value pairs. The party uses `oZKS` to commit to label-value pairs where the labels are all unique.

Initial work on zero-knowledge sets (e.g. [16], [33], [14]) committed to static data stores. The `aZKS` primitive extended this to include insertions only. The recent implementation of an append-only `oZKS` [6] extends the `aZKS` to include a strict ordering on when elements are inserted. Thus, an `oZKS` should support verifiable algorithms for: initially committing to a datastore (`oZKS.CommitDS`), insertions (`oZKS.InsertionToDS`), membership/non-membership queries (`oZKS.QueryMem` and `oZKS.QueryNonMem`). We use the `oZKS` primitive to build our VKD construction in section III-D.

Comparing with the `oZKS` from [6]. Note that the append-only property of the implementations of both [6]’s `oZKS` and [15]’s `aZKS` require an ever-growing storage requirement on the server. In large-scale practical applications, never purging obsolete data can be infeasible. The novelty of the `oZKS` construction in our work consists specifically of the mechanism for *enabling secure deletions*. We present this as a middle ground solution between fully mutable auditable data structures that require users to be always online (e.g. CONIKS), and append-only auditable data structures that lead to an unreasonable server storage cost for long-running systems (e.g. prior and concurrent work on `oZKS` and SEEMless).

We introduce the secure deletion extension here for completeness and apply it in section IV. For verifiability, the deletion operation of an `oZKS` with compaction is actually a two-step process: marking nodes as candidates for deletion (a process we call tombstoning, through `oZKS.TombstoneElts`) and deletion of tombstoned elements (`oZKS.DeleteElts`). In appendix B, we provide a formal definition of all of these algorithms and in appendix C, we describe our `oZKS` construction.

Our `oZKS` implementation. Similarly to [15]’s `aZKS`, our `oZKS` instantiation uses a Merkle Patricia Trie (MPT) to commit to label-value pairs. However, instead of using a complicated persistent data-structure, which requires storing all states of every MPT node, for the `oZKS` data structure, we simply include the epoch a leaf was inserted as part of the value committed for a node. This allows us to store one MPT which mutates over time and old states of nodes in this tree can be garbage collected, instead of persisting forever. Also, note that for privacy, the MPT-based `oZKS` implementation

computes leaf labels using a *verifiable random function* (VRF) (defined in appendix C), which is a deterministic function computable only by the holder of a secret key for a publicly known public key, PK. Any party can, however, check the correct computation of a VRF using PK. For privacy of the actual value associated with a label, we use a hiding commitment scheme. Concretely, this means that to add label with value val to the datastore and update the corresponding commitment at epoch t , the owner of the datastore adds a leaf with label $VRF(label)$ and value $(com(val); t)$ to the MPT. Note that this constitutes only a simple oZKS, without the tombstone or deletion algorithms. We discuss tombstoning and deletion in section IV. See figs. 1 and 2 for simple examples of the underlying data structures of our oZKS and the aZKS of [15].

Soundness and privacy. The oZKS without the `TombstoneElts` and `DeleteElts` algorithms is said to be sound if, given at least one honest auditor in every epoch, the server cannot delete any elements which were previously committed. I.e., the oZKS without tombstones or deletions is append-only. At first glance, the terms “append-only” and “with deletion” for an oZKS may seem like contradictions. What we intend to capture is a mechanism to commit to a set, in an append-only manner, but with the additional ability to remove very old values which may no longer be needed by the calling application. Hence, the soundness property of this data structure, defined in appendix B, is that values may only be marked for deletion if they are older than an epoch permitted by a system parameter. Other than the fixed epochs for tombstoning old-enough nodes, and correspondingly for deleting nodes marked as tombstoned, the data structure should only allow insertions for labels not present in the datastore already. As in the privacy definition for the aZKS of [15], we require all functions for the oZKS to be zero-knowledge, with a well-defined leakage function.

Leakage for our implementation. Our oZKS construction, when committing to an initial data store (`oZKS.CommitDS`), leaks the size of the datastore. The `oZKS.InsertionToDS` also leaks the size of the datastore before and after the update. For each inserted element, the adversary also learns whether it queried for this element before and if it did, this tells the adversary when this element was added. `oZKS.TombstoneElts` and `oZKS.DeleteElts` leak when the tombstoned (resp. deleted) elements were inserted and the size of the datastore before and after each call. The responses to `oZKS.QueryMem` and `oZKS.QueryNonMem`, in addition to the actual response, also leak the size of the datastore and for `oZKS.QueryMem`, it leaks when this element was added.

D. Revisiting the SEEMless VKD Construction

In this section, we construct a VKD (without compaction) using an oZKS, instead of an aZKS and also summarize the concrete implementation strategy for our VKD implementation. This replacement of an aZKS with an oZKS results in functionality which is equivalent to that of SEEMless [15]

with significant space savings. The formal description of this construction is in the full version.

Concrete implementation. As mentioned in section III-C, our concrete implementation of the oZKS consists of a Merkle Patricia Trie (MPT), used to commit to a leaf’s value and bind it to the epoch in which it was inserted. This MPT-based oZKS is an important component of our VKD construction. In addition to this oZKS, the VKD requires a database to store actual username to value mappings as well as when each value was added. Suppose a user with username Alice first joined the system at a time t with public key value val_1 . Thus, val_1 is the first version of Alice’s public key and the server adds the label ‘Alice|1’ with value val_1 to the oZKS at epoch t . If at a later epoch t' , Alice’s key is updated from version i to $i + 1$ with new value val' , the server inserts the labels ‘Alice| i |stale’ with value equal to the empty string, and ‘Alice| $(i + 1)$ ’ with value equal to val' to the oZKS. At a lower level, this means that at epoch t' , the server adds to the MPT the leaves whose labels are ‘VRF(Alice| $(i + 1)$)’ and ‘VRF(Alice| i |stale)’, with values $(com(val'); t')$ and $(com(\epsilon); t')$, respectively.

Soundness and privacy. The soundness definition of this VKD construction is identical to that of [15]. At a high level, as long as for epochs up till epoch t , (1) a client with label label checks the states of key using the `VKD.VerifyHistory` algorithm, (2) the `WitnessAPI` is honest, and (3) at least one honest auditing party verifies each update using `VKD.Audit`, then the identity provider could not have output a diverging view of the val associated with label at any epoch less than t . In other words, in the presence of auditing and witnesses, Alice and Bob must always agree on the view of val_{Alice} —the value associated with the label Alice. The leakage functions of this construction match those of [15]’s construction exactly.

Space efficiency improvements. The aZKS construction of [15] allows reconstructing the state of verification data for a data store at any time, by storing all intermediate states ever generated. Both our work and [15] implement a compressed version of a Merkle Patricia Trie, where, for a random set of leaves, the expected depth of a leaf in a tree with n leaves is $\log(n)$. Thus, if a new node is added to the MPT of [15], this results in adding $\mathcal{O}(\log(n))$ new states to persistent storage. Even with batching, this means that the space complexity of their implementation depends on the number of epochs, in addition to the number of leaves added. However, the space complexity of our oZKS implementation only depends on the total number of leaves added. The impact of this difference is quite significant, as shown in Section VI.

E. Other Practical Considerations

In addition to the compaction extension, we discuss the following additional considerations for practical deployments of key transparency for messaging applications.

User interfaces. In this work, we assume that the server is trusted for bootstrapping users. For example, if a user loses their phone, there is a mechanism for them to recover their account. We also assume that the client application has

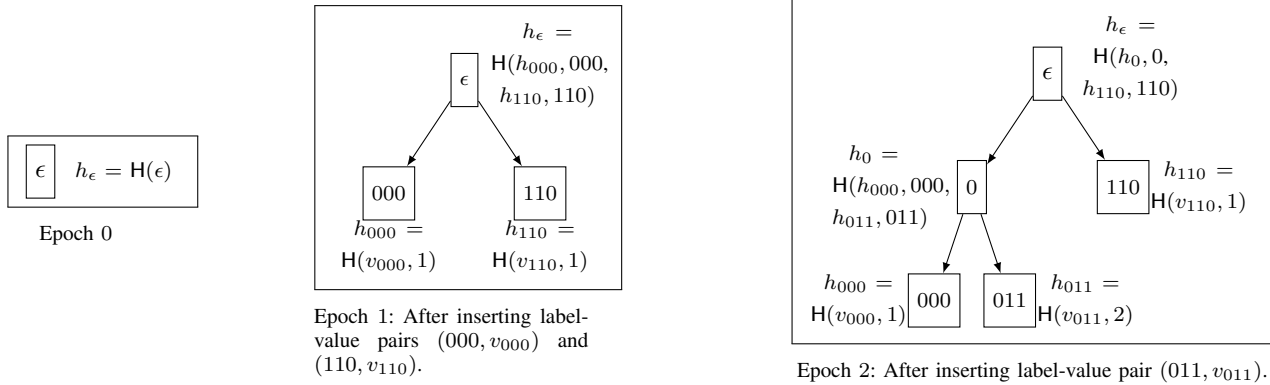


Fig. 1: An example of the evolution of our Merkle Patricia Trie-based oZKS construction with 3-bit labels. The oZKS starts with no entries at epoch 0, which is committed in the tree with a single node, with label ϵ and commitment $H(\epsilon)$. At the first epoch, two new leaves are inserted and at epoch 2, a third leaf is inserted. Note that the tree at epoch 2 can be used to reconstruct any of its previous states.

