

Transparent Key Integrity (TKI): A Proposal for a Public-Key Validation Infrastructure

Tiffany Hyun-Jin Kim, Lin-Shung Huang, Adrian Perrig, Collin Jackson, and Virgil Gligor

July 20, 2012

(revised November 26, 2012)

[CMU-CyLab-12-016](#)

[CyLab](#)

Carnegie Mellon University
Pittsburgh, PA 15213

Transparent Key Integrity (TKI): A Proposal for a Public-Key Validation Infrastructure

Tiffany Hyun-Jin Kim Lin-Shung Huang Adrian Perrig Collin Jackson Virgil Gligor
Carnegie Mellon University
{hyunjin1,adrian,virgil}@ece.cmu.edu {linshung.huang,collin.jackson}@sv.cmu.edu

ABSTRACT

Recent trends in public-key infrastructure research explore the trade-off between decreased trust in certificate authorities (CAs), the level of security achieved, the communication overhead (bandwidth and latency) for setting up a secure connection (e.g., verified via SSL/TLS), and the availability with respect to verifiability of public key information. In this paper, we propose TKI as a new public-key validation infrastructure, where we reduce the level of trust in any CA and increase the security by achieving increased robustness in the case of CA key compromise. Compared to other proposals, we reduce the communication overhead associated with certificate validation during the existing SSL/TLS connection handshake and provide site owners with an optional time window to review potentially malicious key changes. Our design deters CA misbehavior by using a public log that records all certificate events, thereby enabling CAs' accountability for their actions. TKI will help reduce the trust in the hundreds of currently trusted CAs, reduce exposure to CA compromise, and enhance the security of SSL/TLS connection establishment.

1. INTRODUCTION

Secure connection establishment on the Internet through SSL and TLS has been a tremendous success, as it is globally used for practically all secure web-based communication. Given that the security of the majority of network-based financial or commercial transactions relies on SSL/TLS, one would hope that its security is commensurate with its proliferation and importance.

Unfortunately, numerous attack vectors against SSL/TLS exist and recently several high-profile attacks have demonstrated its vulnerability in practice. The main weakness lies in the fact that current browsers trust hundreds of root Certificate Authority (CA) certificates, and a security breach of a single CA can compromise the security of sites protected by any one of the other CAs, powerfully illustrating the concept of weakest-link security. In fact, a maliciously issued certificate for a site can be used by an adversary to mount man-in-the-middle attacks on connections to that site. To demonstrate the large extent of the SSL universe in which the weakest link could occur, EFF's SSL Observatory reports that Microsoft IE and Mozilla Firefox trust 1482 different CA public keys held by 651 organizations located throughout the world [17].

An often-advanced but unjustified argument is that since CAs are in the business of managing cryptographic keys, they have secure processes in place to protect their keys. Regrettably, recent events have highlighted the inability of many CAs to keep their keys and certificate issuance processes secure. We list a few high-profile cases:

- VeriSign issued two false Microsoft ActiveX certificates in Jan-

uary 2001 [24]: "According to Microsoft, someone posing as a Microsoft employee tricked VeriSign, which hands out so-called digital signatures, into issuing the two certificates in the software giant's name on Jan. 30 and Jan. 31."

- According to a U.S. Securities and Exchange Commission (SEC) filing in October 2011, VeriSign was hacked successfully and repeatedly, in 2010 [27], potentially enabling attackers to impersonate any online entities.
- In March 2011, in an attack on a Comodo reseller, fake certificates were issued for: mail.google.com, www.google.com, login.yahoo.com, login.skype.com, addons.mozilla.org, and login.live.com [15, 29]. Comodo suggested that the attack originated from an Iranian IP address.
- In August 2011, news broke that DigiNotar, a Dutch CA, improperly issued a certificate for all Google domains to an external party [32]. It was claimed that as many as 250 false certificates for an unknown number of domains were released. It was reported that these certificates were used by the Iranian government to spy on Iranian citizens' communications with Google email during the month of August 2011.
- For the Stuxnet malware, two Taiwanese CAs' private keys were compromised, which the Stuxnet developers used to sign their malware [18].

Not all CA vulnerabilities become public; for example, without the SEC filing, the VeriSign breach may have been concealed [27]. Also, DigiNotar's problem was noticed only because a vigilant user posted the issue to a Gmail help forum.¹ Thus, an even larger number of CA-breaches may have occurred.

These examples demonstrate that CA breaches can result in real-world attacks. Besides CA-based attacks, SSL/TLS has other vulnerabilities, for example the fact that users will click through browser warnings in case of self-signed certificates [16]. In Syria, such an attack was used to mount a man-in-the-middle attack against Facebook [14], supposedly by the Syrian Telecom Ministry.

Addressing these problems is very challenging, as several seemingly conflicting requirements need to be satisfied. On one hand, adversarial events such as CA private key compromise or domain private key compromise need to be addressed. On the other hand, legitimate events such as switching to different CAs or key re-creation after private key loss need to be supported. For example, legitimate re-creation of a key pair and certificate after private-key loss may appear to be an impersonation attempt. Also, legitimately switching to a new CA to cease using a compromised CA that signs fraudulent certificates may also appear as a security problem instead of a solution. Hence, we aim to create a certificate infrastructure that can prevent adversarial attacks yet gracefully handle

¹<http://productforums.google.com/forum/#!category-topic/gmail/share-and-discuss-with-others/3J3r2JqFNTw>

legitimate key and certificate management events.

We combine several key observations to address these challenges. The first observation is that by letting a domain define which CAs it trusts, we can greatly reduce the immense trust base of the current PKI, where hundreds of global organizations need to be trusted without any tangible evidence of trustworthiness [9, 12, 19]. Next, we leverage globally visible directories (known as public log servers) that enable public integrity validation for directory information. Such public validation provides accountability for CA’s actions, and thus creates a deterrent against fraudulent CA activities. Finally, we observe that unplanned key changes are relatively infrequent, and we thus admit a domain-selectable time period for interested parties to review keys that change due to a private key loss or compromise of a domain.

To evaluate the security, availability, and efficiency of our public-key validation infrastructures and to enable comparison between different systems, we propose a set of new metrics. Specifically, we propose Duration of Compromise (DoC), Duration of Unavailability (DoU), and several efficiency metrics. The DoC metrics provide insight into the security of a system, measuring the impact of a compromise or loss of various credentials, such as the private key of CAs or domains. The DoU metrics measure availability, again depending on various events such as key compromise or loss. Finally, the efficiency metrics measure the overhead of operating the public-key validation infrastructure and the overhead of secure session establishment. We present these metrics in detail in Section 9.

We find that based on these metrics, currently proposed public-key validation infrastructures are inadequate, as they suffer from excessive DoC or DoU, or have unacceptable performance impact. In this paper, we propose the TKI (short for Transparent Key Integrity) system to address these challenges.

2. PROBLEM DEFINITION

The core problem we aim to address is the design of a new public key validation infrastructure that reduces the amount of trust placed in any one infrastructure component (e.g., CA), yet reduces the system’s attack surface and single points of failure. Additional goals are efficiency, and incentives for deployment for all associated parties. In this section, we first describe a list of desired properties, followed by assumptions, and the adversary model. Detailed metrics for evaluation are described in Section 9.

