MIT Libraries | DSpace@MIT

MIT Open Access Articles

Speranza: Usable, Privacy-friendly Software Signing

The MIT Faculty has made this article openly available. ***Please share*** how this access benefits you. Your story matters.

Citation: Merrill, Kelsey, Newman, Zachary, Torres-Arias, Santiago and Sollins, Karen. 2023. "Speranza: Usable, Privacy-friendly Software Signing."

As Published: https://doi.org/10.1145/3576915.3623200

Publisher: ACM|Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security

Persistent URL: https://hdl.handle.net/1721.1/153143

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution

Massachusetts Institute of Technology

Speranza: Usable, privacy-friendly software signing

Kelsey Merrill
MIT CSAIL
Cambridge, MA, USA
kmerrill@csail.mit.edu

Zachary Newman
Chainguard, Inc.
Brooklyn, NY, USA
research@z.znewman.net

Santiago Torres-Arias
Purdue University
West Lafayette, IN, USA
santiagotorres@purdue.edu

Karen R. Sollins
MIT CSAIL
Cambridge, MA, USA
sollins@csail.mit.edu

ABSTRACT

Software repositories, used for wide-scale open software distribution, are a significant vector for security attacks. Software signing provides authenticity, mitigating many such attacks. Developer-managed signing keys pose usability challenges, but certificate-based systems introduce privacy problems. This work, Speranza, uses certificates to verify software authenticity but still provides anonymity to signers using zero-knowledge *identity co-commitments*.

In Speranza, a signer uses an automated certificate authority (CA) to create a private identity-bound signature and proof of authorization. Verifiers check that a signer was authorized to publish a package without learning the signer's identity. The package repository privately records each package's authorized signers, but publishes only *commitments* to identities in a public map. Then, when issuing certificates, the CA issues the certificate to a distinct commitment to the same identity. The signer then creates a zero-knowledge proof that these are identity co-commitments.

We implemented a proof-of-concept for Speranza. We find that costs to maintainers (signing) and end users (verifying) are small (sub-millisecond), even for a repository with millions of packages. Techniques inspired by recent key transparency systems reduce the bandwidth for serving authorization policies to 2 KiB. Server costs in this system are negligible. Our evaluation finds that Speranza is practical on the scale of the largest software repositories.

We also emphasize practicality and deployability in this project. By building on existing technology and employing relatively simple and well-established cryptographic techniques, Speranza can be deployed for wide-scale use with only a few hundred lines of code and minimal changes to existing infrastructure. Speranza is a practical way to bring privacy and authenticity together for more trustworthy open-source software.

CCS CONCEPTS

• **Security and privacy → Key management**; **Digital signatures**; **Pseudonymity, anonymity and untraceability**; *Security protocols*; *Software and application security*; **Usability in security and privacy**.

This work is licensed under a Creative Commons Attribution International 4.0 License.

CCS '23, November 26–30, 2023, Copenhagen, Denmark
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0050-7/23/11.
https://doi.org/10.1145/3576915.3623200

KEYWORDS

code signing, privacy, key management, usable security

ACM Reference Format:
Kelsey Merrill, Zachary Newman, Santiago Torres-Arias, and Karen R. Sollins. 2023. Speranza: Usable, privacy-friendly software signing. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS '23), November 26–30, 2023, Copenhagen, Denmark.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3576915.3623200

1 INTRODUCTION

Open source software is of vital importance to today's society, as almost all codebases use it [38]. However, the nature of open source software's development makes it especially hard to secure, due to both the sheer numbers of maintainers and their varying nature. Each of these maintainers' accounts, along with the repository itself, is a potential entry point for an attacker.

Attacks on software distribution via repositories can have enormous impact [20, 23]. In such attacks, malicious actors compromise software repository credentials [18], repository infrastructure [16], or distribution channels [13] to publish malware masquerading as packages from maintainers that downstream developers implicitly trust. This has led to tens of thousands of malicious downloads and hundreds of real-world incidents [42], including the United States government in the SolarWinds attack [39], resulting in an Executive Order in May 2021 charging NIST with increasing the integrity of the software supply chain [45].

In response to these sorts of attacks, organizations such as Microsoft and the Center for Internet Security (CIS) recommend software signing as one tool for greater package repository security [10, 36]. In a software signing scheme, maintainers digitally sign their software artifacts to assure users that these artifacts were actually produced by the expected maintainer. However, uptake has been low due to usability concerns [52]. To address the usability issues with classic digital signatures, Newman et al. [40] proposed Sigstore, which allows users to generate digital signatures without managing a key pair themselves, using their email addresses as identities rather than a long-lived key pair.

The privacy issue of exposing email addresses has limited uptake of Sigstore; proposals to add Sigstore signing to RubyGems [26] and Crates.io [2, 28] have both stalled in part due to privacy concerns. npm has also expressed the importance of protecting maintainer personally identifiable information [17]. Effectively protecting the vital resource that is open source software requires a software

signing solution that is easy to use but still protects maintainer privacy.

Speranza provides the same authenticity guarantees as the state of the art while hiding personally identifiable information from the public. The package repository manages a public record of ownership for each package, dictating who is authorized to sign that package. However, rather than storing a public key for each owner (as in a traditional package signing scheme) or the identity of each owner, such as an email address (as in a certificate-based scheme like Sigstore), Speranza stores a *commitment* to the identity of the owner (e.g., a Pedersen commitment [44]). Then, the certificate authority provides a certificate with a commitment to the identity of the requester, rather than the identity itself, as its subject. The package maintainer then publishes a package, a signature over the package, a certificate containing the verification key for the package signature, and a zero-knowledge proof [7] that the subject of the certificate and the package's public ownership record contain *identity co-commitments*—commitments to the same identity.

To summarize, the contributions of this work are:

- *Identity co-commitments*: a privacy-friendly technique for using third-party identity providers in security systems like package signing.
- Speranza, a system for easy-but-private package signing:
  - anonymized certificates with identity co-commitments.
  - techniques for efficiently managing an anonymized authorization record.
- An implemention and evaluation of Speranza:
  - signing and verifying microbenchmarks.
  - a measure of server costs and bandwidth utilization.
  - proof-of-concept implementation in Sigstore
- Analysis of data from PyPI on package ownership changes

Like most privacy systems, Speranza trades off between transparency, privacy, usability, and performance. Because certificates no longer store cleartext identities, or even consistent pseudonyms for cleartext identities, clients can no longer monitor for every time their identity is used. Furthermore, Speranza does not provide complete privacy, and maintainer identity is revealed to the certificate authority and the package repository. Future work (Section 9.3) might address these limitations.

In the remainder of the paper, we begin with background in Section 2 followed in Section 3 with enumeration of the elements of the underlying model, the goals of the attackers, and the goals of our system to mitigate those attacks. With identity co-commitments presented in Section 4, we then discuss signing and verifying with them in Section 5. Section 6 introduces the authorization record that reflects the authorization and policy for delegation to sign. This is followed in Section 7 with a security analysis, in turn followed by our implementation and evaluation in Section 8, further discussion and related work (Section 9) and conclusions (Section 10).

2 BACKGROUND

In this section, we present background on package repository security with an emphasis on digital signing, including schemes that allow signers to avoid managing long-lived key material.

A *package repository* is a service for software distribution, allowing *maintainers* to publish named *packages*, and *end users* to look up packages and download them. Almost all codebases (96 % in a 2022 estimate [54]) contain open-source software, often from package repositories like PyPI [11] and RubyGems [46].

Because package repositories are widely used, they are a valuable target for attackers, who can deliver malware to many end users at once. This work focuses on attack distribution, where an end user requests a specific package and receives malware that the authorized maintainer of the package did not publish. These attacks can happen for many reasons, including the compromise of a maintainer's credentials on the package repository or the compromise of repository itself.

Digital software signing provides the basis for significant mitigation of this problem. A signature over a package allows a user to then verify that the given package was published as-is by a maintainer they trust. However, the introduction of cryptographic keys raises several problems, including key management and revocation in the case of compromised or lost keys [47, 59]. One solution to developer-managed key systems is to create a public key infrastructure (PKI) linking digital identities to cryptographic keys via digital certificates. The most popular deployment of PKI is the web PKI that secures communications between clients and web servers. This allows clients to store a small root-of-trust (hundreds of entries) but communicate with web sites hosted under hundreds of millions of different domains.

The Sigstore project [40] creates an automated "PKI" for software signing. Maintainers run a command to sign their package, a browser window opens asking to authenticate with their email account, and the Sigstore infrastructure automatically issues a certificate for that email, creating a keypair linked to the email address, along with a timestamp. When a user goes to download the artifact, they check that the signature matches the identity they are expecting, and they check the published certificate. They also check that the signature was created while the certificate was valid. Sigstore can support any digital identity whose provider supports OpenID Connect (OIDC) [43], an authentication protocol based on OAuth 2.0, such as a Google, Microsoft, or Facebook account.

Sigstore solves many of the usable security issues with traditional cryptographic key pairs. By outsourcing authentication to an OIDC provider, developers only have to manage one account. They also have account recovery provided by the OIDC provider. However, uptake of the Sigstore system has been limited in part due to privacy complaints from maintainers [2, 17, 26, 28]. In order to use Sigstore, the developer's OIDC identity must be public. For example, if using an email account, the developer's email address must be published. Many developers wish to remain anonymous and find this unacceptable [2, 26, 28]. Providing a system that maintains developer privacy but is still usable would encourage the use of digital signatures on artifacts. This is where Speranza steps in.

Current software signing solutions also have the problem of "key transparency" concerning digital signatures. Classic digital signatures rely on the assumption that end users can securely learn the public verification keys associated with a particular package. In practice, this is nontrivial, because the repository itself is not a reliable source of those keys. If end users simply query the package repository for the key associated with a package, this provides little extra protection on top of the checks a repository performs on package publication.

The Update Framework (TUF) [49] systematically addresses many of these issues by separating different responsibilities into different *roles* with corresponding keys, and then recording the mapping between the roles and packages as a large list the user can download. Keys for important roles can be stored offline, reducing the risk of compromise and allowing recovery from the compromise of online keys. These roles support *thresholds* for greater security. In subsequent work, Diplomat [22] introduces *delegations* and a scheme for locking in delegations such that even the package repository itself cannot reverse them.

In TUF, forking/equivocation/split-view attacks are still possible; a TUF repository can present different views to different end users. Further, an attacker who compromises a repository could "rewind" to undo evidence of their attack. To help detect compromise, transparency logs [24], a technique originating in the web public key infrastructure, require all operations to be posted to a public, tamper-proof log. The log can be "audited" by third party *monitors* which communicate amongst themselves [30, 31]. This in turn allows for the detection of both split-view attacks and general unauthorized activity that might indicate a compromise. Speranza, originally proposed in [33], brings this work together and formalizes it in a state machine model to provide a solution to the key transparency problem.