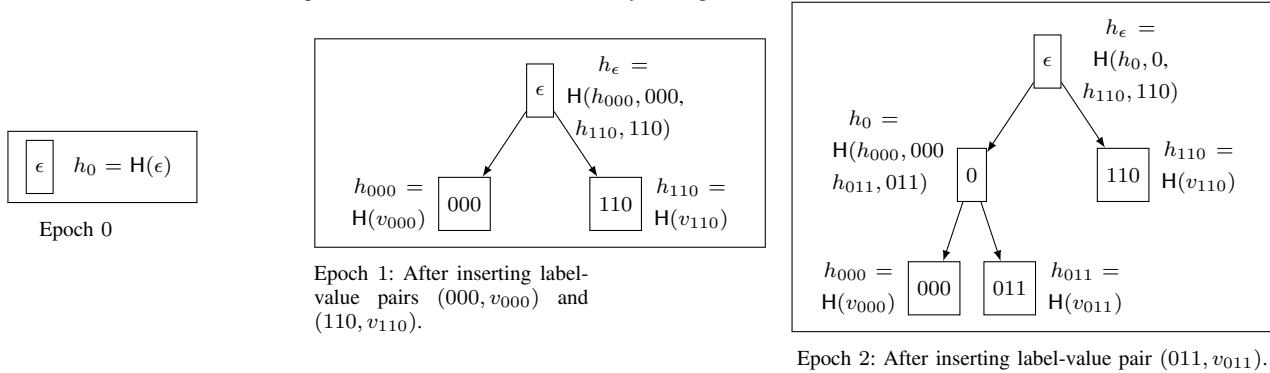


Fig. 2: An example of the evolution of the SEEMless Merkle Patricia Trie-based aZKS construction with 3-bit labels. The aZKS starts with no entries at epoch 0, which is committed in the tree with a single node, with label ϵ and commitment $H(\epsilon)$. At the first epoch, two new leaves are inserted and at epoch 2, a third leaf is inserted. Note that the values inserted in this tree correspond to the entries inserted in Figure 1. However, in contrast to the oZKS construction, the tree at epoch 2 as is cannot be used to reconstruct any of its previous states. SEEMless optimizes the ability to get values from previous states using storage compression.

interfaces to inform users about various events including: if witnesses are offline for extended periods of time, if a proof failed to verify, or if a public key they requested to add to their account did not take effect within a specified time period. [36] discusses some of the design considerations for a user-facing transparency system. We leave further study to future work.

Multiple devices per user. A common property of most mainstream messaging applications is the ability to associate one user account with multiple devices owned by the user (e.g. computer, mobile phone), each with a distinct E2EE keypair. This means that a VKD must support a set of public keys belonging to a user, as opposed to just a single key. This simply involves using the set of public keys as the values in the VKD, ensuring that validating client queries involves a membership check instead of an exact match of public key.

Multiple updates per epoch. Depending on the length of the epochs, clients may submit multiple updates to their public key values within the span of a single epoch. Only publishing the latest key received for a user within an epoch can lead to issues with consistency for clients which may have been messaged within the fast key updates. Instead, when sequencing these updates, the server can designate the ordering by adding an ordered list of public keys as opposed to a single key per

epoch. Then, clients can check for list membership when querying for proofs, and the server can use the ordering to ascertain the latest key after the next epoch begins.

Together, these adjustments to client public keys result in an ordered list (multiple updates per epoch) of sets (multiple devices) of public keys being hashed in our VKD construction.

F. Storage API

As discussed in Section I, a large-scale VKD implementation requires a separate storage solution which we call StorageAPI.

Warmup. As a first attempt at defining StorageAPI, we may want only two operations on top of a simple database:

- $\text{val}/\perp \leftarrow \text{StorageAPI.GetFromStorage}(\text{key})$: This call takes as input a key (key). If the key is in the database, it returns the associated value (val). Otherwise, it outputs \perp .
- $1/\perp \leftarrow \text{StorageAPI.SetToStorage}(\text{key}, \text{val})$: This call takes as inputs a key key and corresponding value val . It outputs 1, if this value is successfully in the database and \perp otherwise.

We could use this simple API for all data types by defining the storage keys for each type as an encoded binary vector with a prefix to demarcate types. The implementer of StorageAPI determines the underlying storage layer to handle each type.

This also maintains unique keys for individual records without changing the storage key type signature.

Handling storage latency. Of course, in a practical implementation, storage will have to consist of a set of geographically spread-out, duplicated nodes. This means that one of the biggest bottlenecks in timely update of the server’s verifiable data structures, as well as client queries, is memory latency. For large-scale end-to-end encrypted messaging identity providers, what may *not* be a problem, however, is bandwidth. Hence, we add batched storage APIs to leverage the high-bandwidth to reduce latency. We define these APIs as:

- $\{\text{val}_i\}_i \leftarrow \text{StorageAPI.BatchGetFromStorage}(\{\text{key}_i\}_i)$: This call should take as input a set of keys key_i and return values val_i stored in the database associated with the corresponding key. If label_i is not in the database, $\text{val}_i = \perp$.
- $1/\perp \leftarrow \text{StorageAPI.BatchSetToStorage}(\{(\text{key}_i, \text{val}_i)\}_i)$: This call takes as inputs a set of key-value pairs $(\text{key}_i, \text{val}_i)$, and outputs 1, if all values are successfully set and \perp if any errors occur. Note that this implicitly implies atomacity of a single `BatchSetToStorage` operation.

Batching-friendly oZKS. Note that batching storage operations is not enough on its own, unless the algorithms we implement are also amenable to batched memory accesses. In particular, underlying our implementation of an oZKS is a compressed Merkle Patricia Trie, whose construction satisfies the following invariant:

Each node is the longest common prefix of its children.

This implies that, for example, just looking at the label for a leaf being inserted does not automatically allow the server to know which nodes (and corresponding key_i ’s) to use with `BatchGetFromStorage`. To solve this problem, our implementation starts with assuming that the storage solution has caching capabilities. Under the assumption of a large enough cache, if we can preload nodes for a batched operations, say, `VKD.Publish`, the operation itself can be done as if in RAM. Later, the cache can be flushed to persistent memory as a single transaction operation. oZKS operations requiring tree-traversals are called with a leaf label as input, batched versions of these operations are called with a set of leaf labels `leaves`. To optimize remote persistent storage accesses, at a high level, we implement procedures which operate following two steps: (1) compute a set `prefixes` of all prefixes of values in `leaves`, (2) starting at the tree’s root node, begin a breadth-first search for nodes with labels in the set `prefixes`, batch-fetching labels at the same depth. This ensures that all of the nodes required to run algorithms for the labels in `leaves` will be loaded at the end of the preload operation with only one access to persistent storage per layer of the Merkle Patricia Trie, without overwhelming the cache by having to load the entire tree. In fact, the set of nodes retrieved to cache is exactly the set required for the original oZKS batched operation.

So far, we have omitted any discussion of batching for set (write) storage operations. This is easier, since batching writes is equivalent to flushing an in-memory cache of the updated nodes, an operation our storage interface requires.

The only writing operation is `VKD.Publish`, and once it has completed its changes, we can commit all of the changes in the transaction cache as a single atomic operation (assuming the storage layer supports this). Since all of the modified data will be in the cache, flushing the cache in a timely manner is all that is needed to ensure up-to-date views of storage.

Other storage considerations. Even with our cache-based solution, there is an underlying assumption that the storage layer can support a cache-sized atomic write operation. For very large systems, we would like to further loosen this requirement, by adding support for the situation in which the storage layer cannot atomically (or at least efficiently) commit the entire transaction of changes in a single operation. We make some simple modifications to the data-structure layer to prevent inconsistencies in read operations during an update.

All proof generation operations use an epoch value stored in the oZKS data-structure to determine the latest epoch, denoted (`LatestEp`). `LatestEp` must be updated last in order to ensure that the data corresponding to `LatestEp` has been written for all parts of the data-structure, before any reading is permitted. Once the oZKS’s latest epoch is updated, all operations will take the new epoch as truth and access consistent values.

With this in mind, we need to change the stored record for a tree node to store two values at a time: the previous and the current value. For atomicity, this means storing a dual-value struct, each value with an associated epoch. Read and write operations read the epoch stored at these values to determine which to designate as current and which as previous. The read operation always reads the value whose epoch is less than or equal to `LatestEp`, and the write operation overwrites values whose epoch is less than `LatestEp`. This allows us to update tree nodes without impacting proof generation operations.

IV. VKD WITH COMPACTION

Previous constructions, such as [15] and ([42], [32]) seem to fall into two extreme categories: (1) works that assume the server’s storage to be linear in the number of updates, and (2) works that assume users are either always online or can retroactively check the server’s view of their keys for any epochs they may have missed. This means that the client has to do linear work in the number of server epochs to monitor their keys. Also, at production scale, we must account for the storage implications of being able to serve users who come online infrequently—e.g., in case (2), the server would need to store a large trove of history, perhaps even all of its data.

Even the soundness of our oZKS-based VKD construction depends on the append-only property of the oZKS. In our implementation, we instantiated the oZKS described above using the Merkle Patricia Trie-based SA, as in [15]. For each version added by the user, the server adds a label of the form `(uname|i)` to the MPT and for each version retired by the user, it adds a label of the form `(uname|i|stale)`. So, the space complexity of `st` for our VKD grows linearly in the total number of updates, i.e. $|\text{st}| \in \mathcal{O}(\sum_{t=0}^n |S_t|)$ where S_t is the input to the t th call to `VKD.Publish`. This poses two problems:

- the labels committed for a given user may never be deleted, so it becomes impossible for a user to request that *all* data associated with their account be deleted. This is (arguably) in conflict with “right to be forgotten” regulations, such as those laid out in the GDPR [4].
- eventually, traversing a tree which is monotonically growing leads to unreasonable blowups in proof sizes as well as the time taken to respond to queries and make server updates.

This motivates the need for a process of secure (and transparent) compaction of storage on the VKD server. Naturally, the best way to do this would be to delete very old data stored on the VKD server, should it no longer be useful.

Overview of our paradigm. To address the issue of ever-growing storage, we extend the VKD functionality to a VKD with compaction, denoted *cVKD*. The compaction consists of two phases: the *tombstone phase*, a special epoch when some of the server’s data is marked for garbage collection, followed, after a period of time, by, the *compaction phase*, a special epoch where values marked as tombstoned are garbage collected i.e. deleted. All other epochs are expected to function as before, i.e., internal data structures remain append-only. The auditors contribute global correctness checks to the tombstone and compaction processes by verifying that only data which is “old enough” is tombstoned (*VKD.VerifyTombstone*), that only tombstoned values are garbage collected (*VKD.VerifyCompact*) and, that in all epochs which are not used for tombstoning or deletion, the server updates its commitments correctly, as before. To ensure that values are properly garbage collected, we require that users monitor their own key history after each tombstone phase, prior to the following compaction phase. Through *VKD.KeyHistory*, the user checks any deletions were correct.

Soundness of a VKD with compaction. We provide a formal definition of a *cVKD* in appendix B. The soundness definition for a VKD is more general than a single tombstone or deletion epoch, but rather, refers loosely to the following intuition. We start with some requirements on the user for verifying her own key. Specifically, we require that between any consecutive tombstone and deletion, a user checks her key at least once. Then, “*if* a user Alice verifies her key according to the requirements and believes her own key at an epoch t was PK , *then*, any other user must also believe that Alices key at epoch t was PK ”. Note that Alice checking her keys between all “tombstone-deletion epochs induces a mapping from epoch t to the freshest public key and version number at t . For this mapping to be unambiguous, and strictly increasing, a version number cannot be deleted and reinserted without detection.

