

KeyForge: Mitigating Email Breaches with Forward-Forgeable Signatures

Michael Specter
MIT

Sunoo Park
MIT & Harvard

Matthew Green
Johns Hopkins University

Abstract

Email breaches are commonplace, and they expose a wealth of personal, business, and political data that may have devastating consequences. The current email system allows any attacker who gains access to your email to prove the authenticity of the stolen messages to third parties — a property arising from a necessary anti-spam / anti-spoofing protocol called DKIM. This exacerbates the problem of email breaches by greatly increasing the potential for attackers to damage the users’ reputation, blackmail them, or sell the stolen information to third parties.

In this paper, we introduce *non-attributable email*, which guarantees that a wide class of adversaries are unable to convince any third party of the authenticity of stolen emails. We formally define non-attributability, and present two practical system proposals — KeyForge and TimeForge — that provably achieve non-attributability while maintaining the important protection against spam and spoofing that is currently provided by DKIM. Moreover, we implement KeyForge and demonstrate that that scheme is practical, achieving competitive verification and signing speed while also requiring 42% *less* bandwidth per email than RSA2048.

1 Introduction

Email is the world’s largest and most ubiquitous messaging scheme, spanning industry, government, and personal use. Email content comprises a vast trove of sensitive information; it is common knowledge that email breaches can have devastating consequences.

Making matters worse, email protocols today ensure that if a malicious party gains access to your email, then she can cryptographically prove the authenticity of the stolen emails to any third party. This greatly increases the intelligence value of stolen email: attackers can (even anonymously) damage the reputation of users by publicly disseminating sensitive messages, or sell the messages to interested third parties, all with high credibility. This property of email (“attributability”) arose as an unintended side-effect of the DomainKeys Identified Mail (DKIM) standard,¹ whose purpose is to fight spam and spoofing attacks.

The consequences of this design are far from hypothetical. For example, during the 2016 U.S. election season, leaked emails from the Democratic National Committee (DNC) and other operatives were released to the public via Wikileaks [50]. While the original senders did not intentionally authenticate the emails, third parties were able to verify the authenticity of emails due to the presence of cryptographic signatures from Google’s email server [49]. In some cases this authentication directly contradicted disavowals by party officials [36].

Regardless of whether one believes the 2016 DNC leaks were ultimately for the better or worse, it is alarming that an unintended side effect of a widely used protocol has created a broader

¹As noted by a DKIM RFC author [15].

ecosystem for credible propagation and misuse of illicitly obtained private communication. That DKIM materially supports such an ecosystem is clear. Wikileaks has a webpage about DKIM authentication [49] indicating that the organization relies heavily on non-repudiability to enhance the credibility of its material. This is corroborated by recent incidents, e.g., [13], in which email leaks were allegedly salted with falsified data.

This raises a natural question: can we mitigate the potential harms of such an ecosystem, while maintaining the efficiency and strong spam- and spoofing-resistance of email today? It is a particularly pressing question as new reports of data breaches, including email account information, seem to surface every few weeks. The 2013 breach of every one of Yahoo!’s 3 million email accounts [46] is a vivid reminder that not only public figures are impacted by the consequences of compromised email. Though the Yahoo! data has not been publicized, data from other large breaches has been posted publicly online.² Besides election season strategy, attackers’ motives have ranged from financial gain — such as by selling patient healthcare data gleaned from emails [29] — to industrial espionage and monitoring political dissidents and foreign officials [24]. In the absence of any outside breach, such incentives may also prompt malicious insiders to sell or leak email data.

While the property that stolen email can be cryptographically attributed exacerbates email leaks, DKIM itself serves a necessary purpose in the email ecosystem. Concretely, DKIM requires email servers to provide a cryptographic signature on each outgoing email’s contents and metadata (including the sender’s identity). This assures the receiving server that the sending domain actually did intend to send the message, providing inter-domain accountability in case of spam, and moreover ensuring that spoofing is detectable.

An initial intuition may be that attributability of stolen email is a necessary side-effect of any spam- and spoofing-resistant email system. Indeed, it is unclear how a recipient can be certain that an email was legitimately sent without gaining the ability to convince a third party of the same.

In this work, we challenge the above intuition and construct efficient protocols that achieve the important spam- and spoofing-resistance that DKIM provides, while simultaneously *guaranteeing non-attributability* of stolen email. Our protocols provably ensure that a wide class of adversaries are unable to convince any scrupulous third party of the authenticity of stolen emails.

1.1 System Requirements

Requirements on email demand certain unique properties not commonly seen in other messaging systems. For example, email is an any-mesh ecosystem; any domain owner on the internet could set up the appropriate DNS records and become a server that interoperates with any other email server. Similarly, a domain may be responsible for signing and verifying anything from hundreds to millions of emails per day, and administrators are likely to have to contend with increasing throughput requirements over time.

In considering throughput for any domain, while maintaining good constants on computation time for signing and verification of emails is important, it is more important that the service be able to scale — that adding more resources to the system actually provide linear or better performance guarantees. Scalability in terms of interconnection with other mail servers is as important as email throughput.³

²E.g., the recent “Collection 1” breach implicating 773 million email addresses [30]. Note that the “Collection 1” breach consists of email addresses and passwords used to log in to websites, not necessarily email accounts.

³The IETF standard for DMARC [1] explicitly states that pre-sending agreements is a poor scalability choice for this reason, see also [47].

An unintuitive result of these throughput requirements is that certain types of overhead that would be trivial in other messaging contexts, such as requiring servers to communicate prior to a message being sent or requiring round trips between servers on every message, are unlikely to be viable in this environment. For example, it would be difficult to scale a service that requires constant connection to, and maintain constant state with, an unbounded number of potentially malicious servers.

In addition to efficiency and scalability, a practical solution for email non-attributability must address the following considerations:

Untrusted recipients. A naive proposal is to rely on recipient behavior to ensure non-attributability: e.g., having the receiving server delete all DKIM header information upon receipt. Such solutions are inadequate for a threat model whose very purpose is to prevent attacks by the recipient and/or the recipient’s server (or an attacker who compromises the recipient). Ideally, non-attributability for a given sender’s emails should arise from the sender’s behavior alone.

Asynchrony. Email is inherently asynchronous and must operate across many nodes that are unable to interact directly, as discussed above. Thus, schemes based on interactive authentication protocols, such as naturally deniable interactive zero-knowledge proofs — while they might suffice for a synchronous communication setting — are not viable solutions for email. Schemes based on other interactive protocols such as OTR [12] are inappropriate for similar reasons.

Universal forgeability One intuitive definition of email non-attributability could be that DKIM signatures must be *forgeable* by the recipient of an email (“recipient forgeability”). Then, even if the attacker presents validly signed emails a third party, she lacks credibility: had she really broken into the recipient’s email, she could have produced the signatures herself. Recipient forgeability could be achieved, for example, by using replacing DKIM signatures with ring signatures [44] such that a valid signature can be produced by either the sender or the recipient, as proposed by [3].

However, recipient forgeability is a weak notion of non-attributability. A stronger, and more desirable, notion is that DKIM signatures must be forgeable by *anyone* after a small time delay to prevent spam and spoofing attacks. This renders alleged attackers far less credible, since anyone can forge valid-looking signatures without breaking into any email account at all. We call this stronger notion “universal forgeability,” and consider it an essential feature of a satisfactory non-attributability solution for email.

Incremental deployment. Given the vast number of existing email servers and the need for interoperability, we consider the majority of the email ecosystem to be entrenched. For example, it would be difficult to require substantial changes to mail routing, and one cannot expect that every actor immediately switch to some new scheme. We do, however, hold the view that DKIM itself can be replaced incrementally, and that new signing mechanisms represent a less onerous impact on the email ecosystem.

Long-lived keys. One natural approach to short-lived credibility of signatures is to leverage correspondingly short-lived keys and publish each secret key at the end of its lifetime; this sort of approach to achieving forgeability has been mentioned in passing outside of the context of email [12]. Similarly, one may leverage easily breakable keys in tandem with rapid key rotation, a solution originally suggested by Jon Callas, one of the original creators of the DKIM standard [15].

Unfortunately, too-frequent key rotation entails a variety of likely insurmountable practical problems; rotating keys is a manual process that introduces risk of misconfiguration and key theft, and no administrator wants to be responsible for disabling email for an entire organization. Worse, DNS results are often cached, so replacing an individual DNS record with a new one takes time and can yield inconsistent results for different DNS servers on the receiver’s end. Thus, given the impracticability of very frequent key rotation, this simple scheme cannot give a satisfactory non-attributability guarantee.

These observations point to a possible solution involving long-lived *public* keys but short-lived *secret* keys, reminiscent of forward-secure signatures (FSS) [4, 10]. However, FSS were designed with a rather different goal from ours: namely, to allow efficient key updating while preventing derivation of *past* keys from present and future keys. In contrast, our setting requires that *present and future* keys cannot be derived from past keys. The only way to achieve this property using an FFS would be to precompute a long list of keys and then use the secret keys *in backwards order* — this is arguably better than the simplest solution based on short-lived keys, but having to store the whole list of keys is inefficient and unsatisfactory.

Our solution, KeyForge, does in fact leverage a signature scheme with long-lived public keys and short-lived secret keys: we define a new primitive called *forward-forgeable signatures* (FFS) capturing the specific requirements of our application, and give a construction of FFS from any hierarchical identity-based signature scheme, as detailed in Section 4. Our second proposal, TimeForge, requires no exposure of private keys, but achieves a slightly weaker notion of non-attributability than KeyForge. TimeForge is explained in Section 5.2.

1.2 Key Ideas

There are two main ideas underlying the solutions we explore: *delayed universal forgeability* and *immediate recipient forgeability*.

Delayed universal forgeability. The key idea in this approach is to ensure that signatures with respect to past emails “expire” after a time delay Δ and thereafter become forgeable by the general public (i.e., arbitrary outsiders or non-parties). This property ensures that no attribution will be credible after the time delay has elapsed. We call this property *delayed universal forgeability*. As long as Δ is set larger than the maximum viable time for email latency, the signature will still be convincing to the recipient at the time of receipt, thus maintaining the spam- and spoofing-resistance of DKIM.

Signatures that possess delayed universal forgeability retain all of the standard unforgeability properties one would expect of a standard signature scheme, until the set time Δ has passed. Thus in cases where an attacker gains access to email and shows it to a third party within Δ time after the email was sent, a third party will be convinced of the email’s authenticity. Effectively, delayed universal forgeability protects against adversaries that compromise an email account by breaking in and taking a snapshot (“after-the-fact attacks”), but not adversaries that fully control an email account and monitor its email in real time (“real-time attacks”). After-the-fact attacks cover a broad range of realistic attacks, for example, including many data breaches. Next, we discuss how we address real-time attacks.

Immediate recipient forgeability. Suppose that the fact of access to a particular client account implies the ability to forge messages from arbitrary other servers *to that recipient only*: that is, the ability to obtain valid DKIM signatures on email content and metadata of one’s choice. We call this *immediate recipient forgeability*. Importantly, the recipient constraint ensures the inability

to impersonate any other server for the purposes of email addressed to *other* recipients, thus not interfering with DKIM’s spam- and spoofing-resistance.

This undermines the credibility of any attacker who claims to have ongoing access to a particular email account and attempts to convince third parties of the authenticity of emails that she claims were sent to (and from) that account. This is also the case for real-time attacks, which may publish allegedly-incoming emails immediately after they are received.

Combining both notions. While immediate recipient forgeability has the advantage of no time delay, it has the disadvantage that the ability to forge is not universal: rather, it is limited to those who have access to the recipient email account. In contrast, delayed universal forgeability enables anyone at all to forge. Our protocols achieve the “best of both worlds,” in order to leverage both their advantages.

Section 3 defines our threat model, discusses its limitations, and formalizes immediate recipient forgeability and delayed universal forgeability with indistinguishability-based definitions (Definitions 1–2).

1.3 Our Solutions

This paper fully specifies and evaluates our main proposal, KeyForge. We also outline an alternative protocol TimeForge that achieves email non-attributability by rather different techniques assuming the existence of a timestamping authority.

KeyForge Our main proposal is KeyForge, a protocol to replace DKIM and the associated authentication mechanisms. The first idea of KeyForge is to achieve *delayed universal forgeability* by publicizing signing keys after a set time delay Δ (say, by posting them on the internet).