2.1 Desired Properties

- **Checks and balances:** The infrastructure should limit trust in any single party. Also, limited trust should be distributed over multiple parties to prevent a single point of failure. Furthermore, multiple parties should be able to monitor each other to detect misbehavior.
- **Brief compromise period:** Given the compromise of a private key of any trusted party, the time during which an attacker can successfully attack a legitimate client should be brief. This includes the compromise of CAs, public log servers, or domains.
- **Brief unavailability period:** After various events (benign or adversarial), the time during which a legitimate client (who is not under an attack) cannot verify a domain’s certificate should be brief. This includes when a domain’s certificate is newly registered or updated, and when the CAs’, log servers’, and domains’ private keys are compromised.
- **Trust agility:** Users can decide which entities they trust to form their root of trust, and they can modify their trust decisions at any time [25]. Furthermore, such changes should be made without

undue delay. We extend this notion of trust agility to domains, enabling domains to select their roots of trust and to modify their trust decisions at any time.

- **Privacy:** Clients should not reveal which entity/server they establish an SSL/TLS connection to any party other than the server.
- **Efficiency:** One of the most important efficiency requirements is to avoid increasing the latency of a SSL/TLS connection establishment – in particular, avoiding any additional round-trips to external servers. Moreover, no additional infrastructure servers should be needed.

2.2 Assumptions

In our approach of checks and balances, some trusted public entities continuously validate public log servers’ operation and publish/disseminate any misbehavior. We thus assume that these trusted entities do not collude with CAs or public log servers. In the context of the current Internet, the Electronic Frontier Foundation may be one entity that would play the role of a validator. Such validators exhibit the property of *security of the strongest link*, because as long as a single verifier is correct, misbehavior will be detected and publicly disseminated.

We assume that browsers store the authentic public keys of root CAs and the authentic public keys of public log servers.

2.3 Adversary Model

We consider an adversary whose main goal is to impersonate a victim web site that is using HTTPS. To achieve this goal, the adversary may compromise some CAs and log servers that the victim trusts. An adversary may be able to *temporarily gain control of some trusted CAs and log servers*, but can gain *long-term control of untrusted CAs and public log servers*. Gaining control means access to private signing keys.

3. BACKGROUND

Two main families of proposals to reduce trust in CAs exist that are related to our design: 1) certificate observatories, 2) timeline servers that provide public visibility into all certificate operations. We briefly provide an overview of these approaches.

3.1 Certificate Observatories

The first type to reduce the trust in CAs and prevent many of the attacks discussed in Section 1 is to create a public repository of SSL/TLS server certificates, and enable browsers to compare the key they have received (presumably from the server) with the observation of the observatory which is received over an integrity-protected connection. This integrity-protected connection is set up by embedding a root public key of the observatory system into the client, creating a PKI just to authenticate the relatively small number of observatory nodes.

Perspectives. The Perspectives system [3, 34] has globally distributed notary servers that contact known SSL / TLS servers once a day to fetch the current server’s certificate. These notary servers then *store the history of observed certificates and support queries into their database*. Perspectives offers a Firefox plugin, which contacts a random subset of notary servers after an HTTPS connection is opened, to compare the notary certificate observations with the received certificate. A configurable policy decides if the received certificate is presumed valid.

Convergence. Convergence [1] enhances Perspectives in several dimensions, most notably *providing privacy for certificate lookups by including a two-step onion routing approach*, where the first Convergence server redirects the query to a second server, and the second server responds (the first server knows the identity of the

querier, but not the web site queried, and the second server only knows the query but not the querier).

Certificate Catalog. Google Certificate Catalog [2] is a similar approach implemented in the Chrome browser, however, not yet enabled. It is a database of observed certificates, that can be queried through DNS requests.

SSL Observatory. EFF’s SSL Observatory [17] also collects global certificate information. However, it does not support any online queries, as far as we are aware of.

Discussion. The main advantage of these certificate observatories is that *server operators need not be aware of this approach*. Hence, no additional steps are necessary on their part. The approach even works to validate self-signed certificates. These systems are also effective to prevent numerous CA-based attacks, for example, all the attacks in Iran and Syria would have been prevented, as the illegitimate certificates would have been detected as different from the legitimate server certificates.

The main disadvantage of these systems is that they require additional connections to query the observatories, resulting in higher latency for establishing an HTTPS connection, and thus reduce overall performance whenever HTTPS connections need to be set up frequently. A major disadvantage is a period of unavailability for new certificates and when certificates change. However, this can be remedied with a scheme where a new certificate is offered ahead of time to the observation servers, but this approach negates one of the main advantages of these systems in that they do not require assistance from the server operators.

3.2 Certificate Log Servers

Sovereign Keys [13] and Certificate Transparency [22, 23] are two recent proposals that suggest a public log to record all certificate transactions.

Sovereign Keys. In the Sovereign Keys model, timeline servers act similar to timestamping servers [20], where certificate owners register their certificates and obtain a non-repudiable statement that their certificate has been added to the read-only and append-only log. The property achieved is that *the first registration of a certificate binds the key to the domain name, preventing any subsequent registrations for the same name*. Only the legitimate owner of the certificate has the private key, enabling revocation and re-registration of the name with a new public key.

The browser can contact the timeline server to inspect if the received server certificate is indeed the correct certificate registered in the timeline server’s log. More specifically, the browser requests all events within a time period pertaining to the server whose certificate is to be validated. *Certificate correctness means that the current certificate has not been revoked and can be derived through a chain of authorized key updates from the first registered certificate*. To reduce the load on the timeline servers, distributed mirror servers store a copy of the timeline server database and provide more efficient information dissemination.

Certificate Transparency. In Certificate Transparency [22, 23], the authors construct an append-only log by using an append-only Merkle hash tree structure, enabling efficient validation of each certificate that was added to the log. The general approach is similar to Sovereign Keys described above, but several details differ as we describe below and in Section 9.

Discussion. The main advantage of these approaches is that they prevent the attacks mentioned in the introduction. For example, compelled certificates are simply invalid in this framework, as the legitimate owner did not approve the newly generated key.

Unfortunately, these approaches have some shortcomings. For Sovereign keys, the browser needs to query a mirror server to down-




Figure 1: TKI certificate registration process. This figure depicts A.com registering a certificate (signed by two CAs) to a single ILS, but the domain can acquire multiple certificates from multiple CAs and register to multiple ILSs. Dashed arrows represent occasional communications.

load the server’s certificate history, increasing latency, decreasing availability, and sacrificing privacy. Consequently, the mirror servers need to be high-performance, as their latency directly affects the speed at which the SSL/TLS server certificate can be validated. In Certificate Transparency, the audit proof (i.e., proof of including the certificate in the log) is combined with the certificate.

A major disadvantage of both schemes is that a compromised private domain key cannot be easily revoked as the audit proof continues to be valid until (1) the log server publishes the new log entry, and (2) the client validates that there has been an update to the key since the log entry (which can be challenging in case of Certificate Transparency as the adversary can keep replaying the old certificate validation information). In Section 9, we discuss these systems in more detail.

4. TKI: Transparent Key Integrity

Before we describe the details of TKI, we first provide a high-level overview.

4.1 Overview

The main entities in our system are as follows.

- A **Domain (server)** represents a named entity with a public/private key pair that clients contact to establish secure connections with.
- A **Client (browser)** is an entity that establishes SSL/TLS connections with domains (servers).
- A **Certification Agency (CA)** is similar to a current certificate authority that certifies domains’ public keys, but “agency” instead of “authority” indicates that CAs are not absolute authorities any more in TKI.
- In addition to current systems, **Integrity Log Server (ILS) Operators (ILSO)** are Internet services that operate Integrity Log Servers (ILSs) that log domains’ certificates and make them publicly available. Each ILS maintains an **Integrity Tree**, which is a hash tree of all the registered certificates in lexicographic order. All ILS operations (i.e., log data entries such as registration, updates, and revocation) are digitally signed. We envision large ISPs or widely available Internet-based corporations such as Amazon or Google to operate ILSs.
- **Validators** are entities that monitor ILS operations, by downloading the entire ILS data structure and performing consistency checks. Consistency checks include validation that certificate updates follow the certificate policy or that certificates do not suddenly disappear from the log. Non-governmental Internet governance organizations, such as the EFF, would likely serve as a validator. In the case of ILSO misbehavior, validators can disseminate incriminating information without requiring their own

authentication information, since ILSO misbehavior is self-authenticating due to non-repudiation of ILSO/CA/domain signatures.