A. Construction

We extend the *oZKS*-based VKD construction from section III to include a compaction phase, to allow garbage collection and verification of data which is no longer needed.

The following example and Figure 3 illustrate what kind of data the server may want to delete in order to compact its storage. Recall that when a user with username Alice

first joins the system, with public key PK_1 , the VKD adds the label $Alice|1$ with value PK_1 to the *oZKS*. When Alice updates their key from the first version PK_1 to PK_2 , the server simultaneously adds the labels $Alice|2$ and $Alice|1|stale$ with values PK_2 and ϵ to the *oZKS*. This pattern generalizes to any further updates by the user Alice. Suppose, after several years of joining, Alice is on their 10th key version, and has come online and checked their key history a few times in the interim. The data for $Alice|1$, PK_1 may not be useful and the server may want to reduce storage costs by deleting such entries.

We motivate our final *cVKD* construction with preliminary attempts to support compaction directly from a VKD.

Attempt 1. The most straightforward attempt at adding a compaction functionality to the VKD built using an *oZKS* involves simply allowing the deletion of arbitrary, “old enough” *oZKS* entries by introducing a *oZKS.DeleteElts* functionality. We define “old enough” as a system parameter called *StaleParam*, with the requirement that any entries deleted by the server must have been inserted at least *StaleParam* epochs ago. That is, if *oZKS.DeleteElts* is called in some epoch t , the only entries the auditors will verify as correctly deleted were added before the epoch $t - \text{StaleParam}$. In our MPT-based *oZKS* implementation, this is easy to support, since the epoch a leaf was added is included in the leaf’s hash. From the auditors’ point of view, as long as a deleted node is old enough, its deletion was valid. For privacy, the auditors should not access the plaintexts used to compute the VRF for the MPT.

What if certain users check their key history infrequently? Suppose a user Alice is at version 20 of their public key and version 10 was added a long time ago. The server could mount an attack where it deletes a label $Alice|10|stale$ but not $Alice|10$. If the user Alice does not come online for a while, the server could serve the stale value at version 10 for Alice, in response to *VKD.Query*, then delete $Alice|10$. The server could do all of this before Alice comes online to check their public key. The only way to prevent this seems to be to store all states for a long time, or to have users be always online: both of which significantly degrade efficiency.

Attempt 2. We could try to patch the issues with the *oZKS.DeleteElts* functionality in the previous attempt by introducing a system parameter called *DeletionEpochs* such that an *oZKS.DeleteElts* proof only verifies if it is called in an epoch in the set *DeletionEpochs* with the important invariant that the *oZKS* remains append-only at all non-deletion epochs. We could then rely on the assumption that a user comes online between any two consecutive elements of *DeletionEpochs* to run *VKD.KeyHistory* on their own label. Now, if a user Alice comes online, they can ensure that (1) the epochs for version numbers are correctly ordered, for example, if the label $Alice|9$ was inserted at epoch t_9 and $Alice|10$ was inserted at t_{10} , then $t_9 < t_{10}$, (2) versions are marked stale at the appropriate epoch, in the previous example, this means that $Alice|9|stale$ was inserted at t_{10} , and (3) for an honest IdP, labels of the form $Alice|version$ are always inserted *before* the corresponding label $Alice|version|stale$, they should ensure that




Fig. 3: An example of labels corresponding to user Alice’s key updates. Since Alice has changed their key enough times since their initial entry into the system and epochs before epoch 200 are considered *old enough*, Alice|1, Alice|1|stale can be deleted.

if Alice|version|stale is deleted, then so is Alice|version.

However, an issue still remains: suppose the server has two consecutive deletion epochs del_1 and del_2 , with n non-deletion epochs between them. Suppose, Alice comes online at epoch $\text{del}_1 + 1$ and checks their key history and then again at epoch $\text{del}_2 + 1$ checks their key history. Suppose at epoch $\text{del}_1 + 1$ their latest key’s version number was 10, and the server now inserted 3 fake keys for their between epochs $\text{del}_1 + 2$ and del_2 , i.e. their key’s version right before the next deletion is fallaciously 13. As part of del_2 , the server could delete the labels it added for Alice’s versions 11 through 13 and the label Alice|10|stale. Such an attack could go undetected, unless Alice came online in exactly the epoch $\text{del}_2 - 1$. In the presence of network delays and users who are possibly offline for long periods of time, we consider this requirement too restrictive.

Hence, we propose adding the restriction that StaleParam is at least as large as the space between two deletions. Meaning, auditors would not accept deletion of labels which were inserted more recently than the most recent deletion. In the above example, this would mean that $\text{StaleParam} > \text{del}_2 - \text{del}_1$ and any values inserted between del_1 and del_2 would remain for Alice to check when they comes online at epoch $\text{del}_2 + 1$.

Even with the restriction that $\text{StaleParam} > \max(\text{del}_{i+1} - \text{del}_i)$, if the user does not immediately come online at each epoch $\text{del}_i + 1$, and her key is old enough, the server could temporarily rollback her key’s version. Consider the following case: a user Alice updates their key very infrequently and their most recent key, say, version 10, is old enough to fit the criteria for deletion. That is, Alice|10 was added at an epoch older than StaleParam . At the next deletion epoch del , the server could delete all record of their most recent key, i.e. the labels Alice|10 and Alice|9|stale, effectively rolling back their key to version 9. It could then re-insert the labels Alice|10 and Alice|9|stale before Alice checks their key history for this round of deletions. All of the checks we mentioned above would pass and yet, for a period of time, if another user Bob queries for Alice, they would see a different version and value than what Alice will see later. This is a subtle technical issue but preventing such an attack is integral for a construction to be considered sound!

The auditors could mitigate this issue by storing old labels and ensuring deleted labels were not reinserted. This requires auditors to be stateful and have linear storage complexity in the size of the updates. This would make the auditors storage as large as the servers which, our construction avoids.

We argue that the requirement that the user come online immediately after every deletion epoch is also too strong. In the next construction, we propose a design that fixes this issue.

Final construction. In our final construction, we patch this last issue by including a set of *tombstone epochs*, denoted TombstoneEpochs in the public parameters. A tombstone epoch is an epoch when items are marked for deletion, i.e. *tombstoned*, but not actually deleted. This allows users to check their own key history and ensure that values marked for deletion are appropriate, *before* they are actually deleted.

For example, if the IdP wants to delete the label Alice|10 at a deletion epoch del , it would first have to set the value of Alice|10 to TOMBSTONE at the preceding tombstone epoch TombEp . Between TombEp and del , Alice can run KeyHistory and see if this is appropriate at the VKD level.

Meanwhile, at TombEp , the auditors ensure that the only oZKS elements which are modified at TombEp are *old enough*, according to the parameter StaleParam . At all epochs between TombEp and del , the auditors continue checking append-only proofs. At del , auditors check that the only oZKS elements which are deleted have the value TOMBSTONE .

The final construction requires a client to come online between deletion epochs and check the following set of conditions. If the minimum valid version number it receives is min and its current version number is current , then:

- *Correct tombstoning or deletion.* For all versions version below min , it gets either (a) a non-membership proof for $\text{uname}|\text{version}$, or (b) membership proofs for $\text{uname}|\text{version}$ and $\text{uname}|\text{version}|\text{stale}$, with either both of them or neither of them having the value TOMBSTONE .
- *Correct ordering.* For any version version in the range min to $\text{current} - 1$, that the epoch t_{version} when $\text{uname}|\text{version}$ was added is less than the epoch $t_{\text{version}+1}$, when $\text{uname}|\text{version} + 1$ was added.
- *Correct version changes.* For any version version in the range min to $\text{current} - 1$, that $\text{uname}|\text{version} + 1$ and $\text{uname}|\text{version}|\text{stale}$ were added at the same epoch.
- *Freshness of current value.* For the current version current , it checks for the non-membership of $\text{uname}|\text{current}|\text{stale}$.
- *Non-membership of next few entries.* For any version $\text{version} \in [\text{current} + 1, 2^{\lceil \log(\text{current}) \rceil + 1})$, $\text{uname}|\text{version}$ wasn’t added.
- *Non-membership of much further entries.* For L , the most recent server epoch, for any $j \in [2^{\lceil \log(\text{current}) \rceil + 1}, 2^{\lceil \log(L) \rceil}]$, it checks a non-membership proof for $\text{uname}|2^j$.

Soundness. We give a detailed proof for the soundness of this construction in appendix B. The intuition for soundness is as follows: For the most part, this construction with deletion is identical to the one in section III, since all but a designated set of epochs is append-only. The only divergence occurs at tombstone or deletion epochs, and we show that dishonest server behavior cannot go undetected, even in these epochs.

We require that a user must check their key history once for each tombstone epoch, prior to the following deletion. If a user’s key is at version, the server can show the wrong key for a user in one of the following ways. The server could try to change the public key committed to the label version: this is mitigated by the auditors checking that the only change being made is tombstoning or deletion of already tombstoned values. If the server tombstones the label `uname|version`, the user will detect that it is tombstoned since they will come online before the next deletion. The second option is that the server could show an older key, i.e. show a key for some version $version' < version$ but this should be detected by the one of the key history checks. The third option is that it could try to add a higher version $version' > version$: this is already detected by the construction in section III. The fourth option is that it could try to re-insert an old version number, which was deleted, with a different key: this particular attack is unique to the scenario with deletions and is detected by the key history check, since if the server re-inserts an old version number, the user will see this the next time they comes online, which will be before the server has the opportunity to hide this change.

Assumptions. Recall that previous works, as well as our construction in section III, make two major assumptions. First, they assume that the server’s underlying data structure is append-only, i.e. the identity provider could never delete any verification data. Second, that users are not always online. The guarantee is that if a user checks their key history for a period of time, they will catch the a cheating server. In fact, there is no upper bound on the time the user can wait before checking their key history but still get the same guarantee. This construction changes these assumptions as follows:

- 1) We assume that the system only considers a value stale-enough to tombstone, if it was added *before* the previous tombstone. The ability to tombstone and delete values is already an improvement over previous, append-only verifiable data structures. We argue that the compromise of keeping data for a certain period of time is still a major improvement for long-term storage costs.
- 2) Instead of allowing the user to go for an unbounded amount of time before checking their key history, we require that they come online between each pair of tombstone and corresponding deletion epochs and check their key history to get guarantees for that period. Since large servers only need to run deletions infrequently, while more strict that the previous assumption, this still provides a user with leeway. In fact, this approach also has the added benefit that there is an upper bound on the amount of time the server can cheat for a particular user and go undetected.