We present a new definition, *forward-forgeable signatures (FFS)*, to formalize the requirements of a signature scheme equipped with a method to selectively disclose “forging information” for subsets of past signatures. The public key is required to persist over disclosures: otherwise, the frequency of manual key updates and DNS updates would render fine-grained disclosures inviable. We moreover define *succinctness* of FFS, a measure of how efficiently the disclosures can be made. We construct FFS based on hierarchical identity-based signatures (HIBS) and use it as an essential building block for KeyForge. Our construction achieves logarithmic succinctness.

The second idea of KeyForge is to achieve *immediate recipient forgeability* by having email servers accept “forge requests” requesting that specific emails be sent to the requester in real time. The restriction that the recipient must be the requester ensures that DKIM’s spam- and spoofing-protection remains effective.

TimeForge. Our second proposal TimeForge assumes a *publicly verifiable timekeeper (PVTk)* model, in which a reliable timekeeper periodically issues timestamps whose authenticity is publicly verifiable (e.g., because the authority signs each timestamp). Rather than deploying and scaling a new timekeeping service, we show that viable timekeeping systems can be constructed from existing Internet infrastructure, including OCSP staple responders [41], Certificate Transparency [35], randomness beacons [39], and proof-of-work blockchains [37, 51].

In a nutshell, the idea of TimeForge is to substitute each DKIM signature on a message m at time t with a succinct zero-knowledge proof of the statement $S(m) \vee R(m) \vee T(t + \Delta)$, where: $S(m)$ denotes knowledge of a valid signature by the sender on m , $R(m)$ denotes knowledge of a valid signature by the receiver on m , and $T(t + \Delta)$ denotes knowledge of a valid timestamp for a time later than $t + \Delta$. Effectively, the inclusion of $T(t + \Delta)$ ensures *delayed universal forgeability*, and

the inclusion of $R(m)$ ensures *immediate recipient forgeability*. The primitive just described can be thought of as a forward-forgeable signature scheme in the PVTk model.

1.4 Summary of Contributions

- We provide formal *definitions* of:
 - *Email non-attributability* (Section 3.3)
 - *Forward-forgeable signatures* (Section 4)
- We give *provably secure constructions* of succinct forward-forgeable signatures (FFS) in both the standard and PVTk models. FFS may be of independent interest beyond the present application.
- We present two *system designs* — KeyForge, and TimeForge — that achieve email non-attributability and could be applied to the existing email system with acceptable practical and processing overhead.
- We present a fully *implemented* KeyForge and provide a comprehensive evaluation of KeyForge’s impact on signing, verification, and bandwidth costs.

2 Background on Email




Figure 1: Simplified email routing infrastructure

RFC 5598 [19] defines terminology for processing and delivering mail. This section introduces basic terminology towards discussing our threat model and system requirements.

2.1 Email Routing

Figure 1 illustrates the architecture of email routing: an asynchronous routing protocol built on top of TCP/IP.

Users first establish a relationship with a trusted email service provider, called a Mail Submission Agent (MSA) on the sender side and a Mail Delivery Agent (MDA) on the receiver side. The MDA is responsible for various tasks including verifying the authenticity of incoming messages.

The user’s email client is generically called a Mail User Agent (MUA). Email originates from an MUA, and arrives at the user’s trusted MSA. Depending on the system’s configuration, the MSA may send the email to other Mail Transfer Agents (MTAs) it trusts. Eventually, an MTA performs

a DNS lookup for the receiving domain to check which MTAs are authorized to process emails for that domain. The email is then sent to one of these MTAs. After a number of hops that depend on the sending and receiving organizations' infrastructure, the email arrives at the receiver's MDA, which is responsible for verifying the message for the receiver's MUA.

2.2 Email Authentication

The IETF has developed a number of standards that allow domains to sign and verify incoming and outgoing messages. We provide a stylized overview of the four that have seen appreciable adoption: DKIM, SPF, DMARC, and ARC.

DKIM. DomainKeys Identified Mail (DKIM) is an IETF standard for an MSA to sign outgoing email, and an MDA to verify that email by looking up the MSA's public key in the DNS.

Setup: The MSA generates a key pair and uploads the public key to the DNS.

Send: The MSA adds the location of its public key to the email's metadata (or *header*), as well as metadata allowing verification of the signature, and then signs the email and its headers.⁴

Verify: On receipt, the MDA performs a DNS lookup for the sender's public key as specified in the header, and then cryptographically verifies the email.

SPF The Sender Policy Framework (SPF) ensures that intermediary MTAs are permitted to send and receive messages as a part of the domain. In a way, this solves a somewhat orthogonal problem to DKIM in that SPF provides little guarantee that the message itself has not been modified by an intermediary, but instead provides spoofing protection by limiting what IP addresses are valid accepting MTAs.

DMARC An SPF or DKIM failure as a result of a misconfiguration is indistinguishable from a failure due to an attempted message spoofing, and neither DKIM nor SPF provide mechanisms for alerting the sending domain that there has been a problem. DMARC solves this by adding a DNS TXT record specifying to the receiver what it should do in the case of such failures (such as quarantine, reject, or accept the message despite the failure), as well as providing an email address to send aggregated statistics on such failures.

ARC Authenticated Received Chain (ARC) is an experimental and largely unimplemented IETF standard that aims to address the issues caused by indirect email flows.

Legitimate modification of messages in transit may occur in a number of circumstances. For example, an MTA that is also a virus scanner may remove malicious attachments, or a mailing list may prepend the mailing list's name to the subject of an email. Unfortunately, any alteration of the body or headers invalidates the original DKIM signature.

ARC acts as an attestation by an intermediary that some subset of DKIM, SPF, DMARC, or previous ARC instances were verifiable before content was modified. To do so, the intermediary adds its own signature to the header of the message, along with metadata about what was verified.

⁴This usually includes a hash of the entirety of the message, but the specification does allow for portions of the body of the message to go unsigned but it's unclear how often this is done, and appears to be relatively rare.

2.3 Challenges for Synchronous Authentication Protocols

Email is inherently asynchronous, and any replacement for DKIM must therefore allow for non-interactive verification of emails. The reason for this is twofold:

Third party MTAs. Many organizations choose to use third-party MTAs as an initial hop between the organization’s own mail server and the internet. The reason one may choose to use such MTAs is varied, but a common use case is to provide some protection from spam, malicious attachments, or DDoS attacks. While these intermediaries are allowed to quarantine messages or provide flow control to the MDA, they are not trusted to refrain from modifying or spoofing emails. By scraping DNS results for the Alexa top 150k, we find that 31,615 have an MX record, and, of these, 22% (6,793) are using a confirmed multi-hop third-party MTA provider.

Mail forwarding. Often users will forward email received from an account on one domain to another through mail forwarding. These emails then must be verified by the receiving domain, and there is no way for the receiver to communicate with the original sender.

3 Model and Security Definitions

This section presents threat models and formal definitions of email non-attributability.

Notation “PPT” means “probabilistic polynomial time.” $|S|$ denotes the size of a set S . $[n]$ denotes the set $\{1, \dots, n\}$ of positive integers up to n , and $\mathbb{P}(\cdot)$ denotes powerset. \approx_c denotes computational indistinguishability. $\tau||e$ denotes the result of appending an additional element e to a tuple τ .

3.1 Model

Time We model time in discrete time-steps and assume that everyone has access to a fairly accurate local clock. (This is a reasonable model because our system requires synchrony only at a coarse granularity, such that realistic discrepancies in local clocks will not have any effect.)

Asynchrony Recognizing the asynchronous nature of email, $\hat{\Delta}$ denotes an upper bound on the time required for email delivery. Our parameter settings depend on $\hat{\Delta}$, and our evaluation sets $\hat{\Delta}$ at 15 minutes, as discussed in Section 5.1.1.

DNS Our model assumes all parties and algorithms have access to DNS and can update their own DNS records.

Publication We assume that each party has method of publishing persistent, updatable information that is retrievable by all other parties and algorithms. This could be via DNS or on another medium, such as posting information on a website.

3.2 Threat models

The type of attack with which we are concerned is disclosure of private communications at the *recipient’s* server, whether because the recipient is compromised or because it is malicious. We define two threat models, defined below.

Both our proposals KeyForge and TimeForge achieve security against Threat Model 2, which is the stronger of the two threat models and captures advanced persistent threats. Threat Model 1

may be a meaningful threat model in a scenario where attackers may illicitly gain access to an email server but are not likely to be able to maintain such access for extended periods of time. A basic, simpler version of KeyForge achieves non-attributability for almost all emails (i.e., all but very recent emails) under Threat Model 1.

Threat Model 1. (After-the-fact attacks) *In this model, the recipient is presumed honest at the time of email receipt, but is later compromised by an attacker that takes a snapshot of all past email content.*

Threat Model 2. (Real-time attacks) *In this model, the receiver is presumed malicious at the time of email receipt, i.e., presumed to have ongoing access and immediate intent to disclose the content of received emails to third parties.*

Client-server trust. In the current system, email clients rely heavily on their email servers. A malicious email server could easily and undetectably misbehave in many essential functions: it could drop incoming emails, modify outgoing emails (since typically, emails are not signed client-side), or falsify the content and metadata of incoming emails (since typically, clients do not perform DKIM verification themselves).

Thus, client-server trust is very high in practice. We treat the client and server as a single entity for the purposes of this paper, except where specifically otherwise indicated.

No trusted parties. In a system where credibility is based on a reliable reputation system, a trusted (reputable) party with the ability to eavesdrop on the communication channel would be able to undermine non-attributability by keeping logs of all traffic it sees. For example, an MTA in the email system, if it kept logs of all traffic, could effectively resolve the question of authorship of emails by referring to its internal logs, in a manner convincing to anyone who simply trusts the MTA based on its reputation. Our model assumes mutually distrustful parties: i.e., that no party is taken simply on its word as in the scenario just described. In other words, credibility in our model is evidence-based and not reputation-based.

3.3 Defining Non-Attributability

We define email non-attributability as a game involving an email protocol E , adversary \mathcal{A} , simulator \mathcal{S} , and distinguisher \mathcal{D} . For any email server S with internal state s , we denote by $E_s(S, R, m, \mu, t)$ the information that recipient R receives when S legitimately sends at time t an email message m with metadata μ to R . For simplicity, our notation leaves implicit that parties other than the sender participate in transmission and may affect the information received by R (for example, the MTAs); however, the final received information $E_s(S, R, m, \mu)$ should be thought to contain any modifications made en route between S and R . While this definition refers to “internal state s ” for generality, the state s can for our purposes be thought to contain secret key material.

Intuitively, we require indistinguishability between a legitimate email $E_s(S, R, m, \mu, t)$ and a “fake” email that was created without access to the sending server at all: that is, without knowing s . To model this, we consider a simulator \mathcal{S} whose “goal” is to create an email that looks legitimate without knowing s , and our security definition requires that $E_s(S, R, m, \mu, t)$ is indistinguishable from the output of \mathcal{S} .

Next, we give two formal definitions of non-attributability. The first (Definition 1) considers a simulator that has access to a particular recipient’s email server, and is required to output email from any sender to that recipient. The second (Definition 2) is a stronger definition that considers

a simulator that is required to output email from any sender to any recipient and has access to neither the sender’s nor the receiver’s email server.

In other words, if an adversary publishes an email allegedly authored by an honest party, Definition 1 guarantees that she can credibly argue that, granting that the attacker indeed broke into her or her correspondent’s account, his allegation inherently lacks credibility because he must have had the ability to forge emails with content of his choice (by virtue of having access to one of their email accounts). Definition 2 gives the yet stronger guarantee that *anyone* can forge past emails: so the allegation is even less credible and the email accounts may not even have been compromised. Note that the two definitions are incomparable: neither is strictly stronger.

Definition 1 (Recipient non-attributability). *An email protocol E is non-attributable for recipient servers if there is a PPT simulator \mathcal{S} such that for any sender S and recipient R , for any email message m and metadata μ ,*

$$E_s(S, R, m, \mu, t) \approx_c \mathcal{S}(R, S, m, \mu) ,$$

where s is the internal state of S , t is the time at which \mathcal{S} is invoked, and R denotes the internal state of the recipient R ’s email server.

E is furthermore non-attributable for recipients if there is a PPT simulator \mathcal{S} such that for any sender S and recipient R , for any email message m and metadata μ ,

$$E_s(S, R, m, \mu, t) \approx_c \mathcal{S}^R(S, m, \mu) ,$$

where the superscript R denotes that \mathcal{S} has the capability of sending outgoing mail as R through R ’s email server.