Figure 1 depicts an overview of our TKI architecture. Alice owns domain A.com and wants to obtain a TKI-protected certificate, as she wants to protect herself against compromise of the CAs that signed her certificate and other rogue CAs, and protect her clients against compelled certificates. To define the security properties that she intends to achieve for her domain, Alice defines CAs and ILSOs that she trusts, the minimum number of CA signatures that she recommends her clients for validation, rules for certificate revocation, replacement, updates, etc. Alice includes these parameters with her public key and contacts more than the minimum number of trusted CAs (according to her security policy) to sign her certificate. She then registers the certificate with concatenated CA signatures with one or multiple ILSs. Each ILS adds A.com to its database, by placing it in the Integrity Tree. The ILS then re-computes hash values over all stored certificates for updated verification information.

Alice now supplements her certificate with the verification information that she downloads from every ILS, and sends it to browsers that connect to her web site via HTTPS. For certificate validation, the browser uses the trusted root-CA certificates as in current practice, and uses the pre-installed ILS public key(s) on her browser to validate ILS information.

Before trusting the ILS information received from Alice, the client browser occasionally checks with validators to confirm that the current root hash values of the ILSs are valid.

4.2 TKI Details

We discuss the details of TKI based on the following stages: certificate creation, CA signature acquisition, ILS registration, browser-based validation, ILS tree update, certificate update, certificate revocation and recovery.

Certificate creation. TKI certificates contain several extensions over standard X.509 certificates and feature the following additional fields:

- **Trusted CAs (CA_LIST):** This field contains a list of trusted CAs for creating a new certificate.
- **Trusted ILSs (ILS_LIST):** This field contains a list of trusted ILSs where the certificate is registered.
- **ILS validation proof timeout (ILS_TIMEOUT):** This field indicates how long an ILS proof is acceptable to the browser after the proof creation time. The tradeoff is between efficiency, availability, and robustness. A long timeout requires fewer queries to the ILS for an updated proof, but increases the amount of time until a certificate can be revoked. This parameter typically varies from one hour to one day.
- **Minimum number of CAs to generate a TKI certificate (CA_MIN):** This field indicates the number of CA signatures required to initially register a certificate to ILS and to update a certificate, typically set to 1 or 2.
- **Threshold number of CAs for certificate re-establishment (CA_TH):** This field indicates the minimum number of CA signatures needed to re-establish a certificate in case of a lost private key. In other words, this parameter indicates the number of different CA signatures that can activate the new key for the domain which lost its key. An adversary can register a certificate to an ILS for a domain who is unaware of TKI, and select CA_TH to be high such that the domain can never revoke the adversary's certificate. To prevent such an attack, we set $CA_TH = CA_MIN + 1$.
- **Cool-off period for an unlinked certificate (COP_UNLINKED):** This field indicates the minimum cool-off period




Figure 2: TKI Integrity Tree structure.

for a new certificate that is not linked to the old certificate (i.e., the new public key is not signed by the previous private key of the domain). In TKI, registering a new certificate does not automatically validate it: in case a previous certificate is present, TKI enforces a “cool-off” period until the new certificate becomes valid which is to replace the previous certificate. This enables protection against an adversary who quickly registers a new key following a CA compromise, as the legitimate owner can revoke the new certificate during the cool-off period. An attacker can register a new key for a domain that is unaware of TKI and set COP_UNLINKED to be high to prevent the domain owner from re-acquiring the ILS entry. To prevent such an attack, an upper bound exists for COP_UNLINKED (e.g., 7 days).

- **Cool-off period for a certificate from an untrusted CA (COP_UNTRUSTED):** This field indicates the minimum cool-off period for a new certificate that is signed by a CA that is not in CA_LIST. In case of a lost or compromised private key, an attacker can acquire a certificate signed by some CA (that the domain owner does not trust), and this parameter provides time to enable the legitimate owner to revoke the bogus certificate. We recommend that COP_UNTRUSTED is defined to be longer than COP_UNLINKED. For similar reasons as for COP_UNLINKED, COP_UNTRUSTED has an upper bound (e.g., 10 days).

CA signature acquisition. After creating a certificate with TKI extensions, the domain contacts the CA_MIN number of CAs from CA_LIST to acquire CA signatures. In TKI, the combination of all the CAs’ signatures validates the domain’s TKI-certificate. Hence, care must be taken to include all the CAs’ serial numbers and timestamps, and the X.509 standard will need to be amended to enable such multi-signatures.

ILS registration. The domain then contacts one or more of the trusted ILSs (from ILS_LIST) to register the TKI-certificate.

The ILS data structure to store TKI-certificates is based on a binary hash tree, and we call it *Integrity Tree*, as depicted in Figure 2. All TKI-certificates are placed at the leaf nodes of the binary hash tree, sorted in lexicographic order. TKI uses a sorted hash tree as opposed to a linear list as in previous work [13] for the following reasons:

- 1) The hash tree efficiently represents the current state of all distinct names, and its height only depends on the number of different entries but not on time (i.e., it does not grow taller with revocations/re-establishments, thus removing a source of DoS).
- 2) The height is logarithmic in the number of entries. Hence, a validation of any leaf node can be efficiently represented based on an authenticated root node and a logarithmic number of nodes to re-compute the root node from the leaf node.
- 3) The sorting enables quick verification of the *absence* of an entry, whereas in a linear list, the entire list needs to be searched.

As depicted in Figure 2, the Integrity Tree enables independent validators to check the integrity of the entire data structure. The

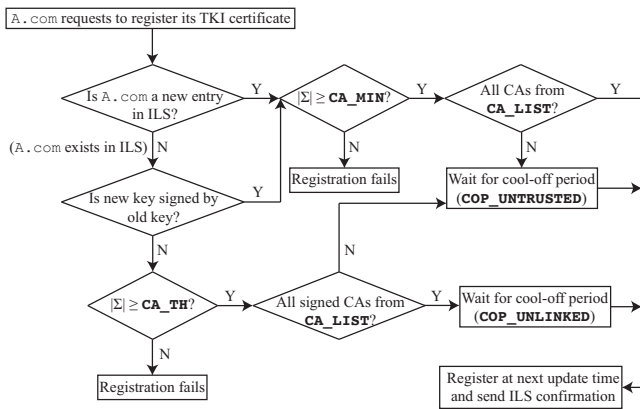


Figure 3: Flowchart of ILS certificate registration, update, revocation and recovery. Upon receiving a certificate registration request from A. com, the ILS follows this flowchart and plans to add the certificates. Σ stands for the CA signatures that A. com obtained.

hash chaining of the trees enables temporal re-construction of all operations, similar to a timestamping service or the timeline server data structure [13].

When adding a new TKI-certificate, the ILS first verifies whether an entry already exists in the Integrity Tree for the same domain name. If the name is indeed new, the ILS schedules the TKI-certificate to be added to the Integrity Tree. The ILS creates a confirmation of the successful addition through a digital signature with its private key, and returns it to the domain. The domain can use this ILS confirmation to start using the new certificate for a limited time, until the ILS generates the next Integrity Tree which will include the new certificate on a leaf node. Figure 3 depicts the certificate registration process.

