



A Formal Model for ACME: Analyzing Domain Validation over Insecure Channels

Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Nadim Kobeissi

► **To cite this version:**

Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Nadim Kobeissi. A Formal Model for ACME: Analyzing Domain Validation over Insecure Channels. [Research Report] INRIA Paris; Microsoft Research Cambridge. 2016. <hal-01397439v2>

HAL Id: hal-01397439

<https://hal.inria.fr/hal-01397439v2>

Submitted on 16 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Formal Model for ACME: Analyzing Domain Validation over Insecure Channels

Karthikeyan Bhargavan¹, Antoine Delignat-Lavaud² and Nadim Kobeissi¹

¹ INRIA

{karthikeyan.bhargavan, nadim.kobeissi}@inria.fr

² Microsoft Research

antdl@microsoft.com

Abstract. Web traffic encryption has shifted from applying only to highly sensitive websites (such as banks) to a majority of all Web requests. Until recently, one of the main limiting factors for enabling HTTPS is the requirement to obtain a valid certificate from a trusted certification authority, a tedious process that typically involves fees and ad-hoc key generation, certificate request and domain validation procedures. To remove this barrier of entry, the Internet Security Research Group created Let’s Encrypt, a new non-profit certificate authority which uses a new protocol called Automatic Certificate Management Environment (ACME) to automate certificate management at all levels (request, validation, issuance, renewal, and revocation) between clients (website operators) and servers (certificate authority nodes). Let’s Encrypt’s success is measured by its issuance of over 12 million free certificates since its launch in April 2016.

In this paper, we survey the existing process for issuing domain-validated certificates in major certification authorities to build a security model of domain-validated certificate issuance. We then model the ACME protocol in the applied pi-calculus and verify its stated security goals against our threat model of domain validation. We compare the effective security of different domain validation methods and show that ACME can be secure under a stronger threat model than that of traditional CAs. We also uncover weaknesses in some flows of ACME 1.0 and propose verified improvements that have been adopted in the latest protocol draft submitted to the IETF.

1 Introduction

Since the dawn of HTTPS, being able to secure a public website with SSL or TLS required obtaining a signature for the website’s public certificate from a certificate authority [1] (CA). These certificate authorities had to be recognized by all major operating system and browser vendors to be legitimate entities that could attest for a reasonable link between a certificate and identity of the server or domain it claims to represent.

For example, all major operating systems ship with Symantec’s root certificate signing key as built-in and trusted. This symbolizes a permission by

these major players for Symantec to then act as a CA, which allows Alice, the owner of `AliceShop.com`, to ask Symantec to attest that some SSL certificate being served by this website is indeed identifying the legitimate server behind `AliceShop.com`. After Alice pays Symantec some verification fee, Symantec performs some check to verify that Alice and her web server indeed have the authority over `AliceShop.com`. If successful, Symantec then signs a certificate intended for that domain. Since the aforementioned operating systems already trust Symantec, this trust now extends towards Alice's certificate as being authentically representative of `AliceShop.com`.

The security of this trust model has always relied on the responsibility and trustworthiness of the CAs themselves, since a single malicious CA can issue arbitrary valid certificates for any website on the Internet. Each certificate authority is free to engineer different user sign-up, domain validation, certificate issuance and certificate renewal protocols of its own design. Since these ad-hoc protocols often operate over weak channels such as HTTP and DNS, with no strong notion of cryptographic client authentication, most of them can be considered secure only under relatively weak threat models, reducing user credentials to a web login, and domain validation to an email exchange.

The main guidelines controlling what type of domain validation CAs are allowed to apply are the recommendations in the CA/Browser Forum Baseline Requirements [2]. These requirements, which are adopted by ballot vote between the participating organizations, cover the definition of common notions such as domain common names (CNs), registration authorities (RAs) and differences between regular domain validation (DV) and extended validation (EV).

These guidelines have not proven sufficient for a well-regulated and well specified approach for domain validation: Mozilla was recently forced to remove WoSign [3] (and its subsidiary StartSSL, both major certificate authorities) from the certificate store of Firefox and all other Mozilla products due to a series of documented instances that range from the CA intentionally ignoring security best-practices for certificate issuance, to vulnerabilities allowing clients to obtain a signed certificate for any website of their choosing.

The lack of a standardized protocol operating under a well-defined threat model and with clear security goals for certificate issuance has so far prevented a systematic treatment of CA security using well-established formal methods.

In 2015, a consortium of high-profile organizations including Mozilla and the Electronic Frontier Foundation launched Let's Encrypt [4], a non-profit effort to specify, standardize and automate certificate issuance between web servers and certificate authorities, and to provide certificate issuance itself as a free-of-charge service. Since its launch in April 2016, Let's Encrypt has issued more than 12 million certificates [5] and has been linked to a general increase in HTTPS adoption across the Internet.