Observe that *non-attributability for recipients* is a stronger than (i.e., implies) *non-attributability for recipient servers*. KeyForge achieves non-attributability for recipients, whereas TimeForge achieves non-attributability for recipient servers. See Remark 3 for further discussion of the differences between the two variants of recipient non-attributability.

Definition 2 (Δ -universal non-attributability). *For $\Delta \in \mathbb{N}$, an email protocol E is Δ -strongly non-attributable if there is a PPT simulator \mathcal{S} such that for any sender S (with internal state s) and recipient R , for any email message m , metadata μ , and timestamp t , the following holds at any time $\geq t + \Delta$:*

$$E_s(S, R, m, \mu, t) \approx_c \mathcal{S}(S, R, m, \mu, t) .$$

The two indistinguishability requirements defined above serve to ensure that no attacker can credibly claim to a third party⁵ that he is providing her with the true content of an email account: the third party is in the role of distinguisher.

Definition 2 is inviable if $\Delta < \hat{\Delta}$. Otherwise, the spam- and spoofing-resistance provided by DKIM would be undermined, since any outsider could use the simulator in real time to send spam email indistinguishable to the recipient from email actually sent by an honest party. This means that the simulator can only begin to generate an indistinguishable distribution after $\hat{\Delta}$ delay. Recall that as time elapses, new information may become available to the simulator, as discussed in Section 3.1.

⁵E.g., the general public (if the allegedly stolen emails are released publicly) or a specific interested party (such as a potential buyer or disseminator of the information).

Relation to the threat models. Δ -universal non-attributability (Definition 2) achieves non-attributability against after-the-fact attacks (Threat Model 1) for all emails sent and received at least Δ before the server is compromised.

The combination of recipient non-attributability and $\hat{\Delta}$ -universal non-attributability (Definitions 1 and 2) yields effective non-attributability for all emails against real-time attacks (Threat Model 2). A real-time attacker with ongoing access to an email server can easily make the fact of his access evident by immediately publishing all emails he sees (in particular, within time $\hat{\Delta}$ of receipt), but will nonetheless be unable to convince a third party of any particular emails since the fact of his access to the email server allows him to forge emails in real time, under Definition 1. With respect to allegedly compromised emails from more than $\hat{\Delta}$ ago, an attacker’s credibility is even more diminished, since for such past timestamps *anyone with internet access* can generate seemingly validly signed emails, even without breaking into any email server at all, under Definition 2.

Limitations of our model. A practical consequence of recipient non-attributability is that a recipient R ’s email server can, *unknown to R* , create fraudulent messages that appear to be legitimate emails from any sender to R , and deliver them to R . As discussed Section 3.2, the level of client-server trust necessitated by the current email system is very high. In this context, recipient non-attributability does not meaningfully increase the trust a client places in her email server. For example, email servers in the current system could (and often do) omit DKIM headers when delivering emails to clients: this effectively implies the ability to deliver fake messages to their clients.

A trivial and uninteresting way to achieve Definitions 1 and 2 is not to sign emails at all. Of course, this is undesirable as it would undermine the spam- and spoofing-resistance for which DKIM was designed. The concurrent requirement of spam- and spoofing-resistance is implicit throughout this paper. It follows that any simulator \mathcal{S} satisfying Definition 1 must use the recipient R ’s secret state r : that is, in order to prevent spam, real-time forgery must be limited to messages whose recipient is the forger herself. This suggests that recipient non-attributability would not be meaningful in an extreme situation where every single use of r can be monitored and attested to, since then an attacker could prove that \mathcal{S} was never invoked on r . This might be plausible assuming secure hardware, e.g., by generating and monitoring all uses of r within a secure enclave (as suggested in [28]) — but even then, such an attack would likely only be feasible by the unlikely attacker who has designed her recipient email server with this unlikely configuration from its very setup. We consider such attacks to be outside our threat model.

Finally, we note that our definitions allow for attackers to convince others of the fact that they have access to a particular email account. Our guarantee is that even so, they cannot make credible claims about the content of the emails, since they gain the ability to falsify emails by the very fact of their access. That attackers can prove access is unavoidable given that universal forgeability is not compatible with spam resistance for too small Δ : e.g., consider a real-time attacker that obtains a secure timestamp (via an outside timestamping service) for all incoming emails immediately on arrival. Under Definition 1, this does *not* prove authenticity of the email content, but just demonstrates access to the recipient’s account. In general, we treat such attacks as being outside our threat model.

4 Forward-Forgeable Signatures

4.1 Definition

Definition 3 formalizes *forward-forgeable signatures (FFS)*. They are a new primitive that this paper introduces, and are an essential building block of our main protocol KeyForge. Informally, FFS are signature schemes equipped with a method to selectively “expire” past signatures by releasing *expiry information* that makes them forgeable. More precisely: in an FFS, each signature is made with respect to a *tag* τ , which can be an arbitrary string. Expiry information can be released with respect to any tag or set of tags.

In our context of wishing to make past signatures forgeable after a time delay, the tag can be thought to be a timestamp: for example, each email is signed with respect to the current time τ , and at some later time $\tau + \Delta$, the signer may publish expiry information for τ .

The *correctness* requirement of FFS is the same as that of standard signatures. The *unforgeability* requirement is modified to include an *expiry oracle*: that is, unforgeability of signatures w.r.t. non-expired tags must hold even in the presence arbitrary, adversarially chosen expirations. The *forgeability on expiry* requirement is a feature of FFS that has no analogue in standard signatures:

Definition 3 (FFS). A forward-forgeable signature scheme (FFS) Σ is implicitly parametrized by message space \mathcal{M} and tag space \mathcal{T} , and consists of five algorithms

$$\Sigma = (\text{KeyGen}, \text{Sign}, \text{Verify}, \text{Expire}, \text{Forge})$$

satisfying the following syntax and requirements.

SYNTAX:

- $\text{KeyGen}(1^\kappa)$ takes as input a security parameter⁶ 1^κ and outputs a key pair (vk, sk) .
- $\text{Sign}(sk, \tau, m)$ takes as input a signing key sk , a tag $\tau \in \mathcal{T}$, and a message $m \in \mathcal{M}$, and outputs a signature σ .
- $\text{Verify}(vk, \tau, m, \sigma)$ takes as input a verification key vk , a tag $\tau \in \mathcal{T}$, a message $m \in \mathcal{M}$, and a signature σ , and outputs a single bit indicating whether or not σ is a valid signature with respect to vk , m , and τ .
- $\text{Expire}(sk, T)$ takes as input a signing key sk and a tag set $T \subseteq \mathcal{T}$, and outputs expiry info η .
- $\text{Forge}(\eta, \tau, m)$ takes as input expiry info η , a tag $\tau \in \mathcal{T}$, and a message $m \in \mathcal{M}$, and outputs signature σ .

REQUIRED PROPERTIES:

1. Correctness: For all $m \in \mathcal{M}, \tau \in \mathcal{T}$, there is a negligible function ε such that for all κ ,

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{KeyGen}(1^\kappa) \\ \sigma \leftarrow \text{Sign}(sk, \tau, m) \\ b \leftarrow \text{Verify}(vk, \tau, m, \sigma) \end{array} : b = 1 \right] \geq 1 - \varepsilon(\kappa).$$

⁶Technically, all five algorithms take 1^κ as an input, and \mathcal{M} and \mathcal{T} may be parametrized by κ . For brevity, we leave this implicit except in KeyGen.

2. Unforgeability: For any PPT \mathcal{A} , there is a negligible function ε such that for all $\kappa \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{KeyGen}(1^\kappa) \\ (\tau, m, \sigma) \leftarrow \mathcal{A}^{\mathcal{S}_{sk}, \mathcal{E}_{sk}}(vk) \\ b \leftarrow \text{Verify}(vk, \tau, m, \sigma) \\ b' = \tau \notin Q'_E \wedge (\tau, m) \notin Q_S \end{array} : b = b' = 1 \right] \leq \varepsilon(\kappa),$$

where \mathcal{S}_{sk} and \mathcal{E}_{sk} respectively denote oracles $\text{Sign}(sk, \cdot, \cdot)$ and $\text{Expire}(sk, \cdot)$, Q_S and Q_E denote the sets of queries made by \mathcal{A} to the respective oracles, and $Q'_E = \bigcup_{T \in Q_E} T$.

3. Forgeability on expiry: For all $m \in \mathcal{M}, T \subseteq \mathcal{T}$, for any $\tau \in T$, there is a negligible function ε such that for all κ ,

$$\Pr \left[\begin{array}{l} (vk, sk) \leftarrow \text{KeyGen}(1^\kappa) \\ \eta \leftarrow \text{Expire}(sk, T) \\ \sigma \leftarrow \text{Forge}(\eta, \tau, m) \\ b \leftarrow \text{Verify}(vk, \tau, m, \sigma) \end{array} : b = 1 \right] \geq 1 - \varepsilon(\kappa).$$

Generalized expiry policies Definition 3 allows for more complex *expiry policies* than simply expiring a sequence of timestamps in order. For example, it may be desirable to expire certain types of email more quickly than others (say, email internal to an organization which is likely to be delivered faster, or email with particularly sensitive content). Alternatively, it may be desirable to expire email not only according to its timestamp but also based on other metadata (say, a recipient whose email server is revealed to have been compromised). Our FFS definition would support expiry on the basis of any attribute which is recorded in the tag: for example, tags might be of the form $\tau = (t, \delta)$ where t is a timestamp and δ is the delay until expiration, or alternatively, $\tau = (t, R)$ where R is the recipient. Moreover, even when tags have the form $\tau = (t, \cdot)$, our construction still allows for efficiently expiring based solely on timestamp: that is, expiring “all tags that begin with t .”

FFS in the publicly verifiable timekeeper model Recall the *publicly verifiable timekeeper* (PVTk) model, in which a reliable timekeeper periodically issues timestamps whose authenticity is publicly verifiable (e.g., because the authority signs each timestamp). We show that in this model, expiration of signatures may occur “automatically” over time: that is, it is not necessary for the signer to publish any additional expiry information in order for the signature to become forgeable after a delay. In this case, the algorithm `Expire` is unnecessary, and `Forge` need not take η as input. Our construction of FFS in the PVTk model is given in Section 5.2.

Difference with forward-secure signatures Both FFS and FSS yield a system of short-lived secret keys all corresponding to one long-lived public key. However, the definitions differ in two main ways, described below and depicted in Figure 2.

1. Forward-secure signatures require that past keys cannot be computed from future keys, whereas forward-forgeable signatures require that future keys cannot be computed from past keys.
2. Forward-secure signatures are specifically designed to protect against compromise of all secret key material, so the each key must be derivable based solely on the previous short-lived key. Forward-forgeable signatures, in contrast, may have persistent “master secret key” material that is used to generate each short-lived key.




Figure 2: Forward-secure vs. forward-forgeable signatures

4.2 Succinctness

Next, we define *succinctness* of FFS, a measure of the efficiency of disclosure in terms of the size of expiry info per tag expired. KeyForge uses a construction of FFS based on hierarchical identity-based signatures (Section 4.3), which achieves logarithmic succinctness.

Definition 4. Let $z : \mathbb{N} \rightarrow \mathbb{N}$. Let $S \subset \mathbb{P}(\mathcal{T})$ be a set of sets of tags. A forward-forgeable signature scheme Σ is (S, z) -succinct if for any $T \in S$, there is a negligible function ε such that for all κ ,

$$\Pr_{(vk, sk) \leftarrow \text{KeyGen}(1^\kappa)} \left[|\text{Expire}(sk, T)| \leq z(|T|) \right] \geq 1 - \varepsilon(\kappa) .$$

Remark 1. Definition 4 is a worst-case definition: it guarantees the size of expiry information with overwhelming probability. In certain applications, an average-case definition may be appropriate instead, i.e., defining succinctness by bounding the size on *expectation*. We use a worst-case definition since it is stronger than an average-case definition, and our construction achieves it.

4.3 FFS Construction from (Hierarchical) IBS

We first outline a simple FFS construction BasicFFS based on identity-based signatures (IBS) [45], as a stepping stone to our main construction from hierarchical IBS (HIBS). The next paragraph assumes familiarity with standard IBS terminology; readers unfamiliar with IBS may skip to the main construction which is explained formally with explicit syntax definitions.

Let tags in the FFS correspond to identities in the IBS. BasicFFS.KeyGen outputs IBS master keys. The BasicFFS signing and verification algorithms for tag τ respectively invoke the IBS signing and verification algorithms for identity τ . BasicFFS.Expire outputs the secret key for each input tag $\tau \in T$, and BasicFFS.Forge uses the appropriate secret key from the expiry information to invoke the IBS signing algorithm. The simple solution has linear succinctness. By leveraging hierarchical IBS (HIBS), our main construction achieves logarithmic succinctness, as described next.

