

Analyzing proposals for improving authentication on the TLS-/SSL-protected Web

Christopher W. Brown¹ · Michael Jenkins²

Published online: 3 February 2016
© Springer-Verlag (outside the USA) 2016

Abstract “Secure” Web browsing with HTTPS uses TLS/SSL and X.509 certificates to provide authenticated, confidential communication between Web clients and Web servers. The authentication component of the system has a variety of weaknesses, which have led to a variety of proposals for improving the current environment. In this paper, we survey, analyze, compare and contrast five prominent proposals. To do this, we attempt to systematically capture the properties one might require of such a system: authentication properties, forensics/privacy properties, usability properties and pragmatic properties. Enumerating these properties is an important part of understanding these proposals and the nature of the authentication problem for the secure Web. Finally, we offer a few conclusions and suggestions pertaining to these proposals and possible future directions of research.

Keywords Web security · Authentication · TLS · HTTPS · Certificates

1 Introduction

This article provides a summary and analysis of five proposals for improving the authentication component of the current environment for the TLS-/SSL-protected Web (HTTPS)—specifically the client’s authentication of the server. As a part

of this analysis, we identify the properties the overall system would, ideally, satisfy and describe the various proposals in terms of the degree to which they do or do not provide/have these properties. The paper’s real contributions, hopefully, are a clearer picture of what the authentication problem for the TLS-/SSL-protected Web really is, and a framework for evaluating new proposals both individually, and in combination with one another.

The current environment for secure Web browsing is based on TLS (see RFC5246 [1]). TLS provides a mechanism by which a Web client (browser) and a Web server establishing a connection between one another make use of public key cryptography to agree on a shared secret key, which is then used to encrypt communication using symmetric encryption for the rest of that session. Typically the process begins with the client being furnished a domain name by the user, which is then translated to an IP address via DNS name resolution, after which the client sends the Client Hello TLS message to the server presumed to be listening on port 443 at that IP Address. The client then receives messages from the server, one of which contains a public key. The client chooses a secret key (more accurately, a value that determines the secret key), encrypts it with the public key the client received and sends the resulting ciphertext back to the server. At this point, both parties have the same secret key, and encrypted communication can commence.

This process guarantees that the client and the owner of the public key—i.e., the entity in control of the associated private key—are the only ones who know the secret key, assuming that the private key is kept private. It does not, however, guarantee anything about the identity of the owner of the private key. There is no assurance that the party the user of the client wanted to contact is the one with whom the client now has a secure connection. Routing to the IP address could have gone wrong. DNS resolution could have produced the

✉ Christopher W. Brown
wcbrown@usna.edu

Michael Jenkins
mjjenki@tycho.nsc.mil

¹ National Security Agency / U. S. Naval Academy, Annapolis, MD, USA

² National Security Agency, Fort Meade, MD, USA

wrong IP address. The domain name itself might be wrong—e.g., amazon.com instead of amazon.com. Thus, there is a critical authentication problem to be dealt with!

The current environment for secure Web browsing generally handles this authentication problem in the following (highly simplified!) way: (1) It is the user's responsibility to ensure that the domain name provided to the browser is in fact the correct domain name for the entity they are trying to contact. (2) The Web server sends its public key to the client in an X.509 certificate (see RFC5280 [2]) as part of a certificate chain in the TLS Server Certificate message, and it is the browser's responsibility to validate the certificate chain and verify that it chains to one of the certificates in the browser's list of trust anchors. (3) It is the responsibility of the CA that issued the trust anchor at the end of the certificate chain to ensure that the public key in the certificate really belongs to the entity that owns the given domain name.¹

It is assumed that the reader of this paper is familiar with this process, including details such as certificate revocation lists, self-signed certificates and X.509 validation not covered in the very brief description above. It is also assumed that the reader is familiar with the manifold problems inherent in this system—for instance, that every CA trusted by the browser represents a single point of failure for the whole system or realities of how users usually bypass the whole system [3, 4]. The key point is that the authentication component of the current secure Web-browsing environment has problems both in principle and in practice.

The many problems that exist in the current environment have prompted a wide variety of proposals for improvements. The proposals considered in this paper are limited to ones with a reasonable level of pragmatism (full-scale replacement of the current environment with something new is extremely unlikely!) and a reasonable degree of visibility or momentum behind them. User interface improvements, while important, are also outside of the scope of this analysis. The proposals examined here are: DANE, Certificate Transparency, HTTP Pinning (HPKP), TACK and Perspectives. We looked at the Sovereign Keys proposal closely, but it is not yet mature enough to really analyze to the same extent as the above. Projects similar to Perspectives, such as Convergence, Google's now defunct Certificate Catalog project and the Berkeley ICSI project, are interesting as well, but space precludes covering them here.

This paper consists of three parts. The first describes the set of properties that might be desirable for the authentication component of the "secure Web." The second provides summaries of the five proposals, along with some commentary and analysis—all of which is done in terms of the properties

from the previous section. The last focuses on comparisons of different proposals and their potential to work in combination with one another.

2 Desirable properties for the authentication component of the "secure Web"

It is easier to understand and compare these various proposals if we first describe what it is we really want. In other words, what properties do we really desire of the authentication component of the "secure Web"? What follows is a list of such properties, grouped into four categories. The fourth category is more accurately described as a list of properties we would like to see in a proposal for improving on the current state of affairs, for example that the proposed improvements could be phased in incrementally, as opposed to expecting the whole Internet infrastructure to change all at once. These properties were in large part deduced from reading a number of proposals and commentary on those proposals.

*Authentication Properties*²

continuity: that when a client connects to a host with name X, one can be sure that, in some meaningful sense, it is communicating with the same entity as it was communicating with the last time it connected to a host with name X. Note: It could make sense to think of continuity on two levels, the individual client level and the community level. Pinning proposals are about individual clients observing continuity. Notary proposals such as Perspectives or the Berkeley ICSI, and gossiping proposals are about a community of clients observing continuity.

domain name authentication: that the client is connected to the server authorized, intended or allowed to run under that name-and-port by the legitimate owner of that domain name.

higher-level authentication: that the client is connected to the server authorized, intended or allowed to run under that name-and-port by an entity described by some notion of identity beyond merely ownership of a given domain name (e.g., Southwest Airlines, the US Postal Service). There are actually different classes of certificates—EV (extended validation) [5], OV (organization validation), DV (domain validation)—that seek to provide some higher-level authentication. Note: Many companies have made their domain name their identity, e.g., amazon.com, so that domain name identity and higher-level identity are one and the same.

attribute authentication: that the client is connected to the server authorized, intended or allowed to run under

¹ The responsibility for managing trust anchors falls to some combination of browser vendors, OS vendors, users and, perhaps, IT support departments.

² This article is only concerned with *server* authentication, so client-authentication properties are not addressed.

that name-and-port by an entity with certain attributes (e.g., FDIC insured bank, NASDAQ listed company).

Forensics and Privacy Properties

client privacy: that third parties cannot, without the willing participation of the client or the server, deduce what Web sites a given client has been connecting to (without special access to the client machine).

impostor discoverability: that the legitimate owner of a given domain name should be able to discover what servers are presenting themselves as belonging to that name.

server privacy: that third parties cannot, without the willing participation of the client or the server, deduce the existence of a given server without actually attempting a connection themselves.

local privacy: that someone with access to the client machine after a Web site has been visited cannot deduce facts about what sites have been visited by the client. Of course, if a proposal includes client-side data, a user should be able to “clear history” as they can with current mechanisms like cookies. However, this presumably degrades the improvements to authentication. So this property refers to privacy concerns assuming that this clearing of authentication history has not happened.