Let's Encrypt also introduces ACME [6], an automated domain validation and certificate issuance protocol that gives us for the first time a protocol that can act as a credible target for formal verification in the context of domain validation. ACME also removes almost entirely the human element from the process

of domain validation: the subsequently automated validation and issuance of millions of certificates further increases the necessity of a formal approach to the protocol involved.

In this paper, we formally specify, model and verify ACME using the automated protocol verifier ProVerif [7]. Against a classic symbolic protocol adversary, ACME achieves most of its security goals. Notably, it prevents many more attacks than traditional CAs through stronger cryptographic notions of user identity, achieved through automated client signatures and strong binding between the clients identity and the validated domain. In comparison, we show that ACME’s design allows it to resist a substantially stronger threat model than the traditional ad-hoc protocols of other CAs that rely on bearer tokens (passwords, cookies, authorization strings) for authentication and domain validation.

Nevertheless, we still discover issues and weaknesses in ACME’s domain validation and account recovery features, potentially amounting to user account compromise. We attempt to address in this paper what seem to be open questions regarding ACME: how does ACME compare to the existing security model of the actual top real-world certificate authorities? How can we most fruitfully illustrate and formally verify its security properties, and what can we prove about them?

Contributions We present an outline of our contributions in this paper:

- **A Survey of Existing Domain Validation Practices on Major CAs** In §2, we survey some of the most-used certificate authorities both in terms of protocols and infrastructure. We argue that all of the top 10 traditional CAs operate under an unrealistic threat model.
- **A Threat Model for Domain Validation on the Internet** In §3, we specify a high-level threat model for traditional CAs as well as ACME. We link this threat model to our network topology analysis in §2. In §4, We demonstrate that ACME resists a stronger threat model than ad-hoc protocols.
- **Formally Specifying and Verifying ACME** In §4, we formally specify the ACME protocol and verify it using ProVerif. Although ACME is shown to be more resistant to attacks than ad-hoc CAs, we also discover weaknesses in ACME’s domain validation and account recovery and suggest countermeasures.

2 Current State of Domain Validation

A goal of this paper is to establish a relationship between current domain validation practices in the real world and a more formal threat model on which we base our security results. We begin by taking a closer look into the network infrastructure, user authentication and domain validation protocols currently in use by ad-hoc CAs.

Basing ourselves on earlier studies of Web PKIs [8,9], we chose to investigate the domain validation mechanisms of top traditional certificate authorities due

to their high usage and relevance for web certificate issuance. With each CA, we attempted to obtain a regular, one-year, single-domain certificate signature for a domain name that we own.

§3.2.2.4 of the CA/Browser Forum’s Baseline Requirements allow for domain validation to occur in ten different ways, including over postal mail and by validating a random number over a TLS certificate. Of these methods, only three are in popular use: validation via email, the setting of an arbitrary DNS record, or serving some HTTP value on the target domain.

2.1 Domain Validation Mechanisms

With ad-hoc CAs, user C authenticates its identity C_{pk} to CA A as a simple username/password web login, with an option for account recovery via email. C can then request that A validate some domain $C_{wx} \subset C_w$. A ’s flow with the various domain validation channels proceeds thus:

- *HTTP Identifier* A sends to C a nonce A_{URIC} via an HTTPS channel that C must then advertise at some agreed-upon location under C_{wx} . A then accesses C_{wx} using an unauthenticated, unencrypted HTTP connection to ensure that it can retrieve A_{URIC} . This identifier depends on honest DNS resolution query responses and an HTTP connection that is resistant to tampering.
- *DNS Identifier* A sends to C a nonce A_{DNSC} via an HTTPS channel that C must then advertise at some agreed-upon TXT record under the DNS records of C_{wx} . A then queries C_{wx} ’s name servers using to ensure that it can retrieve A_{DNSC} . This identifier is dependent on honest DNS resolution query responses.
- *Email Identifier* A sends to C a URI A_{URIC} via email that C must then visit. Visiting this URI satisfies A , which then issues the certificate for C_{wx} . This identifier is dependent on a private email channel and honest DNS resolution query responses.

Once one of the above identifiers succeeds in validating C ’s ownership of C_{wx} to A , A issues the certificate and the protocol ends.

2.2 User Authentication and Domain Validation

While CAs are required to document their certificate issuance policies in Certificate Practice Statements [10–17], a real-world survey found that these statements are not always accurate. Most ad-hoc CAs in our study restricted their validation capabilities to that of Email Identifiers. Unlike HTTP and DNS Identifiers, Email Identifiers effectively offer C a *read*-based challenge instead proof of some write access. In §3, we discuss how Email Identifiers are the weakest available form of identification given our threat model. In §4, we elaborate on a weakness in ACME affecting both account recovery and domain validation. While this weakness is also generalizable to traditional certificate authorities, ACME offers an opportunity for a stronger fix.