3 SETTING AND THREAT MODEL

With this background in mind, we consider our setting and model. We first introduce the parties involved in the system, and then consider the nature of the attackers. We then consider our system goals, as well as some key assumptions.

3.1 Parties and Roles

The system has the following parties and roles:

- **Signers**: individuals wishing to sign an artifact
- **Verifiers**: individuals wishing to verify the authenticity of an artifact.
- **Package repository**: independent service hosting artifacts for download. The repository also maintains an *authorization record*, a data structure managed by the package repository that maps artifacts to those identities with the authorization to sign them.
- **Identity provider**: entity vouching that an individual controls an identity.
- **Certificate authority**: entity that verifies identity tokens and issues certificates to signers.
- **Monitors**: third parties auditing the package repository for correctness and consistency.

For example: Alice (signer) wishes to sign artifact `Foo` attesting to its authenticity. She generates a key pair, authenticates with Google (identity provider), and receives a certificate from the certificate authority. She then uses her key pair to sign `Foo`. Bob (verifier) wishes to download `Foo` and check its authenticity. He verifies that the signature on `Foo` is valid and the identity in the certificate is the same as the identity stored under `Foo` in the authorization record.

3.2 Attacker Capabilities and Goals

We assume that an attacker may compromise the maintainer's package repository account. We examine how security in our system degrades under further compromise of the package repository or the certificate authority in Section 7.4. Maintainers' identity provider passwords or cryptographic material belonging to others, such as secret keys, are assumed to be inaccessible to attackers. We also assume a one-time trusted setup phase, perhaps conducted by a standards-setting body, that no attacker may compromise. An attacker may have the following goals:

3.2.1 Get users to run malicious code. An attacker may try to get unsuspecting users to download and run malicious code. They might do this by modifying an existing package to include their code or replacing an existing package with their code. The ability to run arbitrary code on someone else's system enables an attacker to achieve any number of goals including using computing resources, introducing malware into the user's system, or introducing security vulnerabilities for later exploit.

3.2.2 Rollback attacks. An attacker may try to change a package back to a previous version with known security vulnerabilities. This allows the package to largely function as usual and likely remain undetected by the user but still introduces security vulnerabilities that the attacker can use for later exploit.

3.2.3 Violate the privacy of maintainers. An attacker may try to learn the identity of the maintainer of a given package, determine that the same maintainer is publishing multiple given packages, or confirm or deny a guess about the maintainer of a given package. A maintainer's OIDC identity likely includes personally identifiable contact information. An attacker could use this contact information to contact the maintainer themselves to try to influence the maintainer socially, conduct marketing campaigns, or engage in harassment. An attacker could also collect many such identities and attempt to sell this personally identifiable information for advertising purposes.

3.3 System Goals

We have the following goals for Speranza: correctness, authenticity, privacy, transparency, and deployability. We introduce security goals informally here, and we define these more formally in Section 7.1.

3.3.1 Correctness. A signature generated by an honest signer should always be accepted by the verifier.

Correctness means that the system has no "false positives" when signalling malicious behavior. When honest parties are using Speranza, the user should be able to download artifacts successfully.

3.3.2 Authenticity. Verifiers should be able to confirm that an authorized signer (as defined by the authorization record) attested to the validity of the artifact.

Authenticity means that Speranza is secure from the user's point of view. When users download artifacts using Speranza, they are convinced that the artifact is valid according to the rightful maintainer of that artifact.

3.3.3 *Privacy.*

Only the signer, the certificate authority, and the package repository should know the identity of the signer. Verifiers and other signers should not be able to determine anything about a given signer's identity. They should not learn the identity, learn which artifacts are signed by the same identity, or have the ability to confirm or deny guesses about a signer's identity.

Privacy means that Speranza is secure from the maintainer's point of view. It means that maintainer identities will not be made public when publishing artifacts with Speranza.

Note that Speranza aims to provide privacy from *the public*, not *the system*. Under this privacy goal, maintainer identities may be exposed to the package repository and the certificate authority. Also note that this privacy goal allows for maintainer linkages across the *same* package, but not across different packages. A user can see that some maintainer x signed package `Foo` multiple times, but they cannot tell if maintainer x is also signing package `Bar`.

3.3.4 *Transparency.*

All changes to the authorization record should be public and verifiable. Thus, malicious tampering with the authorization record should be detectable.

"Tampering" can be divided into two categories - consistency and correctness. Consistency means that the server cannot equivocate about the authorization record, and all clients see the same record. Correctness is the idea that all mappings in the authorization record are correct and as they should be, i.e., there are no mappings between a package and a non-owner of that package (like an adversary). Our transparency goal requires that both types of tampering be detectable.

Transparency means that we can respond to compromise of the authorization record. If compromise can be detected, users, maintainers, and other parties can then respond to that event, work to recover from compromise, and use alternate methods in the meantime to protect themselves.

This transparency goal excludes the identity provider, the certificate authority, and the repository as a whole from compromise detection. These tradeoffs in transparency are made in order to gain privacy. More discussion of the privacy-transparency trade-off inherent in this system can be found in Section 9.2.

3.3.5 *Deployability.*

This system should be easily usable and deployable with minimal changes to existing infrastructure.

3.4 Cryptographic Assumptions

We assume the existence of a collision-resistant hash function (CRHF) and a commitment scheme with efficient zero-knowledge arguments of commitment equality (see Section 4); we realize such a scheme using the Chaum-Pedersen construction [7] in a group where the discrete logarithm problem is assumed hard. We also require a digital signature scheme. Specifically, we use SHA2-512 as a CRHF, the Ristretto255 [15] group over Curve25519 [4], and Ed25519 signatures.

4 IDENTITY CO-COMMITMENTS

Identity co-commitments incorporate identity into a system while respecting privacy. This section describes identity co-commitments and the cryptographic techniques they use.

4.1 Cryptographic Commitments

Identity co-commitments rely on a cryptographic commitment scheme and zero-knowledge proof of commitment equality. Cryptographic commitments have the following algorithms:

- $\mathsf{Generate}(\lambda) \rightarrow \mathsf{pp}$: generate public parameters.
- $\mathsf{Commit}(\mathsf{pp}, m) \rightarrow (c, r)$: create commitment c to message m with randomness r.
- $\mathsf{Verify}(\mathsf{pp}, m, c, r) \rightarrow$ yes/no: check the commitment.

They satisfy three security properties (formal definitions can be found in the tech report version of this work [34]) :

- Correctness: a commitment to some value x can always be opened to value x.
- Hiding: a commitment cannot be "inverted" to reveal the preimage without the random key needed to "unlock" it.
- Binding: a commitment to some value x cannot be opened to a different value x'.

A proof of commitment equality is a zero-knowledge proof-of-knowledge convincing a verifier that (a) these two commitments are commitments to the same message and (b) the prover knows what that message is using the following algorithms:

- $\mathsf{GenerateEq}(\lambda) \rightarrow \mathsf{pp}$: generate public parameters for proofs of equality.
- $\mathsf{ProveEq}(\mathsf{pp}, m, c_1, r_1, c_2, r_2) \rightarrow \pi$ or $\perp$: prove c_1 and c_2 are commitments to the same identity with keys r_1 and r_2 respectively.
- $\mathsf{VerifyEq}(\mathsf{pp}, c_1, c_2, \pi) \rightarrow$ yes/no: check an equality proof.

They have three security properites:

- Completeness: the verifier will always accept a proof from an honest prover.
- Zero knowledge: the verifier cannot learn anything about the value committed to from the proof.
- Knowledge soundness: in order for the verifier to accept the proof, the prover must know the value committed to.

We construct identity co-commitments using Pedersen commitments and Chaum-Pedersen proofs of commitment equality (constructions detailed in [34]). However, this technique does not actually require Pedersen commitments specifically. It only requires a commitment scheme with the following properties:

- hiding and binding (all commitment schemes)
- efficient zero-knowledge proof of commitment equality

Pedersen commitments have a particularly simple proof of commitment equality, making them a good choice for this scheme.

4.2 Identity Co-Commitments

Identity co-commitments are useful when pseudonyms are desirable, but long-term pseudonyms are not sufficient. Long-term pseudonyms (like verifiable credentials [51]) allow someone to be linked across uses of that pseudonym. Even if the underlying identity is hidden, it is known that one identity is doing all the actions associated with that given pseudonym. Identity co-commitments provide *selectively linkable* pseudonyms. They allow one party to link pairs of signatures under their identity together when and how they wish.

4.2.1 *Parties and roles.*

Identity co-commitments have the following roles:

- **Prover**: an individual holding a given identity that wishes to keep that identity private.
- **Verifier**: an individual wishing to verify a given identity.
- **Authority**: a third party that knows the identity of the prover and keeps a public and private record of this identity.

4.2.2 Methods. Identity co-commitments have the following methods, implemented in terms of Pedersen commitments:

- $\mathsf{CoCo.Generate}(\lambda) \to \mathsf{pp} = (\mathsf{pp}_{\mathsf{pdrsn}}, \mathsf{pp}_{\mathsf{nizk}})$: create public parameters for co-commitments.
  (1) Return $(\mathsf{Pedersen.Generate}(\lambda), \mathsf{Pedersen.GenerateEq}(\lambda))$.
- $\mathsf{CoCo.Commit}(\mathsf{pp}, \mathtt{id}) \to (c, r)$: create a commitment.
  (1) Return $\mathsf{Pedersen.Commit}(\mathsf{pp}_{\mathsf{pdrsn}}, \mathtt{id})$
- $\mathsf{CoCo.Prove}(\mathsf{pp}, \mathtt{id}, c_1, r_1, c_2, r_2) \to \pi$ or $\perp$: prove c_1 and c_2 are co-commitments.
  (1) Return $\mathsf{Pedersen.ProveEq}(\mathsf{pp}_{\mathsf{pdrsn}}, \mathsf{pp}_{\mathsf{nizk}}, \mathtt{id}, c_1, r_1, c_2, r_2)$
- $\mathsf{CoCo.Verify}(\mathsf{pp}, \mathtt{id}, c, r) \to$ yes/no
  (1) Return $\mathsf{Pedersen.Verify}(\mathsf{pp}_{\mathsf{pdrsn}}, \mathtt{id}, c, r)$
- $\mathsf{CoCo.VerifyEq}(\mathsf{pp}, c_1, c_2, \pi) \to$ yes/no
  (1) Return $\mathsf{Pedersen.VerifyEq}(\mathsf{pp}_{\mathsf{pdrsn}}, \mathsf{pp}_{\mathsf{nizk}}, c_1, c_2, \pi)$

We also define FillGraph, which creates a "co-commitment graph": a graph where nodes are commitments to the same identity, and edges are proofs of equality between them.