ILS update period (ILS_UP) is the interval between two tree updates; at every ILS_UP, an ILS finalizes and commits the next Integrity Tree. At this point, the domain contacts the ILS to request the signed root node ($\{Root\}_{K_{ILS}^{-1}}$) and the hash tree verification nodes (h) that are needed to validate its certificate as depicted in Figure 2, where ILS_UP is set to one hour. In practice, ILS_UP is set to one or two hours, to enable quick certificate revocation. ILS verification information is combined with the TKI-certificate to enable client browsers to validate the ILS information without the need to contact an ILS during connection setup.

Browser-based validation. The browser receives the ILS information together with the server/domain’s TKI-certificate during the second phase of the SSL/TLS protocol. The CA signatures are validated using the browser’s CA root certificates, and the ILS information is validated using the ILS public keys stored in the browser. The ILS_TIMEOUT field in the TKI-certificate is validated to ensure that the ILS information is sufficiently recent depending on the domain’s preferences specified in the certificate.

A potential security vulnerability is due to incremental deployment: if browsers validate the ILS information only in TKI-certificates then no security would be gained since an adversary could simply create a traditional certificate without the TKI extensions. Consequently, browsers will need to contact trusted ILS servers for traditional certificates to prevent such attacks. In the absence of an ILS response, the browser needs to abort the connection. While this adds considerable latency to connection setup and reduces availability, it actually represents a positive incentive for TKI deployment: an ILS-registered certificate will result in a considerably faster connection setup.

ILS tree update. Periodically at a well-specified time, the ILS up-

dates its Integrity Tree by purging TKI-certificates that have been revoked or expired without renewal. The ILS also activates certificates that have passed their cool-off periods.

We envision update intervals of an hour up to a day. Hourly updates enable more fine-grained certificate revocation but increase overhead, as servers/domains need to frequently query the new signed root value to ensure that their name remains unchanged.

Certificate update. Before a TKI-certificate expires, the domain creates a new private key, and requests the trusted CAs in CA_LIST to sign the new key. The domain also signs the new key with its previous private key. After gathering CA_MIN number of CA signatures, the domain combines all signatures and other relevant information into a TKI-certificate. The domain then sends the TKI-certificate with an update request to the ILS, which will readily accept the new TKI-certificate since it is signed with the domain’s old key and the update request is signed by both new and old keys. (Requiring a signature with the new key confirms possession of the new private key.) There is no cool-off period in this case, and the new TKI-certificate is added when the ILS finalizes the next Integrity Tree. Hence, the new key can be readily used.

Certificate revocation and recovery. In case a key needs to be prematurely removed, a certificate revocation message needs to be sent to the ILS. Either the private key corresponding to the certificate’s public key is used to sign the revocation message, or a special revocation key can be used, for which the public key is included in the certificate. The point of using a different revocation key could speed up recovery for the case where the main private key is compromised, as a shorter cool-off period can be used if the new public key would be signed by the revocation key.

The cool-off periods (COP_UNLINKED, COP_UNTRUSTED) in the TKI-certificate specify the amount of time that needs to elapse before the new certificate becomes active. In case of private key compromise (and potentially private revocation key compromise), the COP_UNLINKED and COP_UNTRUSTED values enable the legitimate owner to react and revoke a fraudulent certificate that was potentially registered by the adversary.

Since some domains may not have the best key secrecy and availability practices in place, we need to consider the case of catastrophic key compromise and loss when only the adversary is in possession of all secrets. In that case, we need recovery mechanisms where the legitimate owner can re-gain control of its domain. By contacting CA_TH number of CAs and obtaining signatures on a fresh key, the legitimate owner can eventually re-gain control. However, the adversary will be able to use the key until a valid revocation message arrives.

Figure 3 depicts the ILS checks for certificate registration, update, revocation and recovery.

4.3 Checks and Balances among Parties

In this section, we describe how TKI achieves checks and balances among CAs, ILSs, validators, domain owners, and clients to reduce trust and prevent misbehavior by any party. Figure 4 illustrates what each party monitors and how each party reports.

Validation by CAs. Once the domain owner acquires signatures from trusted CAs for the certificate, the CAs monitor the ILS for any malicious changes in the domain’s ILS entry. If the ILS makes a potentially invalid update (e.g., updated certificate without any of the trusted CAs’ signatures), the CAs immediately inform the domain owner.

Validation by validators. Validators maintain a list of revoked ILSs that are detected for misbehavior possibly due to compromise. Validators disseminate the revoked ILSs, especially to the domain owners who are registered to those revoked ILSs, in which case the domain owners attempt to register with other valid ILSs. Thanks to




Figure 4: Checks and balances among TKI entities. In TKI, each entity monitors other entity for misbehavior detection and reports to other entities.

the fact that all ILS operations are signed, the validator can easily demonstrate misbehavior in case the signed records are inconsistent with the ILS’s state. In the absence of a compromised ILS private key, a validator cannot perform a slander attack, as it cannot forge signatures that would incriminate the ILS for malicious behavior.

Validation by domain owners. Prior to initial registration of a certificate to an ILS, the domain first ensures that a CA has not created a bogus certificate for that domain by checking the ILS as follows: the domain owner queries the ILS for entries that can be close to itself once added to the Integrity Tree (e.g., E.com can query for the directly preceding and succeeding entries for its domain to ensure that no unaccounted certificates are registered at that location of the tree. After confirming that no entry exists for the domain, the domain owner registers the new certificate to the ILS.

The domain owner can occasionally query for the leaf nodes that are adjacent to its certificate in the Integrity Tree to ensure that there is no equivocation in the ILS. Equivocation is detected if the ILS uses two different hash tree roots for two different replies. This behavior can be easily detected if the domain keeps track of the received hash-tree root values. If any equivocation is found, the domain owner contacts the validators to blacklist the ILS.

Validation by clients. A client browser occasionally checks with validators to see if the information received from the ILS server is valid. For this, the validator distributes the observed integrity tree root value for each ILS. If the client’s obtained ILS integrity tree root value is inconsistent, it will inform the validator and ignore that ILS server. The validator can then use that information to blacklist the ILS.

5. SECURITY ANALYSIS

In this paper, we provide an informal analysis of the TKI architecture. The formal analysis is presented in an extended version of this paper.

The main security property we aim to achieve is the prevention of successful impersonation of a victim server. More concretely, given a domain S with a certificate C_S , an adversary M attempts to impersonate S to a client C during the SSL/TLS connection establishment. The attack succeeds if C has a connection with M while believing that the connection is with S .

We assume that S uses a TKI-certificate, and that C uses a TKI-enabled browser. We show that TKI is resilient against attackers that compromise different entities’ keys.

M compromises S ’s private key. Assuming that S detects the compromise, it can immediately revoke the key. However, M can use the key during the ILS_TIMEOUT period, as specified in the TKI-certificate. This timeout is typically selected on the order of the ILS’s Integrity Tree update period. Thus, if the update period is

one hour, the timeout may be selected at 2 hours. Shorter revocation times could be achieved through an online validation protocol such as OSCP or Sovereign Keys, however, paying dearly in terms of latency and practicality [8].

To re-instantiate a key after compromise, there is unfortunately an unavailability period (COP_UNLINKED as specified in the certificate) in case the domain lost access to its private key. However, if the domain owner still has access to the private key, it can obtain trusted CA signatures for its new key, sign it with its old key, and immediately obtain an ILS confirmation that will enable use of the new key.

M compromises CAs’ private keys. As long as fewer than the CA_TH number of keys of trusted CAs in CA_LIST are compromised, there is no impact on browsers who contact the trusted ILSs. The CA_LIST evicts untrusted CAs from the set of potential weak links. Given a small well-selected list of trusted CAs, it is highly unlikely that more than a threshold number are compromised.

