

SEEMless: Secure End-to-End Encrypted Messaging with *less* trust

Melissa Chase
Microsoft Research
melissac@microsoft.com

Apoorvaa Deshpande
Brown University
apoorvaa_deshpande@brown.edu

Esha Ghosh
Microsoft Research
esha.ghosh@microsoft.com

Harjasleen Malvai
Cornell University
hm553@cornell.edu

Abstract—End-to-end encrypted messaging (E2E) is only secure if participants have a way to retrieve the correct public key for the desired recipient. However, to make these systems usable, users must be able to replace their keys (e.g. when they lose or reset their devices, or reinstall their app), and we cannot assume any cryptographic means of authenticating the new keys. In the current E2E systems, the service provider manages the directory of public keys of its registered users; this allows a compromised or coerced service provider to introduce their own keys and execute a man in the middle attack.

Building on the approach of CONIKS (Melara et al, USENIX Security '15), we formalize the notion of a *Privacy-Preserving Verifiable Key Directory* (VKD): a system which allows users to monitor the keys that the service is distributing on their behalf. We then propose a new VKD scheme which we call SEEMless, which improves on prior work in terms of privacy and scalability. In particular, our new approach allows key changes to take effect almost immediately; we show experimentally that our scheme easily supports delays less than a minute, in contrast to previous work which proposes a delay of one hour.

Index Terms—Privacy-preserving verifiable directory service, Zero knowledge sets, Accumulators, Persistent Patricia Trie, History Tree, PKI, Transparency, Security definitions

I. INTRODUCTION

A number of popular messaging apps such as iMessage, WhatsApp and Signal have recently deployed end-to-end encryption (E2EE) in an attempt to mitigate some of the serious privacy concerns that arise in these services. E2EE is a system of communication where only encrypted messages leave the sender's device. These messages are then downloaded and decrypted on the recipient's device, ensuring that only the communicating users can read the messages. But E2EE relies on a Public Key Infrastructure (PKI); this in practice requires the provider of the messaging service (such as Apple, Facebook, Microsoft etc.) to maintain a centralized directory of the public keys of its registered users. To be accessible to average users, these systems assume the user will store her secret key on her personal device, and do not assume she has any other way to store long term secrets. When a user loses her device (and thus her secret key), she will need to generate a new (secret key, public key) pair and replace her old public key stored in the PKI with the newly generated public key.

Such a system naturally places a lot of trust in the service provider – a malicious service provider (or one who is compelled to act maliciously, possibly because of a compromise) can arbitrarily set and reset users' public keys. It might, for

example replace an honest user's public key with one whose secret key it knows, and thus implement a man-in-the-middle attack without the communicating users ever noticing. Ironically, this defeats the purpose of E2EE. Without some way of verifying that the service provider is indeed returning the correct keys, E2E encryption does not provide any protection against malicious (or coerced) service providers.

Some service providers provide a security setting option to notify a sender when a recipient's public key changes. In WhatsApp, the sender can scan a QR code on the recipient's device to verify the authenticity of the new public key. This option, however, is turned off by default to provide a seamless user experience. Moreover, the communicating users will be able to verify each other's QR codes only if their devices are physically close to each other, which is most often not the case. And these features are something few users use, as is evidently uncovered in this attack [wha].

To enable E2EE with real security, we need to keep the inherent constraints of the system in mind. To begin with, the primary objective of any E2EE messaging system is to provide a *secure and seamless* communication service between remote users. This problem is made even more challenging by the fact that we must assume that a user can lose her device and along with it all of her secrets. Moreover, in our attempt to reduce trust on the service provider, we must not introduce any new attack surface. For example, if we introduce a mechanism that will enable a user to verify that she is receiving the correct key of the intended recipient, this mechanism should not leak any additional information about the other users and the public keys they have registered with the service provider. Privacy may not be very important in a traditional PKI, where all the players are usually public entities like businesses, but in the context of private messaging, privacy is very important. Hiding usernames may help prevent spam messaging. And the user's update pattern may itself be sensitive information. Users change their keys primarily when they change devices, or detect that their devices/keys have been compromised, either of which may be sensitive. Moreover, if a user rarely changes her key, then compromising her secret key gives the attacker the ability to decrypt large volumes of messages at once, making her a more vulnerable and attractive target. Or if a device is compromised but the user does not update her key, then the attacker knows the compromise has gone undetected.

To address this fundamental problem, we design and pro-

vide a prototype implementation of SEEMless – a privacy-preserving directory service for end-user key verification. We also initiate the study of the privacy and security guarantees of such a key directory service formally. Our major contributions are the following.

Formalizing Key Directory Services. We formalize the security and privacy requirements of a verifiable key directory service, (such as SEEMless, CONIKS [MBB⁺15], EthIKS [Bon16]), in terms of a new primitive that we call Verifiable Key Directories (VKD) (Section II). A VKD consists of three types of entities: an identity provider or server, users, and external auditors. The server stores a directory *Dir* with the names of the users and their corresponding keys. VKD provides different query interfaces to the users to interact with the server: 1. Alice can register her (username, key) to *Dir* and update it at any point* 2. Bob can query for Alice’s key 3. Alice can obtain the history of her key updates. This last query allows Alice to verify that the only key updates that occurred were those that she requested (we assume that Alice roughly remembers when she changed her device or reinstalled her software forcing a key update). Note that we do not assume Alice can remember any cryptographic secret as she can lose her device any time. We discuss the assumptions we make for a VKD system in Section IV.

A key criterion of such a system is that Alice must be able to verify that any user Bob who queried for her key received a value consistent with what the server tells her when she checks her key history. To enable consistency checking, a VKD server periodically signs and publishes a privacy-preserving digest of its latest *Dir*. We assume an additional set of parties which we call auditors who help verify that the server is showing the same view of *Dir* to all its users by periodically auditing these digests. These auditors are not trusted for privacy, so any interested party (privacy conscious users, or privacy advocacy organizations) can serve as an auditor. If a VKD server ever tries to equivocate by issuing multiple keys for a single username, this would require publishing invalid digests or conflicting digests which would provide irrefutable proof of the server’s equivocation.

Soundness of a VKD requires that Bob sees Alice’s latest key as long as Alice has checked her key history and there is at least one honest auditor that audited each of the published server digests. *Privacy* is parameterized by a well specified leakage function \mathcal{L} ; informally, we require that the published digests and the proofs should not leak any information about *Dir* beyond the query answer and $\mathcal{L}(\text{Dir})$. For example, for every query performed by the users or auditors, our construction leaks the total number of key registrations and updates performed so far. In addition, when Bob queries for Alice’s key, he learns when Alice’s key was last updated and how many times she has changed her key so far, but nothing else about her past or future updates, or about any other user.

*We assume that the identity provider can verify that the update request comes from Alice through some out-of-band, non-cryptographic check, like verifying security questions, sending a text to an appropriate phone number, and/or requiring voice or in person communication.

Specifying a precise leakage function leads to a better understanding of the privacy guarantee of any VKD system. As a concrete example, we were able to identify a *tracing attack* in current implementations [MBB⁺15], [Bon16] through which it might be possible to trace the entire update history of a particular user. We explain this attack and the corresponding leakage in Appendix.

Building Blocks. We take a modular approach in designing our VKD system, SEEMless (Section IV). We first define a new primitive, *Append-Only Zero Knowledge Set* (aZKS), that generalizes the definition of a traditional static zero-knowledge set [MRK03], [CHL⁺05] by accounting for updates. We construct SEEMless using aZKS in a blackbox way.

Then we construct aZKS modularly from a *strong accumulator* [CHKO08] and a *verifiable random function* [CLO7]. By providing modular constructions, we simplify the presentation and analysis and allow for independent study of each of these building blocks: an improvement in any of them would directly translate to improvement in SEEMless.

The high level idea in our VKD construction (Section IV) is to use aZKS to commit to the current directory and an append-only Merkle Hash Tree to efficiently maintain server commitments over time. A zero knowledge set allows a server to commit to a set of (label, value) pairs, respond to lookup queries for different labels, and prove that those responses are correct without revealing any additional information. Append-only ZKS allow for additional pairs to be added to the database; that is the only legal update. Since the aZKS is append only and each label can only appear once, we append each username with a version number showing how many times that key has been updated so far.

We use *two* aZKS; one which consists of all the (username | version, key) pairs corresponding to all the updates made so far. The other aZKS has (username | version, key) pairs with all the versions except the latest, i.e. the versions that are now old. The purpose of maintaining two aZKS is to prevent the server from serving stale keys; now when Alice updates her key, instead of deleting the old entry, the server will add it to the “old” aZKS, and when Bob queries for Alice’s key the server will additionally prove that the entry it returns is not “old”. We have a more sophisticated approach for preventing the server from showing Bob a higher version number for Alice’s key, which involves storing a “marker” every 2th update; we defer details to Section IV.

Auditors in the system will verify that (1) the server’s published commitments are linearly ordered and (2) the underlying aZKS are indeed append-only. The first goal can be achieved using a hash chain (or for improved efficiency an append-only Merkle Hash Tree as described above). SEEMless does not take a position on how to ensure that users and auditors eventually see the same commitments. This can be achieved either with a gossip protocol [MBB⁺15] or by posting the commitments on a public blockchain like [TD17], [Bon16].

Persistent Patricia Trie. Recall that SEEMless provides a KeyHistory API c all which allows Alice to audit all her key

changes; the same is true for CONIKS in order to support users who are not always online. Adding this functionality requires the server to maintain the history of the two aZKS starting with the very first epoch. This is obviously very space consuming and wasteful since there might be a significant overlap between the aZKS of consecutive epochs. Motivated by this constraint, we develop a *Persistent Patricia Trie* (PPTr) data structure and use it to build a persistent version of our strong accumulator, i.e. one in which we can also answer queries relative to older versions of the data structure. This directly translates to space-efficient persistent aZKS. PPTr is a space-efficient history-preserving data structure (like history tree [OB09]) which lets the server generate historical proofs (KeyHistory in our case) efficiently. The details of this construction are in Section VI. Our PPTr construction presents a way of maintaining any temporal, append-only, authenticated log of data. We believe PPTr will find applications in several contexts, including tamper-evident logging, distributed identity providers and distributed certificate authorities and may be of independent interest.

Epoch length. One key parameter in a VKD is the length of the epochs between server updates. Note that, when Alice updates her key, the change is not reflected in the server’s digest until the next update. This presents two options: either the server can continue to respond with the old key until the end of the epoch, or the server can respond with a key which is not consistent with the digest. CONIKS proposes this latter approach; it is unclear how the user can verify that this new key is in fact correct. In our solution, we propose making the epochs short enough that we can take the first approach without significantly affecting security or usability of the system. This requires significant efficiency and scalability improvements over previous work.

Privacy and Efficiency Improvement. SEEMless improves upon the existing VKD constructions [MBB⁺15], [Bon16] on privacy and scalability aspects. Our construction is more space efficient for the server, thanks to our PPTr construction. In particular, the space required depends not on the number of epochs, but only on the number of key updates performed, despite the fact that we must be able to answer questions about historical key values. We also significantly improve on the monitoring cost of [MBB⁺15], [Bon16]. CONIKS [MBB⁺15] required that Alice query the server for every epoch to ensure that her key was correct at that time; this meant both that epochs couldn’t be too short without substantial burden on the system, and that a user who was offline for a period would have to perform all the checks for that time on coming back online. ETHIKS [Bon16] improved this, by moving most of the work to the blockchain, but the epoch length is still limited by the blockchain block frequency.[†] In our system, using the KeyHistory API, Alice can monitor her key revisions with cost that depends on the number of times her key has been updated rather than the number of epochs that have passed. Similarly she can be offline for several server epochs and can

[†]It also required Alice to remember her last version number which we feel is less natural to remember.

run KeyHistory when she comes online. This together with the storage improvements mean that we can support much shorter epochs, which in turn leads to a more secure system as discussed above, in that Alice’s key updates can be very quickly reflected in the directory. Finally, our construction provides stronger privacy guarantees compared to [MBB⁺15], [Bon16].