- $\mathsf{FillGraph}(\mathsf{pp}, G, m) \to H$: for graph structure G and message m, returns an instantiated graph H where nodes are commitments to m computed as $\mathsf{Pedersen.Commit}(\mathsf{pp}_{\mathsf{pdrsn}}, m)$, and edges are equality proofs between their connected nodes computed as $e_{ij} = \mathsf{Pedersen.ProveEq}(\mathsf{pp}, m, c_i, r_i, c_j, r_j)$.

4.2.3 Security properties. Identity co-commitments are a particular deployment of regular cryptographic commitments. Thus, they retain all of the security properties of Pedersen commitments discussed in Section 4.1, and they have the following additional properties:

- **Privacy of Identity**: co-commitments do not leak the underlying identity.
  For all $\mathtt{id}_0, \mathtt{id}_1, G$, where G is an undirected graph structure:
  $$(\mathsf{pp}, \mathsf{FillGraph}(\mathsf{pp}, G, \mathtt{id}_0)) \approx_c (\mathsf{pp}, \mathsf{FillGraph}(\mathsf{pp}, G, \mathtt{id}_1))$$
  (where $\approx_c$ indicates computational indistinguishability).
- **Linkability**: if two commitments are co-commitments, even transitively, then they are commitments to the same identity.
  For all connected graphs H, and all PPT adversaries $\mathcal{A}$:
  $$\Pr\left[\begin{array}{l} \mathsf{pp} \leftarrow \mathsf{CoCo.Generate}(\lambda) \\ (H, \{\mathtt{id}_i\}, \{r_i\}) \leftarrow \mathcal{A}(\mathsf{pp}) \\ \text{such that } |\{\mathtt{id}_i\}| > 1 \text{ and} \\ \forall e_{ij}, \mathsf{CoCo.VerifyEq}(\mathsf{pp}, c_i, c_j, e_{ij}) = \text{yes and} \\ \forall \mathtt{id}_i, c_i, r_i, \mathsf{CoCo.Verify}(\mathsf{pp}, \mathtt{id}_i, c_i, r_i) = \text{yes} \end{array}\right] \leq \mathsf{negl}(\lambda).$$
  This means that no adversary should be able to generate a connected graph of co-commitments such that all of the following hold:
  - The set of nodes contain commitments to at least two distinct identities ($|\{\mathtt{id}_i\}| > 1$)
  - All of the commitment equality proofs on the edges verify
  - All of the commitments on the nodes verify

4.3 Construction

Identity co-commitments require the following setup phase:

(1) Trusted setup: $\mathsf{pp} = \mathsf{CoCo.Generate}(\lambda)$ is securely computed and published.
(2) Authority computes $c_a, r_a = \mathsf{CoCo.Commit}(\mathsf{pp}, \mathtt{id})$. They publish c_a in the public record, and they store $(\mathtt{id}, r_a)$ in the private record.

When a prover wishes to use her identity, they take the following steps:

(1) Prover computes $c_p, r_p = \mathsf{CoCo.Commit}(\mathsf{pp}, \mathtt{id})$.
(2) Prover communicates with the Authority to retrieve r_a. This should involve some form of authentication.
(3) Prover computes $\pi = \mathsf{CoCo.Prove}(\mathsf{pp}, \mathtt{id}, c_a, r_a, c_p, r_p)$.
(4) Prover publishes c_p, π.

When a verifier wishes to check the identity co-commitment, they compute $\mathsf{CoCo.Verify}(\mathsf{pp}, c_a, c_p, \pi)$.

4.4 Co-Commitments Security Analysis

The following provides intuition about why identity co-commitments provide privacy of identity and linkability as defined above. Full proofs are omitted here for space; they can be found in our tech report [34].

4.4.1 Privacy of identity. Intuitively, privacy of identity must hold by the non-interactive zero knowledge property and the hiding property of commitments. A graph is made up of commitments for nodes and proofs for edges. Hiding says that the commitments do not reveal any information about the underlying identity, and NIZK says that the proofs do not either.

The proof is a hybrid argument. Starting with a graph of $\mathtt{id}_0$, the edges can one at a time be switched out for simulated proofs without an adversary being able to distinguish by the NIZK property. Then, the graph can be replaced with commitments to $\mathtt{id}_1$ and the relevant simulated proofs by the hiding property. The process can then go in reverse, with simulated edges being replaced with real proofs by the NIZK property, until we arrive at a graph of $\mathtt{id}_1$.

4.4.2 Linkability. Intuitively, linkability must hold by the knowledge soundness property and the binding property of commitments. If a connected graph contains a commitment to a different identity than the nodes it is connected to, then either the proof edges are not sound, or an adversary was able to open a commitment to something other than the identity it was computed with.

The proof is a reduction to the binding game that also uses the extractor from the knowledge soundness definition. If there exists some adversary that can win the linkability game with non-negligible probability, there must also exist an efficient extractor. The extractor can then be used to win the binding game.

5 SPERANZA PROTOCOLS

This section contains a walkthrough of the three protocols that make up Speranza: package registration, signing, and verifying. The system builds on Sigstore [40], and many of the components are similar or identical to that system. We leave the authorization record as a generic mapping data structure from package to identity.

More details on our proposed implementations for the authorization record follow in Section 6.

5.1 Preliminaries

This section contains APIs for other system components outside of identity co-commitments that we will reference below.

Digital signatures. We use standard asymmetric digital signature algorithms:

- $\mathsf{DigSig.Generate}(\lambda) \to (\mathsf{sk}, \mathsf{pk})$: generate key pair.
- $\mathsf{DigSig.Sign}(\mathsf{sk}, m) \to \sigma$: sign message m with secret key sk.
- $\mathsf{DigSig.Verify}(\mathsf{pk}, m, \sigma) \to$ yes/no: check that σ is a valid signature over message m for public key pk.

Security properties for digital signatures can be found in [34].

X.509 certificates. We use the following algorithms for creating and verifying signing certificates (for simplicity, ignoring certificate details like the public key and validity periods):

- $\mathsf{X509.Generate}(\lambda) \to (\mathsf{sk}_{\mathtt{cert}}, \mathsf{pk}_{\mathtt{cert}})$: generate key pair.
- $\mathsf{X509.SignCert}(\mathsf{sk}_{\mathtt{cert}}, \mathsf{sub}, \mathsf{pk}) \to \mathtt{cert}$: sign sub and pk with $\mathsf{sk}_{\mathtt{cert}}$ to produce $\mathtt{cert}$.
- $\mathsf{X509.VerifyCert}(\mathsf{pk}_{\mathtt{cert}}, \mathtt{cert}) \to \mathsf{sub}$ or $\perp$: verify that $\mathtt{cert}$ is valid according to $\mathsf{pk}_{\mathtt{cert}}$.

OIDC. We use a simplified model of an OpenID Connect identity provider:

- $\mathsf{OIDC.Generate} \to (\mathsf{sk}_{\mathtt{oidc}}, \mathsf{pk}_{\mathtt{oidc}})$: generate key pair.
- $\mathsf{OIDC.Issue}(\mathsf{sk}_{\mathtt{oidc}}, \mathtt{id}) \to \mathtt{tok}$: issue $\mathtt{tok}$ to $\mathtt{id}$.
- $\mathsf{OIDC.Verify}(\mathsf{pk}_{\mathtt{oidc}}, \mathtt{tok}) \to \mathtt{id}$ or $\perp$: verify that $\mathtt{tok}$ is valid for $\mathsf{pk}_{\mathtt{oidc}}$, and return the underlying $\mathtt{id}$ if so.

Authorization record. The authorization record maps from packages to owners:

- $\mathsf{AuthRecord.Register}(\mathsf{pkg}, \mathtt{id})$: register pkg to $\mathtt{id}$ in the authorization record
- $\mathsf{AuthRecord.Lookup}(\mathsf{pkg}) \to c$: look up pkg in the authorization record, and return a commitment c to the identity that owns it.

Certificate authority. We use a simplified model of an OIDC-backed certificate authority:

- $\mathsf{CA.Generate}(\lambda) \to (\mathsf{sk}_{\mathsf{ca}}, \mathsf{pk}_{\mathsf{ca}})$: generate key pair.
- $\mathsf{CA.Issue}(\mathsf{sk}_{\mathsf{ca}}, \mathsf{pk}_{\mathtt{oidc}}, \mathtt{tok}) \to \mathtt{cert}$: Issues a certificate to the the identity corresponding to token. Puts a commitment to that identity in the subject instead of the identity itself. See steps below:
  (1) $\mathtt{id} \leftarrow \mathsf{OIDC.Verify}(\mathsf{pk}_{\mathtt{oidc}}, \mathtt{tok})$
  (2) If $\mathtt{id} = \perp$ abort
  (3) $c, r \leftarrow \mathsf{CoCo.Commit}(\mathtt{id})$
  (4) $\mathtt{cert} \leftarrow \mathsf{X509.Issue}(\mathsf{sk}, c)$
  (5) Return $(\mathtt{cert}, r)$

5.2 Setup

Before any of the Speranza protocols can be run, the trusted setup phase for Pedersen commitments and non-interactive zero knowledge proofs must be run. This means securely computing $\mathsf{pp} = \mathsf{CoCo.Generate}(\lambda)$, then publishing pp.

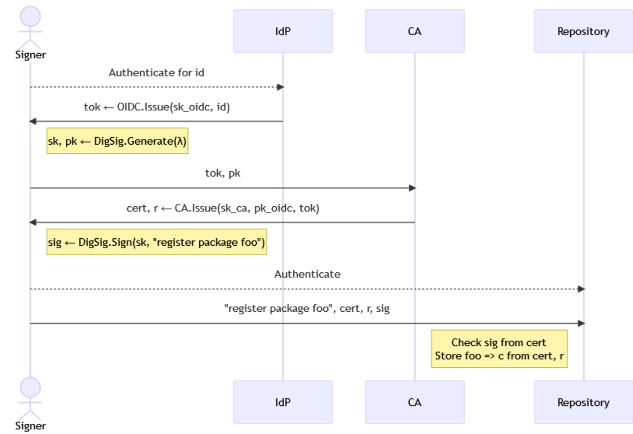

Figure 1: Speranza package registration protocol

5.3 Package Registration

Before a maintainer can sign a given package, they must first register that package with the repository. Registration contains the following steps (see Fig. 1):

(1) Maintainer authenticates with the Identity Provider.
(2) Maintainer receives an OIDC token from the Identity Provider.
(3) Maintainer generates a Digital Signature key pair.
(4) Maintainer requests a certificate from the CA; sends their public key and their OIDC token to the CA.
(5) CA issues a certificate; returns the certificate and the commitment key to the signer.
(6) Maintainer uses their key pair to sign a message requesting that package `foo` be registered.
(7) Maintainer authenticates with the package repository (logs in).
(8) Maintainer sends message, certificate, commitment key, and signature to the package repository.
(9) Package repository checks that the certificate is valid and that the signature over the message is valid. If so, the repository registers package `foo`. The repository privately stores a mapping from `foo` to the commitment key, and it publicly stores a mapping from `foo` to the commitment in the subject of the certificate (see CA.Issue).