CA	Identifiers	Email Recovery	Public Key Auth.	Per-CSR Check
AlphaSSL	Email	✓	✗	N/A
Comodo PositiveSSL	Email	✓	✗	✓
DigiCert	Email	✓	✗	✗
GeoTrust QuickSSL	Email	✓	✗	✗
GlobalSign	HTTP, DNS, Email	✓	✗	✗
GoDaddy SSL	HTTP, DNS	✓	✗	✗
Let's Encrypt (ACME draft-1)	HTTP	✓	✓	✗
Network Solutions	Email	✓	✗	✗
RapidSSL	Email	✓	✗	✓
SSL.com BasicSSL	HTTP, DNS, Email	✓	✗	N/A
StartCom StartSSL	Email	✓	✓	✗

Fig. 1: Popular CAs, their validation methods, whether they permit account recovery via email, whether they allow login via a public-key based approach (such as client certificates) and whether domain validation is carried out once for every certificate request, even for already-validated domain names.

No CA we surveyed offered a login mechanism that was completely independent of email. An exception almost occurs with StartSSL, which uses browser-generated client certificates for web login, but this exception is negated by StartSSL still allowing email-based account recovery in case of a lost certificate private key. Reliance on the security of the email channel can in many cases be even more serious: in many surveyed CAs, simply being able to complete a web login will allow a user to re-issue certificates for domains they had already validated before, without any further validation.

A scan of the DNS MX and NS records of the web's top 10,000 websites (according to Alexa.com) [18] showed that roughly 45% of surveyed domain names used only six DNS providers, of which CloudFlare alone had a 18% share. Meanwhile, Let's Encrypt's infrastructure is hosted almost entirely on Akamai, rendering it a centralized point of failure for ACME and ad-hoc CAs alike. While ACME is a centralization-agnostic protocol, Let's Encrypt operates with a fully centralized infrastructure. A similar centralization of authority exists with email, where the top six providers serve more than 55% of domain names surveyed, with Google alone holding roughly 27% market share (Figure 2.)

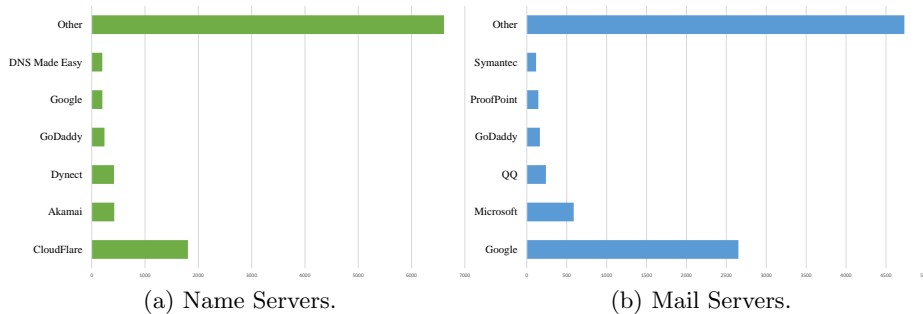


Fig. 2: Provider repartition among the Alexa Top 10,000 global sites, as of October 2016. Notably, CloudFlare and Akamai also provide CDN services to domains under their name servers, allowing them stronger control over HTTP traffic.

These results suggest that the number of actors of which the compromise could affect traditional domain validation is significantly small. This is relevant given how top CAs allow for account recovery, certificate re-issuance and more with simple email-based validation.

3 A Security Model for Domain Validation

The protocols considered in this paper operate between a party C claiming to serve and represent one or more domain names C_w (for which it wants certificates), and it is incumbent upon a certificate issuer A to verify that all domains in C_w are indeed controlled and managed by C . User C authenticates itself to CA A using a public key C_{pk} of a private identity C_k . C can then link identifiers under C_k that prove that it manages and controls domains in C_w .

This and following sections are largely based on our full symbolic model³ of ACME and ad-hoc CA protocol and network flow, which is written in the applied pi calculus and verified using ProVerif. Excerpts of this model are inlined throughout.

3.1 Security Goals and Threat Model

Our security goals are straightforward: for any domain in C_w , A must not sign a certificate for that domain unless C_{pk} is linked with a valid identifier that binds C as the identity that owns and manages this domain. A domain name under C_w is considered to be *validated* under C_w if an identifier for that domain can be linked to C_{pk} .

The network topology, channels and actors are essentially the same for both ACME and ad-hoc CAs. However, the manner in which these actors communi-

³ Full models available at <https://github.com/Inria-Prosecco/acme-model>

cate over the channels is different, and leads to different attempts to establish the same security guarantees.

Channels Intuitively, the channels we want encapsulate the following properties:

- **HTTPS Channel** Intuitively a regular web channel, we treat it as a A -authenticated duplex channel whereupon anyone can send a request to A , only A can read this request and respond, and only the sender can read this response.
- **Strong Identifier Channels** These channels must be assumed to be writable only by C . They are therefore relevant for HTTP and DNS Identifiers.
- **Weak Identifier Channel** Anyone can write to this channel, but only C can read from it. This makes it relevant for domain validation via Email Identifiers.