Experiments We give a prototype implementation of our VKD system, SEEMless and experimentally evaluate the cost incurred for the users, the server and the auditors. We also run CONIKS to fairly compare our performance costs. In our experiments, the cost of lookup query in SEEMless is slightly more expensive than in CONIKS (at most 3×) due to our stronger privacy guarantees, but it is still about 10ms from request to complete verification, which is reasonably fast. However, SEEMless performs and scales well with frequent server updates, as opposed to CONIKS. For the same update frequency and experimental setup, CONIKS cannot handle more than 4.5M users, owing to the inefficiency of their underlying data structure (our system could handle 10M users seamlessly) and hangs with more frequent epochs due to garbage collection overhead. Moreover, users of SEEMless need not monitor their keys as frequently as the server publishes its digest; they can verify their entire key history by downloading around 2MB (even if they are updating their keys often) when they come back online; in contrast, a CONIKS, user will need to download about 577MB in a similar scenario, which makes the monitoring cost quickly infeasible. The details of our experiments and performance comparison with CONIKS is in Section VII.

Organization of the paper: In Section II, we define the primitive of Verifiable Key Directories (VKD) along with the security properties. In Section III, we define append-only Zero Knowledge Sets (aZKS) which is a building block for VKD construction. In Section IV, we describe our SEEMless construction starting with an overview of the same. In Section V, we give concrete instantiations of aZKS. In Section VI we describe our space efficient Persistent Patricia Trie construction. We present our prototype implementation and performance results in Section VII. Finally, in Appendix A we discuss some design considerations, including assumptions we are making about the users, the underlying infrastructure, and the way that this system would be integrated with a messaging app.

II. VERIFIABLE KEY DIRECTORY (VKD)

In this section, we will define the primitive of a *Verifiable Key Directory* (VKD) and formalize its properties. The goal of a VKD is to capture the functionality and security of a privacy-preserving verifiable key directory system such as CONIKS or EthIKS [MBB⁺15], [Bon16].

A VKD consists of three types of parties: an identity provider or server, clients or users and external auditors. The server stores a directory *Dir* with the names of the users (which we call labels) and their corresponding public keys (the values corresponding to the labels). For the ease of exposition, let

Alice and Bob be two users. VKD provides the following query interface to the users. 1) Alice can add her (username, key), i.e., (label=username, val=key) to Dir. 2) Alice can update her key and request that Dir be updated with the new key value. 3) She can query the server periodically to obtain history of her key updates over time (VKD.KeyHistory). 4) Bob can also query for the key corresponding to username Alice (VKD.Query) at the current time.

The functionality VKD.KeyHistory warrants further discussion since this is not a functionality that one usually expects from a key directory service. In a privacy-preserving verifiable key directory service, intuitively, we expect the server to be able to prove to Bob that he is seeing Alice’s latest key without leaking any additional information about the directory. This is trivial to achieve if we assume that Alice can always sign her new public key with her old secret key. But this is a completely unreasonable assumption from an average user who may lose her device or re-install the software, thereby losing her previous secret key; the user will only have access to her latest secret key which is stored on her latest device. It is crucial to *not assume* that Alice or Bob can remember any cryptographic secret. Under this constraint, we need Alice to monitor her key sufficiently often to make sure her latest key is in the server directory. Only then, we can talk about Bob getting Alice’s latest key in a meaningful way. Alice could of course check every epoch to make sure that her key is being correctly reported, but this becomes costly, particularly when epochs are short. Instead, we allow Alice to query periodically and retrieve a list of all the times her key has changed and the resulting values. This is precisely the VKD.KeyHistory interface which Alice users to monitor her own key.

The server applies updates from its users (of type 1 and 2 described above) in batches, and publishes a commitment to the current state of the database com and proof Π^{Upd} that a valid update has been performed (VKD.Publish). The batch updates should happen at sufficiently frequent intervals of time, so that the user’s keys are not out-of-date for long. The exact interval between these time intervals, or epochs has to be chosen as a system parameter. We use time and epoch interchangeably in our descriptions. The server also publishes a public datastructure which maintains information about all the commitments so far.

The auditors in the system keep checking the update proofs and the public datastructure in order to ensure global consistency (VKD.Audit) of Dir. Our definition captures a general notion of audit where independent auditors can audit arbitrary intervals $[t_1, t_n]$. The audit need not be monolithic and can consist of many independent audit steps. As long as some honest auditor executes each audit step, soundness will be guaranteed.

The server also produces proofs for VKD.Query and VKD.KeyHistory. At a very high level, the users verify the proofs (VKD.QueryVer, VKD.HistoryVer) to ensure that the server is not returning an incorrect key corresponding to a username or an inconsistent key history for the keys corresponding to a username. VKD also requires the proofs to be

privacy-preserving, i.e., the proofs should not leak information about any other key (that has not been queried) in Dir. The auditors may not be trusted and hence, the proofs that the server produces as part of VKD.Publish need to be privacy-preserving. Since the auditors do not have to be trusted with any private information, anyone can become an auditor.

Definition 1. A Verifiable Key Directory is comprised of the algorithms (VKD.Publish, VKD.Query, VKD.QueryVer, VKD.KeyHistory, VKD.HistoryVer, VKD.Audit) and all the algorithms have access to the system parameters. We do not make the system parameters explicit. The algorithms are described below.

Periodic Publish:

$\triangleright (\text{com}_t, \Pi_t^{\text{Upd}}, \text{st}_t, \text{Dir}_t) \leftarrow \text{VKD.Publish}(\text{Dir}_{t-1}, \text{st}_{t-1}, S_t)$:

This algorithm takes the previous state of the server and the key directory at previous epoch $t - 1$ and also a set S_t of elements to be updated. Whenever a client submits a request to add a new label or update an existing label from epochs $t - 1$ to t , the corresponding (label, val) pair is added to S_t to be added in the VKD at epoch t . The algorithm produces a commitment to the current state of the directory com_t and a proof of valid update Π^{Upd} all of which it broadcasts at epoch t . It also outputs the updated directory Dir_t and an updated internal state st_t . If this is the first epoch, i.e., no previous epoch exists, then the server initializes an empty directory and its internal state first, then updates them as described above. We will denote $(\text{com}_t, \Pi_t^{\text{Upd}})$ as pub_t .

Querying for a Label:

$\triangleright (\text{val}, \pi) \leftarrow \text{VKD.Query}(\text{st}_t, \text{Dir}_t, \text{label})$: This algorithm takes the current state of the server for epoch t , the directory Dir_t at that epoch and a query label label and returns the corresponding value if it is present in the current directory, \perp if it is not present, a proof of membership or non-membership respectively.

$\triangleright 1/0 \leftarrow \text{VKD.QueryVer}(\text{com}, \text{label}, (\text{val}, \pi))$:

This algorithm takes a commitment with respect to some epoch, a label, value pair and verifies the above proof.

Checking Consistency of Key Updates:

$\triangleright (\{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}}) \leftarrow \text{VKD.KeyHistory}(\text{st}_t, \text{Dir}_t, \text{label})$: This algorithm takes in the server state, the directory at current time t and a label. It outputs $\{(\text{val}_i, t_i)\}_{i=1}^n$ which are all the times at which the value corresponding to label was updated so far, the resulting val’s, along with a proof Π^{Ver} .

$\triangleright 1/0 \leftarrow \text{VKD.HistoryVer}(\text{com}_t, \text{label}, \{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}})$: This algorithm takes the commitment published by the server for the current time t , a label, and $\{(\text{val}_i, t_i)\}_{i=1}^n$, and verifies the above proof.

Auditing the VKD:

$\triangleright 1/0 \leftarrow \text{VKD.Audit}(t_1, t_n, \{\text{pub}_t\}_{t=t_1}^{t_n})$: This algorithm takes the epochs t_1 and t_n between which audit is being done, the server’s published pub for all the epochs from times t_1 to t_n . It outputs a boolean indicating whether the audit is successful. Note that for auditing we assume that all auditors see consistent versions of published commitments and the users see

the same commitments as the auditors. That is, everyone sees the same broadcast values com_t at any epoch t . This can be enforced in different ways. For a discussion on the different implementation mechanisms, please see Section A.

Now we discuss the security properties we require from a VKD. We give the informal descriptions here and defer the formal definitions to the Appendix. The security properties are the following.

- **Completeness:** We want to say that if a VKD is set up properly and if the server behaves honestly at all epochs, then all the following things should happen for any label updated at t_1, \dots, t_n with $\text{val}_1, \dots, \text{val}_n$: 1) their history proof with $\{(\text{val}_i, t_i)\}_{i=1}^n$ and Π^{Ver} should verify at t_n 2) the query proof for the label at any $t_j \leq t^* < t_{j+1}$ should verify with respect to the value consistent with the versions proof at t_j which is val_j and 3) the audit from epochs t_1 to t_n should verify. Note that for KeyHistory and HistoryVer, we consider epochs t_1, t_2, \dots, t_n when the updates have happened for a label. These will be epochs distributed in the range $[t_1, t_n]$. However for Audit, we consider all possible pairwise epochs between t_1 and t_n . For example, for $t_1 = 3$ to $t_n = 10$, there might be updates at 3, 5, 8, 10 but for audit we need to consider all of the epochs 3, 4, 5, 6, 7, 8, 9, 10.

- **Soundness:** VKD soundness guarantees that if Alice has verified the update history of her key till time t_n and if there exists at least one honest auditor whose audits have been successful from the beginning of time till time t_n then, whenever Bob queried before t_n , he would have received Alice's key value that is consistent with the key value reported in Alice's history query. Thus soundness is derived from all of VKD.Publish, VKD.QueryVer, VKD.HistoryVer and VKD.Audit.

Note that the onus is on the user, Alice, to make sure that the server is giving out the most recent and *correct* value for her key. Soundness guarantees that under the circumstances described above, Bob will always see a key consistent with what Alice has audited. But Alice needs to verify that her key as reported in the history query is consistent with the actual key that she chose.

- **Privacy:** The privacy guarantee of a VKD system is that the outputs of Query, HistoryVer or Audit should not reveal anything beyond the answer and a well defined leakage function on the state of the directory. In other words, the proofs for each of these queries should be simulatable given the output of the leakage function and the query answer.

III. APPEND-ONLY ZERO KNOWLEDGE SET (aZKS)

In this section we introduce a new primitive, *Append-Only Zero Knowledge Set* (aZKS) which we will use to build SEEMless. Zero Knowledge Set [MRK03], [CHL⁺05] (ZKS) is primitive that lets a (potentially malicious) prover commit to a static collection 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 collection

2) the prover can prove statements about membership/non-membership of labels (from the domain of the labels) in the committed collection with respect to the succinct commitment 3) the proofs are efficient and do not leak any information about the rest of the committed collection. Our primitive, *Append-Only Zero Knowledge Set* (aZKS) generalizes the traditional zero-knowledge set primitive by accounting for append-only updates and characterizing the collection with a leakage function.

Here it is worth pointing out that the notion of soundness one would expect from updates in a ZKS is not obvious. For example, if the expectation is that updates leak absolutely no information about the underlying sets or type of updates (inserts/deletes), then there is no reasonable definition of soundness of updates: any set the prover chooses will be the result of some valid set of updates. In [Lis05], Liskov did not define any soundness notion for updates. In our context, we want to be able to define an append-only ZKS, which makes the expectation of update soundness clear: it should ensure for any label, its value never gets modified and in particular, it never gets deleted.

Here we describe the primitive and informally define its security properties. The formal security definition is deferred to the full version.

Definition 2. Append-Only Zero Knowledge Set is comprised of the algorithms (ZKS.CommitDS, ZKS.Query, ZKS.Verify, ZKS.UpdateDS, ZKS.VerifyUpd)[‡] described as follows:

▷ $(\text{com}, \text{st}_{\text{com}}) \leftarrow \text{ZKS.CommitDS}(1^\lambda, D)$: This algorithm takes the security parameter and the datastore to commit to as input, and produces a commitment to the data store and an internal state to pass on to the Query algorithm. Datastore D will be a collection of (label, val) pairs.