Definition 5. A hierarchical identity-based signature scheme HIBS is implicitly parametrized by message space \mathcal{M} and identity space $\mathcal{I} = \{\mathcal{I}_\ell\}_{\ell \in \mathbb{N}}$, and consists of four algorithms

$$\text{HIBS} = (\text{Setup}, \text{KeyGen}, \text{Sign}, \text{Verify})$$

with the following syntax.

- Setup(1^κ) takes as input a security parameter⁷ and outputs a master key pair (mvk, msk) .

⁷Technically, all four algorithms take 1^κ as an input, and \mathcal{M} and \mathcal{I} may be parametrized by κ . For brevity, we leave this implicit except in Setup.

- $\text{KeyGen}(sk_{\vec{id}}, id)$ takes as input a secret key $sk_{\vec{id}}$ for a tuple of identities $\vec{id} = (id_1, \dots, id_\ell) \in \mathcal{I}_1 \times \dots \times \mathcal{I}_\ell$ and an additional identity $id \in \mathcal{I}_{\ell+1}$ and outputs a signing key $sk_{\vec{id}, id}$ where $\vec{id}' = (id_1, \dots, id_\ell, id)$. The tuple may be empty (i.e., $\ell = 0$): in this case, $sk_{()} = msk$.
- $\text{Sign}(sk_{\vec{id}}, m)$ takes as input a signing key $sk_{\vec{id}}$ and a message $m \in \mathcal{M}$, and outputs a signature σ .
- $\text{Verify}(mvk, \vec{id}, m, \sigma)$ takes as input master verification key mvk , tuple of identities \vec{id} , message $m \in \mathcal{M}$, and signature σ , and outputs a single bit indicating whether or not σ is a valid signature with respect to mvk , \vec{id} , and m .

A depth- L HIBS is a HIBS where the maximum length of identity tuples is L , i.e., the identity space is $\mathcal{I} = \{\mathcal{I}_\ell\}_{\ell \in [L]}$.

Definition 6. For an identity space $\mathcal{I} = \{\mathcal{I}_\ell\}_{\ell \in \mathbb{N}}$, we say \vec{id} is a level- ℓ identity if $\vec{id} \in \mathcal{I}_1 \times \dots \times \mathcal{I}_\ell$. For any $\ell' > \ell$, let \vec{id} is a level- ℓ identity and \vec{id}' is a level- ℓ' identity. We say that \vec{id}' is a sub-identity of \vec{id} if \vec{id} is a prefix of \vec{id}' . If moreover $\ell' = \ell + 1$, we say \vec{id}' is a immediate sub-identity of \vec{id} .

Deriving subkeys Observe that given a master secret key of a HIBS, it is possible to derive secret keys corresponding to level- ℓ identities for any ℓ , by running KeyGen ℓ times. By a similar procedure, given any secret key corresponding to a level- ℓ identity \vec{id} , it is possible to derive any “subkeys” thereof, i.e., secret keys for sub-identities of \vec{id} . For our construction, it is useful to name this procedure and refer to it directly: we define HIBS.KeyGen^* in Algorithm 1. We write the randomness ρ_1, \dots, ρ_ℓ of HIBS.KeyGen^* explicitly.

Algorithm 1 HIBS.KeyGen^*

Input: $sk, \ell, \vec{id} = (id_1, \dots, id_\ell)$ \triangleright Require: $\ell \leq \ell'$
Randomness: $\rho_1, \dots, \rho_{\ell'}$
for $j = \ell + 1, \dots, \ell'$ **do**
 $sk \leftarrow \text{HIBS.KeyGen}(sk, id_j; \rho_j)$
end for
return sk

HIBS security requirements Due to space constraints, Definition 5 gives only the syntax and not the security guarantees of a HIBS. Informally, an HIBS must satisfy the following security guarantees. For a formal security definition, see, e.g., [25].

- *Correctness:* For any identity tuple \vec{id} , an honestly produced signature w.r.t \vec{id} must verify as valid w.r.t. \vec{id} .
- *Unforgeability:* For any PPT adversary \mathcal{A} that has access to a $\text{KeyGen}(msk, \cdot, \cdot)$ oracle, the probability that \mathcal{A} outputs a valid signature w.r.t. an identity $\vec{id} \notin Q$ must be negligible, where Q is the set of all sub-identities of identities \mathcal{A} has queried to the oracle.

Succinctly representing expiry information Given any set T of tuples of identities, the simplest way to make signatures with respect to T forgeable would be release the secret key corresponding to each $\vec{id} \in T$, much as in BasicFFS :

$$\eta = \left\{ sk_{\vec{id}} = \text{HIBS.KeyGen}^*(msk, 0, \vec{id}) \right\}_{\vec{id} \in T} . \quad (1)$$

Algorithm 2 Compress

Input: $\mathcal{I} = \{\mathcal{I}_\ell\}_{\ell \in [L]}$, $T \subseteq \mathcal{I}_1 \times \cdots \times \mathcal{I}_\ell$
// Remove redundant sub-identities
for all $\tau \in T$ **do**
 if $\exists \tau' \in T$ s.t. τ' is a prefix of τ **then**
 $T = T \setminus \{\tau\}$
 end if
end for
// Replace identities with prefix identities where possible
for $\ell = L - 1, \dots, 1$ **do**
 for all $\vec{\tau} = (\tau_1, \dots, \tau_\ell) \in \mathcal{I}_1 \times \cdots \times \mathcal{I}_\ell$ **do**
 // X represents all level- $(\ell + 1)$ sub-identities of $\vec{\tau}$
 $X = \{(\tau_1, \dots, \tau_\ell, \tau')\}_{\tau' \in \mathcal{I}_{\ell+1}}$
 if $X \subseteq T$ **then**
 $T = T \setminus X$
 $T = T \cup \{(\tau_1, \dots, \tau_\ell)\}$
 end if
 end for
end for
return sk

However, leveraging the hierarchical nature of HIBS, η can often be represented more succinctly than (1). Based on the fact that Algorithm 1 allows the derivation of any subkey, we make two optimizations. First, before computing (1), we can delete from T any $\vec{id} \in T$ that is a sub-identity of some $id' \in T$. Secondly, if there is any $\vec{id}' = (id_1, \dots, id_\ell) \in \mathcal{I}_1 \times \cdots \times \mathcal{I}_\ell$ such that every immediate subkey of \vec{id}' is in T , i.e.,

$$\forall id_{\ell+1} \in \mathcal{I}_{\ell+1}, (id_1, \dots, id_\ell, id_{\ell+1}) \in T,$$

then all sub-identities of id' can be removed from T and replaced by id' before computing (1). Such replacement is permissible only when *every* possible subkey of id' is derivable from T : otherwise, adding id' to T would implicate additional subkeys beyond those originally in T .

These two optimizations lead us to an algorithm **Compress**, which takes as input the identity space of a HIBS and a set of identity tuples T , and outputs a (weakly) smaller set of identity tuples T' such that knowledge of the secret keys corresponding to T' suffices to produce valid signatures with respect to exactly the identity tuples in T . Next, we describe the operation of **Compress** using a tree-based representation of identity tuples.

Tree representation It is convenient to think of identity tuples represented graphically in a tree. A node at depth ℓ represents a tuple of ℓ identities (considering the root node to be at depth 0). The set of all depth- ℓ nodes corresponds to the set of all ℓ -tuples of identities. The branching factor at level ℓ is $|\mathcal{I}_{\ell+1}|$. Given a secret key for a particular node (i.e., identity tuple), the secret keys of all its descendant nodes are easily computable using HIBS.KeyGen^* . (The secret key for the root node is the master secret key.) In this language, **Compress** simply takes a set T of nodes and returns the smallest set T' of nodes such that (1) all nodes in T are descendants of some node in T' and (2) no node not in T is a descendant of any node in T' . Figure 3 gives an illustration of the **Compress** algorithm on a small example tree.

Lemma 1 (Logarithmic succinctness). *Let HIBS be a depth- L HIBS with message space \mathcal{M} and identity space $\mathcal{I} = \{\mathcal{I}_\ell\}_{\ell \in [L]}$. For each $\ell \in [L]$, let \preceq_ℓ be a total order on \mathcal{I}_ℓ . Let $\mathcal{T}_L = \mathcal{I}_1 \times \cdots \times \mathcal{I}_L$. For $i \in |\mathcal{T}_L|$, let τ_i denote the i th element of \mathcal{T}_L in the lexicographic order induced by $\{\preceq_\ell\}_{\ell \in [L]}$. Construction 1 instantiated with HIBS is $(S, 2z)$ -succinct and also (S_1, z) -succinct, where*

$$S = \{\{\tau_i\}_{j' \leq i \leq j} : j, j' \in [|\mathcal{T}_L|]\} \quad \text{and} \quad S_1 = \{\{\tau_i\}_{1 \leq i \leq j} : j \in [|\mathcal{T}_L|]\} .$$

and $z(\cdot) = B \cdot \log_B(\cdot)$ where $B = \max_{\ell \in [L]} \{|\mathcal{I}_\ell|\}$. We assume B is constant.

Proof. Fix any $j, j' \in [|\mathcal{T}_L|]$ and any set $T = \{\tau_i\}_{j \leq i \leq j'}$. By the definition of succinctness, it suffices to show that the output of **Compress** on T , is a set of nodes of size at most $2B \cdot \log(|T|)$.

For any identity tuple $\iota \in \mathcal{I}$, let Sub_ι be the set of all level- L identities of which ι is a prefix. We say that T *covers* ι if $\text{Sub}_\iota \subseteq T$. Let Cover_T be the set of all identities covered by T . We say T *subsumes* an identity ι if ι is a descendant of some $\iota' \in \text{Cover}_T$ such that $\iota' \neq \iota$. By construction of Algorithm 2 (**Compress**), any identity *subsumed* by T will not be in the output set of **Compress**(T) (specifically, it will be removed in the innermost for-loop of Algorithm 2).

For any $\ell \in [L]$, consider any consecutive sequence of s level- ℓ identities covered by T . By definition of lexicographic ordering, there are fewer than $2B$ level- ℓ identities in the sequence that are not subsumed by T (these identities will be at the beginning and/or end of the sequence). Moreover, if the sequence begins at the smallest level- ℓ identity in the lexicographic order, then there are fewer than B identities in the sequence that are not subsumed by T (these identities will be at the end of the sequence).

It follows that for each $\ell \in [L]$, there are fewer than $2B$ level- ℓ identities in the output set of **Compress**(T). The number of levels is $L \leq \log_B(|\mathcal{T}_L|) \leq \log_B(T)$. Therefore, the size of the output set is at most $2B \log_B(T)$. Moreover, if $T \in S_1$, then the output set size is at most $B \log_B(T)$. \square

5 Our Protocol Proposals

Sections 5.1 and 5.2 respectively describe our two proposed systems KeyForge and TimeForge.

5.1 KeyForge

KeyForge consists of three components:

1. Replace the digital signature scheme used in DKIM (currently, RSA) with a succinct forward-forgeable signature scheme. Details are in Section 5.1.1.
2. Email servers periodically publish expiry information. Details are in Section 5.1.2.
3. A *forge-on-request protocol*. Details are in Section 5.1.3.

Note that the first two components already achieve *delayed universal forgeability* (discussed in Section 1.2).

5.1.1 FFS configuration for KeyForge

KeyForge uses the succinct FFS given by Construction 1 instantiated with a HIBS. Our implementation uses the proposed by Gentry and Silverberg [25] which is based on their ‘‘BasicHIDE’’ hierarchical identity-based encryption scheme, hereafter denoted by FFS. This [25] construction has the advantages of simplicity and low computation costs over many other schemes, and allows us to vary the height of the tree post-implementation.

KeyForge is based on an L -level tag structure, corresponding to identity space $\mathcal{I} = \{\mathcal{I}_\ell\}_{\ell \in [L]}$ where the level- L identities represent 15-minute time chunks spanning a 2-year period. We use the following intuitive 4-level configuration for the purpose of exposition, but as discussed Section 6, it is preferable for efficiency purposes to keep $|\mathcal{I}_\ell|$ equal for all $\ell \in [L]$.