Even in case more than the CA_TH number of CAs are compromised, a newly registered key will have to cool off during the prolonged period ($CA_UNLINKED$) as the fraudulent certificate is not linked to the previous one (as we assume that S ’s private key was not compromised in this case). Such an extended time should leave sufficient time for S to detect and react to the impersonation attempt, without suffering any compromise (thanks to trusted CAs who watch out for ILS entry changes). If S ’s private key was compromised in addition, M can impersonate S during the entire $CA_UNLINKED$ period. However, this case is exceedingly unlikely, as several well-selected CAs and the domain’s private key all need to be compromised at the same time.

An adversary can also contact a different ILS to register a TKI-certificate for a victim domain whose TKI-certificate is already registered at a legitimate ILS. For example, an attacker can contact a Pakistan-ILS to register a forged citibank.com TKI-certificate. Since ILSs coordinate to cross-verify that a domain name is only registered with a single, consistent TKI-certificate, such an attack becomes visible. Furthermore, validators will detect such inconsistencies with high probability.

M compromises ILSs’ private keys. M can create a different Integrity Tree for a compromised ILS. In this case, two different Integrity Tree root values are active in a given time period. If M succeeds at completely suppressing the legitimate ILS, then the validators detect if a certificate was replaced without the proper revocation and certificate re-issuance policy as specified in S ’s certificate.

If M attempts to perform equivocation (i.e., create a shadow Integrity Tree with malicious entries and then answer queries from either tree depending on the querier), then clients and validators can readily detect this since the ILS would have signed two different root values for different Integrity Trees in a given time period, which is a visibly malicious action. The non-repudiation of the signature enables incrimination without permitting slander attacks.

The fact that the Integrity Tree root value prevents equivocation even helps in the case when M compromises CAs’ private keys in addition to the ILS’s. If M attempts to re-register a new key, S will immediately detect this behavior (as we discuss in Section 4.2) and raise an alarm. If M attempts to provide different answers to S ’s queries, it would need to create different Integrity Trees within one time period, which can be detected as described in the previous paragraph.

Another attack for M would be to attempt to create two different entries for S.com at different places in the Integrity Tree, one for the legitimate S.com and the other for a fraudulent S.com. In this scenario, M would provide the legitimate response to S ’s queries, and a fraudulent certificate for other queries. Fortunately, this case

is easily detectable by the validators, as the leaf nodes would not be in sorted order. Placing the two leaf nodes next to each other will be detected by S , when it also queries for the leaf nodes that are adjacent to its certificate in the Integrity Tree.

Another case is where M misuses the ILS's compromised private key to sign the ILS registration confirmation. As we describe in Section 4, such a confirmation would enable a freshly generated and initially registered TKI-certificate to be immediately used and trusted by TKI-enabled browsers (without contacting any ILSs). During the entire lifetime of a name, such an ILS confirmation can only be used during the initial period of registration, in practice for only one hour out of a multi-year lifetime. It would be thus highly suspicious if such a confirmation were used with names that are already part of the ILS trees: the browser could thus report the suspicious certificate to validators, as it is clear from the browser history that the site was accessed in the past over HTTPS.

6. DISCUSSION

TKI vs. CA. One may question about TKI's difference from current Certificate Authorities. The main differences are that all ILS operations are public and that compelled certificates can be easily detected. CAs in TKI have strong accountability for their actions, which cannot be circumvented. Consequently, trust in individual CAs is greatly reduced in TKI.

Censorship resilience. Corporations/governments may want to eavesdrop on all employees/citizens' communication. More specifically, corporations/governments can set up their own CA and ILS that create fake certificates. In such a case, users can opt out by installing legitimate CAs and ILSs as roots of trust.

Absence of ILS information. Similar to EV certificates [10], absence of ILS information may not raise any suspicions. To prevent an attack where a non-TKI certificate is used to attack a domain that is using a TKI-certificate, we require browsers to contact ILSs to validate the absence of TKI information. Note that no additional latency is required for deploying sites, since they provide the ILS information during the SSL handshake; only legacy domains and attackers have additional latency. Hence, this additional latency provides a positive incentive for deployment.

Furthermore, ILSs can cooperate to provide proofs of non-existence as follows: given that site $E.com$ maintains the TKI-information with ILS_j that browsers do not trust as much as ILS_1, \dots, ILS_i , these trusted ILSs can provide the absence proof for $E.com$ by providing the authenticated Integrity Tree leaf nodes before and after the point where $E.com$ would be located at (since all the nodes are sorted in lexicographic order).

In some legitimate environments, ILSs may not be reachable, for example paywalls at airports or hotels do not permit any external connections until the user has authenticated, paid, or accepted the terms of service. A challenge then is: how can the browser verify the non-TKI certificate of the paywall service without access to the ILSs? In this case, geographically-linked certificates [21] can be used, or the paywall obtains a TKI-certificate.

Globally consistent registration. Ideally, all global ILSs coordinate registration and provide one global name space, preventing the same name to be registered at different ILSs with different certificates. However, global coordination is cumbersome to implement in practice, and we can achieve global consistency by detecting and resolving short-term inconsistencies.² An example of an inconsis-

²In any case, ILSs could implement a minimal validation that the same domain name has not already been registered elsewhere. In particular, if the registration originates from an entity outside the legal region of the domain owner, or was signed by CAs outside the legal region of the domain owner, then the TKI-certificate can be listed but become valid only after a cool-off period.

istency is where a rogue CA issues a bogus certificate for $A.com$, presumably the CA and the requester of $A.com$ are in a different legal region from the legitimate owner of $A.com$ preventing the conflict to be locally resolved through legal means. In such cases of inconsistent registrations, the CAs local to the registered name of $A.com$ obtain precedence in determining the correct certificate. If the foreign ILS does not unregister the conflicting entry, it loses its credibility and will be subsequently ignored. It is the task of the validators to document, store, and disseminate such incriminating ILS information.

To detect inconsistent registrations, validators can inspect the global ILSs and inform domain owners in case of inconsistent registrations (i.e., registrations with a different public key). In addition, CAs can offer a service to their clients to watch over potential misuse or inconsistencies of their domain name. Finally, domains themselves can also inspect ILSs' operations and detect misbehavior. Since all CA and ILS operations are non-repudiable (since every operation is digitally signed), misbehavior does not need to be further authenticated by validators, CAs, or domains. Consequently, slander attacks that plague reputation systems are averted. **Usability.** Prior work has shown that users ignore and click through certificate warnings [33]. However, TKI can identify real attacks and completely block users from proceeding without an option to click through.

7. REALIZATION IN PRACTICE

To demonstrate the feasibility of TKI in a real-world setting, we built a prototype as a proof-of-concept system. For testing, we created a CA with OpenSSL. We pre-installed the CA and ILS root certificates on our servers and clients. We implemented an ILS server in Python that maintains a TKI Integrity Tree. The Integrity Tree node hashes were computed with SHA-256, and the root node was signed with RSA-2048. We used Coordinated Universal Time (UTC) to define precedence for domain to key mappings and to audit timeline integrity.³ We configured a stock Nginx HTTP server to serve our TKI-certificates, which are basically X.509 certificates with custom TKI extension fields (described in Section 4.2). We implemented our TKI client in the Chromium web browser.

ILS proof stapling. To deliver fresh ILS proofs for TKI-certificates to clients, we explored the following options. One option is to let CAs embed the ILS proof in the certificate itself, by inserting it into a certificate extension. However, once the certificate file is updated with a time-bounded ILS proof, the hash of the updated certificate would not match the original certificate recorded in the ILS Integrity Tree. To avoid this issue, another option is to let servers send the ILS proofs over the TLS handshake, utilizing a TLS extension. An alternative is to provide the ILS proof in a separate dummy certificate, appended to the leaf of the server's certificate chain. In our prototype, we sent ILS proofs via TLS extensions, given that Nginx currently supports the `TLS Certificate Status Request` extension (primarily used for OCSP stapling). This allows our server to deliver ILS proofs as a stapled response over the TLS handshake to clients without modifying Nginx. The server could use a side-loaded script to periodically fetch fresh ILS proofs and load them into Nginx. We modified the Chromium browser to extract the embedded ILS proofs via the `TLS Certificate Status Request` extension, and validate the ILS proofs.