▷ $(\pi, \text{val}) \leftarrow \text{ZKS.Query}(\text{st}_{\text{com}}, D, \text{label})$: This algorithm takes the state output by ZKS.CommitDS, the datastore and a query label and returns its value (\perp if not present) and a proof of membership/non-membership.

▷ $1/0 \leftarrow \text{ZKS.Verify}(\text{com}, \text{label}, \text{val}, \pi)$: This algorithm takes a (label, value) pair, its proof and a commitment by ZKS.CommitDS and verifies the above proof.

▷ $(\text{com}', \text{st}'_{\text{com}}, D', \pi_S) \leftarrow \text{ZKS.UpdateDS}(\text{st}_{\text{com}}, D, S)$: This algorithm takes in the current server state st_{com} , the current state of the datastore and a set $S = \{(\text{label}_1, \text{val}_1), \dots, (\text{label}_k, \text{val}_k)\}$ of new (label, value) pairs for update. It outputs an updated commitment to the datastore, an updated internal state and an updated version of the datastore and proof π_S that the update has been done correctly.

▷ $0/1 \leftarrow \text{ZKS.VerifyUpd}(\text{com}, \text{com}', \pi_S)$: This algorithm takes in two commitments to the datastore before and after an update and verifies the above proof.

[‡]The original ZKS definition also included a setup algorithm run by a trusted party to generate public parameters used in all the algorithms. Our construction does not need such a set up (we show security in the random oracle model), so we omit it here.

We require the following security properties from an *append-only* ZKS:

Soundness: For soundness we want to capture two things: First, a malicious prover \mathcal{A}^* algorithm should not be able to produce two verifying proofs for two different values for the same label with respect to a com. Second, since the aZKS is append-only, a malicious server should not be able to change or delete an existing label.

Zero-Knowledge with Leakage: We generalize the definition ZKS [MRK03], [CHL⁺05] by introducing leakage functions in the classical definition. The goal of our privacy definition is to capture the following: we want the query proofs and update proofs to leak no information beyond the query answer (which is a val/ \perp in case of query and a bit indicating validity of the update operation). But often, it is reasonable to tolerate a small leakage to gain more efficiency. To capture this sort of leakage formally, we parameterize our definition with a leakage function. If the leakage function is set to null, then our definition reduces to the classical ZKS definition.

Append-Only Strong Accumulator: Finally, we remark that the primitive *strong accumulator* (SA) [CHKO08] can be extended to a new primitive of *append-only strong accumulator* (aSA) trivially from the definition of aZKS. aSA is essentially a aZKS with completeness and soundness and *without* privacy requirement. In Section V, we will first construct an efficient aSA and then construct aZKS using aSA in a blackbox way.

IV. SEEMLESS CONSTRUCTION

In this section, we will describe SEEMless: our construction of VKD from *append-only* zero-knowledge sets (aZKS). We first give an informal overview of the construction and then describe the construction more formally.

System Setup: The high level idea is to have two aZKS that are updated every epoch: one is what we call the “all” aZKS which consists of the entries for all the updates of a label so far. Here label is the username and the value is the user’s public key. The other aZKS that we will refer to as “old” aZKS has entries with all the versions of the label except the latest, versions that are now old. The purpose of maintaining two aZKS is that when a user Alice updates her key, instead of deleting the old entry, it will be added to the “old” aZKS. This means we can use an append operation to capture key updates, while relying on the append-only property to guarantee that the server maintains an accurate record of past updates. In both of these aZKS, the server stores each (label, value) pair along with its version number, so when an item is initially added, we add (label| 1, val) indicating that this label is on version 1. The “all” aZKS will also include *marker entries* i for version 2^i of each label. Intuitively, the markers will help limit the checks that Alice needs to make when verifying her key history. It will help her make sure that the server is not responding to queries with version numbers much higher than the correct version. The server also stores an internal table T with a list for each label of all of the epochs when it has been updated. The server also maintains a hash chain with each pair

of “old” and “all” aZKS commitments. We refer to the head of this hashchain at time t as com_t .

Remark 1. We describe our construction with a hash chain to avoid confusion with the Merkle tree used in our aSA construction in Section V-B. However, we note that this could instead be replaced with a Merkle tree built over the list of all commitments to date. This would result in slightly higher update and audit costs (adding a new entry to the end of this list would require up to a logarithmic number of hashes), but would significantly reduce the cost of history queries (from linear in the number of epochs to logarithmic). We discuss this in Section VII (Update, Gethistory and Audit experiments).

Periodic Publish: 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 α , else sets $\alpha = 0$. It adds a new entry (label | $\alpha + 1$, val) to the “all” aZKS and also adds (label | α , val_{old}) to the “old” aZKS for the label’s previous value val_{old} if it exists i.e. for $\alpha > 0$. If the new version $\alpha + 1 = 2^i$ 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, and 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 $\text{com}_{\text{all}, t-1}$, $\text{com}_{\text{all}, t}$ and $\text{com}_{\text{old}, t-1}$, $\text{com}_{\text{old}, t}$ and the corresponding aZKS update proofs. However, the auditors can get the update proofs from the server by explicitly querying it, i.e., the server *need not broadcast* them along with the new hashchain head com_t . We do not require any non-tamperability or privacy guarantees from the distribution mechanism of the update proofs.

Querying for a Label Value: 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: First is the membership proof of (label | α , val) in the “all” aZKS. Second is the membership proof of the most recent marker entry (label | mark | a) for $\alpha \geq 2^a$. And third is 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.

Querying for the History of Key Updates: Now we are explicitly maintaining versions in the construction and hence one label will have multiple entries associated with it depending on the number of times it was updated. Perhaps a malicious server can give out an older version of the label or it could create an entry corresponding to a later version on its own and give that whenever someone queries. In order to prevent these behaviors we incorporate the some checks in the key history queries. The purpose of each of these checks are described in the algorithm description.

Auditing: 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 timestamp 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. For security, it is sufficient to have at least one honest auditor perform audits over each adjacent pair.

Privacy: The main insight for the gain in privacy is that in our construction we always add new entries since we are using aZKS, we never delete or update existing entries. This hides which entries are being updated and hides the lifespan of each key. However, our construction does leak some benign information that we make explicit in each algorithm description (Query, History and Audit).

VKD CONSTRUCTION: In the construction we assume that the 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.

▷ VKD.Publish($\text{Dir}_{t-1}, \text{st}_{t-1}, S_t$) :

$t = 0$: This is when we are setting up the directory for the first time and let $\text{Dir}_0 = (D_{\text{all},0}, D_{\text{old},0})$ as initialized before. Compute $(\text{com}_{\text{all},0}, \text{st}_{\text{all},0}) \leftarrow \text{ZKS.CommitDS}(\text{st}_{\text{all}}, D_{\text{all},0})$ and $(\text{com}_{\text{old},0}, \text{st}_{\text{old},0}) \leftarrow \text{ZKS.CommitDS}(\text{st}_{\text{old}}, D_{\text{old},0})$. Set the hash chain head to $\text{com}_0 = H(H(\text{com}_{\text{all},0}, \text{com}_{\text{old},0}), 0)$. Initialize an update set and a marker set for next epoch $S_0 = \phi$ and $M_0 = \phi$. Also, initialize a table T which will later store the usernames (labels) and their times of updates for all epochs until the current.

Output $(\text{com}_0, \text{st}_0 = (\text{st}_{\text{all},0}, \text{st}_{\text{old},0}, T, S_0, M_0))$ and $\text{Dir}_0 = (D_{\text{all},0}, D_{\text{old},0})$.

$t > 0$: Parse $\text{Dir}_{t-1} = (D_{\text{all},t-1}, D_{\text{old},t-1})$. The directory Dir_{t-1} includes $(\text{label} \mid i, \text{val})$ pairs and marker entries that are already part of it before time t and S_t is the set of updates to be added to the directory at epoch t . Let $S_t = \{(\text{label}_1, v_1), \dots, (\text{label}_k, v_k)\}$ be the (label, value) pairs to be added to the VKD at epoch t . These could be new additions or updates to existing labels.

- For each $\text{label}_i \in S_t$, append t to the list of times (if it exists) or add (label_i, t) for new labels in the table T . If label_i is a new addition, v_i^0 will not exist. Let α_i be the version corresponding to each label_i which is the number of entries in T for label_i , 0 if no entry exists. Based on the version number, we will create the (label, value) pairs to be added in the “all” aZKS.
- If for any label_i , $\alpha_i + 1 = 2^y$ for some $y \geq 0$ then add the following (label, value) pair to the marker set M_t : $(\text{label}_i \mid \text{mark} \mid y, \text{“marker entry } 2^y \text{ for label}_i\text{”})$.
- Compute new update set to do ZKS.UpdateDS : $S'_t = \{(\text{label}'_i, \text{val}_i) \mid (\text{label}_i, \text{val}_i) \in S_t \wedge \text{label}'_i = \text{label} \mid \alpha_i + 1\}$ Compute the update on the “all” aZKS for the set $S'_t \cup M_t$: $(\text{com}_{\text{all},t}, \text{st}_{\text{all},t}, D_{\text{all},t}, \pi_{S_t}) \leftarrow \text{ZKS.UpdateDS}(\text{st}_{\text{all},t-1}, D_{\text{all},t-1}, S'_t \cup M_t)$.
- Form a new set of (label, value) pairs to be added to the “old” aZKS. Let S_t^o be the list of entries in S_t that are updates to existing labels. For each $\text{label} \in S_t^o$, concatenate it with its version before the update. Hence, let $S_t^{o'} =$

$\{(\text{label}'_i, \text{val}_i) \mid (\text{label}_i, \text{val}_i) \in S_t^o \wedge \text{label}'_i = \text{label} \mid \alpha_i\}$. For the “old” tree, compute $(\text{com}_{\text{old},t}, \text{st}_{\text{old},t}, D_{\text{old},t}, \pi_{S_t^o}) \leftarrow \text{ZKS.UpdateDS}(\text{st}_{\text{old},t-1}, D_{\text{old},t-1}, S_t^{o'})$.

- Update the hash chain: $\text{com}_t = H(H(\text{com}_{\text{all},t}, \text{com}_{\text{old},t}), \text{com}_{t-1})$.

Output $\text{com}_t, \Pi_t^{\text{Upd}} = (\pi_{S_t}, \pi_{S_t^o}), \text{st}_t = (\text{st}_{\text{all},t}, \text{st}_{\text{old},t}, T, S_{t+1}, M_{t+1}, \text{Dir}_{t-1})$ and $\text{Dir}_t = (D_{\text{all},t}, D_{\text{old},t})$.

▷ VKD.Query($\text{st}_t, \text{Dir}_t, \text{label}$) : Retrieve latest version number α for queried label from table T (by counting the number of time entries for label). Let β be the largest power of 2 less than α such that $\beta = 2^b$. Compute the following proofs:

- $(\pi_1, \text{val}_1) \leftarrow \text{ZKS.Query}(\text{st}_{\text{all},t}, D_{\text{all},t}, \text{label} \mid \alpha)$: This gives a proof of membership of the latest version of label in the “all” aZKS and its corresponding value.
- $(\pi_2, \text{val}_2) \leftarrow \text{ZKS.Query}(\text{st}_{\text{all},t}, D_{\text{all},t}, \text{label} \mid \text{mark} \mid b)$: This gives a proof of membership of the marker entry right before the current version α .
- $(\pi_3, \text{val}_3) \leftarrow \text{ZKS.Query}(\text{st}_{\text{old},t}, D_{\text{old},t}, \text{label} \mid \alpha)$: This gives a proof of non membership of the latest version in the “old” aZKS making sure that the claimed “latest” version is not outdated.

Output $\Pi = (\pi_1, \pi_2, \pi_3)$ and $\text{val} = (\text{val}_1, \text{val}_2, \perp)$ and α .