Usability Properties

minimal false positives: that the client seldom refuses a connection or presents the user with an error message when the server it is connected to is indeed the right one. Ideally the client never decides to disbelieve that the entity with which it is communicating has the proper identity when, in fact, it does. This is the ideal. In reality, we can only try to minimize the number of false positives. As noted elsewhere, encountering many false positives conditions users to simply override-and-accept. Conversely, users that err on the side of caution are actually discouraged from connecting to legitimate servers. Note that this is described as a “false positive” because of the analogy to medicine — the test says there is a problem when, in fact, everything is OK.

maximal actionable information: that the client should have good, actionable information to present to the user, information that allows for good decisions. This includes information that allows for a reasonable plan of action in case the user decides not to override-and-accept.

minimal user trust: that the trust users must place in other entities is minimized. Note: This is a complex issue, since it involves not only how many entities need to be trusted, but also to what extent they are trusted, and how bad things could be if some number of them colluded.

user control: that the user is allowed to override decisions, determine who or what to trust, etc. This means the system must allow for it. Particular clients might not.

minimal server trust: that the trust server operators must place in other entities is minimized.

minimal server roadblocks: that setting up a TLS server is not overly burdensome. Already, lots of people and organizations have difficulty doing it right. Ideally improved authentication mechanisms should not set up too many barriers, technical, logistical or financial, to organizations that want to set up Web sites that use them.

Implementation/Infrastructure/Pragmatics Properties

incremental: that there are benefits (progress toward other properties) for those that opt-in without requiring that everyone participates.

minimal-impact: that big changes to the Internet architecture are not required. Client-only solutions, for instance are very low impact, as are server-only. When both have to change, when third parties are involved and so forth, the impact grows.

no-break: that things that currently work relying on relatively common practices (e.g., local network TLS connections, changing domain ownership, Web-hosting services, user-level TLS servers on hosts, etc.) do not break under the new proposal.

scalable/maintainable/robust: that the system works at Internet scale and can function over a long period of time. For example, if keys are lost or no longer secure: Can they be changed? The system should function reasonably in the face of outages and attacks.

resource-friendly: that adopting the proposal does not slow communication too much, or require too much CPU time or memory. Resource-constrained mobile devices must be able to participate.

realistic: that the proposal does not make unreasonable assumptions or demands on individuals or society. For example, expecting organizations with no suitable financial or other incentive to run big servers might not be realistic.

In each of the following five sections, we briefly describe a well-known proposal for improving authentication on the secure Web and analyze that proposal in terms of the properties described above. Each property is listed, along with commentary on the proposal in terms of that property. A small “gauge” icon accompanies each property to give a quick indication of whether the proposal positively or negatively affects a certain property, or has no substantial impact. For the first three categories—Authentication, Forensics and Privacy and Usability — the gauge value is to be understood as relative to the current secure Web environment. For the last category

— Implementation/Infrastructure/ Pragmatics—gauge values are to be understood as relative to all the other proposals for improvements. The gauge values are simply manifestations of qualitative judgments, not true quantitative data. The hope is that when combined in tables as in the following sections, they are a concise means to provide insight into the strengths, weaknesses and purpose of any given proposal. However, it is the accompanying commentary that provides the actual analysis.

3 DNS-based authentication of named entities (DANE)

DNS-based authentication of named entities (DANE) uses DNSSEC to make assertions that constrain valid certificate chains. These assertions can specify the end-entity certificate or public key, the trust anchor certificate or public key, or an intermediate certificate or public key. By using DNSSEC to distribute these assertions, clients can guarantee that the assertions really belong to the domain name in question. Thus, DANE is a mechanism that provides very strong *domain name authentication*. DANE depends on DNSSEC, and DNSSEC adoption seems to be proceeding very slowly. DANE (see RFC6698 [6]) adds a new record type to DNS, the TLSA resource record, which allows the nameserver that is authoritative for a given domain name to make assertions tied to a pairing of the name with a port-and-protocol. These assertions are of the following form:

(usage, selector, matching type, certificate association data)

where $usage \in \{0, 1, 2, 3\}$, $selector \in \{0, 1\}$, matching type $\in \{0, 1\}$ and the “certificate association data” can be a certificate, the Public KeyInfoField of a certificate or a hash value.

The semantics of DANE assertions are essentially this: The certificate chain is constrained in usage 0 to contain a given certificate/public key and in usage 1 to start with a given certificate/public key (which specifies the end entity). For both, the client is still required to validate the certificate chain up to a client-trusted anchor. Usages 2 and 3 are similar, except that the client’s role in determining trust is eliminated. The certificate chain is constrained in usage 2 to contain a given certificate/public key as trust anchor, and that trust anchor must be trusted by the client for purposes of that validation, regardless of the client’s current trust store. The certificate chain is constrained in usage 3 to have a given end-entity certificate/public key, and if it does, the client must accept the connection without doing certification path validation of the chain.


Because it is tied to DNSSEC, DANE’s pragmatic outlook is tied to DNSSEC adoption. It is worth clarifying the sense

in which DANE offers something more than DNSSEC alone would offer. A client relying on DNSSEC to resolve a given hostname to an IP address has a strong guarantee that the IP address it uses for that domain name is correct. However, it has no guarantees that the entity it connects to is a really a host properly associated with that IP address. After all, the attacker could be corrupting routing rather than DNS. With DANE, however, the client has stronger guarantees that the host it is connected to is properly associated with the given domain name.


Usage 2 and 3 assertions are potentially problematic. With usage 3, certification path validation does not occur, i.e., if the end-entity certificate presented to the client matches the certificate in the DANE assertion, the connection is accepted. A usage 2 DANE response mandates a certain trust anchor for validation, and mandates that it be trusted—regardless of whether it is currently in the client’s trust store. Both deny the user (or system administrator) the opportunity to make trust decisions. The security issue here is that without usages 2 and 3 both DNSSEC/DANE and the CA system have to be defeated in a successful attack, while with them an attacker that can successfully subvert DNSSEC can successfully pull off a man-in-the-middle attack. Defense in depth has been lost.

 **continuity:**—


 **domain name auth.:** This is DANE’s strength.

 **higher-level auth.:** Usages 2 and 3 response bypasses certification path validation, so information in certificates is less trustworthy than in the current system.


 **attribute auth.:**—


 **client privacy:** Generally, having to contact a third-party server is a client privacy concern. However, clients would almost always be contacting a DNS server for name resolution anyway, so it is not really a concern here.




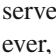



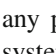
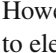

 **impostor discoverability:**—

 **server privacy:** Normally a DNS server would store the IP address associated with a given name and nothing more. DANE records include a port as well as IP, so the fact that a particular host is running a TLS server at a particular port number is then known to the DNS server.

 **local privacy:**—

 **minimal false positives:** DANE provides mechanisms (usages 3 and 4) by which a certificate that does not chain to a trust anchor would be accepted without any error or warning, which reduces false positives, although it can also defeat authentication, as pointed out above.

 **maximum actionable info:** If a site uses DANE and the client issues an error, the DANE assertion itself provides extra information about what public key / certificate should have been expected.