5.4 Signing

Once a maintainer wishes to sign an artifact, they run the following protocol (see Fig. 2):

(1) Signer authenticates with the Identity Provider.
(2) Signer receives an OIDC token from the Identity Provider.
(3) Signer generates a Digital Signature key pair.
(4) Signer requests a certificate from the CA; sends their public key and their OIDC token to the CA.
(5) CA issues a certificate; returns the certificate and the commitment key to the signer.
(6) Signer uses their key pair to sign the package.
(7) Signer authenticates with the repository.

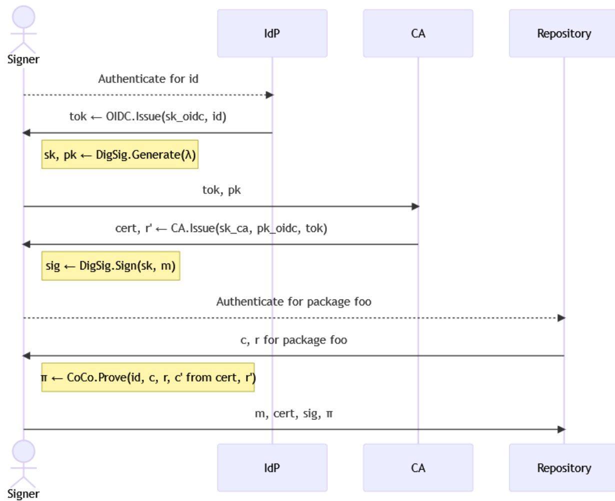

Figure 2: Speranza signing protocol

(8) Signer receives the repository's stored commitment and commitment key for the package.
(9) Signer computes an equality proof for the commitment on the certificate and the commitment returned from the repository (creating an identity co-commitment).
(10) Signer sends package, certificate, signature over the package, and the equality proof to the repository for publishing.
(11) Signer can destroy their ephemeral key pair and commitment keys.

5.5 Verification

(1) Verifier retrieves the package, the certificate, the signature over the package, and the equality proof from the package repository.
(2) Verifier retrieves the commitment associated with the package from the authorization record.
(3) Verifier checks that the certificate is valid; if not, reject.
(4) Verifier checks that the signature over the package is valid; if not, reject.
(5) Verifier checks that the identity co-commitments is valid; if not, reject. If so, accept.

Once all of the relevant pieces have been retrieved from the package repository, verifying can be thought of as an algorithm: $\mathsf{Speranza.Verify}(\mathsf{pk_{ca}}, \mathsf{pp}, \mathsf{pkg}, \sigma, \mathtt{cert}, \pi)$:

(1) $\mathsf{sub} \longleftarrow \mathsf{X509.VerifyCert}(\mathsf{pk_{ca}}, \mathtt{cert})$; If $sub = \bot$, return no
(2) $\mathsf{c_{pub}} \longleftarrow \mathsf{AuthRecord.Lookup}(\mathsf{pkg})$
(3) If not $\mathsf{DigSig.Verify}(\mathtt{cert}.\mathsf{pk}, \mathsf{pkg}, \sigma)$, return no
(4) Return $\mathsf{CoCo.Verify}(\mathsf{pp}, \mathsf{c_{pub}}, \mathsf{sub}, \pi)$

6 AUTHORIZATION RECORD

Up to this point, we have assumed a trusted authorization record. We have also not explained how the server is prevented from misrepresenting the state of authorization to users or equivocating on artifact record, nor have we described how the map can represent more complicated ownership, delegation, and publishing policies than a one-to-one artifact-signer mapping. In this section, we explore different models of trust and methods for preventing the server from acting dishonestly. We also provide methods for implementing complex delegation policies, such as those present in TUF [49], as well as methods for greater efficiency for clients using the authorization record.

For clarity, we first present the authorization record without any concern for privacy. In Section 6.3, we describe how the same techniques from the signature scheme in Section 5 can be used to anonymize it.

6.1 State Machine Model

We begin by introducing a model of an authorization record as a state machine, with allowed and disallowed transitions. Note that there is no notion of "security" in this section. The authorization record is assumed to be held in entirety by all parties.

6.1.1 Basic Authorization Record. An authorization record must support the following algorithms:

- $\mathsf{Initialize}(\mathsf{artifacts}, \mathsf{signers}) \longrightarrow \mathsf{AuthRecord}$: takes in a list of artifacts and a corresponding list of signers and returns an authorization record.
- $\mathsf{Lookup}(\mathsf{artifact}) \longrightarrow \mathsf{Signer}$: takes in an artifact and returns the signer associated with that artifact.
- $\mathsf{Register}(\mathsf{artifact}, \mathsf{signer})$: sets the signer for a given artifact (without authorization).
- $\mathsf{Update}(\mathsf{artifact}, \mathsf{new_signer}, \mathsf{proof})$: takes in an existing artifact and the new signer as well as a proof that the update is authorized (for instance, a digital signature by the previous signer over the new signer). The transition only occurs if the proof verifies.

If there are no concerns for bandwidth efficiency, the authorization record can be backed by a simple hashmap of $(\mathsf{artifact}, \mathsf{signer})$ mappings. However, this requires that clients download the entire authorization record. In order to ensure that state changes the server makes are honest (calls to Update), the client will need to replicate all state changes in their own copy of the authorization record and also verify any calls to Update by checking the proof themselves. This is highly inefficient.

Still lacking from this model is support for more complicated ownership policies. The following sections will discuss support for these policies as well as techniques for greater client efficiency.

6.1.2 Ownership policies. In order to implement more complex ownership policies, we introduce an object, Policy, for describing ownership, delegation, and publishing rules. A policy needs the following algorithms:

- $\mathsf{CheckPublish}(\mathsf{artifact}, \mathsf{proof}) \longrightarrow \mathsf{yes/no}$: checks whether the artifact is authorized, per this policy. The type of proof varies by policy.
- $\mathsf{CheckPolicyChange}(\mathsf{update}, \mathsf{proof}) \longrightarrow \mathsf{yes/no}$: checks whether a requested change to this policy is authorized. The types of update and proof vary by policy.

Algorithms CheckPublish and CheckPolicyChange are both public and computable by the client as well as the server. Transitions in the authorization record state machine are now:

- $\mathsf{Initialize}(\mathsf{artifacts}, \mathsf{policies}) \rightarrow \mathsf{AuthRecord}$
- $\mathsf{Lookup}(\mathsf{artifact}) \rightarrow \mathsf{Policy}$
- $\mathsf{Register}(\mathsf{artifact}, \mathsf{policy})$
- $\mathsf{Update}(\mathsf{artifact}, \mathsf{new_policy}, \mathsf{proof})$

Note that all instances of signer identities have been replaced with policies. A policy contains relevant signer identities.

Remaining goals. With this construction of the authorization record (dictionary mapping $(\mathsf{artifact}, \mathsf{policy})$), the authorization record now supports arbitrary artifact ownership structures and changes in the artifacts themselves as well as in ownership. The remaining issue is client efficiency. In this model, the client has to download the entire authorization record. In the following sections, we explore techniques for making these operations more bandwidth-efficient for the client while also providing authenticity guarantees.

6.2 Compressing the Authorization Record for Users: Third-Party Monitoring

Instead of sending the entire authorization record to users, the repository can use a Merkle tree and send only a *digest* of the authorization record. Specifically, a basic authorization record is a dictionary from artifacts to policies. We can use a Merkle prefix tree (as in [32]) to store the record. The server publishes the entire Merkle tree in the clear, and it also publishes the Merkle digest (constant-sized—typically around 32 B). With this structure, the client only needs to download the constant-sized digest and a logarithmic-sized "lookup proof" (1.5 KiB for 1 million packages) to verify that the server has acted honestly on a lookup call.

We rely on third-party monitors to verify both the correctness and consistency. A monitor can replay *each* change to the tree, verify its correctness according to the state machine model (Section 6.1), and sign the new digest. If a client sees signatures from the monitors for a digest, they know that the monitors have checked the record for that digest.

This model now has a notion of security concerning the consistency of the server's responses to Lookup calls with the true authorization record. This security is acheived from the security properties of a Merkle tree. When the Merkle digest is published, the server is *bound* to a particular version of the authorization record, and all Lookup responses must be consistent with this one version or they will not verify.

6.3 Anonymizing the Authorization Record

So far in this section, we have ignored the privacy concerns motivating our proposed system. All signer privacy is lost if signer identities are published in the clear in the authorization record. This subsection will discuss how the same techniques used in signature scheme described in Section 5 can be used to anonymize the authorization record.

Recall that our current model of the authorization record has the following algorithms:

- $\mathsf{Initialize}(\mathsf{artifacts}, \mathsf{policies}) \rightarrow \mathsf{AuthRecord}$
- $\mathsf{Lookup}(\mathsf{artifact}) \rightarrow (\mathsf{Policy}, \pi)$
- $\mathsf{Register}(\mathsf{artifact}, \mathsf{policy})$
- $\mathsf{Update}(\mathsf{artifact}, \mathsf{new_policy}, \mathsf{proof})$

In this model, signer identities could be exposed within policies. Recall that a policy has the following algorithms:

- $\mathsf{CheckPublish}(\mathsf{artifact}, \mathsf{proof}) \rightarrow$ yes or no:
- $\mathsf{CheckPolicyChange}(\mathsf{update}, \mathsf{proof}) \rightarrow$ yes or no:

Also recall that these are public algorithms, and any state contained in a Policy is also public. This state likely includes a list or tree of signer identities that are authorized to publish a artifact of make changes to its policy. The proof input likely includes signatures providing authenticity with publish or policy change requests.

As in Section 5, any place where a signer identity would be used can instead be replaced with a commitment to that identity to provide privacy. In the same way, a commitment equality proof can be used to prove a match with the commitments stored inside a policy. The signature scheme described in Section 5 can then be used out of the box to provide signatures for the proof input.

Concrete examples of how verifying a publish event or changing a artifact policy might look with an anonymized authorization record can be found in Appendix A.

7 SECURITY ANALYSIS

In this section, we first more formally define the security properties of Speranza introduced in Section 3.3. We then discuss how Speranza meets these security properties. Next, we discuss how these security properties protect against attacker goals in an honest system. Lastly, we describe how the security of Speranza degrades under system compromise.