A shared consideration between ACME and ad-hoc CAs involves the critical importance of DNS resolution: if the attacker can simply produce false DNS responses for A resolving a domain request for any domain in C_w , it becomes impossible to safely carry out domain validation under any circumstances. As a sidenote, this allows us to argue that since the DNS channel must be trusted, it could also be considered as the safest channel on which to carry out domain validation using DNS Identifiers since that would allow C to avoid needlessly involving other channels.




Fig. 3: Channels overview.

In formally describing our network model in ProVerif, we simulate simultaneous requests from Alice, Bob and Mallory as independent clients C . We also simulate two independent ACME CAs, which interchangeably assume the role of A . For each C , we specify a triple of distinct channels:

$$(C_{HTTP}, C_{EMAIL}, C_{DNSTXT})$$

Each channel represents access to a different domain validation mechanism. While C is given complete access over these channels, a write transformation $w(channel) \rightarrow channel$ is applied to C_{EMAIL} before it's handed to A . Similarly, a read transformation $r(channel) \rightarrow channel$ is applied to C_{HTTP} and C_{DNSTXT} .

A routing proxy is then specified in order to model the transportation across these channels by executing the following unbounded processes in parallel⁴:

⁴ We also specify a fully public channel named `pub`.

$$\begin{aligned}
& in(w(C_{EMAIL}), x); out(r(C_{EMAIL}), x) \\
& in(pub, x); out(r(C_{EMAIL}), x) \\
& in(w(C_{HTTP}), x); out(pub, x); out(r(C_{HTTP}), x) \\
& in(w(C_{DNSTXT}), x); out(pub, x); out(r(C_{DNSTXT}), x)
\end{aligned}$$

Threat Model We assume that the adversary controls parts of the network and so can intercept, tamper with and inject network messages. As such, an attacker could make requests for domains they do not own, intercept and delay legitimate certificate requests, and so on. Our adversary has full access to pub , $w(C_{EMAIL})$, $r(C_{HTTP})$ and $r(C_{HTTP})$. We also publish Mallory’s channels and C_k over pub . As such, the attacker controls a set of valid participants (e.g. M) with their own valid identities (e.g. M_k , M_{pk}). The attacker may advertise any identity for its controlled principals, including false identities, and may attempt to obtain a certificate for domains not legitimately under M_w .

The adversary also has at his disposal certain special functions:

- *poisonDnsARecord*, which takes in a domain C_{wx} and allows the attacker to poison its DNS records to redirect to a server owned by M . Using this function triggers the *ActiveDnsAttack*(C_{wx}) event.
- *manInTheMiddleHttp*, which allows the attacker to write arbitrary HTTP requests as if they were emitting from C_{HTTP} by disclosing C_{HTTP} to the attacker. Using this function triggers the *ActiveHttpAttack*(C_{wx}) event.

3.2 Events and Queries

Queries under our model are constructed from sequences of the following events, each callable by a particular type of actor:

- **Client** The client is allowed to assert that they own some domain by triggering the event *Owner*(M, C_{wx}). Once C receives a certificate $C_{wx_{cert}}$ for C_{wx} from A , they also trigger *CertReceived*($C_{wx}, C_{wx_{cert}}, C_{pk}, A_{pk}$)
- **Server** The server (ACME instance or CA) triggers the event *HttpAuth*(C_{pk}, C_{wx}), *DnsAuth*(C_{pk}, C_{wx}) and *EmailAuth*(C_{pk}, C_{wx}) depending on the type of domain validation used. Once A issues a certificate $C_{wx_{cert}}$ for C_{wx} to C , they also trigger *CertIssued*($C_{wx}, C_{wx_{cert}}, C_{pk}, A_{pk}$)
- **Adversary** As noted above, the adversary may trigger the events *ActiveDnsAttack*(C_{wx}) and *ActiveHttpAttack*(C_{wx}). In addition, the adversary is allowed to masquerade as M in order to assert that they own some domain by triggering the event *Owner*(M_{pk}, C_{wx}).