Leakage: The leakage is the version number of the queried key, the epoch at which it was last updated, and the total number of key registrations and of key updates (across all users) up until the current time.

▷ VKD.QueryVer($\text{com}_t, \text{label}, \text{val}_i, \pi_t, \alpha$) : The client checks each membership or non-membership proof, and the hash chain. Also check that version α as part of proof is less than current epoch t .

▷ VKD.KeyHistory($\text{st}_t, \text{Dir}_t, \text{label}$): The server first retrieves all the update times t_1, \dots, t_α for label versions $1, \dots, \alpha$ from T , the corresponding $\text{com}_{\text{all},t_1-1}, \text{com}_{\text{all},t_1}, \dots, \text{com}_{\text{all},t_\alpha-1}, \text{com}_{\text{all},t_\alpha}$ and $\text{com}_{\text{old},t_1}, \dots, \text{com}_{\text{old},t_\alpha}$ and the hashes necessary to verify the hash chain: $H(\text{com}_{\text{all},0}, \text{com}_{\text{old},0}), \dots, H(\text{com}_{\text{all},t}, \text{com}_{\text{old},t})$. 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 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 $\text{com}_{t_i} \text{ com}_{t_i-1}$. Also output the values necessary to verify the hash chain: $H(\text{com}_{\text{all},0}, \text{com}_{\text{old},0}), \dots, H(\text{com}_{\text{all},t}, \text{com}_{\text{old},t})$.
- 2) **Correct version i is set at time t_i :** For each i : Membership proof for $(\text{label} \mid i)$ 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 $(\text{label} \mid i-1)$.
- 4) **Server couldn’t have shown version i before time t_i :** For each i : Non membership proof for $(\text{label} \mid i)$ in “all” aZKS with respect to com_{t_i-1}

- 5) **Server can't show any version from $\alpha + 1$ to $2^{\alpha+1}$ at time t or any earlier time:** Non membership proofs in the “all” aZKS with respect to com_t for $(\text{label} \mid i+1), (\text{label} \mid i+2), \dots, (\text{label} \mid 2^{\alpha+1} - 1)$.
- 6) **Server can't show any version higher than $2^{\alpha+1}$ at time t or any earlier time:** Non membership proofs in “all” aZKS with respect to com_t for marker nodes $(\text{label} \mid \text{mark} \mid a+1)$ up to $(\text{label} \mid \text{mark} \mid \log t)$.

Leakage: The leakage is the total number of key registrations and of updates (across all users) in each epoch.[§]

▷ $\text{VKD.HistoryVer}(\text{com}_t, \text{label}, \{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}})$: Verify each of the above proofs.

Remark 2. *The client software runs VKD.HistoryVer to monitor the history of the user keys. This software either downloads the entire proof from the server each time it runs VKD.HistoryVer (when the software is re-installed or the user installs it on a new device) or cache parts of the proof to from the first run of VKD.HistoryVer to use in the subsequent verifications. We experimentally evaluate the performance for VKD.HistoryVer with and without caching in Section VII.*

▷ $\text{VKD.Audit}(t_1, t_n, \{\text{pub}_t\}_{t=t_1}^{t_n})$: Recall $\text{pub}_t = (\text{com}_t, \text{acc}_t, \Pi_t^{\text{Upd}})$. Do the checks $\text{ZKS.VerifyUpd}(\text{com}_{\text{all},t}, \text{com}_{\text{all},t-1}, \Pi_{\text{all},t}^{\text{Upd}})$ and $\text{ZKS.VerifyUpd}(\text{com}_{\text{old},t}, \text{com}_{\text{old},t-1}, \Pi_{\text{old},t}^{\text{Upd}})$ for all consecutive epochs from t_1 to t_n . Also check the the hash root is correctly updated.

Leakage: The leakage to the auditors is the total number of new registrations and of key updates in each epoch.

Remark 3. *In our implementation, the marker entry doubles up as the normal key entry for that version, where the version number is $\log(k)$, k being the count of number of updates for a specific user. We use prefix 1 for marker entries and prefix 0 for non-marker entries. For example, the 3rd key for Alice will be labeled using the value $0 \mid \text{Alice} \mid 3$ and the 4th entry will be saved with $1 \mid \text{Alice} \mid 2$, hence only two new entries are made for any user update request. This also reduces the leakage and is captured in our leakage functions above.*

Discussion: We discuss some of the design considerations and assumptions that SEEMless makes about the users and the deployment infrastructure in Appendix A. We also discuss how we envision our system to be integrated on top of current E2E apps to provide a seamless user experience. Here we give a brief summary of the discussion.

Having a smooth user-experience while interacting with the functionality of SEEMless is crucial for successful adoption of this system. We envision a user interface where the client software would periodically run KeyHistory in the background and only notify a user if the verification fails.

In designing SEEMless, we make the following assumptions: 1) a user device can store a cryptographic key and the server has a non-cryptographic means of authenticating

[§]In fact it is slightly less, but we overestimate, which our definition allows for, to make the statement cleaner.

its users 2) the participating parties have approximately synchronized clocks 3) the infrastructure provides a mechanism that ensures all parties have consistent views of the root commitments published by the server.

Finally, even though we present SEEMless as a single *logical* server, it can be implemented using a distributed network of servers. Our system also supports multiple devices per user with keys for each device being stored on the server.

V. aZKS INSTANTIATIONS

In this section we will give a concrete instantiation for the aZKS used for SEEMless and the primitives used in the construction of the aZKS.

A. Cryptographic Primitives

Collision Resistant Hash Function (CRHF): A hash function H is collision resistant if it is hard to find two inputs that hash to the same output; that is, two inputs x and y such that $H(x) = H(y)$, and $x \neq y$. We will in some of our building blocks require a stronger hash function, which we will treat as a random oracle. The random oracle heuristic proves security when the hash function is replaced with a random function. This roughly provides security against attackers who use the hash function as a black box.

INSTANTIATION. For our experiments we use SHA384 in our VRF construction and SHA256 everywhere else.

Simulatable Commitment Scheme (sCS) : A simulatable commitment scheme consists of algorithms $(\text{CS.Commit}, \text{CS.Open}, \text{CS.VerifyOpen})$: (1) $\text{com}_\sigma \leftarrow \text{CS.Commit}(1^\lambda, m; r)$: takes security parameter λ and produces commitment com_σ to message m using randomness r . (2) $\tau \leftarrow \text{CS.Open}(\sigma, m, r, \text{com}_\sigma)$: outputs a decommitment value corresponding to commitment com_σ for message m and randomness r . (3) $1/0 \leftarrow \text{CS.VerifyOpen}(\sigma, \text{com}_\sigma, m, \tau)$: accepts or rejects the decommitment of com_σ to message m in terms of the decommitment value τ .

An sCS satisfies the standard requirements of commitment schemes with respect to hiding and binding: commitments hide the committed message, and it is hard to open a commitment to two different messages. In addition, sCS requires that there exist a hypothetical “simulator” with special powers (either a trapdoor or in our case control of the random oracle) who could form commitments that can later be opened to any value m . This then shows that the commitment and opening give the adversary no additional information.

INSTANTIATION. We can construct this commitment scheme from a hash function which we model as a random oracle: $\text{Commit}(\sigma, m; r)$ samples random string $r \leftarrow \{0, 1\}^\lambda$ and outputs $H(m, r)$.

Simulatable Verifiable Random Function (sVRF) [CL07]: A Verifiable Random Function (VRF) is similar to a pseudo-random function, with the additional property of verifiability: corresponding to each secret key SK , there is a public key PK , such that, for any $y = \text{VRF.Eval}(\text{SK}, x)$, it is possible to verify that y is indeed the value of the VRF seeded by SK evaluated on input x . A simulatable VRF (sVRF) is a VRF for

which this proof can be simulated, so a hypothetical simulator with special powers (e.g. controlling the random oracle) can fake a proof that the value of $\text{VRF.Gen}(\text{SK}, x)$ is any y . This guarantees that the proof does not compromise the pseudorandomness of this or any other output. An sVRF is comprised of the following algorithms: 1) $(\text{PK}, \text{SK}) \leftarrow \text{sVRF.KeyGen}(1^\lambda)$: takes security parameter λ and outputs the public key PK and secret key SK. 2) $y \leftarrow \text{sVRF.Eval}(\text{SK}, x)$: takes the secret key SK and x and outputs the sVRF evaluation of x as y 3) $\pi \leftarrow \text{sVRF.Prove}(\text{SK}, x)$: takes the secret key SK and x and outputs a proof for the sVRF evaluation of x 4) $1/0 \leftarrow \text{sVRF.Verify}(\text{PK}, x, y, \pi)$: verifies the proof.

INSTANTIATION. We describe an efficient constructions of sVRFs based on the DDH assumption as proposed in [MBB⁺15]. The proof for pseudorandomness follows from DDH as observed in [MBB⁺15]. Let \mathbb{G} be a DDH group of order q and let g be a generator. Let $H_1 : \{0, 1\}^* \mapsto \mathbb{G}$ and $H_2 : \{0, 1\}^* \mapsto \mathbb{Z}_q^*$ be two hash functions. The sVRF construction is the following. (1) sVRF.KeyGen: Choose $k \xleftarrow{\$} \mathbb{Z}_q^*$ and output $\text{SK} = k$ and $\text{PK} = g^k$ (2) sVRF.Eval: Output $y = (H_1(x))^{\text{SK}}$ (4) sVRF.Prove: The proof is the Fiat-Shamir transformation of Schnorr protocol for proving common exponent; proving that prover knows $\text{SK} = k$ such that $\text{PK} = g^k$ and $y = h^k$ for $h = H_1(x)$. Prover chooses $r \xleftarrow{\$} \mathbb{Z}_q^*$ and outputs proof as $\pi = (s, t)$ for $s = H_2(x, g^r, h^r)$ and $t = r - sk$. (5) sVRF.Verify: Parse π as (s, t) . Check that $s = H_2(x, g^t(\text{PK})^s, h^t y^s)$ for $h = H_1(x)$

B. Append-Only Strong Accumulator (aSA) Construction

Here, we give a construction of an *append-only* strong accumulator over a data collection of (label,value) pairs. The high level idea is to build a Patricia Trie (Ptr) [Knu98] over the labels. Ptr is a succinct representation of the labels such that each child has a unique suffix string associated with it and the leaf nodes constitute the actual label values. See Fig 6 is in Appendix for an illustrative example. Our aSA construction is built on a Ptr. We use a collision-resistant hash function $H : \{0, 1\}^* \mapsto \{0, 1\}^m$ in our construction.

\triangleright SA.CommitDS($1^\lambda, D$) : Datastore $D = \{(l_1, v_1), \dots, (l_n, v_n)\}$, a collection of label-value pairs. Choose a constant $k_D \xleftarrow{\$} \{0, 1\}^\lambda$. Let $\{y_1, \dots, y_n\}$ be a lexicographic ordering corresponding to $\{l_1, \dots, l_n\}$. Build a Patricia trie on $\{y_i\}$ and output $\text{com} = (h_{\text{root}}, k_D)$. For the different nodes in the tree, hashes are computed as follows.

Leaf nodes: For a node y_i , compute: $h_{y_i} = H(k_D | y_i | v_i)$. For example, in the tree in Figure 6, $h_2 = H(k_D | 0100 | v_2)$.

Interior nodes: For an interior node x , let $x.s_0$ and $x.s_1$ be the labels of its children. Compute: $h_x = H(k_D | x | h_{x.s_0} | h_{x.s_1} | s_0 | s_1)$. For example in Figure 6, $h_6 = H(k_D | 0 | h_1 | h_7 | 010 | 1)$

\triangleright SA.Query(D, l) : If $l \in D$, output value v associated to it. Let h_l be the hash value of the node. Give the sibling path for h_l in the Patricia trie along with the common prefix x at each sibling node and the suffixes s_0, s_1 which form the nodes $x.s_0$

and $x.s_1$. For example, proof for 0100 will be its value and $[h_2, (h_3, 01, 00, 11), (h_1, 0, 010, 1), (h_8, \epsilon, 0, 1)]$ If $l \notin D$, let z be the longest prefix of l such that z is a node in the Patricia tree. Let zu_0, zu_1 be its children. Output $z, h_z, u_0, u_1, h_{zu_0}, h_{zu_1}$ along with sibling path of z . For example, proof for 1010 will be \perp and $[1, 1000, 1100, h_8, h_4, h_5, (h_6, \epsilon, 0, 1)]$.