$\mathcal{I}_1 = \{1, 2\}$	representing a 2-year time span
$\mathcal{I}_2 = \{1, \dots, 12\}$	representing months in a year
$\mathcal{I}_3 = \{1, \dots, 31\}$	representing days in a month
$\mathcal{I}_4 = \{1, \dots, 92\}$	representing 15-minute chunks of a day

A tag $\tau = (y, m, d, c) \in \mathcal{I}_1 \times \mathcal{I}_2 \times \mathcal{I}_3 \times \mathcal{I}_4$ thus corresponds to a particular 15-minute chunk of time. The 15-minute chunks are contiguous, consecutive, and disjoint, so that any given timestamp is contained in exactly one chunk. For a timestamp t , we write $\tau(t)$ to denote the unique 4-tuple tag (y, m, d, c) that represents a 15-minute chunk of time containing t . We write $t \sqsubset \tau$ if τ represents a chunk of time containing timestamp t .

KeyForge requires each DKIM signature at time t to be produced with respect to a tag representing timestamp $t + \hat{\Delta}$ (except when DKIM signatures are produced within the forge-on-request protocol, as discussed in Section 5.1.3). The tag is sent alongside the email, and is used for verification at the receiving server. The signing and verification algorithms are specified more precisely in Algorithms 3–4.

Algorithm 3 KeyForge.Sign

Input: sk, m, Δ
 $t = \text{CurrentTime}()$
return FFS.Sign($sk, \tau(t + \Delta), m$)

Algorithm 4 KeyForge.Verify

Input: vk, τ, m, σ
 $t = \text{CurrentTime}()$
return $t \sqsubset \tau$ **AND** FFS.Verify(vk, τ, m, σ)

Why 15-minute chunks? The time period that a leaf node represents — i.e., a “chunk” — corresponds to the maximum granularity at which expiry information can be released, as also discussed in Section 5.1.2. $\hat{\Delta}$ is a practical lower bound on the chunk size: since $\hat{\Delta}$ represents the time for an email to be delivered, publishing expiry information more than once per $\hat{\Delta}$ time does not make sense. 15 minutes is a conservative estimate of $\hat{\Delta}$.

Why a 2-year period? Rotating signing keys is good practice; the Messaging, Malware, and Mobile Anti-Abuse Working Group (M3AAWG) recommends rotation every 6 months [34]. Having signing keys that explicitly correspond to time periods, as in KeyForge, has the convenient side effect of encouraging regular key rotation.

For an FFS, rotation means generating a fresh master key pair. Thus, the leaves corresponding to any master key need not exceed the rotation time span. Recognizing that realistically, key rotation times often exceed 6 months, we run our evaluation assuming a 2-year period.

How many levels? We evaluate performance for different hierarchy depths; our evaluation is discussed in Section 6. We find that the best value of L depends on a tradeoff between computation time and succinctness of expiry information; the optimal value may thus depend on the specific application.

5.1.2 Publishing expiry information

KeyForge requires email servers to publish expiry information at regular intervals. A natural option is to publish expiry information every 15 minutes: that is, to publish the expiry information corresponding to each chunk c at the end of the time period that c represents.

Publishing every 15 minutes yields the finest granularity of expiry possible under the basic four-level tag structure. Based on an email server’s preference, it could release information at longer intervals (e.g., days). In case of an attack, an adversary would be able to convince third parties of the authenticity of all emails in the current interval (e.g., the current day), so the risk-averse approach is to prefer shorter intervals. If an email server prefers to use intervals even shorter than 15 minutes, the basic tag structure could easily be adapted for this.

Flexible expiry policies The basic tag structures described above can be augmented easily. For example, an additional level \mathcal{I}_ℓ might represent the “sensitivity” of an email: this would allow expiry information to be released with respect to highly sensitive emails more quickly. Alternatively, it may be desirable that certain emails expire more slowly or not at all (e.g., an email clearing an employee for a certain privilege for a year, or emails in the course of an important contract negotiation), and that is expressible in the sensitivity paradigm too: sensitivity can be expressed as the desired delay until expiry.

KeyForge is very configurable: beyond the first four levels, different email servers’ policies need not be consistent. The verification algorithm refers only to the first four levels of the tag (when checking $t \sqsubseteq \tau$), so a sending server can add more levels beyond the basic four without sacrificing compatibility, to match its desired expiry policy.

5.1.3 Forge-on-request protocol

The final component of KeyForge is a protocol by which email servers accept real-time requests for specified email content to be sent to the requester. Any client with the capability of sending outgoing mail can thus initiate a request.

The recipient of the requested email is required to be the requester: this is essential so the expiry request protocol cannot be used for spoofing and spam. To enforce this, receiving servers only respond to requests signed by the requester’s server, and address the response to the originating client. The request content that is signed by the requester’s server includes the identity of the originating client.¹⁰

Note that a forge-on-request protocol achieves a stronger guarantee than the use of ring signatures (i.e., signing with respect to both the sending and receiving servers’ public keys). A forge-on-request protocol enables any recipient with the ability to send mail from an email server to forge mail from any sender to herself. The ring signature approach enables her to do this only if she has the ability to sign fraudulent mail with the receiving server’s secret key.

¹⁰We note that a malicious server could, in principle, unauthorizedly sign requests for any client account it controls. This is outside our threat model, and is equally possible for outgoing mail under DKIM; see “Client-server trust” under Section 3.2 for further discussion.

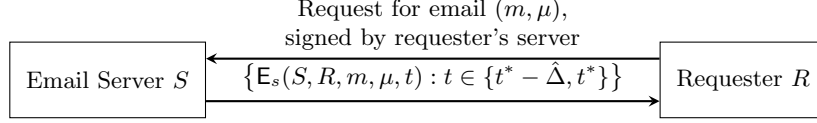


Figure 4: Forge-on-request protocol

The notation $E_s(S, R, m, \mu, t)$ is as defined in Section 3.3, and denotes an email (m, μ) from S to R at time t . t^* denotes the current timestamp when the reply is sent. Protocol messages are sent via email.

Theorem 2. *KeyForge is non-attributable for recipients (Definition 1) and $\hat{\Delta}$ -universally non-attributable (Definition 2).*

Proof. Follows directly from Lemmata 2 and 3. □

Lemma 2. *KeyForge is non-attributable for recipients (Definition 1).*

Proof. Recall from Definition 1 that we must show that there is a PPT simulator \mathcal{S} such that for any sender S and recipient R , for any email message m and metadata μ ,

$$\text{KeyForge}_{sk}(S, R, m, \mu, t) \approx_c \mathcal{S}^R(S, m, \mu), \quad (2)$$

where sk is the (master) secret key of S , t is the time at which \mathcal{S} is invoked, and the superscript R denotes that \mathcal{S} has access to the recipient's email server. We construct \mathcal{S} in Algorithm 5.

Algorithm 5 Simulator \mathcal{S} for recipient non-attributability

Input: S, m, μ
 $t = \text{CurrentTime}()$
send forge request (m, μ) to S
receive answer $\{e_0, e_1\}$
parse e_0, e_1 as emails w.r.t. tags τ_0, τ_1 respectively
if $\tau(t) = \tau_0$ **then return** e_0
else if $\tau(t) = \tau_1$ **then return** e_1
end if

By definition of $\hat{\Delta}$, we know S received the request at some time $t' \leq t + \hat{\Delta}$. Thus, by construction of KeyForge, the emails e, e' must be signed with respect to the tags $\tau(t' - \hat{\Delta}), \tau(t')$ (say, respectively). It follows that $\tau(t) = \tau(t' - \hat{\Delta})$ or $\tau(t) = \tau(t')$. Therefore, at least one of the if-conditions in Algorithm 5 must be satisfied, and \mathcal{S} always produces an output. By construction of the if-statements and the forge-on-request protocol, the output of \mathcal{S} is an email signed by S for a tag corresponding to timestamp t , as required by (2). Indeed, we achieve equality of distributions, rather than just indistinguishability. □

Lemma 3. *KeyForge is $\hat{\Delta}$ -universally non-attributable (Definition 2).*

Proof. Recall from Definition 2 that we must show that there is a PPT simulator \mathcal{S} such that for any sender S (with secret key sk) and recipient R , for any email message m , metadata μ , and timestamp t , the following holds at any time $\geq t + \hat{\Delta}$:

$$\text{KeyForge}_{sk}(S, R, m, \mu, t) \approx_c \mathcal{S}(S, R, m, \mu, t). \quad (3)$$

Let $\text{KeyForge}_{sk^*}^*$ be defined as identical to KeyForge_{sk} except that whenever KeyForge_{sk} invokes $\text{KeyForge}.\text{Sign}(sk, \cdot)$, $\text{KeyForge}_{sk^*}^*$ instead invokes $\text{HIBS}.\text{Sign}(sk^*, \cdot)$. Using this notation, we construct \mathcal{S} in Algorithm 6.

Algorithm 6 Simulator \mathcal{S} for $\hat{\Delta}$ -strong non-attributability

Input: S, R, m, μ, t
retrieve published expiry information η for S
for all $(\vec{\tau}, sk_{\vec{\tau}}) \in \eta$ **do**
 if $t \sqsubset \vec{\tau}$ **then**
 $sk^* \leftarrow \text{HIBS}.\text{KeyGen}^*(sk_{\vec{\tau}}, \ell, \tau(t); \vec{O}(\vec{\tau}))$
 return $\text{KeyForge}_{sk^*}^*(S, R, m, \mu, t)$
 end if
end for

Since \mathcal{S} is invoked at a time $\geq t + \hat{\Delta}$, and *KeyForge* prescribes publication of expiry information at the end of each time chunk of duration $\hat{\Delta}$, it holds that the expiry information η retrieved by \mathcal{S} includes expiry information with respect to time t . Therefore, the if-condition in Algorithm 6 will be satisfied for at least one element of η .¹¹

Recall that $\text{KeyForge}.\text{Sign}$ invokes $\text{FFS}.\text{Sign}$, which in turn invokes $\text{HIBS}.\text{Sign}$. By definition, if $t \sqsubset \vec{\tau}$ and η is expiry information with respect to sk , then sk^* as computed in Algorithm 6 is the same key used to invoke $\text{HIBS}.\text{Sign}$ (within $\text{FFS}.\text{Sign}$, which is within $\text{KeyForge}.\text{Sign}$) at time t . Therefore, the output distributions of $\text{KeyForge}_{sk^*}^*$ and KeyForge_{sk} are identical. It follows that \mathcal{S} satisfies (3). Again, we achieve equality of distributions, not just indistinguishability. \square

5.2 TimeForge

The main limitation of the *KeyForge* approach is that it requires signers to continuously release key material. Once this key material has been revealed and widely distributed, any party can formulate apparently-valid signatures. The need to widely distribute this key material can pose a practical challenge for users, who must depend on their provider to perform this task reliably. Failures to distribute key material may in practice limit the degree of deniability that a system can offer.

The *TimeForge* protocol uses a different approach, one that is designed to ensure that *any* party can forge signatures after the time period Δ has expired, even if the signer fails to release key material. Rather than depending on key release, *TimeForge* employs a new service that we term a *publicly verifiable timekeeper* (PVTk). This service, which can in practice be realized using various extant Internet systems, does not share any secret key material with the signing party. Instead, it is a global service that maintains a monotonically-increasing clock. The essential property of a PVTk is that, at any clock time t , any party with access to the service can obtain a publicly verifiable proof π_t that the current time is at least t . Simultaneously, the system ensures that it is infeasible for an attacker to forge such a proof at an earlier period.

Given a PVTk service, the intuition behind the *TimeForge* protocol is straightforward. Let M be some email message sent at time period t . The sender first signs each message using a standard SUF-CMA signature scheme to produce a signature σ . However, rather than using this signature

¹¹In fact, it will be satisfied for exactly one element of η , by construction of *Compress* which ensures that no timestamp is represented by more than one element.

directly, she instead authenticates the message using a witness indistinguishable and non-interactive proof-of-knowledge (PoK) of the following informal statement:

I know a valid sender signature σ on M
 OR
 I know a valid PVTk proof π_{t+d} , for some $d \geq \Delta$.

Assuming a trustworthy PVTk service, this proof authenticates the message during any time period prior to $t + \Delta$. Once a PVTk proof $\pi_{t+\Delta}$ becomes generally available, the knowledge proof becomes trivial for any party to formulate. The witness indistinguishability property of the proof system ensures that a proof formulated by the valid signer cannot be distinguished from a forgery constructed using a revealed PVTk proof at a later time period.

Publicly verifiable timekeeping. As a building block for our main result, we define a publicly verifiable timekeeping (PVTk) scheme.