Performance cost. TKI induces no round trip latencies (no extra network requests) to the TLS handshake. However, TKI increases the TLS handshake message size by roughly a kilobyte due

³UTC is a well-known, independent external reference, and in theory, many synchronization problems should reduce to accurate timekeeping.

to ILS proof stapling (assuming millions of domain names registered to the ILS). The ILS proof is composed from the following constituents: a list of authentication node hashes (32 bytes per node), a timestamp (4 bytes), and a root node signature (256 bytes). Further, the server’s certificate size is slightly increased (by roughly 40 bytes) due to the additional custom X.509 extension.

We measured the client validation processing time in Chromium on a machine with 2.26 GHz dual-core CPU and 4 GB RAM, tested with a million domains registered to ILS. The overall TKI processing time averaged 990 μ s (median = 936 μ s). Specifically, the RSA verification step averaged 880 μ s (median = 831 μ s), while the Merkle verification step only averaged 95 μ s (median = 87 μ s). The overall TKI processing time is relatively small, especially compared to other approaches, such as Perspectives and Sovereign Keys that require several network round-trips to communicate with servers.⁴ Another advantage of using Integrity Trees that are relatively infrequently updated means that RSA verifications on the client side are amortized, as we envision that many domains will use the same set of ILSs, hence the root of the Integrity Tree will remain the same for numerous sites. Consequently, clients mostly only perform efficient hash tree verifications and only rarely perform signature verifications.

8. RELATED WORK

Several proposals have been made to provide a web of trust for SSL, such as Monkeysphere Web-of-Trust for SSL [4], the EFF’s SSL observatory [17], and Certificate Patrol [5]. These proposals make it easy to detect key changes, but it is difficult to distinguish legitimate key changes from attacks.

Langley et al. [6] implemented a public key pinning mechanism in Google Chrome. The browser vendor maintains a list of trustworthy public key(s) associated with each site. Public key pinning provides similar security benefits to TKI by preventing certificates signed by rogue CA from being accepted by the browser. Typically the keys of trusted CAs are pinned, allowing for an orderly transition from one certificate to the next. To address the scalability challenges of a browser vendor maintained database, the Public Key Pinning Extension for HTTP [7] generalizes this mechanism to an HTTP header that allows a server to declare the keys that can be used in the future for that domain name. Choosing a pin duration that is too long risks a lengthy period of unavailability for the site. Furthermore, if the user is visiting the site for the first time on a device or the pin has expired, no protection is provided. By contrast, TKI provides protection on the first visit to the site.

Marlinspike and Perrin propose Trust Assertions for Certificate Keys (TACK) which pins public keys generated by the domain owners themselves [26]. More specifically, a server generates a TACK key pair, and use the TACK private key to sign the TLS public key. The TACK public key and the signature form the TACK, which clients can see in the TLS extension field, and clients “pin” the domain’s TACK public key after observing the consistent TACK multiple times. Although TACK aims at removing complete trust in CAs, TACK relies on frequent visit patterns by clients to pin the domain’s public key, resulting in long initial unavailability period for every server. Furthermore, if a certificate becomes compromised and the pin is still inactive, the client must delete the observed TACK information. In contrast, TKI provides no initial unavailability period for any servers, providing protection on the first visit to the server.

Huang et al. propose short-lived certificates [8] in conjunction

⁴ A previous study indicates that round-trip latencies for OCSP lookups cost a mean of 497 ms with a median of 291 ms in real-world deployments [31].

with browser vendor maintained Certificate Revocation Lists (CRLs) to mitigate the impact of key compromise. Servers provide certificates with a short validity lifetime and update them from the CA on a daily basis. Short-lived certificates provide similar security benefits to OCSP while eliminating the need for an online check during the HTTPS handshake. However, unlike TKI, they rely on browser vendors to somehow detect certificates that are issued by compromised CAs and block them using a browser vendor maintained blacklist.

DNS-based Authentication of Named Entities (DANE) securely binds certificates with domain names using Domain Name System Security Extensions (DNSSEC), enabling domain holders to assert certificates without reference to CAs [12]. However, the security of DANE heavily relies on the security of DNS operators.

In the following section, we perform an in-depth comparison of all the closely related certificate validation infrastructures.

9. THEORETICAL COMPARISON

In this section, we compare TKI with other proposals with respect to security, availability, and efficiency metrics. One of the contributions of this paper is to establish a set of metrics for comparison, which we present in the following subsection.

9.1 Evaluation Metrics for Comparison

Security metrics. The main security metric is Duration of Compromise (DoC): given the compromise of a private key,⁵ for how long can a domain be impersonated? This metric can be specified into the following:

- DoC after a trusted CA’s private key compromise: This case also covers compelled certificates [30].
- DoC after untrusted CA’s private key compromise: This case is important for TKI, where a domain defines trusted and untrusted CAs.
- DoC after trusted public log server’s private key compromise: To avoid a proliferation of cases, we do not consider untrusted log server’s private key compromise, as it is a strictly weaker attack scenario.
- DoC after domain’s private key compromise: This metric measures the DoC, for how long an adversary can misuse the captured private key. We define the DoC as the duration of when a key is revoked or the domain bootstraps a new key which invalidates the old key, whichever is earlier.

Security guarantees of new systems can sometimes be circumvented due to required compatibility issues with legacy systems. For example, Extended Validation (EV) certificates or OCSP information in a certificate are both optional extensions, and their absence does not raise any suspicions. Therefore, even if an entity obtains an EV certificate and uses OCSP, an adversary can still obtain a fraudulent non-EV certificate without OCSP extensions that will enable MitM attacks. To measure the security of public key validation infrastructures during incremental deployment, we propose the following metric:

- Protection during incremental deployment: This is a binary measure to characterize whether any security is offered while compatibility with legacy systems needs to be ensured.

Finally, we measure privacy of client requests.

- Connection privacy: information about a client is not leaked to entities other than the contacted domain.

⁵ In a private key compromise, the key is disclosed to the adversary. Depending on the attack, the legitimate owner may still possess the key.

Table 1: Comparison of different public-key validation infrastructures based on the security, availability, and efficiency metrics. Entries in bold red font indicate major disadvantages of the corresponding scheme. Server* stands for the ILS, DNS, Notary, or OCSP responder server, depending on which scheme is used. Δ_U corresponds to the public log servers’ update interval, which is in practice on the order of one hour. Section 9.2 describes our methodology for filling in the entries.