7.1 Formalizing Security Definitions

Correctness. For all $\mathtt{pkg}, \mathtt{id}$,

$$\Pr\left[\begin{array}{l} \mathsf{pp} \leftarrow \mathsf{CoCo.Generate}(\lambda) \\ \mathsf{pk_{ca}} \leftarrow \mathsf{CA.Generate}(\lambda) \\ \sigma \leftarrow \text{Speranza registration for } \mathtt{pkg}, \mathtt{id} \\ \mathtt{cert}, \pi \leftarrow \text{Speranza signing for } \mathtt{pkg}, \mathtt{id} \\ \text{s.t. } \mathsf{Speranza.Verify}(\mathsf{pk_{ca}}, \mathsf{pp}, \mathtt{pkg}, \sigma, \mathtt{cert}, \pi) = \text{yes} \end{array}\right] = 1.$$

Authenticity. For all $\mathsf{pk_{ca}}, \mathtt{id}, \mathtt{pkg}$, there does not exist PPT adversary $\mathcal{A}$ that can win the following game with probability greater than $\mathsf{negl}(\lambda)$:

(1) $\mathsf{pp} \leftarrow \mathsf{CoCo.Generate}(\lambda)$
(2) Perform package registration for $\mathtt{pkg}$ and $\mathtt{id}$
(3) $\mathcal{A}$ receives pp and oracle access to $\mathsf{OIDC.Issue}(\mathsf{sk_{oidc}}, \mathtt{id}')$ for $\mathtt{id}' \neq \mathtt{id}$, and $\mathsf{CA.Issue}(\mathsf{sk_{ca}}, \mathsf{pk_{oidc}}, \cdot)$.
(4) $\mathcal{A}$ outputs $\sigma, \mathtt{cert}, \pi$.
(5) $\mathcal{A}$ wins if $\mathsf{Speranza.Verify}(\mathsf{pk_{ca}}, \mathsf{pp}, \mathtt{pkg}, \sigma, \mathtt{cert}, \pi) = \text{yes}$.

Privacy (selective linkability). Let G_0 and G_1 be disjoint graph structures. For all G_0, G_1, for all $\mathtt{id}_0, \mathtt{id}_1$ such that $\mathtt{id}_0 \neq \mathtt{id}_1$, and for all PPT adversaries $\mathcal{A}$:

$$\Pr\left[\begin{array}{l} \mathsf{pp_{pdrsn}} \leftarrow \mathsf{CoCo.Generate}(\lambda) \\ \mathsf{pp_{nizk}} \leftarrow \mathsf{CoCo.GenerateEq}(\lambda) \\ H_0 \leftarrow \mathsf{FillGraph}(\mathsf{pp_{pdrsn}}, \mathsf{pp_{nizk}}, G_0, \mathtt{id}_0) \\ b \xleftarrow{\$} \{0,1\} \\ H_1 \leftarrow \mathsf{FillGraph}(\mathsf{pp_{pdrsn}}, \mathsf{pp_{nizk}}, G_1, \mathtt{id}_b) \\ b' \leftarrow \mathcal{A}(\mathsf{pp_{pdrsn}}, \mathsf{pp_{nizk}}, H_0, H_1) \\ b = b' \end{array}\right] \leq \frac{1}{2} + \mathsf{negl}(\lambda).$$

Though not intuitively linked to the Speranza protocol, this definition captures the necessary privacy requirements. As maintainers use the system, they create disjoint graphs of commitments, one graph for each package. This definition states that no matter the structure of those graphs, privacy is maintained.

Transparency. We get our transparency property by applying existing techniques from key transparency to the package repository setting. Formal analysis is left to the prior work that introduced these techniques (see [30]).

7.2 Meeting Security Definitions

The following section provides some intuition as to how Speranza meets our security definitions. More formal arguments have been omitted for space; they can be found in our tech report [34].

Correctness. This follows trivially by inspection from the correctness and completeness properties of identity co-commitments.

Authenticity. In order for an adversary to break the authenticity property, they would have to produce a Speranza signature for an invalid identity that still passes all of the checks in Speranza.Verify. In order for all of the checks to pass, the adversary must violate identity co-commitments' linkability property: they have linked two unequal identities in one co-commitment. Thus, authenticity must hold.

Selective linkability. Selective linkability reduces to the privacy of identity property of identity co-commitments. The only difference between the privacy of identity property and this selective linkability property is that selective linkability allows for two different graph structures. The graph structure does not reveal any identity information, so selective linkability holds.

Transparency. The authorization record is simply an authenticated dictionary. Transparency is the ability to detect tampering with this authenticated dictionary. This is exactly the security property of an authenticated dictionary - they will not verify if there has been tampering relative to the last published digest.

Clients require that monitors sign off on the authorization record digests. The monitors will only sign a digest after verifying *all* updates to the authorization record. Further, clients can request the history of their own packages, or any packages they depend on, and monitor those specifically by storing those policies locally. Thus, no adversary can tamper with the authorization record without detection.

7.3 Subverting Attacker Goals in an Honest System

Here, we discuss attacks an attacker might attempt in pursuit of the motivations described in Section 3.2 and the ways Speranza prevents these attacks. For this section, we assume all the system components are honest and uncompromised, and an outside attacker is attempting these attacks. In Section 7.4, we discuss the security of the system with different compromised components.

7.3.1 Violate the maintainer/package mapping. An attacker may attempt to violate the maintainer/package mapping in pursuit of running malicious code on user machines or conducting rollback attacks. This looks like signing a package on behalf of an identity the attacker does not hold. The authenticity property states that there does not exist an adversary that can publish under an identity they do not hold, so it protects against this kind of attack.

7.3.2 Violate Speranza privacy guarantees. The selective linkability security property provides protection from this sort of attack. Selective linkability states that no matter how (poly-)many identity co-commitments are generated in whatever graph structure, parties that only have access to published graphs cannot determine what identity they belong to.

The Speranza signing protocol exposes only this graph of identity co-commitments and a signature over the package to the public. Selective linkability means that the graph tells the public nothing. Digital signatures are generated completely independently from identities, so it is impossible for them to reveal any identity information, and adding them onto the graph structure provides no help to an adversary. Thus, outside adversaries cannot violate maintainer privacy.

7.4 Security Under System Compromise

The following section describes how security in this system degrades if the attacker is able to compromise elements of the system itself: the maintainer's identity provider login, an identity provider, the certificate authority, or the repository.

7.4.1 Maintainer's identity provider login/identity provider. As described in Section 3.2, this attack is out of scope, and we provide no security guarantees against an attacker that is able to compromise a maintainer's identity provider login or the identity provider itself. This is equivalent to an attacker recovering the secret key in a standard cryptographic system. We recommend implementing threshold policies for publishing in order to provide safeguards against this kind of compromise.

7.4.2 Certificate authority. If an attacker is able to compromise a certificate authority, Speranza's privacy guarantees fall. The attacker can see requests for certificates from maintainers and thus learn maintainer's identities. They can then learn which maintainer publishes which package by watching which certificate gets published with which package.

If the attacker has not also compromised the maintainer's package repository account, then they are still prohibited from publishing: the maintainer must authenticate with the package repository in order to publish. However, if the attacker has compromised this account *and* the CA, and the attacker knows the identity of the maintainer, the attacker will be able to publish as that maintainer undetected by generating certificates to that identity and then authenticating with the package repository.

Note that there is weaker security in Speranza under certificate authority compromise than in Sigstore because the privacy properties of Speranza have negated the transparency properties present in the Sigstore system. In the Sigstore system, certificate authority compromise can be detected through the transparency log. However, in our system, even if all issued certificates are recorded in a tamper-proof log, there is no way to tell if those are valid or not because the subjects are all opaque. Future work is needed to provide better security under certificate authority compromise.

7.4.3 Repository. If an attacker is able to compromise the repository, Speranza's privacy guarantees get weaker. Since the repository maintains a mapping of package to commitment key, an attacker can now "guess and check" maintainers' identities; they can try to open the commitments in the authorization record using their guess identity and the stored commitment key. However, the attacker does not learn the maintainer's identity outright.

Even in the event of a compromised repository, an attacker cannot publish on behalf of another maintainer. In order to publish, the attacker would need a certificate from the CA with that maintainer's identity. This would require authenticating with an identity provider on behalf of that maintainer, which is outside of the attacker's capabilities.

Repository compromise can be detected and dealt with if the attacker attempts to alter the authorization record due to the transparency property. However, there is no way for privacy violations by compromised repositories to be detected in Speranza as long as the repository continues behaving as normal. Further work is needed to address this problem.

8 IMPLEMENTATION AND EVALUATION

We implemented signing and verifying for the software signing application[1]. The implementation was done in Rust version 1.63.0. SHA2-512 was used as the hash function for all instances of hashes in the system. For implementing Pedersen commitments, we used the Ristretto group [15] over Curve25519 [4]. Ristretto is a variant of Decaf [14] that allows for a prime-order group from an elliptic curve group that actually has cofactors up to 8. The implementation totals to 3700 lines of Rust code, including benchmarking harnesses.

The maintainer map was implemented with a Merkle binary prefix tree, similar to the implementation described in CONIKS [32]. This is consistent with the trust model and the maintainer map described in Section 6.2. All measurements were performed on a desktop computer running a 3.9 GHz AMD Ryzen 5 5600G processor with 32 GiB RAM running Linux. We report the median over at least 10 trials unless otherwise mentioned.

8.1 Signing and Verification Costs

Operation	Create (µs)	Verify (µs)
Ed25519 signature	15	40
Pedersen commitment	82	82
Pedersen equality proof	188	272
Co-commitment	307	362

Table 1: Microbenchmarks for cryptographic operations.

We find that, for a maintainer to create a signature using identity co-commitments for a repository with 3.2 million packages requires 404 µs; verification requires 372 µs. These numbers are consistent for repositories of varying sizes: the primary cost that changes with repository size is verifying the Merkle proof for looking up the authorization; this is small relative to the cost of verifying identity

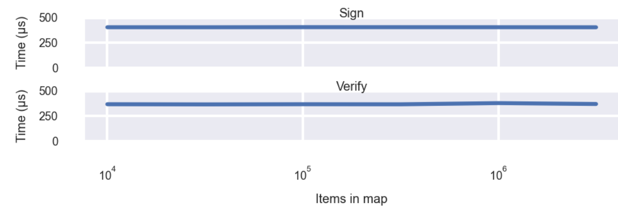

Figure 3: End-to-end costs for maintainers during package publishing (sign) and users during package download (verify).

co-commitments, and scales logarithmically with the size of the repository. We measured, for repositories of varying size, the end-to-end signing and verification costs for maintainers and users when using the Merkle BPT-based authorization record (Fig. 3).

We report timing for the cryptographic primitives and identity co-commitments used in Speranza in Table 1.