\triangleright SA.Verify(com, l, v, π) : Parse π as the hash values and the auxiliary prefix information at each node. Compute h_l according to leaf node calculation and verify the given value. Compute the hash values upto the root with help of the proof. Let this hash value be h . Verify that $h = h_{\text{root}}$. In case of a non-membership proof, additional verification is required. Let z_l and z_r be the labels of the left and right child (respectively) of the returned node z . Verify that 1) z is the longest common prefix of z_l and z_r 2) z_l and z_r are distinct.

\triangleright SA.UpdateDS(D, S): First we check that S is a valid set of updates which means that for all $(\text{label}_i, \text{val}_i) \in S$, $\text{label}_i \notin D$. Initialize sets $Z_{\text{new}}, Z_{\text{old}}, Z_{\text{const}}$ to empty. For all $\text{label}_j \in S$, compute: $h_{\text{label}_j} = H(k_D | l_j | \text{val}_j)$ and add h_{label_j} to the appropriate position in the trie and change the appropriate hash values. Add the old hash values to Z_{old} and the updated to Z_{new} and those that remain unchanged after all updates to Z_{const} . Output the final updated root hash h'_{root} as com' , output the final st' and $D' = D \cup \{(\text{label}_j, \text{val}_j)\}$ and output $\pi_S = (Z_{\text{old}}, Z_{\text{new}}, Z_{\text{const}})$ and set S .

\triangleright SA.VerifyUpd($\text{com}, \text{com}', S, \pi_S$): Parse $\pi_S = (Z_{\text{old}}, Z_{\text{new}}, Z_{\text{const}})$. Compute the root hash from all the values in Z_{old} and Z_{const} and check that it equals com . Similarly, compute the root hash from all the values in Z_{new} and Z_{const} and check that it equals com' . Let Z_S be the set of roots of the subtrees formed by labels in S . Check that $Z_{\text{old}}, Z_{\text{new}}$ are exactly the hash values of all nodes in Z_S . Output 1 if all checks go through.

Security This aSA construction is secure assuming H is a collision-resistant hash function. We defer the formal proof to the full version.

C. aZKS Instantiations

Our construction builds upon append-only strong accumulators which directly give us completeness and soundness. To get zero-knowledge we will use an sVRF and an sCS. At a high level, the sVRF helps map an element in the datastructure to a random position and the commitment helps in hiding the exact value that is being accumulated.

Let $(\text{sVRF.KeyGen}, \text{sVRF.Eval}, \text{sVRF.Prove}, \text{sVRF.Verify})$ be a simulatable VRF and let $(\text{CS.Commit}, \text{CS.Open}, \text{CS.VerifyOpen})$ be a simulatable commitment scheme as described before. Given $D = \{(\text{label}_1, \text{val}_1), \dots, (\text{label}_n, \text{val}_n)\}$, ZKS.Commit will first generate an sVRF key pair $\text{SK}, \text{PK} \leftarrow \text{sVRF.KeyGen}(1^\lambda)$, build a new $D' = \{(l_1, v_1), \dots, (l_n, v_n)\}$ where $l_i = \text{sVRF.Eval}(\text{SK}, \text{label}_i)$, $v_i = \text{CS.Commit}((\text{label}_i, \text{val}_i); r_i)$ for random r_i . It will then build an append-only SA on D' and output that together with PK. ZKS.Query will return the

appropriate l_i along with the sVRF proof, the opening to the commitment, and the strong accumulator membership/non-membership proof for (l_i, v_i) . ZKS.UpdateDS will compute the l_i, v_i pairs for the new datastore entries and then run SA.UpdateDS. We defer the formal proof of security to the full version.

Leakage: In our aZKS construction, ZKS.Query leaks the size of the datastore, D and ZKS.UpdateDS leaks the size of the datastore before and after an update.

VI. PERSISTENT PATRICIA TRIE

In our construction, SEEMless, the server maintains two aZKS at every epoch. In Section IV, we described how to construct a aZKS using a aSA, which in turn, we implemented using Patricia Trie. Recall that, in SEEMless, the server needs to store the entire history of the key directory and the two corresponding aSA (used in the aZKS construction) to be able to answer key history queries. Naively, the server can maintain all the aSA from the beginning of time, but this will blow up the server’s storage significantly and hinder scalability of the system. This is particularly wasteful, when in fact, there might be a significant overlap between the aSA of consecutive epochs. To address this problem, we build a persistent data structure that retains information about every epoch, while being space efficient. We call this persistent data structure *Persistent Patricia Trie* (PPTr). We believe this data structure will find a wide number of applications (such as the tamper-evident logging application of [OB09]).

Challenges: The idea is similar to that of history tree (HT) [OB09], built on Sparse Merkle Tree. However, we can not naively use the same technique to build a Persistent Patricia Trie (PPTr). Recall that Patricia Trie (PTr) is compressed for efficiency — there is no empty node in the tree. Compression introduces several subtle challenges in building its persistent version, i.e., PPTr. For example, unlike in HT [OB09], nodes in a Patricia Trie do not have a fixed position. A node with a certain label may be at lower depth at an earlier epoch (root being at depth 0) and fall to a higher depth at a later epoch. A node with a given label may change depth several times from its birth epoch until the latest epoch. Please see Fig VI for an illustrative example. Therefore, a parent to child pointer does not remain fixed throughout the life time of a PTr. In the example, the node 00 pointed to 0000 as its left child at time t_2 , and to node 000 at time t_4 . The HT construction in [OB09] crucially relies on parent to children pointers being fixed throughout the lifetime of the tree. We will build our final construction of PPTr gradually.

Attempt 1: First, let us observe some invariants in a append-only PTr (aPTr): (1) The depth of a PTr is defined with respect to the root node of the tree, therefore the root node is always at depth 0 (2) Any node in a append-only PTr can only fall to higher depths over epochs, it never goes to lower depths. (3) Any node with k bit label can be at depth at most k in a PTr.

Now let us attempt to build a PPTr. Every node in a PTr is identified by its label with a binary string. For every node u , we store some information corresponding to every epoch




Fig. 1. Append only Patricia Trie: final view at epoch t_4

when the hash value at node u changed. Let us denote this information as *nodestate*. These epochs and the corresponding *nodestates* are stored in a hashtable. Corresponding to every epoch t when the hash value at node u changed, *nodestate_t* stores the following information.

- 1) h_t : hash value of node u at epoch t
- 2) *leftskip*: a binary string indicating how many levels have to be skipped to get the left child of u at epoch t . Let the left child label at epoch t is l
- 3) *leftepoch*: the key to lookup l 's hashtable and get the corresponding *nodestate*
- 4) *rightskip*: a binary string indicating how many levels have to be skipped to get the right child of u at epoch t . Let the right child label at epoch t is r
- 5) *rightepoch*: the key to lookup r 's hashtable and get the corresponding *nodestate*

Since we are constructing a aPTr, we need to support only two operations efficiently on the data structure:

Generate proof for a historical query: Recall (non-) membership proof for certain node u at epoch t consists of the siblings of the nodes on the path from the leaf (or an internal node with both childrens labels non-prefixes of label u) to the root. For any previous epoch t , this information can be easily extracted traversing the tree from the root and following at each node u the (*leftskip*, *leftepoch*) and (*rightskip*, *rightepoch*) pointers from hashtable entry of epoch t . Thus, the cost of generating membership and non-membership proofs for any epoch in the past is proportional to the tree height, and is independent of the number of epochs.

Insert new node in aPTr: When a new node gets inserted to a aPT, all nodes on its path from the root gets updated. But note that, each node on this path has a hash table (of (epoch, *nodestate*)) associated with it. For each of these entries, the skip entries might need to get updated while updating the path. Hence, the cost of updating each node on the path is not constant, it is proportional to the size of the hashtable at that node. Consequently, the cost of inserting a new node is not proportional to the tree height.

Attempt 2: Our next attempt is to retain the cost for generating historical proofs while bringing down the cost of insertion.

Notice that The skip entry at the parent of node u gets updated if and only if a new node with label v is inserted to the tree where v is a proper prefix of u . v becomes the new parent of u . We keep a direction flag with every node v that indicates whether u became the left or the right child of the newly inserted node v . Let us denote this flag as *dir*. Along with this, at every node we store the following additional information: *birthepoch*: the epoch at which the label u was inserted.

Now we do not need to keep the skip fields any more, because *dir* will help us get to the correct node. More specifically, for a historical query (u, t) , we will go down from the root following as long as $\text{birthepoch} \geq t$ guided by the *dir* bit. For example, in Figure VI, at t_1 , 0000 was the leftchild of the root but at time t_2 00 replaced 0000 and 0000 became the left child of 00. So the *dir* flag at t_2 is set to left. The *birthepoch* of 00 is t_2 which equals the queried epoch t_2 . Now, if we query $(0000, t_2)$, we will get to node 00 whose node state entry t_2 will have $\text{leftepoch} = t_1$ and *dir* is set to Left. So we go down the leftchild of 00 and find the entry 0000 at epoch t_2 .

Now insertion cost is proportional to the height of the tree since cost of insertion on every node on the path is constant. But we have introduced yet another problem that arises at the time of historical query. The proof needs the sibling path and we may need to go down several levels from a parent node (following *dir*) to get the correct sibling for the queried epoch. *Final construction*: To overcome this problem, we keep at the nodestate of every node, the label and the hashes of its current children. This way, we get the required sibling directly from the parent node. Nodestate now contains the following.

- 1) h_t : hash value of node u at epoch t
- 2) leftlabel: the label of the left child of node u at time t .
Let l denote the label
- 3) lefthash: the hash value of node l at epoch t
- 4) leftepoch: the key to lookup l 's hashtable and get the corresponding nodestate
- 5) rightlabel: the label of the right child of node u at time t . Let r denote the label
- 6) righthash: the hash value of node r at epoch t
- 7) rightepoch: the key to lookup r 's hashtable and get the corresponding nodestate

With this new nodestate, combined with the additional information of *dir*, *birthepoch*, we have a persistent Patricia Trie data structure where both historical query and insertion works in time proportional to the height of the Trie.

VII. EXPERIMENTS

A. Experimental Setup

Our experiments were conducted on a Linux VM running on an 2.30GHz Intel Xeon E5-2673 CPU with 125GB of memory and a 64GB heap allocated to the JVM for the server-side experiments for both SEEMless and CONIKS. All client experiments are done on a 2.8 GHz Intel Core i7 laptop. We used Google's protocol buffers [Var08] to implement communication between the client and the server, due to their efficiency and compression. All sizes given below are computed using the `getSerializedSize()` API call for

protocol buffer messages, which returns the number of bytes used to encode the message.

We implemented the VRF using the *Secp384r1* curve using the Bouncy Castle library, Icart's function [Ica09] for hashing to the curve and SHA384. We use the technique of [FFS⁺13] of hashing the input twice (using $\text{SHA384}(0||\text{input})$ and $\text{SHA384}(1||\text{input})$) and applying Icart on both these hashes, since applying Icart once does not provide a distribution indistinguishable from random.

B. Performance Evaluation

Server updates: The graph in Fig. 2 shows the time taken for server update in the following experiment: we increased the total users by 100k at each epoch and made 1k new user registrations and 1k key updates by existing users. For a server with 10M registered users, it takes 0.28s on average to update the authentication structure (the two PPTrs) and its directory. As expected, the update time is proportional to the log of total number of registrations and updates in the system. The server adds the aZKS commitments in a Merkle Tree (instead of a hashchain) as discussed in Section IV. We estimate this cost using the sparse Merkle Tree data structure used in CONIKS (which is an over-estimate). The time to insert a aZKS commitment in the Merkle Tree of aZKS commitments is $61\mu\text{s}$ after 10M server epochs have passed.