-  **minimal user trust:** The user does not need to trust the CAs as much, but they put more trust in DNS.
-  **user control:** With usages 2 and 3, there are situations in which a connection will be accepted regardless of the user's trust anchor settings.
-  **minimal server trust:** With DANE, the Web server operator puts even more trust in the DNS; however, CAs do not need to be trusted as much.
-  **minimal server roadblocks:** To use DANE, the Web server operator requires the cooperation of the DNS administrator. To also ensure that clients that do not support DANE are not locked out, a certificate from a CA would also be required.
-  **incremental:** Both Web clients and sites (though not really Web servers) need to change for DANE to work, any pair that supports DANE gain the security of the system, regardless of whether it is adopted elsewhere. However, it is not enough for just one of the two parties to elect to participate.
-  **minimal-impact:** It is not enough for client and server to change, the infrastructure has to change if DNSSEC is not already available.
-  **no-break:**—
-  **scal./maint./robust:**—
-  **resource-friendly:** DNSSEC has some overhead (see [7]) and more signed info will be sent when DANE is used than would be sent using DNSSEC solely for name resolution.
-  **realistic:** Given that DNSSEC is in use, DANE is quite realistic. The question is whether it is realistic to expect significant DNSSEC adoption anytime soon.

4 Certificate Transparency

Certificate Transparency is described in RFC6962 [8]. Its primary purpose is to provide “impostor discoverability.” The basic idea is this: If there was a public logfile of all certificates issued, then domain name holders could view the public logfile to root out bogus certificates for their domain names and, as the proposal puts it, “invoke existing business mechanisms for dealing with misissued certificates.” If TLS clients all agree to reject any certificate not recorded in the public logfile, attackers would be forced to record their forged certificates in the logfile where, hopefully, server/domain owners would observe the bogus certs and do something about it. While this basic idea is straightforward, realizing it in a secure way is non-trivial. The Certificate Transparency proposal is somewhat complicated in terms of the number of entities involved: In addition to servers and clients, there are logfile servers, trusted auditors and logfile monitors (see

Fig. 1). So the description that follows adds these various pieces in small steps.

Step 1: We first consider how clients determine whether a certificate record is in the logfile. Of course, the client could contact the logfile server and ask. Even if that could be done in an utterly secure and authenticated manner, there are still two issues: (1) contacting the third party has a performance and availability concern and (2) letting the logfile server know every domain name you want to contact has privacy concerns. Therefore, the proposal calls for a different approach. Clients have preloaded/out-of-band-received public keys for the logfile server. The TLS server is supposed to send a “signed certificate timestamp” (SCT) along with the certificate, which is essentially a hash of the certificate concatenated with a timestamp, signed by the logfile server. This gives the client something that it can verify quickly, without any third-party communication, so it addresses both concerns 1 and 2. IANA has issued a value for the TLS SCT extension. For technical reasons beyond the scope of this overview, the SCT is issued before the certificate is logged. However, the SCT contains an additional field with a value called the Maximum Merge Delay (MMD). Implicit in the SCT is a promise by the logfile server that the time between when the SCT was issued and when it is logged will not exceed that MMD value.

Step 2: If the logfile is misbehaving or if it has been compromised or its private key stolen or broken, clients could get forged SCTs. In other words, they could be accepting certificates that were not actually logged. To address this, the proposal calls for “trusted auditors” that clients are supposed to submit SCTs to, in order to keep tabs on the logfile server and make sure it really is reporting the submitted SCT as part of the log. The RFC mentions having the client do this (asynchronously, so as not to take the performance hit), but (1) that has all the same privacy concerns and (2) the logfile server could systematically lie to that one host. So it makes more sense to introduce trusted auditors into the system. It is unclear who or what auditors are notifying in case they detect a misbehaving logfile server, nor is it clear what the plan would be were a logfile server discovered to be misbehaving. Recovering from that situation could be quite challenging.

Step 3: Another way that a logfile server could misbehave is by modifying past entries in the log. For instance, maybe a bogus certificate gets a real SCT from the logfile server and is in the logfile (so the auditor does not see any trouble) but then after the attack the logfile entry gets erased, and all this happens before the domain name owner gets a chance to check for any new entries in the logfile for his domain name. This would defeat the whole purpose of Certificate Transparency. Therefore, the log-

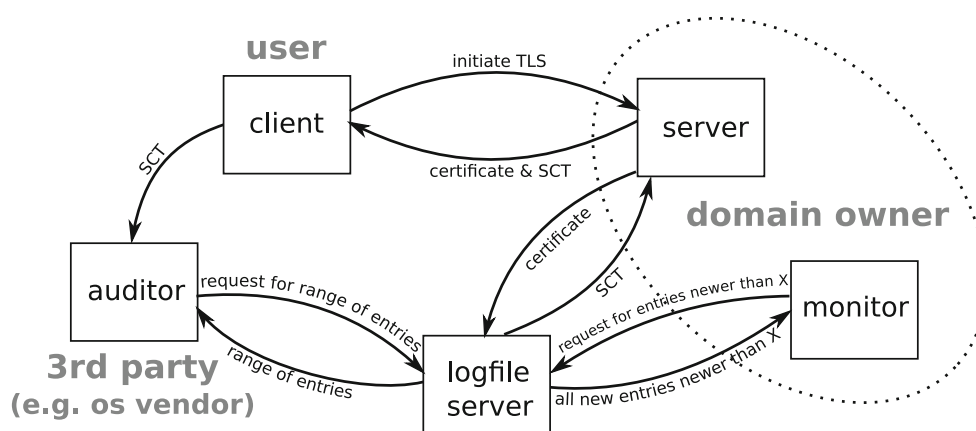


Fig. 1 Diagram illustrating CT. A full picture would show multiple logfile servers

file is *append-only*—once an entry is there, it is there forever. This is done with Merkle trees, which provide a mechanism whereby anyone observing the logfile server could detect modifications or erasures of past entries. Of course, some entity has to bother to make these checks, so the proposal calls for “logfile monitors,” which would periodically query the logfile server and check that it was really operating in append-only mode. These might do double duty by also checking for new logfile entries for domains the host is interested in.

Step 4: Finally, the proposal envisages not one, but multiple logfile servers. To protect against denial-of-service attacks, in which the attacker floods a logfile server with bogus certificates to be logged, the proposal suggests that each logfile server would publish a list of root CAs, and it would only log entries that validate via a chain up to one of the CAs in the list. The proposal also calls for “gossiping” to root out misbehaving logfile servers. However, no details on the gossiping protocol are given.

Certificate Transparency provides impostor discoverability. This is a benefit to server/domain name owners, but only secondarily a benefit to users. It provides no benefit for the initial targets of attacks, but it does offer a potential benefit to the larger user community, in as much as a vigilant server/domain name owner may notice the attacker and take steps to shut him down. The proposal takes great pains to ensure that entities that care to do so can monitor the activities of logfile servers in order to ensure that they are being honest, so that logfile operators do not need to be trusted. There are, however, some issues to consider.











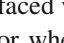
The purpose of the logfile monitors is to ensure that the logfile servers behave properly. However, once again, it is unclear how to deal with a logfile server that has misbehaved. It could be blacklisted somehow, but it is not clear what to make of the SCTs it had previously issued. Web

servers would be sending them out, potentially unaware that it was no longer trusted. Perhaps a bigger question is what to do with logfile servers that are not purposefully acting badly, but fail to meet an obligation—for example a logfile server that does not get the SCT into the log within the window specified by the MMD because of an attack, or a simple programming or administrative error. Simply blacklisting such a server seems highly undesirable. An alternative would be to rollback the log to the point of the error, but that is a problem because all the legitimate SCTs that had been issued in between issuing this SCT and noticing that the merge deadline had been missed would then be invalidated. Some mechanism is required to deal with this gracefully.

The proposal does not address how logfile public keys are distributed and updated. It seems that we end up in a similar situation as with CAs, namely that some arbitrary list of trusted logfile servers is preloaded into the browser. There is then the potential for even more certificate-related error messages, since a client could receive a certificate that is actually OK, but receive an SCT along with it that refers to a logfile server that is not trusted by the client (or for which the public key stored in the client is too old or too new).