Leakage. The leakage of the subset of algorithms of this cVKD construction, which are inherited from the definition of a VKD is identical for this construction. For `cVKD.TombstoneElts` and `cVKD.Compact`, they both leak the size of the cVKD at the start of the operation. Also, `cVKD.TombstoneElts` leaks the number of tombstoned values and `cVKD.Compact` leaks the updated size of the cVKD.

Space efficiency. A compressed MPT is constructed with the invariant that any node with only one child is deleted and replaced by its child. This means that when a new leaf is inserted, somewhere in the tree an additional new node will have to be inserted to accommodate this leaf. When a leaf is deleted, its immediate parent can also be deleted. Hence, for each leaf deleted from this MPT-based VKD, the equivalent of 2 MPT nodes worth of data can be erased from the server.

V. THE PARAKEET CONSISTENCY PROTOCOL

Existing transparency overlays typically fall in two categories:

- **Reactive Client-side Auditing:** These protocols [32] offload the consistency protocol to the clients. Clients are expected to gossip among themselves about their views of the system and construct fraud proofs showing that the IdP has misbehaved. Unfortunately, this class of protocols is too optimistic for a security-critical infrastructure and too heavy for the end user. As a result, there is a high risk that users do not eagerly gossip proofs or that the adversary launches targeted eclipse attacks [22] to isolate targeted users. Finally, gossip relies on always-connected nodes forming a connected graph, perpetually routing messages. This is impractical if the many of nodes in the gossip network are small devices with intermittent network connectivity and rely on battery power.
- **Using a black-box blockchain as a trust anchor:** This approach [40], [15] guarantees consistency as long as the blockchain is secure, i.e., it does not have forks. This alleviates any risks for clients who only need to trust the black-box blockchain, however, (i) it introduces a significant extension of the trust assumptions, which might not be correct given the multitude of attacks blockchains suffer [10], (ii) it limits the update speed of the system to the finality speed of the blockchain (which is tens of minutes in Bitcoin), and (iii) it doesn’t solve the problem of eclipse attacks against clients as they remain a challenge in blockchains. Also, no existing protocols has provided a detailed security analysis for the interactions with any existing blockchain.

One way to avoid these issues is having auditors run consensus on each update of the IdP, essentially replacing the trust assumptions of the blockchain with a custom-made blockchain just for the transparency layer. However, this is simply an overkill. As a final contribution, Parakeet shows a consensus-less protocol that achieves all the desired properties for defending against split-view attacks. Section V-C further extends this protocol to provide censorship resistance and read-freshness.

A. Consensus-less Strong Consistency

The Parakeet protocol consists of network messages exchanged between the participants (Section II-A). The users communicate with the identity provider (IdP) and witnesses. The IdP communicates with the witnesses and users, but the witnesses *need not communicate directly with each other*.

Updating the state. All updates to the state of the IdP start with an *update request* sent by a user to the IdP. The request contains a user identifier `id` and the new key `pknew` associated with `id`. The IdP collects several user requests and runs the




Fig. 4: Illustration of the key update protocol.

key-update protocol illustrated in Figure 4. The IdP broadcasts an *update notification* to the witnesses to notify them that the state of the key directory has been updated (❶). This notification contains the following fields: (i) a commitment to its state com_t , (ii) the epoch t , (iii) a proof π^{Upd} that com_t is a valid commitment and that the state is a valid update of the previous state, and (iv) a signature by the identity provider over this data. Each witness locally verifies that the epoch number t has been incremented by 1 and that the proof π^{Upd} is valid; it then counter-signs the update notification (❷). Finally, the IdP collects at $2/3f + 1$ witnesses' signatures into a *certificate* (❸) which is attesting that a quorum of witnesses verified each past update and the IdP could not have equivocated.

Reading the state. The user sends a *read request* to the IdP to request the key associated with a specific identifier. This request contains the identifier key to query. The identity provider replies to read requests with a message containing the following fields: (i) the key key associated with the identifier, (ii) a certificate from the witnesses over its latest update request, (iii) a proof that key is included in the certified state, and (iv) a signature by the identity provider over this data.

B. Proofs Sketches

We provide the intuition for the proofs of consistency, validity, and termination of the protocol defined in Section II.

Consistency. We prove consistency by contradiction. Let's assume two correct users output different commitments com_t and $\text{com}_{t'}$ for the same epoch t . Then com_t is signed by $2f + 1$ witnesses out of which $f + 1$ are assumed honest. Similarly, $\text{com}_{t'}$ is signed by $2f + 1$ witnesses out of which $f + 1$ are honest and did not sign com_t . But then there should be $f + 1 + f + 1$ honest and f malicious witnesses. But $n = 3f + 1 < 3f + 2$, hence a contradiction.

Validity. Validity directly follows from the integrity property of the signature scheme used by the IdP and the soundness of the proof π^{Upd} . Honest witnesses only counter-sign update notifications if they are correctly signed by the IdP and the proof π^{Upd} verifies. As a result, there cannot exist a valid certificate over an invalid update.

Termination. Assuming all messages are eventually delivered and there exist $2f + 1$ honest witnesses then the honest IdP will send a single commitment for each epoch t which all honest witnesses will counter-sign and produce a valid certificate for.

C. Anti-Censorship and Freshness Subprotocol

We provide an optional enhancement of the consistency protocol presented in Section V providing censorship resistance from malicious IdP. This protocol allows users to tie the liveness of their update requests with the liveness of the entire system effectively defending against selective censorship attacks. The protocol works as follows:

- Users whose updates keep not appearing in the IdP's state send their update request to the witnesses.
- Every epoch witnesses collect these requests and forward them to the IdP.
- Once the witness has forwarded a user update request, it stops signing any future IdP's update notification that do not come with a proof that the user update has been included in the state.

As a result the moment the censored client has sent their update request to $f + 1$ honest witnesses, the liveness of the protocol is tied to the update being included. Hence the IdP can now only act as crashed and halt the whole system if it wants to censor the update (at which point no other requests can be processed).

A second enhancement of the consistency protocol is that during an update the witnesses send their signature over not only the hash of the tree but also their local timestamp. Then clients are given the option to ask for $2f + 1$ timestamps during a read operation. Given that a median timestamp is robust to f faults and assuming bounded clock-drift clients can deduce the freshness of the state tree received as reply. They can thus choose to ignore the reply if it is too old (e.g., more than a day old).

Given that now the updates are coupled to anti-censorship and clients are aware of the time the state is updated in Parakeet, we guarantee that key updates are included in a timely manner in the state and every client quickly receives the new update (otherwise a malicious IdP is forced to halt the whole system).

VI. IMPLEMENTATION AND BENCHMARKS

We implement our VKD scheme presented in Section III in Rust. We also implement a networked multi-core eventually synchronous Parakeet based on our VKD. It uses `tokio` [7] for asynchronous networking, `ed25519-dalek` [8] for elliptic curve based signatures, and data-structures are persisted using `Rocksdb` [9] (unless otherwise specified). It uses TCP to achieve reliable point-to-point channels, necessary to correctly implement the distributed system abstractions. We have open-sourced both our VKD [5] and consistency protocol [5] implementations (which together form Parakeet), including our Amazon Web Services (AWS) orchestration scripts and measurements data to enable reproducible results [37].

Benchmarks. We performed several benchmarks to thoroughly assess the practicality of Parakeet. To this end, we used two kinds of AWS instances, both with up to 5Gbps of bandwidth, running Linux Ubuntu server 20.04, running on 2.5GHz Intel Xeon Platinum 8259CL machines: (1) `t3.medium`

Measure	Mean (ms)	Std. (ms)
(Witness) Verify notification	68.70	1.50
(Witness) Create vote	0.02	0.00
(Witness) Verify certificate	0.16	0.01

TABLE I: Microbenchmark of single core CPU overhead of the consistency protocol operations of Parakeet on witness machines. Committee of 4 witnesses; notifications batch 1,024 updates of 64 bytes each; average and standard dev. of 10 measurements.

instances with 2 virtual CPUs (1 physical core) and 4GB memory. These machines are relatively cheap due to their low specs and we chose them to assess how performant the witnesses of Parakeet are. (2) `t3.2xlarge` instances with 8 virtual CPUs and 32GB of memory. These machines are used to benchmark the most heavyweight IdP operations, namely publish. We implemented IdP and Witnesses in Rust and for signatures we used `ed25519_dalek` with 32-byte public/private keys pairs.

A. Microbenchmarks

We microbenchmark the overhead of the consistency protocol on the AWS `t3.medium` instance and the IdP capabilities of our implementation on the `t3.2xlarge` instances. We set the committee size to 4 witnesses, the size of keys and values to 32 bytes each (i.e. a key-value pair is 64 bytes), and a batch size of 1,024 key-values per batch. In the tables below, every measure displays the mean, standard deviation over 10 runs.

CPU analysis. Table II shows the results of microbenchmarks for the following VKD operations:

- 1) *Create notification*: The IdP generates a batch of (random) public keys and publishes them. Then it generates an audit (append-only) proof over this publish operation. The audit proof along with the the new tree root and the sequence number are signed by the IdP. This message constitutes a notification. Note that the result in table II only includes the cost of the consistency and audit related operations, which we derive by subtracting the server-side publish cost from the total cost of publishing *and* creating a notification.
- 2) *Verify notification*: A witness verifies a notification. This step consists of the verification of the IdP-generated signature and the audit proof.
- 3) *Create vote*: A witness signs the verified notification and creates a vote.
- 4) *Verify vote*: A witness verifies a vote.
- 5) *Aggregate certificate*: A list of votes are combined to form a certificate. The number of votes needed to create a certificate is $2f + 1$ (see Section V). With a four-witness setup, this means combining 3 votes.
- 6) *Verify certificate*: All 3 votes in a certificate are verified.