- $\text{TK.Setup}(1^\lambda)$ takes an adjustable security parameter λ and outputs a set of public parameters $params$ and a trapdoor sk .
- $\text{TK.Prove}(sk, t)$ takes as input sk and the current time epoch t , and outputs a proof π_t .
- $\text{TK.Verify}(params, t, \pi_t)$ on input $params$, a time period t , and the proof π_t , outputs whether π_t is valid.

Correctness and Security. Correctness for a PVTk scheme is straightforward. For security, we consider an experiment in which an adversary with access to an oracle that provides PVTk proofs for arbitrary time periods t , must not be able to produce a valid proof for some time period $t_{\max} + \Delta$ (except with negligible probability) where t_{\max} is the largest value queried to the oracle during the experiment, and $\Delta > 0$ is a constant chosen as an input to the experiment. If a PVTk scheme satisfies this definition for some $\Delta > 0$, then we refer to it as Δ -PVTk secure.

In practice, we note that for any $\Delta > 0$, a Δ -PVTk scheme can be realized using an SUF-CMA-secure signature scheme, where TK.Prove and TK.Verify are respectively implemented using the signing and verification algorithms on message t . We leave the problem of constructing and proving weaker schemes to future work.

Realizing a PVTk service. A simple PVTk system can be constructed using a single server that maintains a monotonically-increasing clock, and implements the PVTk protocol by signing the current time period using an SUF-CMA signature scheme. While this solution is conceptually simple, practically deploying it at the required scale is likely to be quite costly. Moreover, such a system would be highly vulnerable to denial of service and network-based attacks.

A better approach would be to construct a PVTk system from *existing* Internet services, which have already been deployed at scale. Next, we outline several proposals.

OCSP servers. The Online Certificate Status Protocol in its “stapling” configuration [41] allows TLS servers to obtain a standalone, signed certificate validity message from a Certificate Authority (CA). This message is cryptographically signed by the CA, and contains a timestamp intended to guarantee freshness of the data within a ten day expiry window. OCSP staple generation can therefore be viewed as an organic implementation of a PVTk protocol. To avoid the need for reliance on a single CA’s infrastructure, users can define the proof π_t to comprise *multiple* valid staples, *e.g.*, one each from any k out of N chosen CAs. These parameters, as well as the CA identities, can be selected as part of the setup algorithm.

Centralized services: CT and randomness beacons. The Certificate Transparency protocol consists of a centrally-managed and publicly verifiable log for recording the issuance of certificates [35]. Each CT log entry is signed by the log operator (*e.g.*, Google). While CT is not intended as a timestamping protocol, signed CT log entries contain timestamps and may be re-purposed to implement a centralized PVTk service without the need to deploy new infrastructure. Along similar lines, NIST operates a randomness beacon [39] that also distributes signed records containing a current timestamp. While any single centralized service may be unreliable or subject to attack, a (fault-tolerant) combination of these extant services can be used to construct a “composite” PVTk system.

Proof of work blockchains. A number of cryptocurrencies use *proof of work* blockchains to construct an ordered transaction ledger [37, 51]. In these systems, generating each new block involves substantial computational effort by a large network of computers, using parameters tuned to produce new blocks at some chosen average rate. These ledgers can be used as a restricted form of PVTk system, in which *params* comprises some initial block header B_s , and π_t comprises a list of block headers $\{B_{s+1}, \dots, B_t\}$ drawn from the blockchain. Verification involves checking each proof of work and the correct construction of the chain. This approach does not produce an exact timekeeping service (due to the probabilistic nature of block intervals, and difficulty adjustments), nor does it guarantee cryptographic unforgeability, because chains can be forged by an attacker who controls a substantial fraction of the network’s hash power. However, these threats are unlikely to be a significant problem over the short time intervals used in the TimeForge scheme.¹²

VDFs and puzzles. Cryptographic “puzzles” are mathematical problems that require a known (or statistically predictable) number of computational operations to solve. Typical puzzles include cryptocurrency proof of work systems and timelock encryption schemes based on brute-forcing cryptographic keys [43]. A related primitive, the *Verifiable Delay Function* [11, 42] creates a *sequential* puzzle that requires a precisely-known amount of work to solve, and allows the solver to produce a proof of the solution’s correctness. While puzzles and VDFs do not directly allow for the creation of a PVTk system, they enable a related primitive: at time t a sender may generate a puzzle challenge M (*e.g.*, the contents of an email message) such that a proof $\pi_{t+\Delta}$ can be found by applying a computational process to M in expected time approximately Δ . β

A basic TimeForge signature scheme. Using the PVTk primitive above, we can now formally describe our basic TimeForge signature scheme. This initial scheme provides only one of our main deniability guarantees: Δ -universal non-attributability (Definition 2). That is, it allows forgery by any party after a fixed time period. Below, we discuss how to add *recipient non-attributability* as an additional feature.

The TimeForge scheme consists of three algorithms: TF.Keygen, TF.Sign and TF.Verify, as well as a specialized *forgery* algorithm TF.Forge. We assume the existence of a PVTk scheme with parameters *params* and an SUF-CMA signature algorithm Sig.

- TF.Keygen($1^\lambda, params$). Run Sig.Keygen(1^λ) to generate (pk, sk) and output $PK = (pk, params)$, and $SK = sk$.
- TF.Sign(PK, SK, M, t, Δ). Parse $PK = (pk, params)$. On input a message M and a

¹²A related development, the *proof-of-stake* blockchain [7], employs signatures by currency holders to authenticate new blocks. These techniques might provide an alternative to PoW blockchains for instances where Δ is small.

time period t , first compute $\sigma \leftarrow \text{Sig.Sign}(SK, M||t||\Delta)$ and compute the following witness-indistinguishable (WI) non-interactive PoK:¹³

$$\Pi = \text{NIPoK}\{(\sigma, s, \pi) : \text{Sig.Verify}(pk, \sigma, M||t||\Delta) = 1 \vee (\text{TK.Verify}(params, \pi, s) = 1 \wedge s \geq t + \Delta)\}$$

Note that the prover can produce this proof using (σ, \perp, \perp) as the witness. Output $\sigma_{\text{tf}} = (\Pi, t, \Delta)$.

- $\text{TF.Verify}(PK, M, \sigma_{\text{tf}})$. Parse $PK = (pk, params)$ and $\sigma_{\text{tf}} = (\Pi, t, \Delta)$ and verify the proof Π with respect to the public values t, Δ, pk, M . If the proof verifies, output 1, else output 0.

We now describe the forgery algorithm TF.Forge . This forgery algorithm takes as input a PVTk proof π_s for some time period $s \geq t + \Delta$:

- $\text{TF.Forge}(PK, M, t, s, \Delta, \pi_s)$. Parse $PK = (pk, params)$ and compute the NIPoK Π described in the TF.Sign algorithm, using the witness (\perp, s, π_s) . Output $\sigma_{\text{tf}} = (\Pi, t, \Delta)$.

Defining Security. Security for TimeForge is defined according to the following experiment. This experiment can be considered a variant of (weak) UF-CMA security definition for a signature scheme: an attacker must attempt to forge a TimeForge proof over a message M that she has not previously queried to a signing oracle. To assist in this, the attacker is given access to not one, but two oracles. The first is a signing oracle for the TimeForge signature scheme, and produces valid signatures for tuples of the form (M', t', Δ') . The main difference from the standard UF-CMA experiment is the existence of a second oracle that models the PVTk service. To model this, the attacker additionally obtains PVTk parameters $params$ at the start of the experiment, and may repeatedly query the PVTk oracle on chosen epoch numbers s to obtain PVTk proofs of the form π_s . Let s_{\max} be the largest time period queried to the PVTk oracle at the conclusion of the experiment. We say the attacker *wins* iff she outputs a message M and valid TimeForge proof $\sigma_{\text{tf}} = (\Pi, t, \Delta)$ such that $s_{\max} < t + \Delta$. where (M, t, Δ) was not queried to the signing oracle. We say that a TimeForge scheme is unforgeable under *chosen timestamp attacks* if \forall p.p.t. attackers \mathcal{A} , the adversary has at most a negligible advantage in succeeding at the above experiment.

Remark 2. We note that the definition above does not prevent the attacker from constructing forgeries that are “outside the expiration period.” Specifically, an intermediary can intercept a message embedding (M, t, Δ) where $t + \Delta$ is in the future, and author a new message M', t', Δ' where $t + \Delta$ has already been proved by the PVTk oracle. This is explicitly allowed by TimeForge; indeed, it is a goal of the system.

Theorem 3. *If the PVTk service is constructed using an SUF-CMA signature scheme; the WI proof system is sound (extractable); and the underlying signature scheme used by TimeForge is SUF-CMA, then the basic TimeForge scheme defined above is secure under chosen timestamp attacks.*

Proof. Our proof proceeds by contradiction. Let \mathcal{A} be an attacker that succeeds with non-negligible advantage in the chosen timestamp experiment. We show how to construct a pair of algorithms $\mathcal{B}_1, \mathcal{B}_2$ such that that one of the two algorithms (respectively) succeeds with non-negligible advantage in the SUF-CMA game against (1) the signature scheme Sig , or (2) the underlying PVTk signature scheme. We now describe the operation of each algorithm.

¹³Here we use Camenisch-Stadler notation, where the witness values are in parentheses () and any remaining values are assumed to be public.

An attack on the signature scheme Sig. In this strategy we construct \mathcal{B}_1 , which conducts the SUF-CMA experiment for the underlying signature scheme Sig . \mathcal{B}_1 first obtains a public key pk from the SUF-CMA challenger. It next uses the PVTk signature scheme's key generation algorithm to produce a keypair $(params, sk_{\text{PVTk}})$ for the PVTk service and sends $PK = (pk, params)$ to \mathcal{A} .

Each time \mathcal{A} queries the PVTk oracle on some timestamp s , \mathcal{B}_1 implements TK.Prove by using sk_{PVTk} to sign s and return the signature π_s . Whenever \mathcal{A} queries the TimeForge signing oracle on some tuple (M, t, Δ) , \mathcal{B}_1 first queries the SUF-CMA signing oracle to obtain a signature σ on $M||t||\Delta$. It then constructs a TimeForge signature by constructing the proof described in the TF.Sign algorithm using σ as the satisfying witness, and returns σ_{tf} to \mathcal{A} .

When \mathcal{A} outputs a pair $(M^*, \sigma_{\text{tf}}^*)$ that satisfies the win conditions of the experiment, \mathcal{B}_1 parses σ_{tf}^* to obtain (Π^*, t^*, Δ^*) and runs the extractor on Π^* to obtain the witness (σ^*, s^*, π^*) . (If the extractor fails, \mathcal{B}_1 aborts.) If σ^* is a valid signature on $M^*||t^*||\Delta^*$, it outputs the pair (σ^*, M^*) as an SUF-CMA forgery.¹⁴

An attack on the PVTk scheme. In this strategy we construct \mathcal{B}_2 , which conducts the SUF-CMA experiment against the PVTk signature scheme. This algorithm proceeds as in Strategy 1, except that here we set $params$ to be the public key obtained from the SUF-CMA challenger, and generate the keypair $(pk, sk) = \text{Sig.Keygen}(1^\lambda)$. Queries to the PVTk oracle are answered by forwarding s to the SUF-CMA oracle and returning the resulting signature as π_s , and queries to the TimeForge oracle are answered honestly by running TF.Sign with sk as an input.

As in the previous strategy, if \mathcal{A} succeeds in the experiment we run the extractor to obtain σ^*, s^*, π^* . However, in this case \mathcal{B}_2 verifies that π^* is a valid signature on s^* and, if so, outputs (π^*, s^*) as a forgery for the SUF-CMA experiment.

Analysis. We first observe that, in the view of \mathcal{A} , the simulation produced by \mathcal{B}_1 and \mathcal{B}_2 are distributed identically. Thus \mathcal{A} must succeed with identical advantage when interacting with either simulation. Next, we point out that if the WI knowledge extractor must fail with probability at most negligible in the security parameter under the assumption that the WI proof system is sound. Thus both adversaries will abort with at most negligible probability due to extraction error.