Both CONIKS and SEEMless are implemented such that the VRF for an update or registration are computed between epochs, when a client message arrives. This means, the online update phase only needs to update the authentication data structure and does not incur the VRF cost. For SEEMless, registering a new user requires a single VRF computation (for inserting into the "all" aZKS) and an update requires 2 VRF computations – one for the new key entering the "all" aZKS and the other for the old key entering the "old" aZKS. Therefore the offline VRF computation cost at the server (per client request) is at least 1.3ms and at most 2.5ms , which is relatively low computation between epochs. For CONIKS, the cost is a single VRF computation (for registration, no VRF is computed for update), hence the cost is at most 1.3ms .[¶]

Scalability comparison: We tried to run the same update experiment with CONIKS, on the same machine and saw that the heap filled up by the time we had around 4.5M users. Hence, we decided to run a limited experiment with CONIKS which only measured the time for 1k new users and 1k updates at selected epochs. This illustrates the ability of SEEMless to scale in contrast to the limitation of a CONIKS server to less frequent epochs. CONIKS performs even worse (and hangs) with more frequent epochs due to garbage collection overhead. **Query:** To evaluate the cost of Query, we measure the computation and verification times at the server and the client respectively, and the size of a proof. We vary the number of

[¶]In the current implementation, CONIKS does not cache the VRF, so it has to be re-computed for every update and Query. Keeping a hash map for the VRFs would be a simple optimization. Note that CONIKS did not implement the VRF for their evaluation.

	Prove Query	Verify Query
SEEMless	6.03 <i>ms</i>	10.51 <i>ms</i>
CONIKS	2.01 <i>ms</i>	3.47 <i>ms</i>

TABLE I

SIMULATED TIME TAKEN FOR Query PROOF GENERATION AT THE SERVER AND VERIFICATION AT THE CLIENT IN THE ABOVE EXPERIMENTS AT 10M USERS.

users registered with the server up to 10M. We report the costs averaged over a 100 users with 10 trials for each user.

Time: The time spent in computing and verifying a Query proof has two components: The VRF proof and the authentication path in the authentication structure (which is Patricia Trie in SEEMless and sparse Merkle Tree in CONIKS). The numbers in Table I show the amount of time taken to compute and verify Query proofs including the VRF cost, which is the most expensive operation and dominates the Query cost. We computed the VRF costs as follows. The compute, prove and verify functions for our VRF implementation, each on average takes 1.3, 1.9 and 3.4 *ms* respectively over 5k trials. Since the time for each of these computations is independent of the input, we used them as constants to simulate the VRF costs for the Query experiments (recall that each Query response has 1 VRF proof in CONIKS and 3 in SEEMless).

In our experiments, we saw the cost of authentication path generation and verification is negligible compared to the VRF – it varied between 0.2 – 0.3*ms* for SEEMless and 0.05 – 0.06*ms* for CONIKS¹. Since SEEMless requires 3 VRF proofs, its VRF cost is triple that of CONIKS. However, performance of both systems remains comparable and on the order of 10 milliseconds, which is not a noticeable delay for a user. Also, note that the various authentication paths and VRF computations in SEEMless proof generation and verification in our current implementation are done sequentially. We can gain further speedup by parallelizing these operations.

Size of Proofs: In our implementation, the a proof contains the username and the public key (which are sent in addition to the proofs as part of the Querys) along with a commitment to the public key, its opening, the VRF proof and the authentication path in the Patricia Trie. The average size of the proof is about 8600B when the number of users is 10M. For scale, the length of a Twitter post is famously 140 characters (1,120B) [dev], so, the bandwidth consumed by the proof is less that of reading 8 Twitter posts in plaintext.

The size of the proofs is proportional to the logarithm of the total number of users registered, as expected (Fig 7 in Appendix D). We also measure the size of the CONIKS proofs. Since we have 3 authentication paths per Query proof in our system, whereas CONIKS has one, the size is at most 3× the proof of CONIKS, as expected.

KeyHistory: Here we measure the computation and verification times of KeyHistory query and the size of the proofs. We note that our implementation of KeyHistory functionality

¹Note that the CONIKS Java implementation did not store the commitments to the public keys, instead they stored it in plaintext (which obviously is a crucial privacy omission). This slightly reduces their cost

currently implements some redundant checks (for verification without caching) which we optimized in the protocol description in Section IV.

The cost of KeyHistory is determined by 3 parameters: 1) the number of server side epochs, 2) the latest key version or number of key updates made by the user in question and 3) the total number of updates on the server side (i.e. total number of entries in the history trees). We have already seen how the Query proofs grow with respect to parameter 3, so it is not very instructive to explore this again. We conducted two sets of experiments to measure how the first two parameters affect the cost of the history proofs.

Both our experiments include a set S of 10 ‘special’ users, each of whom not only register, but also update their keys as per the experiments. The set of other users, let us call it the set R of ‘regular’ users, serves as a control: these users register and then update their keys only once. In both experiments, we try to simulate *special* users updating their keys at regular intervals. This means that if there are 200k epochs and the *special* users update their keys 10 times: they register initially and then update every 20k epochs. As a control, the users in R register at a fixed rate. For instance, if $|R| = 1M$, with a total of 200k epochs, 5 users in R register and 5 update their keys at every epoch. We first consider the experiments sans VRF. We do not include the cost of verifying the aZKS commitments in the Merkle Hash Tree (as described in Section IV) in these experiments. We simulate the cost of the Merkle Tree by using the sparse Merkle Tree of CONIKS (which is an over-estimate). The verification time for one leaf is about 49 μ s even for 10M server epochs in this sparse Merkle Tree, which would be close to running the server for several years with half minute epochs. Hence, if a user updated her key on 10 occasions, she would have less than an additional 0.5*ms* of verification time.

Dependence on epochs: To measure dependence of KeyHistory response sizes and verification times on epochs, we fixed the size of R at 1M, the number of updates by the users in S at 10 and varied the number of epochs (adding an equal number of elements in R at each epoch). Note that the number of minutes in a year is 525,600, so 1M is close to the number of epochs which would pass if the length of 1 epoch were half a minute.

In our experiments, the average bandwidth consumed by a user who updates 10 times and checks her history after every update, with a server which has 1M epochs is $\approx 1.09MB$ and not significantly more than at 100k epochs which is $\approx 1.03MB$ (Fig 3). The slight growth in size is due to the increasing number of marker entries to be checked as the number of epochs increases. The verification of the history proofs took the client a total computation time of 0.02s for 1M epochs excluding the cost for the VRF. In contrast, CONIKS requires a user to check her key at each epoch, making the total cost of monitoring her key grow rapidly, as shown in Fig 3. From these experiments, we conclude that of SEEMless is scalable with very short epochs as opposed to CONIKS.




Fig. 2. Mean times to update SEEMless and CONIKS for a new epoch with 1k new users and 1k updates. The x -axis is logarithmic in base 10 and each data point is the mean of 10 trials.




Fig. 3. Mean bandwidth consumed in downloading the history proofs over 10 users (running verification with caching) in S with 10 updates each and 1M users in R as the number of total epochs varies. Note that the growth for SEEMless bandwidth is slow: it is $\approx 1.03\text{MB}$ with 10^5 epochs and $\approx 1.09\text{MB}$ with 10^7 epochs.

Dependence on key updates: The graph in Fig 8 (in Appendix D) shows the total bandwidth consumed by the proofs for various numbers of updates by users (running verification with caching) in S , where the epochs are fixed at $200k$ and $|R| = 1\text{M}$. In the case of our system, even a user who updates her keys almost every week for about five months (20 updates), has to download only an average of 2.01MB to verify her entire key history and the downloaded proof can be verified in an average of less than 0.05s excluding the VRF cost. This includes sending the keys and usernames themselves (at least 1000 bytes per key update), since the user may want to verify the used values. If a user changes her key somewhat less frequently, say every two weeks, then she can monitor the entire history of her keys by downloading about 1.05MB . This does not change significantly as the number of epochs grows, as demonstrated by the previous experiment. On the other hand, a CONIKS user would still need to download a proof every epoch to monitor her key binding. This amounts to almost 576.8MB over the course of as many epochs, hence severely constraining the frequency of epochs.

VRF verification with caching: From the experiments above, we see that the size of the authentication paths in the history proofs are small in our case and so is the time to verify these paths. The dominating cost is the cost of VRF verification, but once a VRF has been verified, it can be cached at the client (without the need to verify the VRF proofs again). More specifically, a user may cache the values taken by the SHA256 hashes of the VRF values at various versions as she verifies them. Figures 4 and 5 show the total computation




Fig. 4. Mean computation time to verify SEEMless history proofs over 10 users in S and 1M users in R and 10 key updates as the number of epochs varies including the time VRF verifications. Note that in CONIKS, the total computation time for a non-caching user depends on the number of epochs and equals over 349s , even at 10^5 epochs.




Fig. 5. Mean computation time to verify SEEMless history proofs over 10 users in S and 1M users in R and 200k epochs as the number of key updates by users in S varies including the time for VRF verifications. In CONIKS, the time would only depend on the number of epochs and equal over 699s for a non-caching user and over 14s with caching at 200k epochs.

times for verifying the history proofs including simulated costs for the VRF verifications with or without caching as well as amortized costs. We defer detailed discussion to Appendix D.

Auditing: The auditors of SEEMless verify 1) the aZKS commitments are correctly added to the hash-chain or Merkle Tree 2) the aZKS are growing in an append-only manner. The cost for checking 1) is $49\mu\text{s}$ and requires downloading about 3KB for a tree on 10M epochs (as already discussed in KeyHistory experiments). To monitor 2), i.e., that a Patricia trie is growing in a append-only manner, its auditors receive parts of the tree which are unchanged and the changed leaves, from which they reconstruct the new version of the tree. We simulate this cost by considering the number of new hashes our aZKS trees have to perform for an update (which we experimentally evaluate). We simulate cost for the following: at the start of the epoch, the SEEMless server has 10M registered users and a total of 100k keys that have been updated by those users. Over the course of the epoch, 1k new registrations and 1k updates are made and the server publishes a new digest and proof. For an auditor, the time to verify the append only property of the updated authentication structure is 0.22s , requires computing under 64k SHA256 hashes and the downloaded proof size is upper bounded by 4.24MB . Our simulated proof size estimates a generous upperbound; in practice, this cost will be much lower. The details of our simulation is in Appendix D.