It is not clear why users/clients would opt to submit SCTs to auditors. Collectively, there is the benefit that attackers could be discovered and eventually dealt with. But for the individual user there is little short-term benefit, and there is definitely a risk to privacy. If the intended model is that browser vendors would run their own trusted auditors, the privacy issue is mitigated, since their users are essentially trusting them completely anyway ... at least for Google, Microsoft and Apple. Less so for open source projects like Firefox, where users may put their faith in the “many eyes” that are supposedly on the source code. How a “trusted auditor” run by the Mozilla foundation is set up would not be subject to all those eyes, making it easier for a single individual or small group to misuse it than to introduce errors

into the Firefox codebase. Below is a summary analysis of CT.³

- + **continuity**: no real first-order effect.
- + **domain name authentication**: no real first-order effect.
- + **higher-level authentication**:—
- + **attribute authentication**:—
- + **client privacy** : The auditor sees every secure site the client connects to.
- + **impostor discoverability** : This is the whole point of CT!
- + **server privacy** : A legitimate server has to announce its presence by submitting to a logfile server.
- + **local privacy**: No additional data are stored locally.
- + **minimal false positives** : With CT, users could be faced with errors when valid certificates are not logged, or when SCTs are sent to clients that do not have that logfile server's public key.
- + **maximal actionable information** : When a cert's trust anchor is untrusted by the client, but the cert is logged, the user at least knows that the cert has been available for scrutiny and for how long. Otherwise, the user will know that it is unlogged (which is more suspicious).
- + **minimal user trust** : On the one hand, CT means users do not need to place so much trust in CAs (that

³ Draft revisions of the RFC address some of the issues raised here. This evaluation notes server privacy as an issue—a legitimate server needs to announce its presence by submitting its certificate to a public logfile server. To address this, draft revision 3 of the RFC includes a mechanism for redacting portions of the domain name in the certificate information submitted to a logfile server. For example, if the domain name in the certificate was super.secret.example.com, the information submitted to the logfile server might be `(PRIVATE).example.com`. Another mechanism added in the draft that addresses this problem is logging a name-constrained intermediate authority, along with a field that explicitly allows the SCT for the intermediate authority to stand in lieu of an SCT for the end-entity certificate. Thus, the situation above might be handled by having super.secret.example.com send an SCT for an intermediate CA constrained to example.com. Concerns raised here regarding the Minimal Impact and Minimal Server Roadblocks Properties are addressed in draft revision 3 by providing a mechanism for including a server's SCT in its certificate. This way, the server/server owner does not necessarily need to change or do anything different in order for CT to be used. Instead, CA's could make sure SCTs are bundled in the certificates they issue, and servers simply send those certificates as they always do. Of course, this creates a chicken-and-egg problem: the CA needs the SCT to create the certificate, but the logfile server needs the certificate to create the SCT. To deal with this, draft revision 3 allows CA's to submit "pre-certificates" to the logfile server, which contain enough information for the logfile server to create an SCT. The SCT gets sent back to the CA, which then can complete the certificate. Because these mechanisms are only described in draft revisions under very active development, we are not including them in our analysis.

is the "transparency"), but since a logfile server could essentially blackball a site by refusing to issue a SCT for it, users have to trust them to behave honorably. If there is one (or few) logfile servers for your client, that could become a problem.

- + **user control**:—
- + **minimal server trust** : Server/domain owners do not need to place as much trust in CAs.
- + **minimal server roadblocks** : Server/domain owners have to submit their certs to a logfile server and have to find one that supports their trust anchor.
- + **incremental** : Adoption is a major issue. If clients do participate, all sorts of legitimate sites will suddenly stop working, and users will get swamped in false positives.
- + **minimal-impact**: Both clients and Web servers need to change in order for CT to work, and a lot of additional infrastructure and new kinds of servers needs to be created.
- + **no-break** : How will this work in local network only situations? Will organizations be forced to run logfile servers inside their local networks? How about enterprise trust anchors? None of the usual logfile servers will support them, of course, so would such an enterprise need to run its own logfile server?
- + **scal./maint./robust** : lots of questions: What happens when logfiles make errors or are found to be acting improperly? How can logfile server keys change? How are logfile server keys distributed to clients? There are some significant maintenance problems!
- + **resource-friendly**: Not a lot of extra burden on client or server; although clients have to report to auditors, they do it asynchronously.
- + **realistic** It is unclear what would motivate operators to stand up logfile servers, monitors or auditors.

5 An HTTP extension for public key pinning (HPKP)

At its most basic, "pinning" just means hard-coding or caching cert/public key (or the hash of the cert/public key) in a client, and requiring the cert/public key received at connection time to match what is currently "pinned". More flexibly, the client might pin a set of certs/public keys, or pin the cert/public key of an intermediate element of the certificate chain, both of which allow the end-entity cert/public key to change in a controlled manner. Essentially, pinning is a commitment that the user would not allow certs/public keys to change. What is interesting is looking at the question of who controls whether, when and what pinning takes place. Pinning could be directed by (1) the user, (2) the client (e.g., hard-

coded pins, or pinning that would be updated by the client “calling home” or calling an external service, or a policy of caching certs/public keys after an initial unpinned connection) or (3) the server (directing pinning for itself or for subdomains). An example of (1) is when a user preloads or chooses to accept an ssh public key. An example of (2) is when applications (like Chrome) have preloaded pins or call back home to get new pinning directives. If a Web site were to send pinning directives to the client that would be an example of (3). Pinning, obviously, is a mechanism for providing authentication *continuity*.

The IETF draft document draft-ietf-websec-key-pinning (at the time of this writing in revision 21) [9] proposes an HTTP extension, HPKP, that allows the server to direct the pinning performed by the client—i.e., it is an example of the category (3) type of pinning described above. In this proposal, the server sends clients HTTP directives (the proposed extension) providing (hash-algorithm, hash-of-public-key) pairs that are to be pinned. The client saves this pin information indexed by the domain name it used in creating the connection. On subsequent connections to the same name, the client then checks whether any hash value in the set of pins is matched by a hash of any of the public keys in the certifying chain. If so, the client continues as normal. If not, there is an error. A hash-algorithm + hash-of-public-key pair must be accompanied by a “max-age” value and may be accompanied by a “report-uri” value. The max-age value instructs the client to keep the pin for a certain time. The report-uri gives a URI that is to be used by clients to report pinning errors for that domain name. An additional assertion may be sent that directs the client to apply the pin not only to the server’s domain name, but to all of its subdomains as well.

The obvious benefit of HPKP is the *continuity authentication* it provides. When a user connects to a server often enough (meaning that the time between visits is less than the max age) with the same client, man-in-the-middle attacks should be detected. Because the server directs the pinning and because sets of pins are allowed and intermediate public keys can be pinned, servers can pull off planned transitions to new public keys gracefully. As will be described in more detail below, the proposal is very good in terms of the Usability Properties. Among the Implementation/Infrastructure/Pragmatics properties, the only real concern is the extra resources required by a participating client—namely that a potentially large number of pins will have to be stored, which could be problematic for resource-constrained clients. There is an especial concern that a malicious site could flood the pin store and use up all the available space. The specification could perhaps be modified to bound the storage given a (non-top-level) domain, or reclaim space from the non-top-level domain with the largest storage footprint. Maintainability questions surrounding unplanned key transitions are answered by requiring servers to pin a “backup

key,” which is a key to be used in case the current public key is compromised and needs to be revoked and its use discontinued.