8.2 Repository Server Costs

The primary costs to the repository server lie in maintaining the full authorization record. In Section 6.2, we described two methods: first, a simple method in which users must download the complete authorization record ("basic dictionary") and a method backed by a Merkle binary prefix trie (BPT). We report the costs for varying operations with repositories with up to 3.2 million packages in Fig. 4. First, we report the cost of initializing the authorization record. This is a one-time cost that must be done when importing existing maintainership records, regardless of the use of identity co-commitments. Second, we have the cost to register a new package, or to update an existing package ("insert package"). Third, we have the cost to create a proof that a given policy is accurate. This can also be done for each package ahead of time, so that these proofs can be served via a CDN or mirror. This does not apply to the basic authorization record. Finally, we have the cost to users to verify a policy. For the basic method, this means looking up the policy in the dictionary they already store; for the Merkle BPT method, this means verifying the Merkle proof.

The other primary cost—and the main motivation for the more-complex Merkle BPT method—is the bandwidth used. We report the bandwidth required and storage costs in Fig. 5. The bandwidth for an initial download for the Merkle BPT method is 64 B—the size of a Merkle root, which does not vary with repository size. For the basic dictionary method, this is the size of the complete authorization record—which can be 300 MiB for a repository of 10 million packages (assuming about 32 B per entry—in practice, could be double or triple the size). The Merkle BPT method does require "lookup proofs" each time a user verifies a signature, to check the policy against the authorization record. These top out around 1.6 KiB per package for a repository of 10 million packages. Finally, we note the required storage costs for the repository server to maintain these records. To maintain the Merkle BPT requires store the entire authorization record in the clear, plus some additional Merkle BPT data (sublinear in the size of the record).

[1] https://github.com/znewman01/speranza

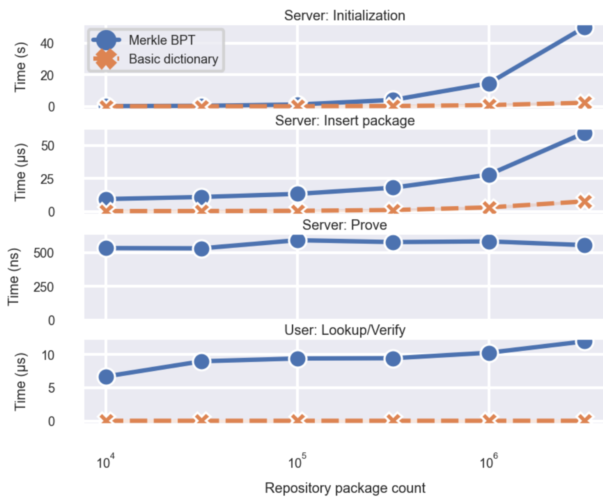

Figure 4: Cost of maintaining the authorization record, using both a Merkle BPT and sending the full authorization records to clients ("basic dictionary").

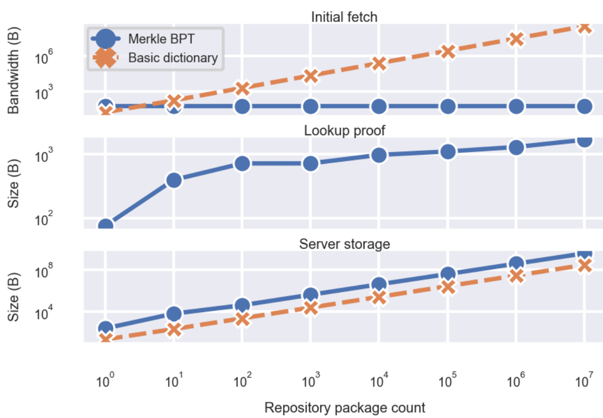

Figure 5: Bandwidth and storage costs in Speranza.

8.3 Proof of Concept: Identity Co-Commitments

While our benchmarks cover the operations for signers, verifiers, and the repository server, we wanted to understand the difficulty of modifying a certificate authority to support identity co-commitments. We added these identity co-commitments to the Fulcio certificate authority in the Sigstore [40] project. Our patch[2] added one dependency, on an implement of Ristretto, and implemented Pedersen commitments and proofs of equality in 115 lines of Go code; to replace email addresses with commitments then required changing 5 lines of code. While a production implementation would require more testing, the changes are overall modest.

[2] https://github.com/znewman01/fulcio/pull/1

9 DISCUSSION AND FUTURE WORK

Our evaluation shows that the costs to repositories, signers, and maintainers are reasonable for repositories of practical sizes (millions of packages and maintainers).

Empirical data about the scale of package repositories allows us to predict how our system will perform in a real-world setting. Table 2 shows the total number of packages for popular package repositories, fetched from the indicated sources on May 4, 2023. In our evaluation, all experiments used a repository size exceeding that of npm, the largest reported repository. We found that, aside from a one-time setup cost (under one minute), the cryptographic server operations were all sub-millisecond—less than a database write would take, and therefore negligible compared to the costs repositories already require in upkeep. The bandwidth requirements for signatures were about 1.6 KiB—far smaller than a typical software package. The cost to a maintainer to sign a package is sub-millisecond—far less than the network roundtrip to the repository to publish a package. Verification times are similar. Even for a project depending on *every* package in a repository like PyPI, verification would take only a couple of minutes, much quicker than downloading these dependencies.

Repository	Packages	Source
npm	"2 million"	npmjs.org
PyPI	451,913	pypi.org
RubyGems	176,365	rubygems.org
Arch User Repository	84,624	aur.archlinux.org
Ubuntu 23.04	35,587	repology.org
Hackage	15,903	hackage.haskell.org
Arch Linux (official)	13,889	archlinux.org

Table 2: Scale of real-world repositories

We obtained data on package registrations and all changes to package ownership from PyPI. The PyPI security team requests that we not publish this data, but encourages interested researchers to reach out to `security@pypi.org`. The rate of changes in ownership affects operations related to maintaining the authorization record, and this data can also help predict future repository growth. 2022 saw about 112,000 total changes to the ownership of packages. Of those, about 103,000 corresponded to new packages registered, leaving about 9000 changes to the ownership of existing packages. This scale (12 role changes per hour) opens the door to more-expensive authenticated data structures backing the authorization record to support client verifiability (see Section 9.2).

9.1 Related Work

Here, we discuss work on software signing and key management.

Packaging signing with long-lived keys. In traditional package signing systems, maintainers use PGP/GPG [12, 55] or similar keys to sign software packages. These systems do respect maintainer privacy, as the public keys are opaque and not linked to their identity. However, key management can be challenging for maintainers [47, 52, 59], and usability concessions like account recovery in the

case of lost keys can reintroduce the same security issues that signatures aim to solve.

Public key infrastructure for software signing. Using public key infrastructure (PKI) for software signing improves usability for both signers and verifiers. Verifiers no longer need to manage a public key for each signer. Instead, they just need to maintain a smaller root-of-trust, which they can use to verify abilities. A certificate authority (CA) issues certificates to signers after a registration authority (RA) checks their identity.

Authenticode [19] is a PKI system for software signing used on Microsoft Windows. Though many commercial software vendors use Authenticode to sign their software, code signing certificates for Authenticode come from a small, trusted list of CAs. These CAs typically charge for certificates, limiting their use among hobbyists and open source maintainers.

Sigstore [40] is a PKI system for software signing featuring an automated CA, inspired by the ACME protocol [3] used by Let's Encrypt [1]. In Sigstore, certificates are linked to identities like a user account with an OpenID Connect [43] identity provider. Sigstore can also issue certificates to machine identities: for instance, a particular build job on a continuous integration/continuous deployment (CI/CD) system. These certificates are free and easy to obtain. Further, because issuance is automated, signers can use a new key pair for each artifact they sign, completely obviating the need for key management.

While these PKI systems handle identity, they do not prescribe which identities should be trusted for particular software artifacts (nor, for that matter, do digital signature schemes without PKI); that requires a package repository policy.

Privacy-friendly credentials. In Verifiable Credentials [51], an issuer issues credentials to holders, who *present* them to verifiers. Instead of showing the credentials directly to a verifier, a presentation can use zero-knowledge proofs of some predicate on the credentials to preserve the privacy of holders. However, the verifier must still know *what* predicate they're verifying: holders would still need to reveal their identity to prove authorization without something like identity co-commitments.

OACerts [27] use Pedersen commitments inside of certificates, and allow holders to prove predicates over their contents. OACerts primarily supports predicates over numerical values, including support equality checks. However, there's no notion of co-commitments—the predicates are assumed to be publicly known.

Package repository policy. A package repository policy specifies, for each package, who must sign a particular software artifact. Centralized package repositories, like the Debian or Red Hat repositories, have a small number of trusted keys used to sign all of their packages. A repository may want to support revocation in case of compromise, support delegation for specific packages to specific maintainers, and address a number of subtle attacks. Naive solutions fall into a trap where verifiers download software from a source, then immediately ask the same source how to verify the software. The Update Framework (TUF) [21, 22, 49] handles these issues, with different roles in the system with different privileges.

TUF presumes a binary model of validation: an artifact is "good" or "bad." The in-toto project [56] allows more nuanced policies, where a package might need a signature from a a maintainer over its source, *and* a signature from a trusted build service tying the ultimate artifact back to that source. Both TUF and in-toto use long-lived keys for package repository security, though enhancement proposals to both [29, 53] propose integrating Sigstore.

CHAINIAC [41] implements a software update policy with an emphasis on build reproducibility and with support for updates to authorization keys. Much of the design of CHAINIAC aims for decentralization, a non-issue for existing package repositories, leading to higher costs, and requires user-managed keys.

Pseudonyms. If maintainers are concerned about privacy, can we identify maintainers by pseudonyms? A privacy-sensitive maintainer could register a new account for their packaging activities. In addition to the extra hassle, however, it represents a security risk: each new account must be managed, leading to security shortcuts (like disabling multi-factor authentication, or reused passwords). Further, this approach does not help in cases (such as harassment) where a maintainer *later* decides they want to obscure their identity.

Parties in this system could provide automated pseudonyms. The OpenID Connect protocol supports pairwise pseudonymous identifiers (PPIDs) [9]. An identity provider, rather than issuing tokens with the account of a user as the subject, instead picks a distinct pseudonym for each *audience*. For example, the provider could use $\mathsf{H}(\mathsf{sk}, \mathsf{user}, \mathsf{aud})$, where $\mathsf{H}(\cdot)$ is a collision-resistant hash function, user is a username, like `user@example.com`, aud is the audience, such as `sigstore.dev`, and sk is a secret, provider-wide salt (preventing brute-force guessing of the user name). Automated pseudonyms require support from upstream identity providers; few major providers implement them.