Queries We evaluate our model against the following queries:

Validation with DNS Identifiers We assert that if DNS validation succeeded, then A must have been able to successfully carry out DNS validation according to spec, *or* an adversary was able to instantiate an active DNS poisoning attack (with no third possible scenario):

$$DnsAuth(C_{pk}, C_{wx}) \implies (Owner(C_k, C_{wx}) \vee DnsAttack(C_{wx}))$$

Validation with HTTP Identifiers We explicitly show that HTTP authentication is weaker than DNS authentication, since it is possible under both cases of DNS poisoning *and* an HTTP man-in-the-middle attack:

$$HttpAuth(C_{pk}, C_{wx}) \implies Owner(C_k, C_{wx}) \vee (HttpAttack(C_{wx}) \vee DnsAttack(C_{wx}))$$

Predictable Certificate Issuance We attempt to verify that all received certificates were issued by the expected CA. This query fails to verify, and leads us to the attack we discuss in §5.2:

$$CertReceived(C_{wx}, C_{wx_{cert}}, C_{pk}, A_{pk}) \implies CertIssued(C_{wx}, C_{wx_{cert}}, C_{pk}, A_{pk})$$

4 Specifying and Formally Verifying ACME

In this section we provide a formal description of the ACME protocol functionality and identify three issues that affect ACME’s security. We also discuss details of how we describe the ACME protocol flow in the applied pi calculus, so that we can verify for certain queries using ProVerif.

4.1 ACME Network Flow

Unlike ad-hoc CAs which are limited to a web login, ACME’s authentication depends on C generating a private value C_k and a public signing key C_{pk} , which are used to generate automated client signatures throughout the protocol.

HTTP Identifier A sends to C a nonce A_{URIC} via the HTTPS channel. C must then advertise, at an agreed-upon location under C_{wx} , the value (C_{pk}, A_{URIC}) . A then accesses C_{wx} using an unauthenticated, unencrypted HTTP connection to ensure that it can retrieve the intended value.

DNS Identifier Since ACME is designed to take advantage of domain validation methods that can be automated and since DNS record management depends on a series of ad-hoc protocols of its own between C and DNS service providers, it is not used by ACME.

Email Identifier This identifier is only used for account recovery, and a resulting weakness is discussed in §5.2.

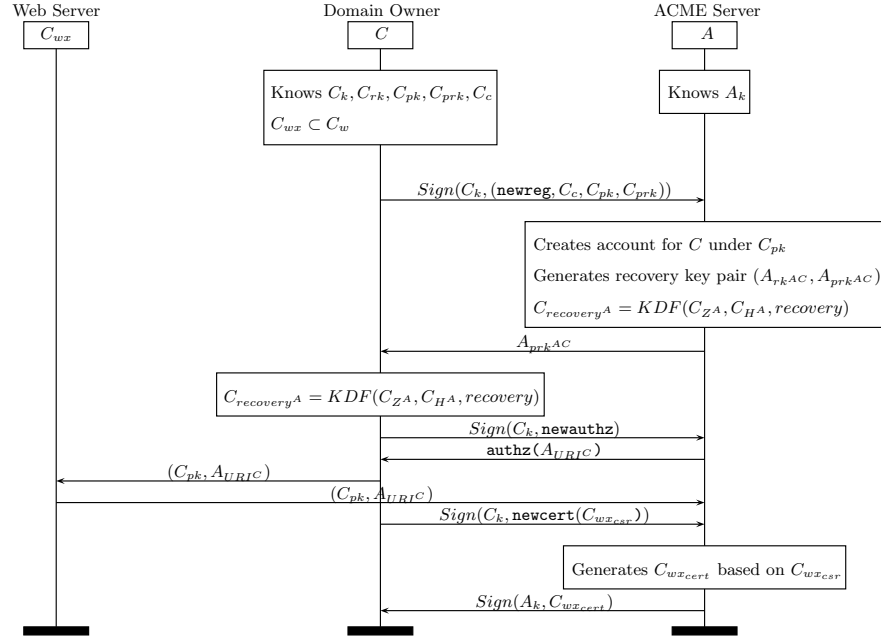


Fig. 4: ACME draft-1 protocol functionality for C account registration, recovery key generation, and validation with certificate issuance for C_{wx} . This chart demonstrates validation via an HTTP identifier. In draft-3 and above, the HTTP challenge (C_{pk}, A_{URIC}) is replaced with $Sign(C_k, (C_{pk}, A_{URIC}))$.

4.2 ACME Protocol Functionality

In this paper we focus on draft-1 of the IETF specification for the ACME protocol. As of October 2016, the draft specification deployed in official Let's Encrypt client and server implementations is somewhere between draft-2 and draft-3⁵. However, draft-1 was adopted after Let's Encrypt's launch and for a majority of 2016. In part due to the issues we discuss in the paper and have communicated with the ACME team, draft-3 (and subsequently draft-4) does away with some features, most notably Account Recovery, and generally is resistant to the issues discussed here.

Preliminaries In some parts of ACME's protocol flow, C and A will need to establish a number of shared secrets, each bound to a strict protocol context, over their public keys. In ACME, this is accomplished using ANSI-X9.63-KDF:

⁵ <https://github.com/letsencrypt/boulder/blob/master/docs/acme-divergences.md>

1. C and A agree on a ECDH shared secret C_{ZA} using their respective key pairs (C_k, C_{pk}) and (A_k, A_{pk}) .
2. A hashing function C_{HA} is chosen according to the elliptic curve used to calculate C_{ZA} : $SHA256$ for $P256$, $SHA384$ for $P384$ and $SHA512$ for $P521$.
3. $C_{labelA} = KDF(C_{ZA}, C_{HA}, label)$, with $label$ indicating the chosen context for this particular key's usage.