Next we consider Forensics and Privacy Properties. What should be another obvious benefit of HPKP is *impostor discoverability*. This is, after all, the point of the reporting mechanism provided by the report-uri directive. But to what degree will HPKP really provide this property? In the case of a man-in-the-middle attack (which is what an “impostor” really has to do), the client will be provided with a domain name X, and the attacker will somehow arrange things so that the client will think it is communicating with the host properly referred to by that name when, in fact, it is communicating with some other host—for example by disrupting routing. If the client receives a certificate chain and it does not match what is pinned for the name X, the client is supposed to send information about the pinning error to the URI in the report-uri directive. However, it seems quite likely that this message will never arrive at its destination given that the attacker is already subverting network traffic to carry out the man-in-the-middle attack. So the only case in which this would actually have its desired effect is when the attacker was unwary enough to allow the reporting message through. We note that this could be remedied by having clients send these reports at exponentially decaying intervals—perhaps until a signed acknowledgment is received. As long as the attack is not permanent, the report should eventually get through. To avoid flooding-style attacks, a carefully analyzed approach that looks at domain relationships and drops multiple error reports from the same (non-top-level) parent domain should be investigated.

There are a variety of ways clients may end up with pins that do not match the public key presented by a legitimate server. A domain name may change hands without the willing cooperation of the party losing the domain. Both primary and backup keys could be lost. An attacker manages to pull off a successful man-in-the-middle attack for a period of time on a site that does not use pinning, and puts a “poison pin” in the browser of all clients that connect during that time, with a very big max age. In all these cases, there is actually a hole in the DNS namespace—a domain name that, for a large number of clients, is unusable for HTTPS connections. This is a potentially serious problem.

Finally, HPKP breaks with the general design principle of separating concerns, and the specific cryptographic principle that different security properties should be safeguarded by different keys (see, for example, Section 5.2 of [10]). Consider the following three scenarios:







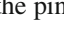

Scenario 1: Suppose a Web server instructs clients to pin the key from the end-entity certificate. Then the key is used both to provide continuity authentication and to provide confidentiality in sending the TLS premaster secret

from the client to the server.⁴ Note in this case that changing the end entity public key is a problem, because the pin becomes invalid—although with some care and sufficient time, this can be dealt with using the “backup key” and max-age directive.


Scenario 2: Suppose instead that the client is instructed to pin the public key from a CA certificate. In this case, changing end-entity certificates will not violate the pin, as long as the same CA signing key is the root of the new end-entity certificate’s validation chain. However, now that signing key is being used to provide both domain name authentication and continuity authentication. Note that in this case the client with the pin is no longer protected against a compromise of the CA key (as it was in the previous case), because that key can sign a fraudulent end entity certificate that will validate with the same certificate chain, thus matching the pin.


Scenario 3: Suppose that the Web server operator, wanting the flexibility to change end-entity public keys but not wanting to have to put all his trust in the CA, gets a signing certificate (hopefully with a limited scope namespace!) and uses it to issue certificates with different public keys as they are needed. Once again, that signing key is being used to provide both domain name authentication and continuity authentication. Note that in this case, the Web server operator is not reliant on the CA for continuity authentication. On the other hand, compromise of the Web server’s private signing key is sufficient to carry out an attack.


In this sense TACK, which is described in the next section, is designed well, because a separate key (the “Tack Signing Key”) provides continuity authentication, and the certificate chain keys provide (as they are supposed to) domain name and higher-level authentication.

-  **continuity:** This is the point of HPKP, although the fact that pins expire limit this property.
-  **domain name auth:**—
-  **higher-level auth:**—
-  **attribute auth:**—
-  **client privacy:** attacks referenced above.
-  **impostor discoverability :** The report-uri directive provides this but, for reasons described above, it is unclear how effectively.
-  **server privacy:**—
-  **local privacy:** There is forensics information in the pins themselves, and simply clearing the pinstore as


you would the browser’s cache is not an attractive option because the user would lose security.


-  **minimal false positives:** Sites the user visits often should not generate false positives, but those visited infrequently might does not address first use.


-  **max actionable information:** For some errors, pinned information shows what public key user should expect to see this can be actionable.


-  **minimal user trust:** The user trusts server X’s pinning directives, but these only pertain to server X itself, so that is a pretty low level of trust. Pinning reduces the trust that must be placed in CAs.

-  **user control:**—


-  **minimal server trust:** A server making use of HPKP needs very little trust in CAs after a given client’s first connection.


-  **minimal server roadblocks:** The server does not need to rely on, or coordinate with, outside entities to use HPKP in a limited way, but using it flexibly so that new keys can be introduced in a reasonable way may require getting a signing cert, which is a much bigger deal.


-  **incremental:** With a conforming client, any participating site is more secure (for the user); however, both client and server must participate.

-  **minimal-impact:** Clients and Web servers need small modifications.

-  **no-break:**—

-  **scal./maint./robust:** Some small concern about how domain names change hands.

-  **resource-friendly:** All the pin information needs to be stored, which could be problematic for memory-constrained clients. There are also concerns about attacks on memory resources via HPKP.

-  **realistic:** HPKP only requires buy-in from browser vendors to get started. Given that this is a Google draft, that buy-in might be there.

5.1 Trust assertions for certificate keys

Trust Assertions for Certificate Keys (TACK) is a 2013 project that proposes a TLS extension to help guarantee continuity through server-directed pinning. The concept is defined in an expired Internet Draft [11]; more information is available at <http://tack.io>. Although its future is unclear, it is an interesting proposal to consider here—especially because with similar goals to HPKP, TACK makes some different design decisions. In TACK, both the client and the server need to participate, but no third parties are involved. The basic idea is that the entity controlling a TLS server, the “owner,” creates a long-term public/private key pair called the Tack Signing Key (TSK) that provides the cryptographic

⁴ Depending on the TLS cipher-suite, the second use of the key may instead be to verify that an ephemeral key belongs to the server.

basis for continuity authentication. A “tack” is an assertion, signed with the TSK private key, that the owner approves of (or revokes, but that would not be covered in much detail here) a particular public key for TLS use. When a client initiates a TLS connection to a host with name *N*, it receives both a tack *T* and a TLS public key *K* (usually in a certificate). The client checks that the TLS public key asserted in the tack matches the received TLS public key *K*, of course, but the important part is what comes next. If the client has not connected to this server before, it “pins” the TSK public key (which is included in the tack) to the name, i.e., it stores the pair (*N*, TSK_pub). If, however, the client has connected to this server before, then there is already a pin for the name *N*, and thus, the client knows the identity of the owner from previous connections to *N*—it is given by the pinned TSK_pub. So the client can verify that the identity of the owner is the same now as it was for previous connections by checking that the TSK public key in the tack matches the TSK public key in the pin. With this mechanism, even if the TLS public key changes, even if the whole certificate chain sent by the TLS server changes, the client still can be confident of “continuity,” i.e., confident that the owner from previous connections approves the TLS public key to be used for this connection.

The actual proposal is trickier than the simplified overview above. And the implications of the mechanism described in the proposal are quite subtle. There may be two pins for a given domain name in the pin store, though no more than two, and each pin can be in one of two states, “active” or “inactive.” Moreover, pins include start and end times that determine the pin’s state (a pin can “age out” and make the transition from active to inaction). Also, there may be two tacks in the ServerHello, though no more than two, and each tack is sent along with a flag labeling it as either “active” or “inactive.” The state of the pin matters: only “active” pins actually count as verifying continuity authentication. The active/inactive status of the tack affects the state of the pin and does it in a way that is intertwined with the current time and the pin’s start and end time. For example, after 30 days without a connection, any active pin ages out to inactive. To top it off, each tack includes a “generation” number and a “max_generation” number, and each pin contains a “max_generation” number. These get used to “revoke” TLS public keys. However, they also interact with the process for determining how tacks affect pin state.