If the identity provider does not support PPIDs, the certificate authority could introduce automated pseudonyms themselves. Creating pseudonyms from a verifiable random function (VRF) [35], as in CONIKS [32], allows someone knowing the cleartext identity, a shared "verifying key" for the CA, and a non-interactive proof to verify pseudonym correctness. However, using a VRF in this manner requires revealing the cleartext identity to any user wishing to verify it (in this setting, just the package repository); identity co-commitments are verifiable *without* knowing the identity. As in Speranza, this limits the ability to monitor the certificate authority for specific identities, though the VRF holder *can* learn their identity's VRF output and scan for it. Further, the fact that a VRF is deterministic for a given key means that a given identity has only one pseudonym for the given ecosystem, allowing correlation across services; multiple VRFs could be used at the cost of extra complexity. A given ecosystem would be tied to a single VRF key, requiring the user to sign with the same CA each time.

Transparency systems. In a transparency system, a centralized party maintains a tamper-proof, public log. These systems are appropriate in "trust-but-verify" settings: unlike in a blockchain, there is a centralized party that can modify the log at will. This centralization avoids the need for expensive consensus mechanisms. Despite centralization, these systems are *accountable*: if the centralized party misbehaves, it can be detected and publicized. These systems use cryptographic techniques to ensure that no data is ever removed from the log and verify.

The transparency system with widespread deployment was certificate transparency [24], which logs all certificates issued in web PKI. Several projects add transparency for software (Go's sumdb [8], Firefox's binary transparency [37]) or signatures (Sigsum [50]). Sigstore [40] features transparency logs for both certificates issued and signatures published.

Key transparency. Several recent works aim to provide a map from identities to public keys (often in support of end-to-end user chat applications), using transparency techniques for accountability. CONIKS [32] uses a verifiable random function [35] to anonymize user identities along with a Merkle prefix tree to map identities to keys. Third-party monitors verify the *consistency* of the data structure and prevent equivocation, but do not check for the correctness of changes. Instead, each user monitors their own key history. The VRF used for privacy in CONIKS inspired identity co-commitments.

Since CONIKS, a number of key transparency systems with interesting properties have emerged. Gossamer [5] manages package ownership on a transparency log, requiring auditors and clients to walk the entire log for updates. Google's Key Transparency [48], since abandoned, is an implementation of key transparency closely following CONIKS. SEEMless [6] improves the performance and scaling characteristics of CONIKS. Verdict [58] proposes a new implementation of a transparency dictionary, which can be used to construct a key transparency system. Parakeet [30] improves on the consistency mechanism required for key transparency and achieves sufficient performance to support billions of users; an implementation powers key transparency for WhatsApp [25].

While the techniques from key transparency systems are useful for package ownership management, there are a few key differences. First, the scale of key transparency systems (Parakeet handles billions of users) is far greater than required for package repository. Second, the correctness of updates to a package's authorization policy can be verified by third parties, while key updates cannot (to allow for lost keys). Third, the privacy considerations are different: package names and package authorization policies are both are public (it is just the identities which are hidden) which enables any client to audit any package. Together, this allows Speranza auditors to verify the correctness of every update to every package, which would be prohibitively expensive for key transparency.

Further, while key transparency systems do support mechanisms for privacy, these mechanisms apply to the *labels* of the directory (usernames or phone numbers). In the package repository setting, these *labels* are public, but the cryptographic identities themselves must be hidden. The techniques for privacy in the key transparency setting do not apply out-of-the box to private signing with certificates, which motivates our identity co-commitments.

9.2 Auditing the Authorization Record

Signatures, even those checked against an authorization record are only useful if that authorization record is correct: a compromised repository could serve a bad authorization record to a user.

Key transparency systems [6, 30, 32] have identified three potential mitigations for such attacks on a different type of directory, which maps identities to public keys. First, a *consistency protocol* ensures that the server cannot equivocate and serve different views of the directory to different clients: clients require signatures from a quorum of distributed witnesses.

Second, updates to the record should be publicly auditable, so that third party *monitors* will notice any attempts to tamper with the history of the record. Third, identity owners should be able to efficiently check their own keys for unauthorized modifications (*correctness*).

Speranza achieves transparency through similar mechanisms, though with slight modifications. The consistency protocol applies exactly as before, and we recommend the use of the protocol used in Parakeet [30]. However, in key transparency systems, the monitors perform only checks of *consistency*, not of the *correctness* of updates. Key transparency systems manage billions of keys, and checking each individual update is infeasible. Except in CONIKS's optional "paranoid" mode [32] which disables replacing lost keys, correctness can *only* be verified by end-users: replacing a lost key looks like unauthorized tampering. In Speranza, it is practical for monitors to check each update (about 12 per hour; see Section 9). Individuals *can* audit the history for packages, but since package maintainers typically only interact with a repository when they publish a release, global correctness monitoring avoids the requirement for each package maintainer to periodically come online. Further, because the *labels* (package names) are cleartext, verifiers can check correctness for packages they would like to use by requesting the history from the repository and verifying the history's consistency.

Client verifiability. In Speranza, monitors audit both the correctness of updates to the authorization map and the consistency of the map itself. That is, the package repository should not be able to tamper with the history of any package, nor should it be able to make changes to a packages ownership inconsistent with its current policy. In the version of Speranza with complete authorization records, clients rely on monitors for consistency. Maintainers and verifiers can, if desired, check the correctness of individual packages they care about by storing the current state for those packages locally.

We note that at the scale of package repositories (using data from PyPI as a test case), *global* correctness checking is possible. First, if clients download the full authorization map, they can request a list of updates to the map and verify each one (at a rate of only about 100,000 per year). To save on bandwidth for the initial download, along with storage costs for the client, we could use *transparency dictionaries* [57, 58], which are exactly the required data structure: each package has its own entry, and the *history* of updates to each package are append-only. Unfortunately, current transparency dictionaries have somewhat higher latency, especially for updating, than would be appropriate for a production deployment. We note the setting where relatively few *updates* to package ownership occur relative to the total number of packages admits an interesting optimization, which we explore in the tech report version of this work [34].

9.3 Future Work

We note the following directions for future work. As discussed in Section 7.4, compromise of either of the certificate authority or repository can cause Speranza's privacy guarantees to fail. Furthermore, detection of this compromise is difficult, and the certificate

authority can issue malicious certificates without detection under Speranza, because identities are no longer cleartext. More work in this space is needed to protect against this kind of event. Redaction mechanisms for transparency logs may help: if information is public by default, but redactable on request, it is transparent (but not private) in real-time but supports privacy later. Further zero-knowledge proofs could support monitoring use cases.

We noted in Section 9.2 that clients could individually audit for the correctness of updates to the policy for a particular package and propose an example architecture supporting this in the tech report version of this work [34]. Future work might search more for specific cryptographic primitives to construct this architecture, and fully evaluate a system supporting global correctness auditing by clients.

10 CONCLUSION

We present Speranza, a system for usable (certificate-based) anonymous software signing. We present identity co-commitments, a technique for one-time-use pseudonyms using Pedersen commitments and Chaum-Pedersen proofs of commitment equality. We then use identity co-commitments and data structure techniques from key transparency to construct our anonymous software signing system. The system is fast, practical, and easily deployable. Future work is needed for better compromise detection and client verifiability of the repository.

ACKNOWLEDGMENTS

This research was partially supported by the NSF under Award No. 2229703 and Cisco Research. We thank the anonymous reviewers for many helpful suggestions, and Nikos Vasilakis for insightful discussions.

REFERENCES

[1] Josh Aas, Richard Barnes, Benton Case, Zakir Durumeric, Peter Eckersley, Alan Flores-López, J. Alex Halderman, Jacob Hoffman-Andrews, James Kasten, Eric Rescorla, Seth Schoen, and Brad Warren. 2019. Let's Encrypt: An Automated Certificate Authority to Encrypt the Entire Web *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 2473–2487. https://doi.org/10.1145/3319535.3363192

[2] Pietro Albini. 2021. Make the authors field optional. https://github.com/rust-lang/rfcs/pull/3052

[3] R. Barnes, J. Hoffman-Andrews, D. McCarney, and J. Kasten. 2019. Automatic Certificate Management Environment (ACME). RFC 8555. https://www.rfc-editor.org/rfc/rfc8555

[4] Daniel J. Bernstein. 2006. Curve25519: New Diffie-Hellman Speed Records. In *Public Key Cryptography - PKC 2006*, Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 207–228.

[5] Marina Sanusi Bohuk, Mazharul Islam, Suleman Ahmad, Michael Swift, Thomas Ristenpart, and Rahul Chatterjee. 2022. Gossamer: Securely Measuring Password-based Logins. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 1867–1884. https://www.usenix.org/conference/usenixsecurity22/presentation/sanusi-bohuk

[6] Melissa Chase, Apoorvaa Deshpande, Esha Ghosh, and Harjasleen Malvai. 2019. SEEMless: Secure End-to-End Encrypted Messaging with Less Trust. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 1639–1656. https://doi.org/10.1145/3319535.3363202

[7] David Chaum and Torben Pryds Pedersen. 1992. Wallet Databases with Observers: Extended Abstract. In *Advances in Cryptology — CRYPTO' 92*. Springer Berlin Heidelberg, Berlin, Heidelberg, 89–105.

[8] Russ Cox and Filippo Valsorda. 2019. Proposal: Secure the Public Go Module Ecosystem. https://go.googlesource.com/proposal/+/master/design/25530-sumdb.md

[9] Curity AB. [n. d.]. Pairwise Pseudonymous Identifiers. https://curity.io/resources/learn/ppid-intro/

[10] Center for Internet Security. [n. d.]. CIS Software Supply Chain Security Guide. https://www.cisecurity.org/insights/white-papers/cis-software-supply-chain-security-guide.

[11] Python Software Foundation. 2023. PyPI - The Python Package Index. https://pypi.org/

[12] Simson. Garfinkel. 1995. *PGP: pretty good privacy.* O'Reilly & Associates, Sebastopol, CA.

[13] Dan Goodin. 2015. Meet "Great Cannon," the man-in-the-middle weapon China used on GitHub. https://arstechnica.com/information-technology/2015/04/meet-great-cannon-the-man-in-the-middle-weapon-china-used-on-github/

[14] Mike Hamburg. 2015. Decaf: Eliminating Cofactors Through Point Compression. In *Advances in Cryptology – CRYPTO 2015 (Lecture Notes in Computer Science)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 705–723.

[15] M Hamburg, H de Valance, I Lovecruft, and T Arcieri. 2017. The ristretto group. (2017).