The slowest operations are the generation and signing of audit proofs, dispatching them to witnesses, *Create notification*, and the verification of the signature and audit proof presented in *Create notification*, i.e. *Verify notification*. Similarly, verification of these two contributed to the second slowest operation *Verify notification*. Comparatively, remaining operations were very fast – completed under one millisecond.

Measure	Mean (ms)	Std. (ms)
Create notification	47.12	9.30
Verify vote	0.07	0.00
Aggregate certificate	0.00	0.00

TABLE II: Microbenchmark of witness CPU overhead of the consistency protocol operations of Parakeet on the IdP machine; notifications batch 1,024 updates of 64 bytes each; average and standard dev. of 10 measurements.




Fig. 5: Memory consumption of Parakeet’s VKD. Measurements at intervals of 100k new users and up to 4.5M users. Total memory consumption is the sum of memory required by the Merkle Patricia Trie data, some metadata and the original database of usernames and public keys.

Storage analysis. To examine the storage costs of Parakeet, as well as SEEMless, we inserted 100,000 users at a time to the respective VKDs with a MySQL-based storage layer, running locally in a docker container. Figure 5 shows the storage cost and its breakdown for Parakeet’s VKD, as the number of users increases. For 5M users, the VKD requires about 4.5GB and this cost grows linearly in number of users. At this rate, we would expect roughly 850GB for 1 billion users – demonstrating the scalability of our solution.

Figure 6, shows the time it takes to run a publish operation on Parakeet’s VKD. The storage layer of this experiment implements a cache and uses a MySQL database for persistent storage. Each of these update operations, as well as the ones in fig. 5 required only one or two persistent storage accesses each, and the storage accesses made up a majority of the time it took for the publish operation to complete. For example, for inserting 100k users on top of a VKD containing 500k existing users, it took about 19 minutes, of which 12 were for MySQL writes. We attempted to run similar benchmarks for SEEMless. Our implementation of an aZKS included all the same caching optimizations as the oZKS implementation we used for Parakeet. With persistent memory, however, this data structure became infeasible at fairly small scale. For example, inserting 100k new users into a set of 300k users and this operation took ≈ 66 minutes. Of this, 51 were spent writing to persistent storage and about 7 were spent on reading from it. This took 47 persistent storage accesses. figs. 7 and 8 show that as the number of epochs increases, for the same set of users, SEEMless sees a large storage size blowup, versus Parakeet, whose storage size remains constant. The analysis used to plot fig. 7 and fig. 8 is included in the full version.

B. End-to-end benchmarks

We evaluate the throughput and latency of our implementation of Parakeet through experiments on AWS. We particularly aim to demonstrate the following claims:




Fig. 6: Time to publish with insertion of a batch of 100k new users into Parakeet’s VKD, with an existing VKD of sizes 0-500k. This graph also includes the time for database reads and writes for each insertion.

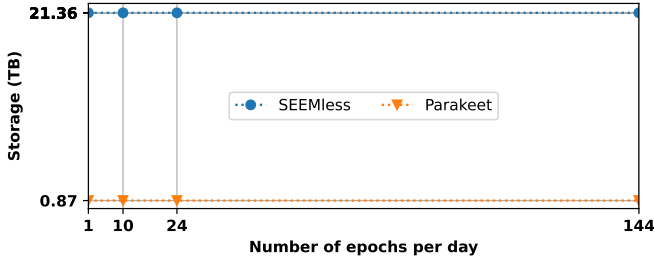


Fig. 7: Comparison of storage costs of Parakeet and SEEMless in the first year with varying number of epochs per day ranging from daily to every 10 minutes with roughly 10 million daily key updates. This graph is meant to illustrate that the storage cost of the Parakeet data structure is independent of the number of epochs and depends only on the total number of key updates.

- (C1) Parakeet achieves enough throughput to operate at planetary scale.
- (C2) Parakeet achieves low latency even under high load, in the WAN, and with large committee sizes.
- (C3) Parakeet runs efficiently on cheap machines with low specs (comparable to common HSMs).
- (C4) Parakeet is robust when some parts of the system inevitably crash-fail. Note that evaluating BFT protocols in the presence of Byzantine faults is still an open question [11].

We deploy a testbed on AWS, using `t3.medium` instances across 5 different AWS regions: N. Virginia (`us-east-1`), N. California (`us-west-1`), Sydney (`ap-southeast-2`), Stockholm (`eu-north-1`), and Tokyo (`ap-northeast-1`). Witnesses are distributed across those regions as equally as possible.

In the following sections, each measurement in the graphs is the average of 3 independent runs, and the error bars represent one standard deviation; errors bars are sometimes too small to be visible on the graphs. Our baseline experiment parameters are 10 honest witnesses, a batch size of 1024, and an update size of 64B. We instantiate one benchmark client (collocated on the same machine as the IdP) submitting transactions at a fixed rate for a duration of 5 minutes. When referring to *latency*, we mean the time elapsed from when the client submits a request to when the IdP receives confirmation that the request is successfully processed. We measure it by tracking sample requests throughout the system.

Benchmark in the common case. Figure 9 illustrates the latency and throughput of Parakeet for varying numbers of witnesses. The maximum throughput we observe is around 800




Fig. 8: Comparison of storage costs of Parakeet and SEEMless in the first five years with 10 minute epochs and 10 million daily updates.




Fig. 9: Throughput-latency performance of Parakeet. WAN measurements with 10, 20, 50 witnesses. No faulty witnesses, 1024 maximum batch size, and 64B updates size.

updates/s while keeping the latency below 3.5 seconds. Based on the system usages estimates for the large-scale end-to-end encrypted messaging service WhatsApp (Section I), we would arrive at the requirement to process around 120 updates/s. Parakeet exceeds by over 6x the throughput required to operate at this scale, and thus satisfies claims (C1) and (C3).

Figure 9 also illustrates that performance do not vary with 10, 20 or even 50 witnesses. This observation concurs with Section VI-A showing that the bottleneck of Parakeet is the IdP. Increasing the number of geo-distributed witnesses up to 50 doesn’t impact the end-to-end performance of the system; Parakeet thus satisfies claim (C2). We however expect that keeping increasing the number of witnesses will eventually make the network to become the system’s bottleneck.

Figure 10 illustrates the performance of Parakeet when varying the batch size form 2^5 to 2^{15} . The maximum throughput we observe for batches sizes of 2^5 and 2^7 is respectively 100 updates/s and 350 updates/s. This is much lower than the 800 updates/s that Parakeet can achieve when configured with a batch size over 2^{10} . Small batch sizes, however, allow Parakeet to trade throughput for latency. Parakeet configured with a batch size of 2^5 can process up to 100 updates/s in under 800ms, setting the batch size to 2^7 allows Parakeet to operate at scale while robustly maintaining sub-second latency.

Benchmark under crash-faults. Figure 11 depicts the performance of Parakeet when a committee of 10 witnesses suffers 1 to 3 (crash-)faults (the maximum that can be tolerated in this setting). It shows that Parakeet’s performance is not affected by (crash-)faults, thus satisfying claim (C4).

Contrarily to BFT consensus systems [26], Parakeet maintains a good level of throughput under crash-faults. The




Fig. 10: Throughput-latency performance of Parakeet. WAN measurements with 10 witnesses. No faulty witnesses, various batch sizes, and 64B updates size.




Fig. 11: Comparative throughput-latency under crash-faults of Parakeet. WAN measurements with 10 witnesses. Zero, one, and three crash-faults, 1024 maximum batch size, and 64B updates size.

underlying reason for steady performance under crash-faults is that Parakeet doesn’t rely on a leader to drive the protocol. This is in sharp contrast with related work (e.g. [15], [12], [40]) that rely on an external blockchain for consistency.

VII. RELATED WORK

Key transparency. As discussed in Section I, this work extends SEEMless [15] by instantiating a light-weight consistency protocol to prevent server equivocation in a practical setting, while also extending SEEMless to handle real-world constraints on storage capacity, efficiency, and scalability to billions of users. While Keybase ([21], [24]) was the first deployment of an auditable public key directory (created as a user-friendly alternative for PGP), CONIKS [32], was the first academic work that formalized the notion of a service that maintains and periodically commits to a public key directory. SEEMless itself can be seen as an extension to CONIKS. CONIKS essentially relies on a synchronous gossip protocol among users in order to detect server equivocation. Unfortunately, this assumption is both hard to scale efficiently for millions of users and easily breakable over the internet.

Since then, multiple works have proposed different ways of publishing directory commitments and serving audit proofs to clients. EthIKS [12] and Catena [40] demonstrate how to leverage a public blockchain (Ethereum and Bitcoin, resp.) as a ledger to provide non-equivocation guarantees and support auditability for the commitment to the CONIKS directory, while minimizing bandwidth overhead for its clients. In this paper we show that the strong assumption of using consensus is unnecessary and instead propose a lightweight consistency protocol. Finally, Mog [31] takes a different approach by proposing a gossip protocol and a verifiable registry that

allows individual clients to perform their own auditing. The gossip protocol still requires synchrony, however, it relies on the assumption of sleepy committees to solve the consensus problem. In theory, this would enable scalability but could hinder liveness for long periods when there are not enough available witnesses or when the network is unstable. In this work, instead, we have shown that consensus is unnecessary for guaranteeing a consistent view of the tree to the clients which enables us to provide consistent views without relying on a good network. Our work could be extended to using sleepy committees however we opted to use highly available witnesses that should be deployed by professional services (e.g., professional blockchain validators) and will provide timely security to the users.

Several recent works on key transparency have focused on improving auditor efficiency. Merkle² [23] proposes a solution which reduces the amount of work required for an auditor to verify a series of key updates from linear to logarithmic, assuming that key updates can be signed by clients through the use of “signature chains”. However, it is unclear what integrity protection the system will provide if the signature keys of the clients are lost. Assuming that the clients can maintain long term cryptographic secret keys is unrealistic, especially in the setting of key transparency, where the focus is on building a PKI for clients who cannot remember cryptographic key material. AAD [39], Aardvark [27], Tomescu et. al. [41], Tyagi et. al. [42], and Verdict [43] have proposed using accumulators (bilinear and RSA) and SNARKs as commitments in order to make auditor verification more efficient. However, the computational overhead incurred from relying on the algebraic assumptions themselves can outweigh the asymptotic improvements over the number of key updates per epoch.