As a protocol, ACME provides the following seven certificate management functionalities (illustrated in Figure 4) between web server C and certificate management authority A :

- *Account Key Registration* In this step, C specifies her contact information (email address, phone number, etc.) as C_c and generates a random private signing key C_k with (over a safe elliptic curve) a public key C_{pk} . A `POST` request is sent to A containing $Sign(C_k, (newreg, C_c, C_{pk}))$. The `newreg` header indicates to A that this is an account registration request. If A has no prior record of C_{pk} being used for an account, and if the message's signature is valid under C_{pk} , A creates a new account for C using C_{pk} as the identifier and responds with a success message.
- *MAC-Based Account Recovery* C may choose to identify an account recovery secret with A . In order to do this, C generates an account recovery key pair (C_{rk}, C_{prk}) and simply includes C_{prk} in an optional `recoverykey` field in its initial `newreg` message to A . A generates the complementary $(A_{rk^{AC}}, A_{prk^{AC}})$ and both parties calculate $C_{recovery^A}$ using their recovery key pairs. A communicates $A_{prk^{AC}}$ in its response to C . Later, if C loses C_k , she can ask A to re-assign her account to a new identity $(C_{k'}, C_{pk'})$ by using $C_{recovery^A}$ as a key to generate a MAC of some value chosen by A .
- *Contact-Based Account Recovery* C can request that A send a verification token to one of the contact methods previously specified in C_c . For example, this could be a URI sent to an email in C_c . If C successfully opens this URI, she becomes free to replace C_{pk} with a $C_{pk'}$ for some arbitrary $C_{k'}$ at A .
- *Identifier Authorization* C can validate its ownership of a domain C_{wx} in C_w by providing one of the identifiers discussed in §3 to A . C must first request authorization for C_{wx} by sending a `newauthz` message. A then responds with the types of identifiers it is willing to accept in a `authz` message. C is then free to use any one of the permitted identifiers to validate its ownership of C_{wx} and allow A to sign certificates for it issued to C and tied to the identity C_{pk} .
- *Certificate Issuance and Renewal* After C ties an identifier to C_{wx} under C_{pk} , it may request that a certificate be issued for C_{wx} simply by requesting one from A . Generally, A will send the signed certificate with no further steps required. The renewal procedure is similarly straightforward.
- *Certificate Revocation* C may ask A to revoke the certificate for C_{wx} by sending a `POST` message containing the certificate in question, signed under either C_{pk} or the key pair for the certificate itself. C may choose which key to use for this signature. A verifies that the public key of the key pair

signing the request matches the public key in the certificate, and that the key pair signing the request is an account key, and the corresponding account is authorized to act for all of the identifier(s) in the certificate.

Given this description of the ACME protocol and the threat model defined in §3, we modeled ACME using the automated verification tool ProVerif [19]. In our model, we involve three different candidates for C : Alice, Bob and Mallory, and two CA candidates as A .

As a result of our automated verification process which an active attacker over the three channels specified in §3, we were able to find the issues discussed in §5.2. The first two are relatively minor; however, the third could lead to account compromise in the case of contact-based account recovery, and potentially to the issuance of false certificate signatures if email-based domain validation were to be implemented in ACME. Furthermore, this third issue is also generalizable to affect traditional certificate authorities, as described in §2.

4.3 Model Processes

Using the modeling conventions we established in §3 which include channels, adversaries, actors and events, we instantiate in our ProVerif model of ACME a top-level process that executes the following processes in parallel:

- *clientAuth* Run simultaneously by Alice, Bob and Mallory assuming the role of C (with a compromised Mallory), this process registers a new account with A and sends the queries illustrated in Figure 4. The events *Owner* and *CertReceived* are triggered as part of this process.
- *serverAuth* Run simultaneously by two independent CAs assuming the role of A , this process accepts registrations from C and follows the protocol illustrated in Figure 4. The events *HttpAuthenticated* and *CertIssued* are triggered as part of this process.

The processes *routingProxy*, *poisonDnsARecord* and *manInTheMiddleHttp*, all described in §3, are also run in parallel with the above.

5 Analysis Results

5.1 Weaknesses in Traditional CAs

Traditional CA dependence on weak channels gives us a threat model where real-world attacks can have a small cost and come with severe consequences.