When a TACK-supporting client sends the TLS Client Hello message, it includes a new TLS extension called “tack” with no associated extension data (and it must also send the domain name using server name indication). This signals to the server that the client wants to support TACK. The ServerHello message sent in response also includes a “tack” extension, but this time there is associated data. It consists of one or two “tacks” and a one byte “activation flags” field. A “tack” is a tuple

$$(\text{public_key}, \text{min_generation}, \text{generation}, \text{target_hash}, \text{signature})$$

where *public_key* is the TSK public key, *target_hash* is the hash of the TLS public key and signature is the signature of the first four fields signed with the TSK public key. The first tack is “active” if bit one of the activation flags is set. The second tack, if it exists, is “active” if bit two of the activation flags is set. A “pin” is a tuple

$$(\text{domain_name}, \text{initial_time}, \text{end_time}, \text{TSK_public_key}, \text{min_generation})$$

A key is “active” if $\text{current_time} < \text{end_time}$. The pin store contains at most two pins with the same values for *domain_name* and never contains two pins with the same values for *domain_name* and *TSK_public_key*.

When a connection is attempted and 0,1,2 tacks are received, a number of checks are performed, and a number of results are possible. Basically, we could get “certificate_revoked,” “contradicted,” “confirmed” or “unpinned” as results relevant to the current connection, and we may have changes to the pin store as well. Pins may be added and deleted, *end_times* may be updated, and *min_generations* may be updated. The precise mechanisms will not be covered in this summary of TACK, as it would require too much detail. Moreover, there are some minor ambiguities in the current proposal.

The primary value of TACK is that it provides continuity authentication. It is essentially orthogonal to the regular certificate mechanism, which is a nice property. It strengthens authentication both for sites that to rely on self-signed certificates and those relying on certificates signed by CAs. TACK and the HTTP Pinning Extension described above both provide server-directed pinning, so it is useful to enumerate some of the key differences between the two proposals.


1. TACK is a TLS extension; thus, it operates on the presentation layer, whereas HTTP Pinning operates on the application layer. This means that not only will Web servers and Web clients need to be modified to participate, but the TLS server and client code needs to change to support TACK. This means a bigger change to the current architecture.
2. TACK does not try to do any kind of back-reporting.
3. TACK TSK’s and TACK signatures are separate from the certificate chain. Conceptually, this is a big improvement over HTTP Pinning, since it follows the principle of using different keys to protect different properties—TSK for continuity authentication, certificate signing keys for domain and higher-level authentication.

The TACK proposal mentions that there are some potential privacy concerns, namely that the pins are another record of where one has browsed and that attackers could throw

fraudulent TSKs at a client and try to deduce based on the client's reactions whether it had connected to the real site previously.


The proposal also mentions that some sort of client gossiping or other kind of mechanism for sharing TSK pins could help with the problem of the initial connection, or connections initiated after a long enough period of time that the TSK pin has expired. If this were to be followed through, it would provide "community continuity authentication", as described earlier in the paper.


TACK suffers from the same potential problem as HTTP Pinning in terms of leaving clients essentially unable to connect to a given domain name because they have a "poison pin," for one reason or another. However, it is mitigated in TACK by the "aging out" policy, which guarantees that a pin that is not "refreshed" by receiving a reinforcing tack within a 30 day window becomes inactive, and therefore, it is as if there was no pin at all.

 **continuity:** The aging out process and the way it is affected by the frequency of visits make TACK's continuity authentication a bit less strong, though it mitigates other issues.

 **domain name auth:**—


 **higher-level auth:**—


 **attribute auth:**—


 **client privacy:** As described above, an attacker may be able to use TACK-based rejection of a connection to deduce that the client has visited a particular site before.


 **impostor discoverability :**—

 **server privacy:**—


 **local privacy:** There is forensics information in the pins themselves, and simply clearing the pinstore as you would the browser's cache is not an attractive option because the user would lose security.

 **minimal false positives:** Sites the user visits often should not generate false positives, but those visited infrequently might does not address first use.


 **max actionable information:** For some errors, pinned information shows what public key user should expect to see this can be actionable.


 **minimal user trust:** The user trusts server X's pinning directives, but these only pertain to server X itself, so that is a pretty low level of trust. Pinning reduces the trust that must be placed in CAs.


 **user control:**—


 **minimal server trust:** A server making use of TACK needs very little trust in CAs after a given client's first connection. Server operators place less trust in CAs


than in HPKP, since with TACK the operator controls the key that safeguards continuity authentication.


 **minimal server roadblocks:** TACK is self-contained (no interaction with third parties) for the server owner. The only real additional burden is the key-management that accompanies the TSKs.


 **incremental:** With a conforming client, any participating site is more secure (for the user); however, both client and server must participate.

 **minimal-impact:** Clients, Web servers and their underlying TLS libraries all need to change.

 **no-break:**—

 **scal./maint./robust:** Just small concerns about switching TSK keys how domain names change hands.

 **resource-friendly:** All the pin info needs to be stored, which could be problematic for memory-constrained clients. There is no "include subdomains" directive like HPKP has; however, there are never more than two public keys stored per name. There are also concerns about using TACK to attack memory resources.

 **realistic:** Adopting TACK requires changes not only to Web clients and servers, but also to underlying TLS libraries.

6 Perspectives

Perspectives is a circa 2008 Carnegie-Mellon project (see [12]). It is included in this survey because it has been influential, and serves as an exemplar of the "notary-based" approach to improving authentication. It is primarily about maintaining meaningful authentication when servers present bare public keys (which includes self-signed certificates). However, it has contributions to make even when CA-signed certificates are used. It seeks to provide what we have characterized as community continuity authentication.

Perspectives starts by assuming we have clients that do straightforward client-side pinning of TLS public keys, i.e., that maintain a map associating with each domain name the last TLS public key accepted for that name. Requiring the public key in the certificate sent by the server to match the pinned public key provides a guarantee of continuity authentication. However, this kind of pinning suffers from two problems:

1. On the first attempt by the client to connect to a given domain name, there is no pin, and therefore, no basis for making a "trust decision" and
2. When the TLS public key presented by a server does not match the pin, the client has no basis for determining whether the server legitimately associated with a particular name now has a different key, or whether the client

is actually communicating with the wrong server. This means that while client-side pinning makes it quite likely that a warning will pop up when an actual attack takes place, it also makes it quite likely that many, many warnings will pop up in benign circumstances. Generating lots of “false positive” warnings like this is a big problem. What Perspectives seeks to provide is a mechanism for determining, in an automated fashion, whether or not the presented public key belongs to the given domain name, so that warnings only make it to users when both pinning and Perspectives fail to validate the name–key association. Thus, the user should see far fewer “false positives,” i.e., fewer warnings in what are actually benign circumstances.

Perspectives is based on third parties called “notaries.” The idea is that there would be a number of “notary servers” out on the Internet connecting to network services just to see what kind of certificates they send to clients. Clients receiving a public key they have not seen before can query one or more notaries to see what keys they have observed for a particular name and use this extra information to make a trust decision. A Perspectives client caches the public keys it receives when it connects to a given server. As long as the public key it receives from a server matches what is cached, the application is happy. When the public key received from the server does not match, the client contacts a notary server and makes a trust decision based on the notary’s response.