Atomic transactions. Atomic transactions [25] allow all-or-none type of execution for a set of operations. In large systems, they are often a necessity since the underlying database can end up in an inconsistent state if operations are not sequentially executed (e.g., withdrawing money from account A and depositing it to account B). To this end, several solutions have been proposed in the literature—[35], [38], [18] to name a few. Although they are a strong primitive for building concurrent applications, transactions come with their cost; locks might leave the systems in a dead-locked state whereas failure in a single operation can cancel a transaction and might require re-execution of the whole set; or alternatively they might not be supported cross-shard [17].

In VKD, we side-step such issues by (1) executing a publish operation by a single writer and (2) preserving the previous value of a node. (1) is needed to ensure that concurrent publish operations not overwrite nodes’ updated states and (2) allows us to allow concurrent reads (e.g., lookup proof generation) and writes (i.e., publish operation). In result of these two properties, we only require that the update order is preserved only between the node updates and the latest epoch, i.e., we allow the nodes to be updated in any order.

The downside of this approach is that the storage cost

is effectively doubled. Yet, we believe this is an acceptable trade-off due to drastic storage reduction compared to existing key transparency solutions such as SEEMless [15] and the flexibility to use any storage as the underlying key directory.

VIII. CONCLUSION

While much recent effort has focused on various aspects of key transparency, large-scale applications on the order of billions of users have not previously been considered. We expose the gaps in purely academic scale implementations and bridge these gaps is our design of Parakeet, a key transparency system with large-scale deployment in mind. Our production-grade implementation of Parakeet shows the feasibility of our approach, which we further demonstrate through experiments.

ACKNOWLEDGMENTS

This work is supported by the Novi team at Meta and funded in part by IC3 industry partners and NSF grant 1943499.

REFERENCES

- [1] “How whatsapp enables multi-device capability,” <https://engineering.fb.com/2021/07/14/security/whatsapp-multi-device/>, [Online; accessed 5-July-2022].
- [2] “Most popular global mobile messenger apps as of January 2022, based on number of monthly active users,” <https://www.statista.com/statistics/258749/most-popular-global-mobile-messenger-apps/>, [Online; accessed 5-July-2022].
- [3] “Whatsapp revenue and usage statistics (2022),” <https://www.businessofapps.com/data/whatsapp-statistics/>, [Online; accessed 5-July-2022].
- [4] <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN#d1e2606-1-1>, 2015.
- [5] <https://github.com/facebook/akd>, 2022.
- [6] <https://github.com/Microsoft/oZKS>, 2022.
- [7] <https://github.com/tokio-rs/tokio>, 2022.
- [8] <https://github.com/dalek-cryptography/ed25519-dalek>, 2022.
- [9] <https://rocksdb.org/>, 2022.
- [10] S. Bano, A. Sonnino, M. Al-Bassam, S. Azouvi, P. McCorry, S. Meiklejohn, and G. Danezis, “Sok: Consensus in the age of blockchains,” in *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, 2019, pp. 183–198.
- [11] S. Bano, A. Sonnino, A. Chursin, D. Perelman, Z. Li, A. Ching, and D. Malkhi, “Twins: Bft systems made robust,” in *25th International Conference on Principles of Distributed Systems (OPODIS 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- [12] J. Bonneau, “Ethiks: Using ethereum to audit a coniks key transparency log,” in *International Conference on Financial Cryptography and Data Security*. Springer, 2016, pp. 95–105.
- [13] C. Cachin, R. Guerraoui, and L. Rodrigues, *Introduction to reliable and secure distributed programming*. Springer Science & Business Media, 2011.
- [14] D. Catalano, D. Fiore, and M. Messina, “Zero-knowledge sets with short proofs,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2008, pp. 433–450.
- [15] M. Chase, A. Deshpande, E. Ghosh, and H. Malvai, “Seemless: Secure end-to-end encrypted messaging with less trust,” in *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, 2019, pp. 1639–1656.
- [16] M. Chase, A. Healy, A. Lysyanskaya, T. Malkin, and L. Reyzin, “Mercurial commitments with applications to zero-knowledge sets,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2005, pp. 422–439.
- [17] A. Cheng, X. Shi, L. Pan, A. Simpson, N. Wheaton, S. Lawande, N. Bronson, P. Bailis, N. Crooks, and I. Stoica, “Ramp-tao: layering atomic transactions on facebook’s online tao data store,” *Proceedings of the VLDB Endowment*, vol. 14, no. 12, pp. 3014–3027, 2021.
- [18] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. J. Furman, S. Ghemawat, A. Gubarev, C. Heiser, P. Hochschild *et al.*, “Spanner: Google’s globally distributed database,” *ACM Transactions on Computer Systems (TOCS)*, vol. 31, no. 3, pp. 1–22, 2013.
- [19] C. Dwork, N. Lynch, and L. Stockmeyer, “Consensus in the presence of partial synchrony,” *Journal of the ACM (JACM)*, vol. 35, no. 2, pp. 288–323, 1988.
- [20] M. J. Fischer, N. A. Lynch, and M. S. Paterson, “Impossibility of distributed consensus with one faulty process,” *Journal of the ACM (JACM)*, vol. 32, no. 2, pp. 374–382, 1985.
- [21] N. Group, https://keybase.io/docs-assets/blog/NCC_Group_Keybase_KB2018_Public_Report_2019-02-27_v1.3.pdf, 2019.
- [22] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, “Eclipse attacks on bitcoin’s peer-to-peer network,” in *USENIX 2015*, J. Jung and T. Holz, Eds. USENIX Association, 2015, pp. 129–144. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/heilman>
- [23] Y. Hu, K. Hooshmand, H. Kalidhindi, S. J. Yang, and R. A. Popa, “Merkle²: A low-latency transparency log system,” in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 285–303.
- [24] Keybase, https://keybase.io/_api/1.0/merkle/root.json?seqno=1, 2014.
- [25] B. W. Lamson, “Atomic transactions,” in *Distributed Systems Architecture and Implementation*. Springer, 1981, pp. 246–265.
- [26] H. Lee, J. Seibert, E. Hoque, C. Killian, and C. Nita-Rotaru, “Turret: A platform for automated attack finding in unmodified distributed system implementations,” in *2014 IEEE 34th International Conference on Distributed Computing Systems*. IEEE, 2014, pp. 660–669.
- [27] D. Leung, Y. Gilad, S. Gorbunov, L. Reyzin, and N. Zeldovich, “Aardvark: A concurrent authenticated dictionary with short proofs,” Cryptology ePrint Archive, Report 2020/975, 2020, <https://eprint.iacr.org/2020/975>.
- [28] D. Malkhi and M. Reiter, “Byzantine quorum systems,” *Distributed computing*, vol. 11, no. 4, pp. 203–213, 1998.
- [29] M. Marlinspike, “Advanced cryptographic ratcheting,” <https://signal.org/blog/advanced-ratcheting/>, 2013.
- [30] —, “Whatsapp’s signal protocol integration is now complete,” <https://signal.org/blog/whatsapp-complete/>, 2016.
- [31] S. Meiklejohn, P. Kalinnikov, C. S. Lin, M. Hutchinson, G. Belvin, M. Raykova, and A. Cutter, “Think global, act local: Gossip and client audits in verifiable data structures,” *arXiv preprint arXiv:2011.04551*, 2020.
- [32] M. S. Melara, A. Blankstein, J. Bonneau, E. W. Felten, and M. J. Freedman, “CONIKS: Bringing key transparency to end users,” in *24th USENIX Security Symposium (USENIX Security 15)*, 2015, pp. 383–398.
- [33] S. Micali, M. Rabin, and J. Kilian, “Zero-knowledge sets,” in *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings*. IEEE, 2003, pp. 80–91.
- [34] K. Nikitin, E. Kokoris-Kogias, P. Jovanovic, N. Gailly, L. Gasser, I. Khoffi, J. Cappos, and B. Ford, “{CHAINIAC}: Proactive software-update transparency via collectively signed skipchains and verified builds,” in *26th {USENIX} Security Symposium ({USENIX} Security 17)*, 2017, pp. 1271–1287.
- [35] D. Peng and F. Dabek, “Large-scale incremental processing using distributed transactions and notifications,” in *9th USENIX Symposium on Operating Systems Design and Implementation (OSDI 10)*, 2010.
- [36] K. Pothong, L. Pshetz, R. Catlow, and S. Meiklejohn, “Problematising transparency through LARP and deliberation,” in *DIS ’21: Designing Interactive Systems Conference 2021, Virtual Event, USA, 28 June, July 2, 2021*, W. Ju, L. Oehlberg, S. Follmer, S. E. Fox, and S. Kuznetsov, Eds. ACM, 2021, pp. 1682–1694.
- [37] A. Sonnino, <https://github.com/asonnino/key-transparency/tree/main/scripts>, 2022.
- [38] Y. Sovran, R. Power, M. K. Aguilera, and J. Li, “Transactional storage for geo-replicated systems,” in *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*, 2011, pp. 385–400.
- [39] A. Tomescu, V. Bhupatiraju, D. Papadopoulos, C. Papamantou, N. Triandopoulos, and S. Devadas, “Transparency logs via append-only authenticated dictionaries,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, pp. 1299–1316.
- [40] A. Tomescu and S. Devadas, “Catena: Efficient non-equivocation via bitcoin,” in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 393–409.

- [41] A. Tomescu, Y. Xia, and Z. Newman, “Authenticated dictionaries with cross-incremental proof (dis)aggregation,” Cryptology ePrint Archive, Report 2020/1239, 2020, <https://eprint.iacr.org/2020/1239>.
- [42] N. Tyagi, B. Fisch, J. Bonneau, and S. Tessaro, “Client-auditable verifiable registries,” *Cryptology ePrint Archive*, 2021.
- [43] I. Tzialla, A. Kothapalli, B. Parno, and S. Setty, “Transparency dictionaries with succinct proofs of correct operation,” *IACR Cryptol. ePrint Arch.*, p. 1263, 2021. [Online]. Available: <https://eprint.iacr.org/2021/1263>

APPENDIX

A. VKD with Compaction Glossary

In order to discuss details about our approach to the compaction of server-side storage, we will need some notation. We gather notation for our construction here:

- **TOMBSTONE**: A special string to replace the value of an oZKS entry which is valid to delete.
- **TombstoneEpochs**: The set of epochs when tombstone marking is permitted. No insertions should be allowed in this epoch, only tombstoning.
- **CompactionEpochs** or **DeletionEpochs**: The set of epochs when entries with value **TOMBSTONE** are permitted to be deleted from the oZKS.
- **StaleParam**: A system parameter for calculating the upper bound for the epochs in which a node must have been inserted in order to be valid for tombstoning.
- **DeletionParam**: A system parameter defining the number of epochs between a tombstone epoch and the subsequent compaction epoch. Assuming **CompactionEpochs** and **TombstoneEpochs** are sorted in increasing order, $\text{CompactionEpochs} = \{t_i + \text{DeletionParam} \mid t_i \text{ is the } i\text{th element of } \text{TombstoneEpochs}\}_i$.