From these observations it remains only to show that at least one of the two algorithms above must succeed with non-negligible advantage in the SUF-CMA experiment. This is shown by observing that the attacker's output (Π^*, t^*, Δ^*) and the extracted witness (σ^*, s^*, π^*) must necessarily satisfy the conditions that (1) for all PVTk queries s made during the experiment, it holds that $s < t^* + \Delta^*$ (by the requirements of the experiment), (2) the message $M^*||t^*||\Delta^*$ has not been queried to the SUF-CMA signing oracle (by the requirements of the experiment), and (3) the following conditions are true (by the soundness of the proof system):

$$\text{Sig.Verify}(pk, \sigma^*, M^*) = 1 \vee (\text{TK.Verify}(params, \pi^*, s^*) = 1 \wedge s^* \geq t^* + \Delta^*)$$

If \mathcal{A} succeeds in the TimeForge experiment with non-negligible advantage, then the extracted witness must contain a valid signature σ^* on a message $M^*||t^*||\Delta^*$ not queried to the signing oracle, **or** it must contain a valid signature π^* on some epoch s^* which was not queried to the PVTk oracle. By these requirements, it must be the case that one of \mathcal{B}_1 or \mathcal{B}_2 succeeds with non-negligible advantage, which concludes our proof. \square

Realizing the TimeForge proof system. TimeForge can be realized using a variety of WI and ZK proof systems, combined with efficient SUF-CMA signature schemes. For example, a number

¹⁴Note that by the restrictions on \mathcal{A} , M^* must represent a message that has not previously been queried to the SUF-CMA oracle.

of pairing-based signature schemes [9, 16, 8] admit efficient proofs of knowledge of a signature using simple Schnorr-style proofs [3]. More recent proving systems such as Bulletproofs [14] and zkSNARKs (e.g., [40, 26]) admit succinct proofs of statements involving arbitrary arithmetic circuits and discrete-log relationships. Using the latter schemes ensures short proofs, in the range of several hundred bytes, in some cases with a small, constant verification cost. Thus, even complex PVTk proofs such as block header sequences, can potentially be reduced to a succinct and readily-verified TimeForge signature.

Adding recipient non-attributability. The basic TimeForge construction above does not provide *recipient non-attributability* (Definition 1). This property allows a message recipient to forge a TimeForge signature prior to the time period $t + \Delta$. The scheme can be altered to achieve *non-attributability for recipient servers* by modifying the algorithms to take as input the recipient’s public key $pk_{\text{recipient}}$. The proof statement then becomes:

$$\Pi = \text{NIPoK}\{(\sigma, s, \pi) : \text{Sig.Verify}(pk, \sigma, M) = 1 \vee \\ \text{Sig.Verify}(pk_{\text{recipient}}, \sigma, M) = 1 \vee \\ (\text{TK.Verify}(params, \pi, s) = 1 \wedge s \geq t + \Delta)\}$$

This new proof statement ensures that the recipient can, at any point, generate a valid TimeForge signature for any message addressed to it, provided that it has the correct signing key.

Remark 3. KeyForge and the above-described version of TimeForge achieve different variants of recipient non-attributability under Definition 1. The recipient non-attributability of TimeForge is weaker than KeyForge in the following sense. In both KeyForge and TimeForge, a recipient client can forge with the help of the recipient server (and a recipient server can forge by itself). In KeyForge, the client needs the server’s help to sign an outgoing email from the client itself and then to receive a reply email — these are both features routinely offered by email servers. However, in TimeForge, the client needs the server’s help to sign an email purportedly authored by the party whom the client designates as the sender of the forged email. Since this is not a feature routinely offered by email servers (and to offer it naively would enable spam and spoofing), it could be difficult for a client to forge without the collusion of a malicious server: realistically, this could substantially undermine the goal of non-attributability. An interesting open question is whether TimeForge’s recipient non-attributability can be strengthened.

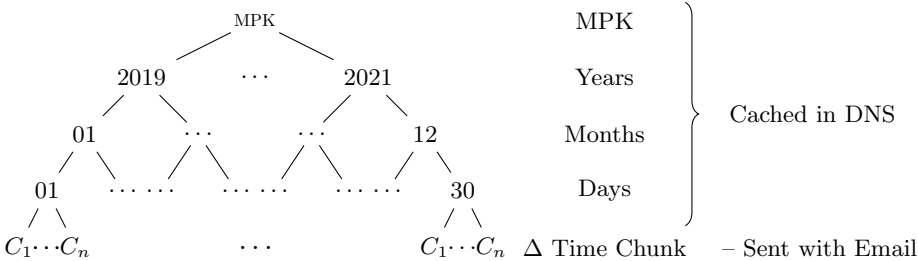


Figure 5: DNS / Email Cached Public Parameters

6 Implementation and Evaluation

We have designed and implemented a prototype of KeyForge suitable for use with common mail (MDA/MSA) servers. The prototype is fully functional and integrates with Postfix and Sendmail. It is designed to be straightforward to extend the system to work with other mail servers, as well as other applications, due to the abstraction of core cryptographic services away from other infrastructure. The entire project consists of roughly 2,000 lines of Go, C, and Python. Our implementation is available to the community open source at <https://github.com/keyforgery/KeyForge>. There are two main parts to KeyForge: a key server, and a mail filter that integrates KeyForge with mail server software, we discuss each in detail below.

Mail Filter: The mail filter is a service that is responsible for ensuring that sent emails are properly formatted, determining if incoming emails are verifiable, and communicating these results to the mail server. The filter works by intercepting emails that the mail server sends and receives, adding required metadata to the sent email’s header, and requesting cryptographic operations from the key server.

When sending an email, the filter adds the expiry time and any other information necessary to verify the signature to the email’s header. The mail filter performs a SHA256 hash of necessary metadata and the message’s content and forwards the hash to the key server to sign, and then adds this signature to the header.

On receipt of a signed email, this service confirms that the message metadata exists, signature’s hash is equivalent to the hash of the message and metadata to be verified. The filter then sends the signature, the sending domain, and the expiry timestamp to the key server for verification. If any step in the verification fails, the filter alerts the server software, and the message is dropped.

Key Server: The key server communicates with the mail filter via JSON RPC [2], a flexible standard that should allow for easy integration into any other mail system. The key server expires old keys by publishing them on a constantly updated website, though this is an implementation detail that could be augmented in any number of ways.

We use the RELIC[5] cryptographic library’s implementation of a 12-degree Barreto-Lynn-Scott curve with $q = 2^{381}$ (BLS12-381). This configuration conservatively yields keys with a 128-bit security level,[6] which puts our implementation on par with the standard 2048-bit RSA seen in current DKIM implementations.¹⁵ RELIC was chosen due to its support for many pairing friendly curves and computational overhead.

6.1 Evaluation


Bandwidth: Table 1 illustrates KeyForge’s bandwidth costs. The second row reflects that email headers and DNS records require that keys be base64 encoded. Our amortized bandwidth per email is 42% *smaller* than a DKIM RSA 2048 signature.

A difficulty with most HIBS schemes is that their signatures can be quite large. For example, in our modified Gentry-Silverberg scheme, a signature created from a particular chunk must include public parameters for each node on the path from that chunk to the Master Public Key. A public

¹⁵We originally considered using a 256-bit Barreto-Naehrig curve, which had a native Go implementation and was close to becoming an IETF standard. Unfortunately, new attacks[32] lowered the curve’s security rating to below the RSA-equivalent 128-bits.

	Secret Keys	Public Params	KeyForge σ (w/ cache)	DKIM RSA2048 σ
Raw Size (B)	64	97	$49 + 97 = 146$	256
Base64 Size (B)	88	132	201	347

Table 1: Amortized bandwidth costs of various parts of the system, as well as DKIM with RSA



(a) Size of expiry data given different tree topologies.

(b) The time it takes to perform a Sha256 hash of particular messages in the Podesta Email Corpus

(c) Time required to sign and verify given whether or not the tree was allowed to pre-cache or is calculating based off of a brand new tree “cold.”

Figure 6: Evaluation of Bandwidth and Computation Speed

parameter in our scheme is 97 bytes, so the parameters required for a signature from a tree structured with Y/M/D/Chunk would wind up being $97 * 4 = 388$ bytes. The total size also includes the signature point (49 bytes), for a total of 437 bytes.

We further decrease the bandwidth and storage costs our implementation by precomputing the path parameters and storing them in DNS along with the Master Public Key (Figure 5). At verification, the server performs a DNS lookup for the appropriate expiry parameters, and stores them in a local cache. This yields a one-time amortized cost of 4kb per server per month.

Computation Costs: We created a number of microbenchmarks to determine the efficiency of our system, as well as the impact of varying the height of the tree. In particular, evaluated the impact of tree depth on signing and verification time. It is important to note that this is a lower bound on throughput – all benchmarks were performed on a laptop with power lower than a common server.¹⁶

To fully capture the best and worst case for signing and verification, we timed execution given a tree where the public key path from the root must be recalculated each time (as would be common in standing up a brand new server, or initially verifying messages from a new domain) and when the tree is allowed cache such parameters; results are illustrated in Figure 6c. We find that signing is largely unaffected by the depth of the tree with caching enabled, which would likely be the case in the majority of settings. Further, we find that increasing the depth of the tree has a linear impact on verification time.

Cached verification with a depth-4 tree (as in our Y/M/D/Chunk tree) takes an average of 0.031579 seconds per verification operation; one laptop alone could verify roughly 2,735,995 messages per day. Cached signing for the same tree takes 0.001388 seconds per signing operation, which means that same computer could sign roughly 62,247,838 messages per day.

Signing and verification times are unaffected by plaintext size, since the key server performs

¹⁶For full details on compilation and machine specifics see appendix A

Table 2: Expiry information size

L	Expiry info size (bytes)			
	1 year		2 years	
	Avg	Max	Avg	Max
1	1121248	2242496	1679814	4485056
2	12700	25344	16934	33792
3	3283	6464	3920	7744
4	1787	3520	2016	3968
5	1275	2496	1408	2752
6	1048	2048	1117	2176
7	859	1664	934	1792

operations only on SHA256 hashes of the message. Since hash-and-sign is the common paradigm between DKIM and KeyForge, both incur these same costs. Figure 6b illustrates hashing time using OpenSSL’s hashing library for different message sizes, using a corpus of real-world leaked emails from Wikileaks’ Podesta Email Corpus [49].¹⁷ We find that the average time to hash for all messages in the corpus is around 0.0102 seconds, indicating that I/O and hashing time are likely as impactful, if not more, than the time it takes to sign and verify.

Optimizing for Expiration Bandwidth: Succinct expiry information may be more or less of a priority depending on the desired application and associated cost in computation time. For example, succinct expiry information would allow for easy publication (even in DNS) or lightweight attachment to emails. The latter idea has the neat feature of guaranteeing that any person who received emails from this domain possesses expiry information.

While the Y/M/D/Chunk tree configuration is easy to intuit, having an equal branching factor across all tree levels yields a large gain in succinctness. Figure 6a illustrates the difference in size between the naive configuration and a depth 4 tree whose branching factor is equal at every level. Table 2 shows the average and maximum size of expiry info of various depth trees with an equal branching factor: e.g., the average expiry size for a full 2-year period is 4.5MB, 4KB, or 1.8KB for depths 1, 4, and 7 respectively. An administrator’s preferred trade-off between expiry size, depth (i.e., verification time), and validity time of the tree likely depends on the application at hand.

Optimizing the expiry time: While expiry time is a configurable parameter of KeyForge (i.e., an administrator may select any appropriate time), there are clear benefits to keeping this time as short as possible. Indeed, the amount of universal forgeability that a user can enjoy depends directly on the expiry time.

Despite that most emails are received very quickly, the SMTP RFC (RFC 5321) [33] has a very lax give-up time of 4–5 days. To get a rough idea of how quickly emails tend to be delivered, we computed the time differences from the first Received header to the last, in the Podesta email corpus, and found that of the 48,246 messages with parseable Received timestamps, 99% (47,349) took less than 15 minutes to be delivered.

We believe that this time can therefore be shortened considerably using any number of tactics. For example, delays are often caused by expected server outages on the receiving end, e.g., by

¹⁷Beyond the dramatic irony, we chose the Podesta email corpus because it was distributed intact with attachments, and thus relatively close to a realistic user’s email distribution, among publicly available datasets.

an administrator using the weekend to update email servers. This could be resolved by using a DMARC-like DNS record to signal that maintenance is happening and to hold messages until later.

Shortening the expiry time in the face of routing delays is made slightly trickier due to the presence of potentially malicious adversaries; providing flow control similar to that of TCP would certainly be systematically possible, but implementers must be careful to ensure that a malicious MTA is not attempting to keep the message unforgeable for undue time. Setting a hard-cutoff maximum here of one to two days would likely be advisable.

7 Related Work and Discussion

A number of works have considered the problem of *deniability* in various settings.