	CA+CRL	CA+OCSP	SLC	Key Pinning	TACK	DANE	Perspectives Convergence Cert Catalog	SK	CT	TKI
Security										
Trusted CA compromise (compelled certificate) DoC	days	days	days	days	N/A	hours	0	0	0	0
Untrusted CA key compromise DoC	days	days	days	0	N/A	0	0	0	0	0
Trusted Server* key compromise DoC	N/A	0	N/A	N/A	N/A	days	days	days	days	hours
Domain key compromise DoC	hours	min	days	days	<month	hours	day	min	∞	ILS_TIMEOUT
Protection during incremental deployment	N/A	N	N	Y	N	N	Y	Y	Y	Y
Connection privacy	Y	N	Y	Y	Y	Y	N/Y/Y	N	Y	Y
Availability										
Initial registration DoU	0	0	0	0	days	hours	days	0	Δ_U	0
Planned key update DoU	0	0	0	0	days	hours	0	0	Δ_U	0
Unplanned key update DoU	0	0	0	0	days	hours	days	0	Δ_U	Δ_U
CA compromise DoU	days	days	days	days	N/A	days	0	0	0	0
Server* compromise DoU	days	days	N/A	N/A	N/A	days	days	days	days	up to 1 day
Domain compromise DoU	min	min	min	min	<month	hours	days	0	Δ_U	COP_UNLINKED
Efficiency										
Number of additional servers required	0	<i>O(C)</i>	<i>O(D)</i>	0	0	0	<i>O(C)</i>	<i>O(C)</i>	<i>O(D)</i>	<i>O(D)</i>
Additional latency for SSL/TLS connection setup	0	RTT	0	0	0	0	RTT	RTT	0	0
Additional bandwidth for SSL/TLS connection setup	0	0.1KB	0	0	0.1KB	0.1KB	10KB	1KB	1KB	1KB

Availability Metrics. The main availability metric we use is Duration of Unavailability (DoU) of a domain’s certificate after various system events. The shorter the DoU, the more available the system. We do not consider DDoS attacks in this paper; thus, we assume the general availability of all servers and communication networks.

- DoU after initial registration: This metric measures the time duration until the registered certificate becomes valid.
- DoU after planned key update: This metric measures the duration when an updated key becomes valid, which was planned to replace the current key.
- DoU after unplanned key update: In case of unplanned events such as losing a domain’s private and backup keys, this metric measures the time duration when the updated key becomes valid.
- DoU after trusted CA’s private key compromise: After a trusted CA’s key becomes compromised, a domain’s certificate may also become invalid. This metric measures the time to acquire a new certificate using the CA’s new key.
- DoU after a trusted log server’s private key compromise: A log server’s public key compromise leads to invalid log entries. This metric measures the time to recover a log server’s private key.
- DoU after domain’s private key compromise: This metric measures the time until the domain’s certificate becomes available with a new private key.

Efficiency Metrics. Below is a list of metrics to measure the efficiency of certificate infrastructures:

- Number of additional servers required: This metric measures how many additional infrastructure servers are required, expressed as an order in the number of new connections established. For example, if C connections are established to D different domains, would we require $O(C) + O(D)$ additional servers, $O(D)$, or even $O(1)$?
- Additional latency to establish a secure connection: Compared to standard SSL/TLS, what additional latency would be required for a secure connection with the proposed scheme?
- Communication overhead: This metric measures the additional network overhead incurred for establishing a secure connection.

In addition, we will evaluate schemes based on their ability for domains to select their trust perimeter (with respect to CAs and

public log servers), as well as providing flexibility for certificate policies, such as specification to achieve different tradeoffs between availability and security metrics as defined above.

9.2 Comparison of Approaches

Based on our metrics, Table 1 compares TKI with the following proposals: CA + CRL [11], CA + OCSP [28], Short-Lived Certificates (SLC) [8], Key Pinning [6], TACK [26], DANE [12], Perspectives (P) [34], Convergence (C) [1], Certificate Catalog (CC) [2], Sovereign Keys (SK) [13], and Certificate Transparency (CT) [23].

We now discuss the methodology we used to fill in the table. For many of the catastrophic failures, such as compromise of a trusted CA or ILS private key, we assume that a software update is required to revoke the old key and setup a new key. We assume that such a software update is secure, and can be completed within a few days for most users.

Security. For the “Trusted CA compromise (compelled certificate) DoC” metric, we assume that it will take days to push out a CA root certificate revocation message through a browser update, which was the method used to revoke DigiNotar’s certificate after the compromise [32]. While some browsers use CRLs to revoke CA keys (e.g., Google Chrome), most browsers still require a software update. OCSP unfortunately does not help in this case, since the CA does not use OCSP to validate the root certificates. Similarly in the case of SLC and DANE, a browser update is required to revoke the CA key. Also in the case of Key Pinning, a browser update is required to remove the pin. Since P/C/CC, and TACK do not rely on CAs, the DoC is 0. Audit-log based schemes also protect from this case, preventing even a trusted CA from registering a new bogus certificate (the security analysis in Section 5 presents this case in more detail for TKI).

For the following properties, we explain the metrics in a less verbose manner. For the “Untrusted CA key compromise DoC” metric, the impact is less than in the previous case. In particular for DANE, the adversary cannot impersonate the domain which was possible in the trusted case.

For the “Trusted Server key compromise DoC,” we consider that the ILS/DNS/Notary/OCSP responder server’s private key is compromised, resulting in a severe disruption for several approaches. Since no additional third parties exist in CA + CRL, SLC, and Key

Pinning, this case is N/A for those schemes. Since a compromised OCSP server’s private key does not enable creation of a fake key, DoC is 0. On the other hand, if the TACK key is compromised, recovery can take up to 30 days, depending on the domain’s parameter setting. In the case of a compromised notary key in P/C/CC, we assume that a software update would require days to be fully deployed, during which time attacks are feasible. SK and CT would also require a software update, requiring days for full deployment. In TKI, a validator can detect ILS misbehavior and disseminate the incriminating information, which may last on the order of hours to reach the majority of clients.

For the “Domain key compromise DoC” metric, we assume that browsers download CRLs every few hours; thus, the DoC for CA + CRL is on the order of hours. For SLC, it may take a few days for the certificate to expire. In TACK, it may require up to a month to have clients switch to a new key. For DANE, it may require hours until DNS entries time out and get replaced by new entries with the updated key information. In P/C/CC, depending on the client configuration, it can take days for an updated key to be consistently observed. Although the online validation of SK revocation is very fast, CT will require more time since stale validation information may be served by the adversary. For TKI, validation information is valid during domain-selected time ILS_TIMEOUT, which is on the order of several hours to one day, until the key is revoked.

For “Protection during incremental deployment,” OCSP, and SLC offer no security, since an adversary can create a legacy certificate without any of these extensions which clients would accept. In TACK, a rollback to a compromised certificate attack is possible at the onset, when the TACK pin is not yet set up. For DANE, DNS responses may be rolled back to non-signed DNS replies. P/C/CC, SK, CT, and TKI all perform an online lookup for the case of a legacy certificate, which will reveal the legacy certificate.

“Connection privacy” is not provided by OCSP, Perspectives, and SK, as the client performs an online lookup for each certificate. Convergence uses a blinding step during lookup, and Certificate Catalog’s lookup via DNS hides client information.

Availability. “Initial registration DoU” requires several days for TACK and P/C/CC to confidently learn a new entry. In DANE, the current DNS entry needs to time out for the updated DNS entry to become available, which we estimate to take hours in the common case. CT requires the log server to update the tree, which we denote with Δ_U , which corresponds to the ILS_UP for TKI (we assume that both CT and TKI use the same log update period).

For “Planned key update DoU,” we consider an optimization we discuss in Section 3, where domains pre-register a key with the notary servers, thus avoiding activation latency. For CT and TKI, Δ_U may be required until a key update becomes active.

For “Unplanned key update DoU,” we assume that P/C/CC use a configured policy where a key has to have been consistently observed for several days for clients to trust the key.

For “CA compromise DoU” and “Server compromise DoU,” we assume that several days are required to recover and roll out new root keys. In key pinning, we assume that one day is required to push out a new software version with a new key. “Domain compromise DoU” indicates the delay required to register a new key.

Efficiency. For the metric “Number of additional servers required,” we specify D for the number of domains and C for the number of connections established per day. For example, $O(D)$ indicates that the number of additional servers needs to be proportional to the number of domains.

For the metric “Additional latency for SSL/TLS connection setup,” we denote a round-trip time to a server by RTT, which includes server processing time. Since P / C / CC, SK, and OCSP also in-

volve additional external connections, they can have a significant time overhead.