REFERENCES

- [ANSF16] Muneeb Ali, Jude Nelson, Ryan Shea, and Michael J. Freedman. Blockstack: A global naming and storage system secured by blockchains. In *2016 USENIX Annual Technical Conference (USENIX ATC 16)*, pages 181–194, Denver, CO, 2016. USENIX Association.
- [BCD⁺17] Foteini Baldimtsi, Jan Camenisch, Maria Dubovitskaya, Anna Lysyanskaya, Leonid Reyzin, Kai Samelin, and Sophia Yakoubov. Accumulators with applications to anonymity-preserving revocation. *IACR Cryptology ePrint Archive*, 2017:43, 2017.
- [Bon16] Joseph Bonneau. Ethiks: Using ethereum to audit a coniks key transparency log. In *International Conference on Financial Cryptography and Data Security*, pages 95–105. Springer, 2016.
- [CHKO08] Philippe Camacho, Alejandro Hevia, Marcos Kiwi, and Roberto Opazo. Strong accumulators from collision-resistant hashing. In *International Conference on Information Security*, pages 471–486. Springer, 2008.
- [CHL⁺05] Melissa Chase, Alexander Healy, Anna Lysyanskaya, Tal Malkin, and Leonid Reyzin. Mercurial commitments with applications to zero-knowledge sets. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 422–439. Springer, 2005.
- [CL07] Melissa Chase and Anna Lysyanskaya. Simulatable vrrfs with applications to multi-theorem nzk. In *Annual International Cryptology Conference*, pages 303–322. Springer, 2007.
- [dev] developer.twitter.com. Counting characters. <https://developer.twitter.com/en/docs/basics/counting-characters.html>. Accessed: 2018-12-25.
- [EK15] Mohammad Etemad and Alptekin Kupcu. Efficient key authentication service for secure end-to-end communications. *Cryptology ePrint Archive*, Report 2015/833, 2015. <https://eprint.iacr.org/2015/833>.
- [EMBB17] Saba Eskandarian, Eran Messeri, Joe Bonneau, and Dan Boneh. Certificate transparency with privacy. *arXiv preprint arXiv:1703.02209*, 2017.
- [FFS⁺13] Reza R Farshahi, Pierre-Alain Fouque, Igor Shparlinski, Mehdi Tibouchi, and J Voloch. Indifferentiable deterministic hashing to elliptic and hyperelliptic curves. *Mathematics of Computation*, 82(281):491–512, 2013.
- [Ica09] Thomas Icart. How to hash into elliptic curves. In *Advances in Cryptology-CRYPTO 2009*, pages 303–316. Springer, 2009.
- [Key] Keybase.io. Keybase is now writing to the bitcoin blockchain. https://keybase.io/docs/server_security/merkle_root_in_bitcoin_blockchain. Accessed: 2018-10-05.
- [Knu98] Donald Ervin Knuth. *The art of computer programming: sorting and searching*, volume 3. Pearson Education, 1998.
- [KZG10] Aniket Kate, Gregory M Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 177–194. Springer, 2010.
- [Lis05] Moses Liskov. Updatable zero-knowledge databases. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 174–198. Springer, 2005.
- [LLK13] Ben Laurie, Adam Langley, and Emilia Kasper. Certificate transparency. Technical report, 2013.
- [MBB⁺15] Marcela S Melara, Aaron Blankstein, Joseph Bonneau, Edward W Felten, and Michael J Freedman. Coniks: Bringing key transparency to end users. In *Usenix Security*, pages 383–398, 2015.
- [MRK03] Silvio Micali, Michael Rabin, and Joe Kilian. Zero-knowledge sets. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium*, pages 80–91. IEEE, 2003.
- [Nam] Namecoin. <https://namecoin.org>. Accessed: 2018-10-05.
- [OB09] Alina Oprea and Kevin D Bowers. Authentic time-stamps for archival storage. In *European Symposium on Research in Computer Security*, pages 136–151. Springer, 2009.
- [Pro] LEAP Encryption Access Project. Nicknym. <https://leap.se/en/docs/design/nicknym>. Accessed: 2018-10-05.
- [TD17] Alin Tomescu and Srinivas Devadas. Catena: Efficient non-equivocation via bitcoin. In *IEEE Symp. on Security and Privacy*, 2017.
- [Var08] Kenton Varda. Protocol buffers: Google’s data interchange format. *Google Open Source Blog*, Available at least as early as Jul, 72, 2008.
- [wha] WhatsApp Security Vulnerability. https://www.schneier.com/blog/archives/2017/01/whatsapp_securi.html. Accessed: 2019-01-25.

APPENDIX

A. Discussion

a) *User Experience*: Having an end-user seamlessly interact with the functionality of SEEMless (or any VKD system) is crucial for successful deployment of this system. Designing a smooth user interface that exposes the VKD functionality to end users without drastically changing their experience of using existing messaging software requires usability studies, which we leave for future work. However, here we discuss a blueprint of what we envision the user interface of SEEMless to be.

The high level goal is to have an end-user’s software run the queries and verification the background *in a timely manner* and alert an end-user only when some verification fails. When Alice first registers for the service, the client software installed on her device generates keys in the background, and makes a KeyHistory query to the server to verify that this is the first key that was issued for her, and a Query to verify that the key is stored correctly. This happens in the background and is invisible to Alice unless verification of one of these fails, in which case she will be alerted of possible misbehavior. Similarly, when Alice requests Bob’s key, the client software will run verification of the proof received from the server in the background and will only notify Alice if the verification fails. In addition to that, we want Alice’s client to check her key history *sufficiently often*. This entails running KeyHistory query in the background periodically and notifying Alice if it is not consistent with the updates she has made.** *keyver* must also be run when Alice changes her device or reinstalls the client software (thereby forcing a key update) in which case cases the software would display to Alice a list of the times when her key was updated.

b) *Multiple devices*: We have described the system as storing only a single key per user, but in reality Alice might use her account on multiple devices, and the system would be used to store a list of public keys, one for each of her devices. Bob would then encrypt under all of these keys, and the service would determine which to deliver. This works without any modification of the VKD system. The only change in user experience is that when Alice’s device runs periodic updates, it might find updates made by other devices: the times for these updates should be displayed to Alice for verification.

c) *What we assume from users*: In this work our goal is to make the system usable by average users who cannot be assumed to remember or correctly store long term cryptographic secrets. We do assume that the user’s device can store a secret (the user’s current key), although that device might

**Determining how often the client software needs to run KeyHistory depends depends how quickly users want to be notified of misbehavior and requires user studies which we defer for future work.

be lost/re-imaged, etc in which case the secret would be lost. We assume that the user has some way of authenticating to the service provider (this could be through a second factor text message, or a phone call to answer security questions, etc); this is already a requirement in existing end-to-end encryption systems or other messaging services so we are not increasing the attack surface. Finally, we assume that the user can recognize the approximate times when she updated her key (enough to be able to identify if extra updates have been included in the list). To help with this we could also have the server store a note from the user about each update, e.g. “bought new iphone” or “re-installed app”.

d) Power users: Our system could of course support options for power users to make KeyHistory queries whenever they want, to explicitly compare public keys with their friends (as in WhatsApp), or to sign their key updates.

e) Assumptions on system: We do assume that Alice, Bob, and the server have clocks that are approximately in sync. We also assume some way of ensuring that all parties have consistent views of the current root commitment or at least that they periodically compare these values to make sure that their views have not forked. A simple way of doing this would be for the server to publish the head periodically to a blockchain (as in Catena-CONIKs [TD17]), but we could also implement it with a gossip protocol as in CONIKs [MBB⁺15]. If we implement this by having the server post the commitment on the blockchain, then this means we assume the client software periodically interact with the blockchain. (Because of the hash chain, even if clients do not check the root on the blockchain every epoch, if two different client’s views of the root diverge, they will eventually detect that when they do contact the blockchain.) Similarly, our epoch length need not be limited by the blockchain block time – it is enough if the server periodically posts the current commitment on the blockchain; a misbehaving server will be caught as soon as the next blockchain post happens. This vulnerability window is a tunable parameter. The bandwidth requirements (for interacting with the blockchain) of log auditors of the system and of thin clients (like mobile phones) can be significantly reduced by using an efficiently-verifiable blockchain witnessing service, like Catena [TD17].

f) Distributing the service: We have described the service in terms of a single server, but in practice, the “server” can be implemented using a distributed network of servers, for reliability and redundancy. Our model captures the server as a *single logical entity*, so it can accommodate a distributed implementation fairly easily. Constructing proofs just requires reading the shared data structures for previous time epochs, so that can easily be done by many processes in parallel. Once all of the updates for a given epoch have been collected, then the authentication data structure needs to be updated, but even that can be parallelized fairly easily because of the tree structure. We can support queries even during the data structure update by keeping a snapshot of the last authentication data structure until the update epoch completes. Once the epoch completes, the snapshot can be discarded (so this does not blow up

memory).

B. Verifiable Key Directory

Notation: A function $\nu : \mathbb{Z}^+ \mapsto \mathbb{R}^+$ is a negligible function if for all $c \in \mathbb{Z}^+$, there exists k_0 such that for all $k \geq k_0$, $\nu(k) < k^{-c}$. An algorithm \mathcal{A} is said to have oracle access to machine \mathcal{O} if \mathcal{A} can write an input for \mathcal{O} on a special tape, and tell the oracle to execute on that input and then write its output to the tape. We denote this oracle access by $\mathcal{A}^{\mathcal{O}}$. For the rest of the sections, we will use (label, val) and (username, public key) interchangeably.

Security Properties: Here we give the formal definitions that we informally described in Section II.

- **Completeness:**

For all possible labels, for all update sets S_1, \dots, S_T such that $(\text{label}, \text{val}_i) \in \{S_{t_i}\}_{i=1}^n$ and $\forall t^*$,

$$\begin{aligned} & \Pr[(\text{PK}_{\text{Gen}}, \text{st}_0, \text{Dir}_0) \leftarrow \text{VKD.Gen}(1^\lambda); \\ & \{(\text{pub}_t, \text{st}_t, \text{Dir}_t) \leftarrow \text{VKD.Publish}(\text{Dir}_{t-1}, \text{st}_{t-1}, \\ & S_t)\}_{t=t_1}^{t_n}; \{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}} \leftarrow \\ & \text{VKD.KeyHistory}(\text{st}_{t_n}, \text{Dir}_{t_n}, \text{label}) ; (\pi, \text{val}) \leftarrow \\ & \text{VKD.Query}(\text{st}_{t^*}, \text{Dir}_{t^*}, \text{label}) : \text{VKD.HistoryVer}(\text{com}_t, \\ & \text{label}, \{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}}) = 1 \text{ for } t = t_n \\ & \wedge \text{VKD.QueryVer}(\text{com}_{t^*}, \text{label}, \text{val}, \pi) = 1 \\ & \wedge (t_j \leq t^* < t_{j+1}) \wedge (\text{val} = \text{val}_j) \wedge \\ & \text{VKD.Audit}(\text{pp}, t_1, t_n, \{\text{pub}_t\}_{t=t_1}^{t_n}) = 1] = 1 \end{aligned}$$

- **Soundness:** Formally, we want to capture that for any label label if versions proofs verifies with respect to $\{(\text{val}_i, t_i)\}_{i=1}^n$ and if the audit verifies from t_1 to t_n then at any time t^* between an interval $[t_j, t_{j+1}]$ for some $j \in [n]$, a malicious server cannot give out a proof for label with a value which is inconsistent with the corresponding versions proof at t_j that is, $\text{val} \neq \text{val}_j$ for t_j . Checking this is enough because if the server has given an incorrect key at time t^* , he will have to introduce an additional update sometime later to potentially fix it and will hence be caught with high probability. Hence a malicious server S^* should not be able to come up with a label, $\{(\text{val}_i, t_i)\}_{i=1}^n$ with versions proof Π^{Ver} , commitments and update proofs Π^{Upd} for all times between t_1 to t_n and query proof (π, val) for some t^* for $\text{val} \neq \text{val}_j$. For all PPT S^* , there exists a negligible function $\nu(\cdot)$ such that for all $\lambda \in \mathbb{N}$:

$$\begin{aligned} & \Pr[(\text{label}, \{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}}), \\ & (\text{PK}_{\text{Gen}}, \{\text{pub}_t\}_{t=t_1}^{t_n}, t^*, j, (\pi, \text{val})) \leftarrow S^*(1^\lambda) : \\ & \text{VKD.QueryVer}(\text{com}_{t^*}, \text{label}, \text{val}, \pi) = 1 \wedge \\ & \text{VKD.Audit}(t_1, t_n, \{\text{pub}_t\}_{t=t_1}^{t_n}) = 1 \wedge \\ & \text{VKD.HistoryVer}(\text{com}_t, \text{label}, \\ & \{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}}) = 1 \text{ for } t = t_n \wedge (\text{val} \neq \text{val}_j) \\ & \wedge (t_j \leq t^* < t_{j+1})] \leq \nu(\lambda) \end{aligned}$$

- **\mathcal{L} -Privacy:**

We will say that a VKD is private if there exists a simulator $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3)$ and a leakage function $L(\cdot)$ such that for any PPT client C^* , the outputs of the following two experiments are computationally indistinguishable:

Real:

$$(\text{PK}_{\text{Gen}}, \text{st}_0, \text{Dir}_0) \leftarrow \text{VKD.Gen}(1^\lambda) :$$

$$C^{*O_P, O_\pi, O_{\Pi\text{Ver}}}(\text{st}_t, -)(\text{PK}_{\text{Gen}}) = 1$$

Simulated:

$$(\text{PK}_{\text{Gen}}, \text{st}) \leftarrow \mathcal{S}_1(1^\lambda) :$$

$$C^{*\mathcal{S}_1(\text{st}'_t, L(S), -), \mathcal{S}_2(\text{st}'_t, L(\text{Dir}), -), \mathcal{S}_3(\text{st}'_t, -)}(\text{st}, \text{PK}_{\text{Gen}}) = 1$$

where, O_P is the publish oracle which on input update set S , outputs $(\text{com}_t, \Pi_t^{\text{Upd}})$ as computed by $\text{VKD.Publish}()$. O_π is the proofs oracle which on query a label label , will output $(\pi, \text{val}) \leftarrow \text{VKD.Query}(\text{st}_t, \text{Dir}_t, \text{label})$. $O_{\Pi\text{Ver}}$ is the key history oracle, which on query, label outputs $(\{(\text{val}_i, t_i)\}_{i=1}^n, \Pi^{\text{Ver}}) \leftarrow \text{VKD.KeyHistory}(\text{st}_t, \text{Dir}_t, \text{label})$.