Email Validation In ad-hoc CAs, C is generally simply sent an email containing a URI to their email inbox, which they're supposed to click in order to validate for their chosen domain. Figure 5 shows an attack rendered possible by this mechanism. A could instead, upon a validation request, redirect C 's browser to

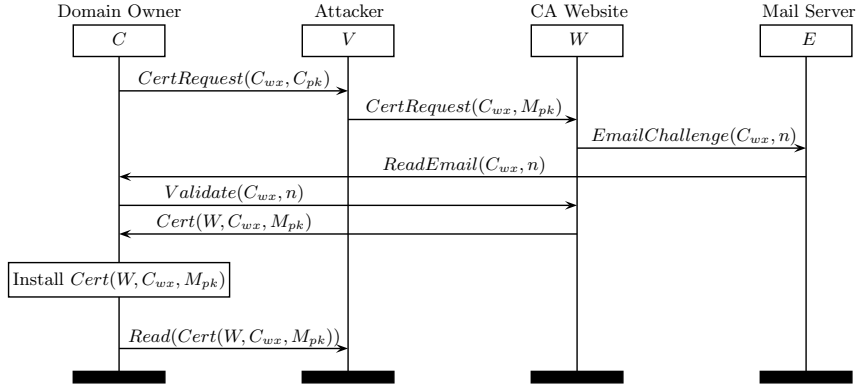


Fig. 5: Attack on Email Validation: Concurrent Request by Active Adversary.

a secret, nonce-based URI A_{URIC} served to C over the HTTPS channel, and independently mail C the value $HMAC(A_{hk}, A_{URIC})$ for some secret A_{hk} held by A . C would need to retrieve this second value and enter it inside the page at A_{URIC} . This approach would largely negate the weakness discussed in §5.2, since an attacker-induced validation email would result in an email that does not include a value matching the URI given by A to C at the beginning of the validation process.




Fig. 6: Active attack on DNS/HTTP/Email Validation when using just nonces.

Usage of Nonces Traditional CAs use random nonces with no special cryptographic properties as the values that they then verify over HTTP, email or DNS.

In addition to this helping caused the attack described above, another more general attack on nonces is shown in Figure 6 in the case of an active attacker. For example, this attack can be used by a compromised CA website to get certificates issued for domain C_{wx} by another (more reputable) CA, hence amplifying the compromise across CAs. None of these attacks would be effective if nonces were tied to some cryptographic properties, such as MACs or even just by deriving them from a hash of the certificate request's public key.

In order to avoid a similar attack, ACME draft-3 and draft-4 require that HTTP identifiers be validated by broadcasting $Sign(C_k, A_{URIC})$ via the web server instead of ACME draft-1's (C_{pk}, A_{URIC}) .

5.2 Weaknesses in ACME

Cross-CA Attacks on Certificate Issuance Suppose an ACME client C requests a certificate from A , and suppose that A 's HTTPS private key is compromised (or that A is malicious, or a man-in-the-middle has compromised the ACME channel). Now, the attacker can intercept authorization and certificate requests from C to A , and instead forward them to another ACME server A' . If A' requests domain validation with a token T , the attacker forwards the token to the client, who will dutifully place its account key K and token T on its validation channel. A' will check this token and accept the authorization and issue a certificate that the attacker can forward to C .

This means that C asked for a certificate from A , but instead received a certificate from A' . Moreover, it may have paid C for the service, but A' might have done it for free. This issue, while not critical, can be prevented if C checks the certificate it gets to make sure it was issued by the expected CA. An alternative, and possibly stronger, mitigation would be for ACME to extend the Key Authorization string to include the CAs identifier.

More generally, this issue reveals that ACME does not provide channel binding, and this appears as soon as we model the ACME HTTPS Channel. We would have expected to model this as a mutually-authenticated channel since the client always signs its messages with the account key. However, although the client's signature is tunnelled inside HTTPS, the signature itself is not bound to the HTTPS channel. This means that a message from an ACME client C to A can be forwarded by A to a different A' (as long as C supports both A and A'). This kind of credential forwarding attack can be easily mitigated by channel binding. For example, ACME could rely on the Token Binding specifications to securely bind the client signature to the underlying channel. Alternatively, ACME could extend the signed request format to always include the server's name or certificate-hash, to ensure that the message cannot be forwarded to other servers.

Contact-Based Recovery Hijacking While the use of sender-authenticated channels in ACME seems to be relatively secure, more attention needs to be paid to the receiver-authenticated channels. For example, if the ACME server uses

the website administrator’s email address to send the domain validation token, a naïve implementation of this kind of challenge would be vulnerable to attack.

In the current specification, the contact channel (typically email) is used for account recovery when the ACME client has forgotten its account key. We show how the careless use of this channel can be vulnerable to attack, and propose a countermeasure. Suppose an ACME client C issues an account recovery request for an account under C_{pk} with a new key $C_{k'}$ to the ACME server A . A network attacker M blocks this request and instead sends his own account recovery request for the account under C_{pk} (pretending to be C) with his own key $M_{k'}$. A will then send C an email asking to click on a link. C will think this is a request in response to its own account recovery request and will click on it. Similarly to the (slightly different) flow described in Figure 5, A will think that C has confirmed account recovery and will transfer the account under C_{pk} to the attacker’s key $M_{k'}$. In the above attack, the attacker did not need to compromise the contact channel (or for that matter, the ACME channel).