We use the terms identity provider and server interchangeably.

B. Formal Definitions

Definition 1. Given a datastore, DS, i.e. a set of tuples $(\text{label}_i, \text{val}_i)$, DS.Labels , denotes the set $\{\text{label} \mid \exists(\text{label}, \cdot) \in \text{DS}\}$.

Definition 2. An ordered append-only zero-knowledge set (with compaction) (oZKS) is a data-structure with a set of public parameters \mathfrak{p} , where \mathfrak{p} includes a set **TombstoneEpochs** of epochs, and parameters **StaleParam**, **DeletionParam**. An oZKS supports the following algorithms (oZKS.CommitDS, oZKS.QueryMem, oZKS.QueryNonMem, oZKS.VerifyMem, oZKS.VerifyNonMem, oZKS.InsertionToDS, oZKS.VerifyInsertions, oZKS.TombstoneElts, oZKS.VerifyTombstone, oZKS.DeleteElts, oZKS.VerifyDel). We defer the details of these algorithms in the full version.

The technical definition for the soundness of an oZKS with compaction is included in the full version.

Definition of a VKD. [15] defined a VKD as a collection of the following algorithms (included here, without description, for completeness): **VKD.Publish**, **VKD.Query**, **VKD.VerifyQuery**, **VKD.KeyHistory**, **VKD.VerifyHistory**, **VKD.VerifyUpd**, **VKD.Audit**.

Definition of a VKD with compaction. Our modified primitive, which we call a VKD with compaction, denoted as cVKD,

is defined below. Note the additional set of public parameters: (**TombstoneEpochs**, **CompactionEpochs**, **StaleParam**). These determine when tombstoning (ordered set **TombstoneEpochs**) or compaction (ordered set **CompactionEpochs**) are permitted, and how old a piece of data needs to be, in order to be eligible for tombstoning (**StaleParam**).

A VKD with compaction (cVKD) is a VKD with public parameters (**TombstoneEpochs**, **CompactionEpochs**, **StaleParam**) and following additional algorithms: **cVKD.TombstoneElts**, **cVKD.VerifyTombstone**, **cVKD.Compact** and **cVKD.VerifyCompact**. We defer the technical details to the full version.

The constraints on **TombstoneEpochs**, **CompactionEpochs** and **StaleParam** for a cVKD are as follows:

- If t_i and t_{i+1} are two epochs in the ordered set **TombstoneEpochs**, then, there exists a deletion epoch $t_{\text{del}} \in \text{DeleteElts}$, between them. That is, $t_i < t_{\text{del}} < t_{i+1}$.
- A piece of data can only be marked as tombstoned in a particular epoch $t_i \in \text{TombstoneEpochs}$ if it was inserted prior to the last tombstone epoch t_{i-1} . That is, $\text{StaleParam} > \min\{|t - t'| : t, t' \in \text{TombstoneEpochs}\}$.

Definition 3. A VKD with compaction is said to be sound if, for any PPT adversary \mathcal{A} ,

$$\begin{aligned} & \Pr[(\text{label}, \{t, \{(\text{val}_k^t, \text{Epoch}_k^t)\}_{k=\min_t}^{\max_t}, \Pi_t^{\text{Ver}}\}_{t \in \{t_1, \dots, t_n\}}, \\ & \{(\text{com}_k, \Pi_k^{\text{Upd}})\}_{k=t_1}^{t_{\text{current}}}, t^*, j, (\pi, \text{val}, \beta) \leftarrow \mathcal{A}(1^\lambda) : \\ & \quad \wedge \{\text{com}_k \leftarrow \text{WitnessAPI.GetCom}(k)\}_{k=t_1}^{t_{\text{current}}} \\ & \quad \wedge \text{VKD.Audit}(t_1, t_{\text{current}}, \{(\Pi_k^{\text{Upd}})\}_{k=t_1}^{t_{\text{current}}-1}) \\ & \quad \wedge \forall t \in [t_1, \dots, t_n], \text{VKD.VerifyHistory}(t, \text{label}, \\ & \quad \{(\text{val}_i^t, \text{Epoch}_i^t)\}_{i=\min_t}^{\max_t}, \Pi_t^{\text{Ver}}) \\ & \quad \wedge \forall t \in \text{TombstoneEpochs}, s.t. \\ & \quad [t, t + \text{DeletionParam}] \cap [t_1, \dots, t_{\text{current}}] \neq \emptyset \\ & \quad \implies \{t_1, \dots, t_n\} \cap [t, t + \text{DeletionParam}] \neq \emptyset \\ & \quad \wedge t_1 \leq t_2 \leq \dots \leq t_n \leq t_{\text{current}} \wedge \text{StaleParam} > \\ & \quad \min\{|t - t'| \mid t, t' \in \text{TombstoneEpochs}\} \\ & \quad \wedge (\\ & \quad (\exists t, t' \in \{t_1, \dots, t_n\}, \gamma, (\text{val}_\gamma^t, \text{Epoch}_\gamma^t) \neq (\text{val}_\gamma^{t'}, \text{Epoch}_\gamma^{t'})) \\ & \quad \vee \\ & \quad ((t_{j-1} \leq t^* < t_j) \wedge (\text{Epoch}_\alpha^{t_j} \leq t^* < \text{Epoch}_\alpha^{t_{j+1}} \vee \alpha = \max_{t_j}) \\ & \quad \wedge \text{val}_\alpha^{t_j} \neq \text{val} \wedge t^* < t'_j \wedge \text{VKD.VerifyQuery}(t^*, \text{label}, \text{val}, \pi)) \\ & \quad)] \leq \text{negl}(\lambda). \end{aligned}$$

C. Constructing an oZKS

Our oZKS construction, first implemented by [6] is very similar to the aZKS construction of [15], based on a strong accumulator (SA), a simulatable commitment scheme (sCS) and a strong VRF (sVRF). We defer the details of this construction to the full version.

1) Efficiency improvements due to oZKS

The major advantage of constructing an oZKS and including the epoch a leaf was inserted in the tree is that if audits are

honestly verified, a SA membership proof includes when the leaf was inserted. This means that only the latest commitment of the SA needs to be verified, even in the case of historical queries about some label. This alleviates the need to store all states of an evolving oZKS for the VKD construction, making space complexity of the VKD's state st linear in the number of leaves, rather than in the number of epochs, unlike the construction using the aZKS of [15]. Additionally, this means that the RAM complexity of multiple, simultaneous history queries can be amortized over the number of queries and the number of proofs π_i^j being proven over all the history queries.

D. cVKD: VKD with Secure Compaction

Here, we present our final construction of a VKD with secure compaction and defer the reader to the full version for intermediate constructions which motivate this final design.

1) Two-phase Compaction

In this section, we present our final construction, which we call a *two-phase compaction*, which allows a VKD to support the algorithm `Compact`, without additional privacy leakage. As stated before, even if compaction were supported with fixed epochs which are demarcated for compaction, since after compaction, there is no record of the changes made, unless a user is online at the compaction epoch, she cannot ensure that any oZKS entries associated with her label were not modified. We weaken this requirement by allowing a *grace period* for a user to check her key history, before a compaction, by marking oZKS entries slated for deletion as tombstoned for a while, before they are actually deleted.

Now, we use the algorithms (`oZKS.TombstoneElts`, `oZKS.VerifyTombstone`, `oZKS.DeleteElts`, `oZKS.VerifyDel`), defined in section III-C to construct a VKD with two-phase compaction as follows.

- $(\text{com}_t, \text{Dir}_t, \text{st}_t, \Pi^{\text{Upd}}, t) \leftarrow \text{VKD.TombstoneElts}(\text{Dir}_{t-1}, \text{st}_{t-1}, t-1, t_{\text{stale}})$: This algorithm initializes $S_t^{\text{oZKS}} = \emptyset$, $S_t^{\text{AVD}} = \emptyset$. For $(\text{label}, L_{\text{label}}) \in \text{Dir}_{t-1}$, for $(\alpha, \text{val}, t_\alpha) \in L_{\text{label}}$: (1) If $t_\alpha \leq t_{\text{stale}}$, it adds $(\text{label}|\alpha)$ to S_t . (2) If $\alpha > 1$, it also adds $(\text{label}|\alpha - 1|'stale')$ to S_t . Then, it obtains the required $\text{st}_t^{\text{oZKS}}$ and DS_t and calls `oZKS.TombstoneElts`($\text{st}_{\text{com}}, \text{DS}_t, S_t, t_{\text{stale}}$). It updates the corresponding state, deletes the corresponding $(\alpha, \text{val}, t_\alpha)$ tuples from L_{label} for each label. Then, the algorithm calls `WitnessAPI.ProposeNewEp`($t, \text{com}_t, \Pi^{\text{Upd}}$). If `WitnessAPI.ProposeNewEp` outputs \perp , the algorithm reverts all its internal changes to st_{t-1} , Dir_{t-1} , and outputs \perp . Finally, this algorithm returns the updated commitments and proof output by `oZKS.TombstoneElts`.
- $0/1 \leftarrow \text{VKD.VerifyTombstone}(\text{com}_t, \text{com}_{t+1}, \Pi^{\text{Upd}}, t_{\text{stale}})$: This algorithm calls `WitnessAPI.GetCom`(t) to obtain $(\text{com}_t, \text{cert}_t)$ and `WitnessAPI.GetCom`($t+1$) to get $(\text{com}_{t+1}, \text{cert}_{t+1})$. Finally, it parses Π^{Upd} to get $\Pi_{\text{oZKS}}^{\text{Upd}}$. If either `WitnessAPI.VerifyCert`($\text{com}_t, \text{cert}_t, p$) or `WitnessAPI.VerifyCert`($\text{com}_{t+1}, \text{cert}_{t+1}, p$) outputs 0, the algorithm outputs 0. Otherwise, it returns the output of `oZKS.VerifyDel`($\text{com}_{\text{oZKS}}, \text{com}'_{\text{oZKS}}, \Pi_{\text{oZKS}}^{\text{Upd}}, t_{\text{stale}}$).