For the simulated game, \mathcal{S}_1 emulates the publish oracle O_P and it gets some leakage on the update set S given by $L(S)$. \mathcal{S}_2 emulates the proofs oracle O_π . It gets the val_i for the label i queried and some leakage $L(\text{Dir})$ and answers the membership queries by producing the proof. \mathcal{S}_3 emulates the key history oracle $O_{\Pi\text{Ver}}$ and outputs proof given $(\text{label}, \{(\text{val}_i, t_i)\}_{i=1}^n)$.

C. Tracing Vulnerability in CONIKS

In [MBB⁺15], we discovered a privacy leakage that can be damaging for targeted user attacks. We call this leakage *tracing vulnerability*. The leakage is the following. In CONIKS, when a user queries for the same label several times, she gets values in the proof which depend on the position of her label in the authentication tree. These values can be positions of other labels and give information about a label that was not queried. Hence if Alice gets the proofs for the same label over time, she can infer about other labels, whether they were updated or deleted. For example, consider a system with 4 users: Alice, Bob, Charlie, Mary with $P(\text{Alice}) = 010$, $P(\text{Bob}) = 011$, $P(\text{Charlie}) = 101$, $P(\text{Mary}) = 110$, $P()$ denotes the position of the label in the tree. The proof for Alice's key will contain 011, being its sibling, which is $P(\text{Bob})$. Since the position is fixed for the entire lifetime of the directory, now Alice can trace when Bob's key changes just by querying for her own key and observing when sibling node changes. While the username is not directly leaked, once Alice queried for Bob's key, she will be able to completely trace when the key changed without ever querying for his key again. This means, even if Bob has deleted Alice from his contact list (and the server exercises access control to not give out Alice's key to Bob after he has removed Alice from his contact list). Alice will still be able to trace when Bob's key changed just by looking at the proof of her own key.




Fig. 6. For the universe of all 4 bit binary strings, a Patricia trie built on subset $P = \{0010, 0100, 0111, 1000, 1100\}$

D. Experiments

Here we discuss the details of caching VRF proofs/labels both at a client and at the server.

Caching VRF labels at the client: Recall, that to verify and key version, i , such that $2^a \leq i < 2^{a+1}$ for some non-negative integer a , her verifications include:

- $i - 1$: She just needs to check for any version that the previous entry is in the “old” aZKS. One additional version is moved to the “old” aZKS for each new update.
- $i + 1, \dots, 2^{a+1} - 1$: This means that unless she updates to a new version i such that $i = 2^k$ for some $k \in \mathbb{Z}$, the VRF values for versions in the set $i + 1, \dots, 2^{a+1} - 1$ have already been verified and are cached and do not need to be verified, just checked for equality. If indeed $i = 2^k$ for some $k \in \mathbb{Z}$, she will have to download the VRF values $i + 1, \dots, 2^{k+1}$ which are 2^k in number. This amortizes her cost to 1 VRF verification per update.
- $2^{a+1}, \dots, 2^{\lceil \log t \rceil}$: This means that only when verifying after some new 2^k th epoch will a user need to verify an additional VRF. For 10 updates at regular intervals, once 1M epochs pass, she will have to verify a total of 20 VRFs, getting an additional amortized cost of 2 VRF verifications, per update.

Once verified, she can just save the 32B Patricia trie labels for the “all” aZKS and the “old” aZKS which, on the client side, results in the following formula

$$\text{client cache} \approx [2i + 2^{\lceil \log i \rceil + 1} + \log t - (\lceil \log i \rceil + 1)] \times 32\text{B}$$

where i is the number of updates by the user and t is the number of server epochs which have passed. In practice, this means that even a user who has 20 updates over 1M epochs only needs to cache 2.8KB. Note that this is much smaller than even a low resolution image take on a flip phone. Figures 4 and 5 show the total computation times for verifying the history proofs including simulated costs for the VRF verifications with or without caching as well as amortized costs. Compare the numbers in figure 4 with a total of over 349s in CONIKS for a user needing to verifying her key every epoch regardless of number of updates when she does not save the VRF locally at just 100k epochs. Alternatively, if she were to store the tree label corresponding to her name, it would require constant monitoring and take a total of over 7s to ensure the server never shows an incorrect key for her.

The numbers for CONIKS would increase significantly when the number of epochs is higher to 3495s when 1M epochs pass if she doesn't save the VRF, and 71s when she does. In the experiment in Figure 5 in Appendix D, the epochs are fixed at 200k so even as the number of updates in CONIKS increases, its computation time would remain fixed at over 699s for verifier without caching and more than 14s for a verifier with caching.

Caching VRF proofs at the server: If the server were to cache the VRF proofs for each user, where it had 10M users and each user had 20 updates, it would correspondingly have to store 28GB of data. However, updates and KeyHistories are relatively infrequent, (such as when a user reinstalls or updates her app), it is reasonable to assume that even about a second of startup time for an app to verify key history is not a barrier to usability. To respond to Query without the need to compute a VRF will require caching the VRF value and proof for the current key of each user. This amounts to $3 \times 48\text{B}$ (the VRF value is 48B) for each registered user, which is 1.44GB over 10M users. Further optimizations can be made, such as saving values for inactive users to disk. We leave such optimizations as future work.

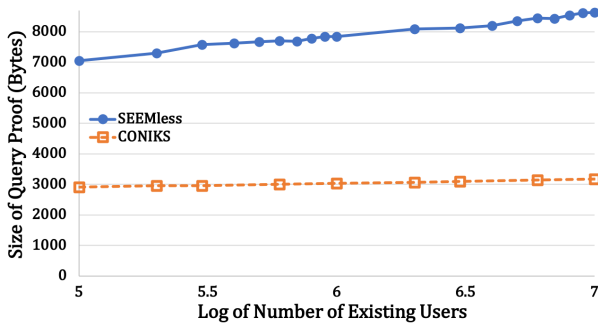


Fig. 7. Mean size of Query proofs for 100 nodes as the number of users in the VKD varies. The x -axis is logarithmic in base 10 and each data point is the mean of 10 trials.




Fig. 8. Mean bandwidth to verify the history proofs over 10 users (running verification with caching) in S and 1M users in R and 200k epochs as the number of key updates by users in S varies and each user checks her key history after each update.

Auditor Cost: To monitor that the Patricia tries are growing in a append-only manner, its auditors receive parts

of the tree which are unchanged and the changed leaves, from which they reconstruct the new version of the tree. We simulate this cost by considering the number of new hashes our aZKS trees have to perform, since our implementation only computes the hashes that changed. For the “all” aZKS, less than 47k new hashes had to be computed and this took 0.17s, and the “old” aZKS computed less than 17k new hashes and it took 0.05s. So the total cost of hashing is 0.22s.

The auditors need to download the respective new leaves and the nodes whose subtrees did not change. For downloading the new leaves, the auditors need to download 1k leaves for the “old” aZKS, 2k leaves for the “all” aZKS, which are 2×256 bits each (leaf label and commitment), totaling 192KB for the two aZKS. To estimate an upper-bound on the number of roots of the unchanged aZKS subtrees, consider the following. Any node which is the root of an unchanged subtree must have a sibling which has changed. If not, then the node as well as its sibling are unchanged, the parent of this node must be unchanged and thus, the parent or some other ancestor must be the root of the unchanged subtree. Therefore, contenders for roots of unchanged subtrees in an aZKS tree are only siblings of nodes along the path from the root to new leaves. Recall that the average depth of a node in a compressed Patricia Trie is $\log n$ (we experimentally confirmed this) where n is the total number of leaves. Hence, the number of unchanged roots of unchanged subtrees is at most $k \times \log n$ where k is the number of new leaves. Corresponding to our experiment in Section VII, this amounts to about $1000 \log 100\text{k}$ unchanged nodes for the “old” aZKS and $2000 \log 10.1\text{M}$ for the “all” aZKS, totaling an upper bound of $1.06\text{MB} + 2.98\text{MB} = 4.04\text{MB}$ (including hash values and labels). Thus, the total proof size that the auditors need to download is less than 4.24MB. Note that, sizes of the labels becomes smaller at each level, but this estimate counts them as 256 bits at every level. This estimate also ignores the double counting of nodes which lie at the intersection of paths to multiple new leaves. So, in practice, this proof size will be much lower.

E. Related Work

Our work broadly falls in the category of building provably secure and efficient privacy-preserving key directory service, particularly relevant in the context of end-to-end secure messaging services. In the recent past, this problem has received significant attention both in the academic community and industry [Bon16], [TD17], [MBB⁺15], [EK15], [Pro], [Key], [ANSF16], [Nam]. This line of work is related to transparency logs [LLK13], [EMBB17], but here we focus on the works that are most relevant to us.

CONIKS [MBB⁺15], a directory service that lets users of an end-to-end encrypted communication system verify that their keys are being correctly reported to all other users. The system is built on ideas similar to *transparency logs* [LLK13] which are public authenticated datastructures for valid SSL/TLS certificates. EthIKS [Bon16] implements CONIKS using the Ethereum for auditing. Catena [TD17]

provides a more generalized infrastructure for managing application specific logs of *append only* statements using the `OP_RETURN` function of the Bitcoin blockchain and implements CONIKS using Catena. Both EthIKS [Bon16] and Catena-CONIKS [TD17] use a global transaction ledger to allow all clients to agree on the same history of the directory. The original CONIKS proposal was to use a different mechanism (gossip) to achieve the same functionality. Apart from that, the core functionality and implementation of CONIKS remains unchanged in [Bon16], [TD17]. But all these works provide weaker privacy than our construction and they lack any rigorous security model. In our SEEMless construction we use append-only strong accumulators (SA) in a way which is somewhat similar to the constructions in [BCD⁺17]. Baldmitsi et al. [BCD⁺17] maintain two accumulators and have an index associated with each element through which they keep track of non membership which is a similar idea to ours. However, they assume that the accumulator manager is trusted, which does not hold for us, since our goal is to detect misbehavior by the server. Considering this stronger setting of an untrusted accumulator adds additional challenges. We added several new ideas to prevent a malicious server from showing arbitrary values on query while still maintaining efficiency. We define and implement a new primitive of *append-only zero-knowledge sets* which generalizes zero-knowledge sets [MRK03], [CHL⁺05] by allowing updates and parameterizing the privacy property with a leakage function. While there have been some attempts to generalize the notion of zero-knowledge sets (e.g., [KZG10], [Lis05]), this is the first attempt that combines both updates and leakage. Our implementation of append-only zero-knowledge sets uses a persistent Patricia Tree data structure which is reminiscent of history trees built on Merkle Trees [OB09]. The construction of [OB09] does not provide any privacy and it is still less efficient than our construction.