[16] Mike Hanley. 2021. GitHub's commitment to npm ecosystem security. https://github.blog/2021-11-15-githubs-commitment-to-npm-ecosystem-security/

[17] Philip Harrison. 2022. Link npm packages to the originating source code repository and build. https://github.com/npm/rfcs/blob/e000b367d9e595bc694893c3d845df269f9b875f/accepted/0049-link-packages-to-source-and-build.md#goals

[18] Henry. 2018. GPostmortem for Malicious Packages Published on July 12th, 2018. https://eslint.org/blog/2018/07/postmortem-for-malicious-package-publishes/

[19] Ted Hudek. 2021. Authenticode Digital Signatures. https://learn.microsoft.com/en-us/windows-hardware/drivers/install/authenticode

[20] Mohit Kumar. 2018. Rogue Developer Infects Widely Used NodeJS Module to Steal Bitcoins. https://thehackernews.com/2018/11/nodejs-event-stream-module.html

[21] Trishank Karthik Kuppusamy, Vladimir Diaz, and Justin Cappos. 2017. Mercury: Bandwidth-Effective Prevention of Rollback Attacks Against Community Repositories. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*. USENIX Association, Santa Clara, CA, 673–688. https://www.usenix.org/conference/atc17/technical-sessions/presentation/kuppusamy

[22] Trishank Karthik Kuppusamy, Santiago Torres-Arias, Vladimir Diaz, and Justin Cappos. 2016. Diplomat: Using Delegations to Protect Community Repositories. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*. USENIX Association, Santa Clara, CA, 567–581. https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/kuppusamy

[23] P. Ladisa, H. Plate, M. Martinez, and O. Barais. 2023. SoK: Taxonomy of Attacks on Open-Source Software Supply Chains. In *2023 2023 IEEE Symposium on Security and Privacy (SP) (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 167–184. https://doi.org/10.1109/SP46215.2023.00010

[24] Ben Laurie. 2014. Certificate Transparency. *Commun. ACM* 57, 10 (sep 2014), 40–46. https://doi.org/10.1145/2659897

[25] Sean Lawlor and Kevin Lewi. 2023. Deploying key transparency at WhatsApp. https://engineering.fb.com/2023/04/13/security/whatsapp-key-transparency/

[26] Roch Lefebvre. 2022. RFC: Proposal for new signing mechanism. https://github.com/rubygems/rfcs/pull/37

[27] Jiangtao Li and Ninghui Li. 2005. OACerts: Oblivious Attribute Certificates. In *Applied Cryptography and Network Security (Lecture Notes in Computer Science)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 301–317.

[28] Ulf Lilleengen. 2023. RFC: sigstore and cargo/crates.io. https://github.com/rust-lang/rfcs/pull/3403

[29] Joshua Lock and Marina Moore. 2021. Ephemeral identity verification using sigstore's Fulcio for TUF developer key management. https://github.com/theupdateframework/taps/blob/c09b344a64076646b11ce4fefe0947a0d7172f31/tap18.md

[30] Harjasleen Malvai, Lefteris Kokoris-Kogias, Alberto Sonnino, Esha Ghosh, Ercan Oztürk, Kevin Lewi, and Sean Lawlor. 2023. Parakeet: Practical Key Transparency for End-to-End Encrypted Messaging. Cryptology ePrint Archive, Paper 2023/081. https://doi.org/10.14722/ndss.2023.24545

[31] Sarah Meiklejohn, Pavel Kalinnikov, Cindy S Lin, Martin Hutchinson, Gary Belvin, Mariana Raykova, and Al Cutter. 2020. Think Global, Act Local: Gossip and Client Audits in Verifiable Data Structures. (2020).

[32] Marcela S. Melara, Aaron Blankstein, Joseph Bonneau, Edward W. Felten, and Michael J. Freedman. 2015. CONIKS: Bringing Key Transparency to End Users. In *24th USENIX Security Symposium (USENIX Security 15)*. USENIX Association, Washington, D.C., 383–398. https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/melara

[33] Kelsey Merrill. 2023. zk-Sigstore: System for Anonymous Certificate-Based Software Signing. https://hdl.handle.net/1721.1/151609

[34] Kelsey Merrill, Zachary Newman, Santiago Torres-Arias, and Karen Sollins. 2023. Speranza: Usable, privacy-friendly software signing. arXiv:2305.06463 [cs.CR]

[35] S. Micali, M. Rabin, and S. Vadhan. 1999. Verifiable random functions. In *40th Annual Symposium on Foundations of Computer Science (Cat. No.99CB37039)*. 120–130. https://doi.org/10.1109/SFFCS.1999.814584

[36] Microsoft. [n. d.]. Open Source Software (OSS) Secure Supply Chain (SSC) Framework Simplified Requirements. https://github.com/microsoft/oss-ssc-framework/blob/main/specification/framework.md.

[37] Mozilla Project. [n. d.]. Security/Binary Transparency. https://wiki.mozilla.org/Security/Binary_Transparency

[38] Frank Nagle, James Dana, Jennifer Hoffman, Steven Randazzo, and Yanuo Zhou. 2022. *Census II of Free and Open Source Software — Application Libraries*. Technical Report. Harvard Laboratory for Innovation Science (LISH) and Open Source Security Foundation (OpenSSF). https://linuxfoundation.org/tools/census-ii-of-free-and-open-source-software--application-libraries/

[39] Lily Hay Newman. 2020. How to Understand the Russia Hack Fallout. https://www.wired.com/story/russia-solarwinds-hack-targets-fallout/

[40] Zachary Newman, John Speed Meyers, and Santiago Torres-Arias. 2022. Sigstore: Software Signing for Everybody. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7-11, 2022*, Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi (Eds.). ACM, 2353–2367. https://doi.org/10.1145/3548606.3560596

[41] Kirill Nikitin, Eleftherios Kokoris-Kogias, Philipp Jovanovic, Nicolas Gailly, Linus Gasser, Ismail Khoffi, Justin Cappos, and Bryan Ford. 2017. CHAINIAC: Proactive Software-Update Transparency via Collectively Signed Skipchains and Verified Builds. In *26th USENIX Security Symposium (USENIX Security 17)*. USENIX Association, Vancouver, BC, 1271–1287. https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/nikitin

[42] Marc Ohm, Henrik Plate, Arnold Sykosch, and Michael Meier. 2020. Backstabber's Knife Collection: A Review of Open Source Software Supply Chain Attacks. In *Detection of Intrusions and Malware, and Vulnerability Assessment*, Clémentine Maurice, Leyla Bilge, Gianluca Stringhini, and Nuno Neves (Eds.). Springer International Publishing, Cham, 23–43.

[43] OpenID. 2022. Openid Connect. https://openid.net/connect/

[44] Torben Pryds Pedersen. 1991. Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing. In *Advances in Cryptology — CRYPTO '91*. Springer Berlin Heidelberg, Berlin, Heidelberg, 129–140.

[45] President Biden. 2021. *Executive Order.* Number 14028. 15 pages.

[46] RubyGems.org. 2023. RubyGems. https://rubygems.org/

[47] Scott Ruoti, Jeff Andersen, Daniel Zappala, and Kent Seamons. 2015. Why Johnny Still, Still Can't Encrypt: Evaluating the Usability of a Modern PGP Client. (2015).

[48] Gary Belvin Ryan Hurst. 2017. Security Through Transparency. https://security.googleblog.com/2017/01/security-through-transparency.html

[49] Justin Samuel, Nick Mathewson, Justin Cappos, and Roger Dingledine. 2010. Survivable Key Compromise in Software Update Systems. In *The 17th ACM Conference on Computer and Communications Security (CCS '10)* (Chicago, IL). ACM, Chicago, IL.

[50] Sigsum. 2021. Sigsum. https://www.sigsum.org/

[51] Manu Sporny, Dave Longley, and David Chadwick. 2022. Verifiable Credentials Data Model v1.1. https://www.w3.org/TR/vc-data-model/

[52] Donald Stufft. 2018. GPG Signatures in the Warehouse UI. https://github.com/pypi/warehouse/issues/3810#issuecomment-405975460

[53] Mikhail Swift. 2022. ITE-7: Signing & Verification With X509. https://github.com/in-toto/ITE/blob/master/ITE/7/README.adoc

[54] Synopsys. 2023. Open Source Security and Risk Analysis Report. https://www.synopsys.com/software-integrity/engage/ossra/rep-ossra-2023-pdf

[55] The GnuPG Project. 2023. The GNU Privacy Guard. https://gnupg.org/

[56] Santiago Torres-Arias, Hammad Afzali, Trishank Karthik Kuppusamy, Reza Curtmola, and Justin Cappos. 2019. in-toto: Providing farm-to-table guarantees for bits and bytes. In *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, Santa Clara, CA, 1393–1410. https://www.usenix.org/conference/usenixsecurity19/presentation/torres-arias

[57] Nirvan Tyagi, Ben Fisch, Andrew Zitek, Joseph Bonneau, and Stefano Tessaro. 2022. VeRSA: Verifiable Registries with Efficient Client Audits from RSA Authenticated Dictionaries. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security* (Los Angeles, CA, USA) *(CCS '22)*. Association for Computing Machinery, New York, NY, USA, 2793–2807. https://doi.org/10.1145/3548606.3560605

[58] Ioanna Tzialla, Abhiram Kothapalli, Bryan Parno, and Srinath Setty. 2022. Transparency Dictionaries with Succinct Proofs of Correct Operation. In *NDSS 2022*. https://www.microsoft.com/en-us/research/publication/transparency-dictionaries-with-succinct-proofs-of-correct-operation/

[59] Alma Whitten and J. D. Tygar. 1999. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. In *8th USENIX Security Symposium (USENIX Security 99)*. USENIX Association, Washington, D.C. https://www.usenix.org/conference/8th-usenix-security-symposium/why-johnny-cant-encrypt-usability-evaluation-pgp-50

A ANONYMIZED AUTHORIZATION RECORD EXAMPLES

This appendix provides concrete examples of what verifying a policy or changing a policy might look like with an anonymized authorization record.

A.1 Example: Verifying with a Threshold Policy

Assume there exists artifact `Foo` that has three signers, Alice, Bob, and Charlie. Their policy is such that in order to publish any changes to `Foo`, at least two of the three signers must sign off on the change.

The following is an example of what CheckPublish might look like for an artifact:

```
struct ThresholdPolicy {
    signers: List<Commit>,
    threshold: int
}
```

The input consists of at least `threshold` commitments to signers and two equality proofs. To check publishing authorization, check, that the inputs validate for at least `threshold` distinct `signers`.

A.2 Example: Changing Policies

Again assume there exists artifact `Foo` with signers Alice, Bob, and Charlie. Their ownership change policy is that Alice is the head signer, and she and she alone must sign off on any changes in ownership.

The following is an example of what the CheckPolicyChange function may look like for this policy:

```
struct HeadsignerPolicy {
    head_signer = Commit(Alice),
    signers = List<Commit>,
}
```

The input consists of a commitment to the head signer, an equality proof, and a new policy. To check authorization, check the identity co-commitment between the input and the `head_signer`. If it holds, overwrite the policy completely with the new policy. Third party monitors can also run these checks themselves to verify that the server has changed the policy correctly and honestly. The security notion for this model includes the consistency notion described in Section 6.2 as well as the privacy property. Further analysis of the security of this model follows in Section 7.