In an ideal world, there would simply be one notary that would keep track through time of what TLS public key every (domain name, port number)-pair (n,p) was serving up. A client who accessed (n,p) for the first time, or which received a different certificate than last time, would contact the one notary and receive a signed response with the public keys observed for (n,p) along with the timespans over which each key was observed. The client would take that information and make a “trust decision.” How it arrives at a decision would be configurable. There are a few potential pitfalls here: (1) an attacker might be able to control some portion of the network surrounding the notary and send it the same forged certificate the victim receives and (2) if that one notary is compromised, all clients lose all their protection. To address these, Perspectives uses a scheme in which there are multiple notaries. A client polls k random notaries, and takes all their responses into account in making the “trust decision.”

Clients need to know what notaries are available and that is something that could easily change over time. To manage this, there is another kind of server, a Notary Authority, whose role is to publish lists of Notary Servers along with their public keys. This list is published periodically, so clients can update their individual lists. The published list is signed by the Notary Authority’s private key, so a client that is pro-


vided with the Notary Authority’s public key can validate the published list for itself.


Perspectives tries to reduce the number of “false positive” warnings and, in cases where warnings must be given, to increase the potentially actionable information available to the user. When a client connects for the first time to a server using a given domain name, notaries may provide the client with sufficient evidence of the trustworthiness of the certificate/public key presented by the server, so that the certificate/public key can be accepted without user interaction (and the resulting warning). For popular sites, the odds are pretty good that the notaries will be able to provide such evidence. If a certificate/public key does change, notaries may (once again) provide sufficient evidence of trustworthiness to clients who have cached the old certificate/public key. When a warning needs to be given, a Perspectives client has the certificate and date/time information from several notary servers, which can be processed to provide (hopefully) useful feedback to users.

One concern is that the Notary Authority represents a single point of failure. The Security Analysis in the paper (Sect. 6) neglects to analyze this case. Essentially, an attacker who gains control of the Notary Authority’s public key can send a client a fraudulent set of Notary Server names and public keys and use them to present the client with fake notary results. Another concern with Perspectives is that clients send notary servers the domain names of the servers they visit, which may violate the client privacy property. Similar concerns exist with the Certificate Transparency proposal. As with TACK and HTTP Pinning, there is a concern that pinning certificates or public keys could be too much of a burden on resource-constrained devices. Moreover, it may be possible for attackers to fill up clients’ pin stores and, in so doing, circumventing the protection of continuity authentication.


Perspectives is predicated on the idea that we should see community continuity in public keys. That is, not only should a single client be seeing the same public key over time as it repeatedly connects using a particular domain name, but any group of clients should see the same key when they use the same domain name. It is possible that servers would send different certificates to different clients based, perhaps, on the client’s IP address or client software version. Such a practice would invalidate a major assumption of the Perspectives project.

Finally, it is interesting to note that there is a big difference between Perspectives and the other proposals considered, in that Perspectives is not about mandating how trust decisions are made, rather it is about providing meaningful, actionable information—how that information gets incorporated into the decision-making process is left to the user or the client. That is a very nice “perspective” on the problem.


-  **continuity:** pinning provides individual continuity, notaries provide community continuity. Perspectives does not suffer from the aging out problem that HPKP and TACK have.


-  **domain name auth:**—


-  **higher-level auth:**—


-  **attribute auth:**—


-  **client privacy:** Collectively, the notaries get to see every server the client connects to.


-  **impostor discoverability :** Site owners could query notaries, thereby identifying impostors—as long as the impostor was discovered by some notary server.


-  **server privacy:** Notary servers could be used as an anonymizing proxy for network reconnaissance; i.e., one could query servers to determine whether a certain host had a server running on a certain port, and the server operator would have no idea that this kind of probing was taking place.


-  **local privacy:** Perspectives assumes client-directed pinning, and there is extra forensics information in the pins themselves.


-  **minimal false positives:** Notary information can allow the client to deal with pin or certificate errors without user intervention


-  **max actionable information:** The pinned values and the notary observations are all available as pertinent information when deciding what to do in response to a warning.


-  **minimal user trust:** Notaries need to be trusted to some extent, the notary server to a greater extent, but much less trust in CAs is required.


-  **user control:** The user/client can implement any policy they want to determine how pin and Perspectives information are combined to make a decision.

-  **minimal server trust:** Required trust in CAs is lessened, but servers must trust notaries not to misrepresent their keys to clients.


-  **minimal server roadblocks:** Servers and server operators do not have to do anything.


-  **incremental:** With a notary authority and some notary servers in place, only clients need to change, and any client that makes use of pinning + perspectives gets all the benefits, even if it is the only one doing it.


-  **minimal-impact:** Only clients need to change.

-  **no-break:** There is concern over how Perspectives clients would deal with internal network resources, since no notary server would be able to query that resource. This needs some thinking. Also a concern is the assumption that, when things are operating properly, for a given domain name every client would be see the same

public key. See, for example, [13] where this assumption is violated. This is the case of an organization rolling out a new public key gradually, or piecemeal.

-  **scal./maint./robust:** Some small concern about how domain names change hands.

-  **resource-friendly:** All the pin info needs to be stored. The delay caused by querying notaries is only encountered in situations where the pin value and the TLS public key do not match, which would happen very infrequently.

-  **realistic:** The main concern is the lack of motivation for providing notary servers. However, the Berkeley ICSI [14] is doing more or less the same thing, so there is apparently some motivation.

7 Comparisons and conclusions

In this paper, we have presented a framework for evaluating proposals to improve the current condition of authentication for the secure Web. Five such proposals were considered and evaluated according to this framework. At a very high level, we might summarize the evaluations as follows:

- DANE offers the prospect of providing strong guarantees of domain name authentication. However, with “usage values” 2 and 3 it eliminates the defense in depth that the system of certificate authorities was supposed to bring. Moreover, DANE is built on top of DNSSEC, and DNSSEC adoption has not progressed very quickly.
- Certificate Transparency offers a mechanism that allows domain/server owners to detect attackers that are impersonating their sites. However, it has a number of pragmatic problems, as detailed above, and may increase the number of “false positive” warnings experienced by users.
- The HTTP Extension for server-directed pinning and TACK are both designed to provide continuity authentication. However, TACK pins a special “Tack Signing Key” rather than one of the public keys from the certificate chain used to deliver the TLS public key. As described above, this is a better, more principled design, although there are some pragmatic advantages to the HTTP Pinning approach. Both TACK and HTTP Pinning suffer from the “poison pin” problem, namely that once the wrong pin gets in a client’s pin store, the client will present the user with a “false positive” error message.
- Perspectives provides continuity authentication, but unlike HTTP Pinning and TACK, it delivers continuity from the community perspective. In fact, Perspectives assumes that the client is using client-directed pinning (i.e., caching end-entity public keys) to provide continuity

authentication on the level of the individual. When that fails, the host queries a number of notary servers to gain the community perspective on what public key should be expected for the given domain name. Perspectives requires a fair amount of trust in the Notary Authority, and there are privacy concerns with the way notary servers see every TLS-protected site visited by a client. Another concern is whether the implicit assumption of Perspectives—that all clients should see the same public key when connecting to the same domain name—is justified.

DANE (ignoring usage 2 and 3), Certificate Transparency, Perspectives and either one of TACK and HTTP Pinning are pretty much orthogonal to one another, meaning that they could be used in combination without interfering with one another or overlapping in what they provide. In fact, used in conjunction we would have stronger domain name authentication (from DANE), continuity authentication in the individual sense (from TACK or HTTP Pinning), continuity authentication in the community sense (from Perspectives) and improved impostor discoverability (from Certificate Transparency).

We finish up by looking beyond the proposals analyzed in this paper and asking whether the analysis suggests new ideas to investigate or has any other interesting implications. The first thing we would like to point out is that instead of viewing proposals like these as trying to “fix” authentication for the TLS-protected Web, we should evaluate a proposal by clearly understanding the authentication or forensics property it is trying to provide and analyzing the extent to which it provides that property along with the positive and negative impacts on the other properties that would result from adopting the proposal.