- $(\text{com}_{t+1}, \text{Dir}_{t+1}, \text{st}_{t+1}, \Pi^{\text{Upd}}, t+1) \leftarrow \text{VKD.Compact}(\text{Dir}_t, \text{st}_t, t, t_{\text{stale}})$: This algorithm parses out $\text{DS}_t^{\text{oZKS}}, \text{st}_t^{\text{oZKS}}$ from st_t and calls `oZKS.DeleteElts`($\text{st}_t^{\text{oZKS}}, \text{DS}_t^{\text{oZKS}}, t, t_{\text{stale}}$), to obtain $(\text{com}_{t+1}^{\text{oZKS}}, \text{st}_{t+1}^{\text{oZKS}}, \text{DS}_{t+1}^{\text{oZKS}}, \pi_S, t+1)$, which it uses to update $\text{com}_t, \text{st}_t, \text{Dir}_t$ to get $\text{com}_{t+1}, \text{st}_{t+1}, \text{Dir}_{t+1}$ and sets $\Pi^{\text{Upd}} = \pi_S$ and returns $(\text{com}_{t+1}, \text{Dir}_{t+1}, \text{st}_{t+1}, \Pi^{\text{Upd}}, t+1)$.
- $0/1 \leftarrow \text{VKD.VerifyCompact}(\text{com}_t, \text{com}_{t+1}, \Pi^{\text{Upd}}, t_{\text{stale}})$: This algorithm simply parses out $\text{com}_t^{\text{oZKS}}, \text{com}_{t+1}^{\text{oZKS}}$ from $\text{com}_t, \text{com}_{t+1}$, respectively, and outputs the output of `oZKS.VerifyDel`($\text{com}_t^{\text{oZKS}}, \text{com}_{t+1}^{\text{oZKS}}, \Pi^{\text{Upd}}, t_{\text{stale}}$).

The `VerifyHistory` algorithm has the following extra checks:

- $0/1 \leftarrow \text{VKD.VerifyHistory}(t, \text{label}, (\text{val}_i, t_i)_{i=\alpha_{\min}}^n, \Pi^{\text{Ver}})$: This algorithm is the same as the original `VKD.VerifyHistory`, except it also ensures that if $\text{val}_i = \text{TOMBSTONE}$, then both the labels $\text{label}|i$ and $\text{label}|i|stale$ have the value `TOMBSTONE`. It also ensures that if $\text{label}|i$ does not correspond to the value `TOMBSTONE`, then neither does $\text{label}|i|stale$. Finally, it ensures that $\text{val}_n \neq \text{TOMBSTONE}$, i.e. the most recent value of this user is not tombstoned. If all checks pass, this algorithm outputs 1, otherwise it outputs 0.

The updated `VKD.Audit` is defined as follows:

- $\text{VKD.Audit}(t_1, t_n, (\Pi_t^{\text{Upd}})_{t=t_1}^{t_n-1})$: If for any of the epochs, $t \in [t_1, t_n - 1]$, t is in `TombstoneEpochs`, the auditor calls `VKD.VerifyTombstone`($\text{com}_t, \text{com}_{t+1}, \Pi_t^{\text{Upd}}, t - \text{StaleParam}$), else, if $t - \text{DeletionParam} \in \text{TombstoneEpochs}$, i.e. $t \in \text{CompactionEpochs}$, call `VKD.VerifyCompact`($\text{com}_t, \text{com}_{t+1}, \Pi_t^{\text{Upd}}, t - \text{DeletionParam} - \text{StaleParam}$). For all other $t \in [t_1, t_n - 1]$, this algorithm calls `VKD.VerifyUpd`($t, t+1, (\text{com}_t, \Pi_t^{\text{Upd}}), \text{com}_{t+1}$). It outputs 1, if all subroutines output 1, otherwise it outputs 0.

Soundness for VKD with Compaction.

Theorem 1. *The construction for a VKD with compaction in appendix D1 satisfies Definition 3.*

We refer the reader to the full version for the proof.

E. Summary of the aZKS Construction from [15]

Append-only Zero Knowledge Set (aZKS) was introduced in [15] (Section 3). We summarize the construction below. Part of the text is copied verbatim from [15].

aZKS is a primitive that lets a (malicious) prover commit to an append-only dictionary of (label,value) pairs (where the labels form a set) such that: 1) the commitment is succinct and does not leak any information about the committed dictionary 2) the prover can prove statements about membership and non-membership of (label,value) pairs with respect to the succinct commitment 3) the prover can prove that for two dictionaries, $D_1, D_2, D_1 \subseteq D_2$ with respect to their respective succinct commitments and 4) the proofs are efficient and do not leak any information about the rest of the committed dictionary.

aZKS has two security properties: *Soundness* and *Zero-Knowledge Privacy (with leakage)*. Soundness ensures that a malicious prover 1) will not be able to produce two verifying

proofs for two different values for the same label with respect to a commitment or 2) should not be able to modify an existing label. Zero-Knowledge privacy property captures that the query proofs and append-proofs leak no information beyond the query answer (which is a value or \perp for membership queries and a bit indicating validity of an append operation) and a well-defined leakage function.

F. Summary of SEEMless construction [15]

In this section, we summarize the construction from [15], Section 4. Part of the text is copied verbatim from [15].

SEEMless assumes that server's identity and public key is known to each user and auditor and all the messages from the server are signed under the server's key, so that the server cannot be impersonated.

SEEMless uses two aZKS: one "all" aZKS to store all versions of all (label, val) pairs, and a second "old" aZKS that stores all of the out of date label versions. SEEMless also uses a *hash chain*. Hash Chain is a classical authentication data structure that chains multiple data elements by successively applying cryptographic hashes, e.g. a hash chain that hashes elements a, b, c is $H(c, H(b, H(a)))$.

- VKD.Publish(Dir_{t-1}, st_{t-1}, S_t): At every epoch, the server gets a set S_t of (label, value) pairs that have to be added to the VKD. The server first checks if the label already exists for some version $\alpha - 1$, else sets $\alpha = 1$. It adds a new entry (label | α , val) to the "all" aZKS and also adds (label | $\alpha - 1$, null) to the "old" aZKS if $\alpha > 1$. If the new version $\alpha = 2^i$ for some i , then the server adds a marker entry (label | mark | i , "marker") to the "all" aZKS. The server computes commitments to both the aZKS, adds them to the hash chain to obtain a new head com_t . It also produces a proof Π^{Upd} consisting of the previous and new pair of aZKS commitments $com_{all,t-1}, com_{all,t}, com_{old,t-1}, com_{old,t}$ and the corresponding aZKS update proofs.
- VKD.Query($st_t, Dir_t, label$): When a client Bob queries for Alice's label, he should get the val corresponding to the latest version α for Alice's label and a proof of correctness. Bob gets three proofs in total: (1) a membership proof of (label | α , val) in the "all" aZKS. (2) membership proof of the most recent marker entry (label | mark | a) for $\alpha \geq 2^a$. (3) a non membership proof of label | α in the "old" aZKS. Proof 2 ensures that Bob is not getting a value higher than Alice's current version and proof 3 ensures that Bob is not getting an old version for Alice's label.
- VKD.VerifyQuery($t, label, val, \pi, \alpha$): The client checks each membership/non-membership proof, the hash chain and that version α as part of proof is less than current epoch t .
- VKD.KeyHistory($st_t, Dir_t, t, label$): The server first retrieves all the update epochs t_1, \dots, t_α for label versions $1, \dots, \alpha$ from T , the corresponding $com_{all,t_1-1}, com_{all,t_1}, \dots, com_{all,t_\alpha-1}, com_{all,t_\alpha}$ and $com_{old,t_1}, \dots, com_{old,t_\alpha}$ and the hashes necessary to verify the hash chain: $H(com_{all,0}, com_{old,0}), \dots, H(com_{all,t}, com_{old,t})$. For versions $i = 1$ to n , the server retrieves the val_i for t_i and version i of label from Dir_{t_i} . Let $2^a \leq \alpha < 2^{a+1}$ for some a where α is




Fig. 12: Schematic of the building blocks for SEEMless.




Fig. 13: Schematic of the building blocks for Parakeet. Note that the shaded components are different from the corresponding components of SEEMless, which are shaded in Figure 12.

the current version of the label. The server generates the following proofs (together called as Π):

- 1) **Correctness of com_{t_i} and com_{t_i-1}** : For each i , output $com_{t_i} com_{t_i-1}$. Also output the values necessary to verify the hash chain: $H(com_{all,0}, com_{old,0}), \dots, H(com_{all,t}, com_{old,t})$.
 - 2) **Correct version i is set at epoch t_i** : For each i : Membership proof for (label | i) with value val_i in the "all" aZKS with respect to com_{t_i} .
 - 3) **Server couldn't have shown version $i - 1$ at or after t_i** : For each i : Membership proof in "old" aZKS with respect to com_{t_i} for (label | $i - 1$).
 - 4) **Server couldn't have shown version i before epoch t_i** : For each i : Non membership proof for (label | i) in "all" aZKS with respect to com_{t_i-1} .
 - 5) **Server can't show any version from $\alpha + 1$ to 2^{a+1} at epoch t or any earlier epoch**: Non membership proofs in the "all" aZKS with respect to com_t for (label | $i + 1$), (label | $i + 2$), \dots , (label | $2^{a+1} - 1$).
 - 6) **Server can't show any version higher than 2^{a+1} at epoch t or any earlier epoch**: Non membership proofs in "all" aZKS with respect to com_t for marker nodes (label | mark | $a + 1$) up to (label | mark | $\log t$).
- VKD.VerifyHistory($t, label, (val_i, t_i)_{i=1}^n, \Pi^{Ver}$): Verify each of the above proofs.
 - VKD.Audit($t_1, t_n, (\Pi_t^{Upd})_{t=t_1}^{t_n-1}$): Auditors will audit the commitments and proofs to make sure that no entries ever get deleted in either aZKS. They do so by verifying the update proofs Π^{Upd} output by the server. They also check that at each epoch both aZKS commitments are added to the hash chain. Note that, while the Audit interface gives a monolithic audit algorithm, our audit is just checking the updates between each adjacent pair of aZKS commitments, so it can be performed by many auditors in parallel.