The key observation here is that on receiver-authenticated channels (e.g. email) the receiver does not get to bind the token provided by A with its own account key. Consequently, we need to add a further check. The email sent from A to C should contain a fresh token in addition to C ’s new account key. Instead of clicking on the link (out-of-band), C should cut and paste the token into the ACME client which can first check that the account key provided by A matches the one in the ACME client and only then does it send the token back to A , or alternatively that the email recipient at C visually confirms that the account key (thumbprint) provided by A matches the one displayed in the ACME client.

The attack described here is on account recovery, but a similar attack appears if we allow email-based domain validation. A malicious ACME server or man-in-the-middle can then get certificate issued for C ’s domains with its own public key, without compromising the contact/validation channel. The mitigation for that attack would be very similar to the one proposed above.

6 Conclusion

In this paper, we have provided the results of an empirical case study that allowed us to describe a real-world threat model governing both traditional certificate authorities and ACME in terms of user authentication and domain validation. We formally modeled these protocols and provided the results of security queries under our threat model, using automated verification. As a result of our disclosures to the ACME team, the latest ACME protocol version (draft-4) has been designed to avoid the pitfalls that make these attacks possible.

Given the weak threat model that ad-hoc CAs are expecting to survive under and the subsequent weaknesses that we describe, we believe that there is a strong need for either larger ACME adoption, or the improvement of the current state of the art for ad-hoc CA domain validation mechanisms. We hope that this work will help act as a bedrock for future formal description of domain validation systems. We also hope to see a wider deployment of related technologies, such

as DNSSEC [20], DANE [21] and SMTPS, that help strengthen the channels involved in domain validation protocols.

References

1. S. Chokhani, W. Ford, R. Sabett, C. Merrill, and S. Wu. Internet X.509 Public Key Infrastructure: Certificate Policy and Certification Practices Framework. RFC 3647, Internet Engineering Task Force, November 2003.
2. CA/Browser Forum. Baseline requirements for the issuance and management of policy-trusted certificates, v.1.1.5, May 2013.
3. Gervase Markham, Ryan Sleevi, Richard Barnes, and Kathleen Wilson. Wosign and startcom.
4. Internet Security Research Group. Let's encrypt overview, 2016.
5. Internet Security Research Group. Let's encrypt statistics, 2016.
6. Richard Barnes, Jacob Hoffman-Andrews, and James Kasten. Automatic certificate management environment (acme), Jul 2016.
7. Bruno Blanchet, Ben Smyth, and Vincent Cheval. Proverif 1.90: Automatic cryptographic protocol verifier, user manual and tutorial, 2014.
8. Antoine Delignat-Lavaud, Martín Abadi, Andrew Birrell, Ilya Mironov, Ted Wobber, Yinglian Xie, and Microsoft Research. Web pki: Closing the gap between guidelines and practices. In *NDSS*, 2014.
9. Olivier Levillain, Arnaud Ébalard, Benjamin Morin, and Hervé Debar. One year of SSL Internet measurement. In *Proceedings of the 28th Annual Computer Security Applications Conference, ACSAC '12*, pages 11–20, New York, NY, USA, 2012. ACM.
10. Comodo CA Ltd. Comodo Certification Practice Statement. Technical report, Comodo CA Ltd., August 2015.
11. DigiCert. DigiCert Certification Practices Statement. Technical report, DigiCert, September 2016.
12. GeoTrust Inc. GeoTrust Certification Practice Statement. Technical report, GeoTrust, September 2016.
13. GlobalSign CA. GlobalSign CA Certification Practice Statement. Technical report, GlobalSign CA, August 2016.
14. Internet Security Research Group. Certification Practice Statement. Technical report, Internet Security Research Group, October 2016.
15. Symantec Corporation. Symantec Trust Network (STN) Certification Practice Statement. Technical report, Symantec Corporation, September 2016.
16. StartCom CA Ltd. StartCom Certificate Policy and Practice Statements. Technical report, StartCom CA Ltd., September 2016.
17. LLC Network Solutions. Network Solutions Certification Practice Statement. Technical report, Network Solutions, LLC, September 2016.
18. Alexa Internet Inc. Top 1,000,000 sites (updated daily), 2013.
19. R. Kusters and T. Truderung. Using ProVerif to analyze protocols with Diffie-Hellman exponentiation. In *IEEE Computer Security Foundations Symposium (CSF)*, pages 157–171, 2009.
20. Giuseppe Ateniese and Stefan Mangard. A new approach to dns security (dnssec). In *Proceedings of the 8th ACM conference on Computer and Communications Security*, pages 86–95. ACM, 2001.
21. Paul Hoffman and Jakob Schlyter. The dns-based authentication of named entities (dane) transport layer security (tls) protocol: Tlsa. Technical report, 2012.