A second point is that once we stop looking for a single, monolithic, universal fix to authentication for the TLS-protected Web, the importance of “orthogonality” of proposals becomes quite clear, by which we mean that the adoption of a proposal would not interfere with existing mechanisms or other proposed improvements. When proposals are orthogonal they can be composed, and that strengthens authentication. DANE’s usages 2 and 3 are unfortunate precisely because they ruin orthogonality. Without usage 2 or 3, DANE and the current certificate infrastructure compose nicely.

A third point is that when we view these various mechanisms as providing evidence for one or more of the four authentication properties, we see each connection attempt as making a case for accepting the purported identity of the server on the other end. It might be reasonable to pin the “shape” of that evidence—i.e., what kind of evidence was presented. So, for example, suppose a user has been using client-directed pinning, and DANE (without usage 2

or 3) is used in conjunction with the usual certificate validation process. The user tries to go to <https://example.com> and there is an error—the public key in the certificate presented by the server does not match what the browser has pinned for the name example.com. However, the DANE record is validated, and the certificate chains to a trust anchor. The decision about whether to trust this server despite the pin mismatch is unclear. Now suppose that on prior connections, example.com had sent the client an EV certificate, and suppose the client had pinned that fact. It would not be at all unreasonable to base the trust decision on whether or not the certificate presented by the server is an EV certificate. Note that pinning the “shape” of the authentication evidence provided by a site has the really nice property that it actually provides increased security to sites that choose to employ stronger authentication evidence. In the example above, the client would have “pinned” the facts that example.com employs DANE and uses an EV certificate. Thus, an attack will generate a warning to the user unless the attacker both subverts DNSSEC and gets a fraudulent EV certificate for example.com. Taking control of DNS records might be enough to get an ordinary certificate for example.com, but it should not be enough to get an EV certificate. This kind of pinning could also make gradual adoption of some of these proposals easier. For example, if the client pinned the fact that a given site used Certificate Transparency (i.e., sent an SCT) in prior connections, then the client could be configured to *require* CT for that site from that point on, but not for other sites. This would eliminate the problem of clients being flooded with false positives. One of the interesting things about Perspectives [12] is that it explicitly presents itself as a mechanism for providing evidence about identity, not as a procedure that proscribes trust decisions. That is a powerful and flexible idea.

The fourth and final point is a suggestion that we feel falls out of this analysis. We start with the observation that if clients were to do client-directed pinning of end-entity certificates and servers would do Online Certificate Status Protocol (OCSP) stapling, then most of the time there would be strong authentication of the server on the basis of those two pieces of evidence alone, and the connection could proceed⁵. The process would be quick and involve very little overhead provided that the certificates match, and the point is that they

⁵ Client-directed pinning would ensure that the certificate had not changed since the client’s last connection to the site, and OCSP stapling would ensure that the certificate had not been revoked—at least as of some reasonably recent point in the past. Online Certificate Status Protocol (OCSP) stapling is a piece of the modern certificate infrastructure. It allows a server to send clients a message, signed by the relevant certificate authority, that asserts that as of a certain point in time, the server’s certificate has not been revoked. This is a nice alternative to contacting OCSP servers to check for revocation, or either pushing or pulling blacklists.

usually would. So the question really is how to deal with the infrequent situation in which the above is unable to confirm authentication. This can happen when a client connects with a given name for the first time, when a different end-entity certificate is sent by the server, or when an end-entity certificate is revoked. Making the right decision in these cases, and doing it to the greatest extent possible without user intervention, is crucial. However, since these situations are infrequent (as well as important), when they do arise it would be acceptable to have the client take substantially longer to make a decision, or to gather information to present to the user in case it is necessary. We suggest research into mechanisms or the development of standards that allow the client to collect a lot of relevant data in order to make a strong case for or against trusting the server. As a very simple example, suppose there was a standard way for a client that was not able to authenticate the server using the pinning and stapling mechanism above, to fetch additional certificates for the server. A client could implement a policy requiring that, in this event, it is able to fetch an additional certificate that contains the same TLS key, the same domain name, and chains to a different trust anchor (without cross-signing). This increases the difficulty of a man-in-the-middle attack significantly, since the attacker would have to obtain fraudulent certificates from two different trust anchors. This is orthogonal to other proposals, and it strengthens all of them. For example, HPKP has potential problems with unplanned key transitions. With a mechanism like this, an organization that is forced to deal with an unplanned key transition could have strong evidence (we have suggested multiple certificates as a possible form) that a client could fetch on that single connection for which HPKP broke. The client could be convinced with overwhelming evidence and accept the TLS connection — without user intervention. The delay caused by fetching and analyzing this extra evidence would only be incurred once; then, HPKP would suffice for subsequent connections.

Acknowledgments The first author's research related to this work was carried out at the National Security Agency under the Sabbatical program administered by the Mathematical Sciences Program. The US Government is authorized to reproduce and distribute reprints notwithstanding any copyright notation herein.

References

1. Dierks, T., Rescorla, E.: The transport layer security (TLS) protocol version 1.1. RFC 5246, RFC Editor (2006)
2. Housley, R., Santesson, S.: Update to directorystring processing in the internet X.509 public key infrastructure certificate and certificate revocation list (CRL) profile. RFC 5280, RFC Editor (2006)
3. Herley, C.: So long, and no thanks for the externalities: the rational rejection of security advice by users. In: Proceedings of the 2009 Workshop on New Security Paradigms Workshop, NSPW '09, New York, NY, USA, pp. 133–144. ACM (2009)
4. Sunshine, J., Egelman, S., Almuhamedi, H., Atri, N., Cranor, L.F.: Crying wolf: an empirical study of ssl warning effectiveness. In: Proceedings of the 18th Conference on USENIX Security Symposium, SSYM'09, Berkeley, CA, USA, pp. 399–416. USENIX Association (2009)
5. CA/Browser Forum. Guidelines for the issuance and management of extended validation certificates (2014). https://cabforum.org/wp-content/uploads/EV-V1_5_2Libre
6. Hoffman, P., Schlyter, J.: The DNS-based authentication of named entities (DANE) transport layer security (tls) protocol: TLSA. RFC 6698, RFC Editor (2012)
7. National Institute of Standards and Technology (NIST) <http://www.dnsops.gov/dnssec-perform.html>
8. Laurie, B., Langley, A., Kasper, E.: Certificate transparency. RFC 6962, RFC Editor (2013)
9. Evans, C., Palmer, C., Sleevi, R.: Public key pinning extension for http. Internet-Draft draft-ietf-websec-key-pinning-21, IETF Secretariat (2014)
10. Barker, E., Barker, W., Burr, W., Polk, W., Smid, M.: Recommendation for key management—part 1: general (revision 3). Technical Report NIST Special Publication 800–57, National Institute of Standards and Technology (2007)
11. Marlinspike, M., Perrin, T.: Trust assertions for certificate keys. Internet-Draft draft-perrin-tls-tack-02, IETF Secretariat (2013)
12. Wendlandt, D., Andersen, D.G., Perrig, A.: Perspectives: improving ssh-style host authentication with multi-path probing. In: USENIX 2008 Annual Technical Conference on Annual Technical Conference, ATC'08, Berkeley, CA, USA, pp. 321–334. USENIX Association (2008)
13. Kaminski, D.: These are not the certs you're looking for. Personal blog (2011). <http://dankaminsky.com/2011/08/31/notnotar/>
14. The ICSI Certificate Notary. <http://notary.icsi.berkeley.edu/>