For “Additional bandwidth for SSL/TLS connection setup,” we list the order of magnitude of additional bandwidth required to set up an SSL/TLS connection. For the case of SK, CT, and TKI, we assume that extra signatures are on the order of 256 Bytes, hash tree values are on the order of 32 Bytes, and that a hash tree has about 30 levels, resulting in about 2 KBytes of additional information, which is on the order of 1KB as listed in the table.

9.3 Observations

As is evident from Table 1, all the newer Certificate Validation Infrastructures handle the case of untrusted CAs or CA key compromise, dramatically increasing the security over the current certificate validation infrastructure.

For practical deployment, it is critical that the SSL/TLS connection establishment does not incur any additional latency. Consequently, the additional RTT incurred by OCSP, P/C/CC, and SK is problematic. Moreover, any system requiring $O(C)$ additional server infrastructure load is likely to incur excessive cost. Performing an online per-connection lookup to an external server also challenges privacy, as it may leak information about the connection to a third-party server.

Another important factor is that certificates become immediately usable after initial registration. However, CT, TACK, and P/C/CC do not support this feature.

Overall, CT and TKI emerge with many desirable features. Since (1) the overhead of TKI is lower due to the different hash tree structure, (2) TKI allows immediate use of initially registered certificates, and (3) TKI can rapidly validate the absence of an entry, TKI provides a more efficient solution in practice.

10. CONCLUSION

Protecting current PKIs against CA root key compromises is becoming a topic of critical importance, as the weakest-link security model of the current PKI system is clearly too weak to provide meaningful security for critical web communication.

We observe that a public integrity log offers a promising approach to prevent the attacks we have recently witnessed. Unfortunately, proposed log-based approaches suffer from several drawbacks that hamper adoption.

With real-world adoption in mind, we propose TKI, a new approach that offers flexibility for entities to select a security policy for their certificates, enabling a tradeoff between availability and security. TKI also provides tangible deployment incentives that we anticipate will drive adoption. In addition, TKI provides a useful point in the design space towards a more trustworthy public key validation architecture.

11. REFERENCES

- [1] Convergence. <http://convergence.io/>.
- [2] Google Certificate Catalog. <http://googleonlinesecurity.blogspot.com/2011/04/improving-ssl-certificate-security.html>.
- [3] Perspectives Project. <http://perspectives-project.org/>.
- [4] The Monkeysphere Project. <http://web.monkeysphere.info/>, 2010.
- [5] Certificate Patrol. <https://addons.mozilla.org/en-US/firefox/addon/certificate-patrol/>, 2011.
- [6] Public Key Pinning. <http://www.imperialviolet.org/2011/05/04/pinning.html>, May

- 2011.
- [7] Public Key Pinning Extension for HTTP. <http://tools.ietf.org/html/draft-ietf-websec-key-pinning-01>, Dec. 2011.
- [8] Towards Short-Lived Certificates. In *Web 2.0 Security and Privacy*, May 2012. To appear.
- [9] A. Birrell, B. Lampson, R. Needham, and M. Schroeder. A global authentication service without global trust. In *Proceedings of the IEEE Symposium on Security and Privacy*, 1986.
- [10] CA/Browser Forum. Guidelines For The Issuance And Management Of Extended Validation Certificates. http://www.cabforum.org/Guidelines_v1_3.pdf, 2010. Version 1.3.
- [11] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. Technical report, RFC 5280 (Proposed Standard) Internet Engineering Task Force, 2008.
- [12] M. M. Correia and M. Tok. DNS-based Authentication of Named Entities (DANE). Technical report, Universidade do Porto, 2011–2012.
- [13] P. Eckersley. Sovereign Key Cryptography for Internet Domains. <https://git.eff.org/?p=sovereign-keys.git;a=blob;f=sovereign-key-design.txt;hb=HEAD>.
- [14] P. Eckersley. A Syrian Man-In-The-Middle Attack against Facebook. <https://www.eff.org/deeplinks/2011/05/syrian-man-middle-against-facebook>, May 2011.
- [15] P. Eckersley. Iranian hackers obtain fraudulent HTTPS certificates: How close to a Web security meltdown did we get? <https://www.eff.org/deeplinks/2011/03/iranian-hackers-obtain-fraudulent-https>, Mar. 2011.
- [16] S. Egelman, L. F. Cranor, and J. Hong. You’ve Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI)*, 2008.
- [17] Electronic Frontier Foundation. SSL Observatory. <https://www.eff.org/observatory>.
- [18] N. Falliere, L. O. Murchu, and E. Chien. W32.Stuxnet Dossier. Technical report, Symantec Corporation, 2011.
- [19] V. Gligor, S.-W. Luan, and J. Pato. On inter-realm authentication in large distributed systems. In *Proceedings of the IEEE Symposium on Security and Privacy*, 1992.
- [20] S. Haber and W. S. Stornetta. How to time-stamp a digital document. In *Advances in Cryptology, CRYPTO*, 1990.
- [21] T. H.-J. Kim, V. Gligor, and A. Perrig. GeoPKI: Converting Spatial Trust into Certificate Trust. In *Proceedings of the 9th European PKI Workshop (EuroPKI)*, Sep 2012.
- [22] B. Laurie and A. Langly. Certificate Authority Transparency and Auditability. <http://www.links.org/files/CertificateAuthorityTransparencyandAuditability.pdf>, 2011.
- [23] B. Laurie, A. Langly, and E. Kasper. Certificate Transparency certificate-transparency-draft. <http://www.links.org/files/sunlight.html>, Mar. 2012.
- [24] R. Lemos. Microsoft warns of hijacked certificates. <http://news.cnet.com/2100-1001-254586.html>, Mar. 2001.
- [25] M. Marlinspike. SSL And The Future Of Authenticity. <http://blog.thoughtcrime.org/ssl-and-the-future-of-authenticity>, Apr 2011.
- [26] M. Marlinspike and T. Perrin. Trust Assertions for Certificate Keys. <http://tack.io/draft.html>, May 2012.
- [27] J. Menn. Key Internet operator VeriSign hit by hackers. <http://www.reuters.com/article/2012/02/02/us-hacking-verisign-idUSTRE8110Z820120202>, Feb. 2012.
- [28] M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams. X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP. Technical report, RFC 2560 (Proposed Standard) Internet Engineering Task Force, 1999.
- [29] P. Roberts. Phony SSL Certificates issued for Google, Yahoo, Skype, Others. http://threatpost.com/en_us/blogs/phony-web-certificates-issued-google-yahoo-skype-others-032311, Mar. 2011.
- [30] C. Soghoian and S. Stamm. Certified Lies: Detecting and Defeating Government Interception Attacks against SSL. <http://files.cloudprivacy.net/ssl-mitm.pdf>, 2010.
- [31] E. Stark, L.-S. Huang, D. Israni, C. Jackson, and D. Boneh. The case for prefetching and prevalidating TLS server certificates. In *Proceedings of the 19th Network and Distributed System Security Symposium*, 2012.
- [32] T. Sterling. Second firm warns of concern after Dutch hack. <http://news.yahoo.com/second-firm-warns-concern-dutch-hack-215940770.html>, 2011.
- [33] J. Sunshine, S. Egelman, H. Almuhamidi, N. Atri, and L. F. Cranor. Crying Wolf: An Empirical Study of SSL Warning Effectiveness. In *Proceedings of the USENIX Security Symposium*, 2009.
- [34] D. Wendlandt, D. G. Andersen, and A. Perrig. Perspectives: Improving SSH-style Host Authentication with Multi-Path Probing. In *Proceedings of USENIX Annual Technical Conference*, June 2008.