Deniable authentication. The cryptographic literature features many works on deniable signatures and authentication, including (but in no way limited to) [31, 22, 44, 38, 20]. Indeed, our constructions can be viewed as a practical instantiation of a deniable signature scheme, with tight systems-based requirements on the deniability properties. A much stronger primitive, *deniable encryption* [18] considers encryption schemes that can survive the later compromise of an encrypting device.

Messaging deniability. In the context of interactive messaging, the OTR protocol [12] uses MACs and key agreement protocols to ensure that message transcripts can be repudiated by either party to a conversation. Along similar lines, Unger and Grubbs considered deniable key exchanges for new messaging systems [48]. Adida *et al.* proposed “lightweight email signatures” for DKIM, based on ring signatures [3], but with more limited properties than we consider in this work. Recently, Gunn *et al.* noted that trusted hardware systems can be used to circumvent deniability in many protocols [28]. In the theoretical realm, Canetti *et al.*’s recent possibility result addressed a yet stronger deniability guarantee for interactive messaging, where transcripts can be “explained” as encrypting any plaintext of the adversary’s choice, even when both parties are coerced [17].

Message franking. Finally, a recent line of works [21, 27] considers the problem of *franking* in encrypted messaging systems. These protocols consider a different problem than we address in this work: they allow attribution of messages even after an arbitrary period, but only with the cooperation of a centralized service provider. One such system, based on a secret HMAC key, has been deployed by Facebook [23].

Complications from ARC. As mentioned in Section 2, ARC is experimental and widely unadopted as a standard, but would pose some minor complications for non-attributability if adopted, because it has third-party MTAs sign emails in transit. KeyForge, though presented as a modification of DKIM, is conceptually straightforward to extend to accommodate ARC, by having third-party signer MTAs to use a FFS for signing, publish expiry information, and offer a forge-on-request protocol, just like the MSA in KeyForge.

Acknowledgements

We are grateful to Jon Callas for his in-depth review of the applicability of our scheme to DKIM, and to Dan Boneh, Daniel J. Weitzner, John Hess, Bradley Sturt, and Stuart Babcock for their feedback. This work was supported in part by the William and Flora Hewlett Foundation grant 2014-1601.

References

- [1] Domain-based Message Authentication, Reporting, and Conformance (DMARC). <https://tools.ietf.org/html/rfc7489>.
- [2] JSON-RPC. <https://www.jsonrpc.org/>.
- [3] Ben Adida, Susan Hohenberger, and Ronald L. Rivest. Lightweight encryption for email. In Dina Katabi and Balachander Krishnamurthy, editors, *Steps to Reducing Unwanted Traffic on the Internet Workshop, SRUTI'05*. USENIX Association, 2005.
- [4] Ross Anderson. Invited lecture at the fourth annual acm conference on computer and communications security (ccs), 1997.
- [5] D. F. Aranha and C. P. L. Gouvêa. RELIC is an Efficient Library for Cryptography. <https://github.com/relic-toolkit/relic>.
- [6] Diego F. Aranha, Laura Fuentes-Castañeda, Edward Knapp, Alfred Menezes, and Francisco Rodríguez-Henríquez. Implementing pairings at the 192-bit security level. In *International Conference on Pairing-Based Cryptography*, pages 177–195. Springer, 2012.
- [7] Christian Badertscher, Peter Gaži, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18*, pages 913–930, New York, NY, USA, 2018. ACM.
- [8] Mira Belenkiy, Melissa Chase, Markulf Kohlweiss, and Anna Lysyanskaya. P-signatures and noninteractive anonymous credentials. In *TCC 2008*, 2008.
- [9] Mira Belenkiy, Melissa Chase, Markulf Kohlweiss, and Anna Lysyanskaya. Compact E-Cash and Simulatable VRFs Revisited. In *Pairing-Based Cryptography '09*, 2009.
- [10] Mihir Bellare and Sara K. Miner. A forward-secure digital signature scheme. In Michael J. Wiener, editor, *Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 1999, Proceedings*, volume 1666 of *Lecture Notes in Computer Science*, pages 431–448. Springer, 1999.
- [11] Dan Boneh, Joseph Bonneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part I*, pages 757–788, 2018.
- [12] Nikita Borisov, Ian Goldberg, and Eric Brewer. Off-the-record communication, or, why not to use pgp. In *Proceedings of the 2004 ACM Workshop on Privacy in the Electronic Society, WPES '04*, pages 77–84, New York, NY, USA, 2004. ACM.
- [13] Thomas Brewster. Russian ‘fancy bear’ hackers tainted their huge leaks with fake data, May 2017. <https://www.forbes.com/sites/thomasbrewster/2017/05/26/russian-dnc-hackers-planted-leaks-with-fake-data>.
- [14] B. Bunz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy (SP)*, volume 00, pages 319–338.

- [15] Jon Callas. [ietf-dkim] DKIM Key Sizes, October 2016. <http://mipassoc.org/pipermail/ietf-dkim/2016q4/017207.html> [<https://perma.cc/H3SR-YB44>].
- [16] Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In *Advances in Cryptology—CRYPTO 2004*, 2004.
- [17] Ran Canetti, Sunoo Park, and Oxana Poburinnaya. Fully bideniable interactive encryption. *IACR Cryptology ePrint Archive*, 2018:1244, 2018.
- [18] Rein Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. Deniable encryption. In Burton S. Kaliski, editor, *Advances in Cryptology — CRYPTO '97*, pages 90–104, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg.
- [19] D. Crocker. Internet Mail Architecture, 2009. <https://tools.ietf.org/html/rfc5598>.
- [20] Mario Di Raimondo and Rosario Gennaro. New approaches for deniable authentication. *Journal of Cryptology*, 22(4):572–615, Oct 2009.
- [21] Yevgeniy Dodis, Paul Grubbs, Thomas Ristenpart, and Joanne Woodage. Fast message franking: From invisible salamanders to encryptment. In *CRYPTO (1)*, volume 10991 of *Lecture Notes in Computer Science*, pages 155–186. Springer, 2018.
- [22] Cynthia Dwork, Moni Naor, and Amit Sahai. Concurrent zero-knowledge. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, pages 409–418, New York, NY, USA, 1998. ACM.
- [23] Facebook. Messenger Secret Conversations Technical Whitepaper. Available at https://fbnewsroomus.files.wordpress.com/2016/07/secret_conversations_whitepaper-1.pdf, July 2016.
- [24] Center for Strategic and International Studies (CSIS). Significant cyber incidents, 2018. <https://www.csis.org/programs/cybersecurity-and-governance/technology-policy-program/other-projects-cybersecurity>.
- [25] Craig Gentry and Alice Silverberg. Hierarchical ID-based cryptography. In Yuliang Zheng, editor, *Proceedings of ASIACRYPT 2002*, volume 2501 of *Lecture Notes in Computer Science*, pages 548–566. Springer, 2002.
- [26] Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016*, pages 305–326, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.
- [27] Paul Grubbs, Jiahui Lu, and Thomas Ristenpart. Message franking via committing authenticated encryption. In *CRYPTO (3)*, volume 10403 of *Lecture Notes in Computer Science*, pages 66–97. Springer, 2017.
- [28] Lachlan J. Gunn, Ricardo Vieitez Parra, and N. Asokan. Circumventing cryptographic deniability with remote attestation. Available at <https://eprint.iacr.org/2018/424>, January 2019.

- [29] HIPAA Journal. United hospital district phishing attack impacts 2,143 patients, 2019. <https://www.hipaajournal.com/united-hospital-district-phishing-attack-impacts-2143-patients/>.
- [30] Troy Hunt. The 773 million record "collection #1" data breach, 2019. <https://www.troyhunt.com/the-773-million-record-collection-1-data-reach/>.
- [31] Markus Jakobsson, Kazue Sako, and Russell Impagliazzo. Designated verifier proofs and their applications. In *Proceedings of the 15th Annual International Conference on Theory and Application of Cryptographic Techniques*, EUROCRYPT'96, pages 143–154, Berlin, Heidelberg, 1996. Springer-Verlag.
- [32] Taechan Kim and Razvan Barbulescu. Extended tower number field sieve: A new complexity for the medium prime case. In *Annual Cryptology Conference*, pages 543–571. Springer, 2016.
- [33] John Klensin. RFC5321: Simple Mail Transfer Protocol, 2008.
- [34] Kurt Andersen. M3aawg DKIM Key Rotation Best Common Practices | M3aawg, December 2013. <https://www.m3aawg.org/documents/en/m3aawg-dkim-key-rotation-best-common-practices>.
- [35] Ben Laurie, Adam Langley, and Emilia Kasper. Certificate Transparency. RFC 6962, June 2013.
- [36] Mike Masnick. The Clinton Campaign Should Stop Denying That The Wikileaks Emails Are Valid; They Are And They're Real. *TechDirt*, October 2016. <https://www.techdirt.com/articles/20161024/22533835878/clinton-campaign-should-stop-denying-that-wikileaks-emails-are-valid-they-are-theyre-real.shtml>].
- [37] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system, 2009.
- [38] Moni Naor. Deniable ring authentication. In Moti Yung, editor, *Advances in Cryptology — CRYPTO 2002*, pages 481–498, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [39] NIST. NIST randomness beacon. <https://www.nist.gov/programs-projects/nist-randomness-beacon>.
- [40] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE, May 2013. Best Paper Award.
- [41] Yngve N. Pettersen. The Transport Layer Security (TLS) Multiple Certificate Status Request Extension. RFC 6961, June 2013.
- [42] Krzysztof Pietrzak. Simple verifiable delay functions. *IACR Cryptology ePrint Archive*, 2018:627, 2018.
- [43] R. L. Rivest, A. Shamir, and D. A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, Cambridge, MA, USA, 1996.
- [44] Ronald L. Rivest, Adi Shamir, and Yael Tauman. How to leak a secret. In *ASIACRYPT*, 2001.

- [45] Adi Shamir. Identity-based cryptosystems and signature schemes. In G. R. Blakley and David Chaum, editors, *Proceedings of CRYPTO '84*, volume 196 of *Lecture Notes in Computer Science*, pages 47–53. Springer, 1984.
- [46] Jonathan Stempel and Jim Finkle. Yahoo says all three billion accounts hacked in 2013 data theft, 2017. Reuters Business News. <https://www.reuters.com/article/us-yahoo-cyber/yahoo-says-all-three-billion-accounts-hacked-in-2013-data-theft-idUSKCN1C8201>.
- [47] Michael Thomas. Requirements for a DomainKeys Identified Mail (DKIM) Signing Practices Protocol. <https://tools.ietf.org/html/rfc5016>.
- [48] Nik Unger and Ian Goldberg. Deniable key exchanges for secure messaging. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, pages 1211–1223, New York, NY, USA, 2015. ACM.
- [49] Wikileaks. WikiLeaks: DKIM Verification, nov 2016. <https://wikileaks.org/DKIM-Verification.html> [<https://perma.cc/H3SR-YB44>].
- [50] Wikileaks. Wikileaks: Search the DNC Database, July 2016. <https://wikileaks.org/dnc-emails/>.
- [51] Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151:1–32, 2014.

A Compilation and Evaluation parameters

We performed all benchmarks on a 2017 MacBook Pro, 15-inch, with an Intel 4-core 3.1GHz processor and 16GB of RAM.

Openssl was compiled with the following flags:

```
clang -I. -I.. -I../include -fPIC -fno-common
-DOPENSSL_PIC -DOPENSSL_THREADS -D_REENTRANT
-DDSO_DLFCN -DHAVE_DLFCN_H -arch x86_64 -O3
-DL_ENDIAN -Wall -DOPENSSL_IA32_SSE2
-DOPENSSL_BN_ASM_MONT
-DOPENSSL_BN_ASM_MONT5 -DOPENSSL_BN_ASM_GF2m
-DSHA1_ASM -DSHA256_ASM -DSHA512_ASM -DMD5_ASM
-DAES_ASM -DVPAES_ASM -DBSAES_ASM -DWHIRLPOOL_ASM
-DGHASH_ASM -DECP_NISTZ256_ASM
```

RELIC was compiled with the following flags:

```
cmake ../ -DALLOC=DYNAMIC -DFP_PRIME=381
-DARITH=gmp-sec -DWSIZE=64
-DFP_METHD="INTEG;INTEG;INTEG;MONTY;LOWER;SLIDE"
-DCOMP="-O3 -mtune=native -march=native"
-DFP_PMERS=off -DFP_QNRES=on
-DFPX_METHD="INTEG;INTEG;LAZYR"
-DEP_SUPER=off -DPP_METHD="LAZYR;OATEP"
```