Identity and Trust beyond EV Certificates

Henrik Biering, Peercraft ApS
v0.7, 2018-10-21

The conflated dual purpose of web site certificates

The main purpose of website certificates is to ensure confidential communication between a website and its users. In their basic form certificates are bound to one or more specific domains over which the provider has documented control. So the server is effectively authenticated to users in terms of a domain name. Recently the provision of such basic domain validation certificates has become fully automated and is offered for free by several issuers.

A problem that has not been solved, however, is the mostly weak identity proofing of domain registrants and operators. To enable businesses to authenticate themselves to their users in terms of their legal identity, some certificate authorities started to offer identity verification as an additional service to their customers.

This verification of the relationship between a domain and a legal entity is principally unrelated to the purpose of domain validation certificates. Nevertheless, the result of this identity verification effort was baked into the certificate offering, referring to simple “organization validation” (OV) and stricter “enhanced validation” (EV) certificates. Unfortunately this conflation with the basic certificate purpose effectively created a mostly useless dead end for identity verification, rather than opening a path for using the identity verification as a basis for further discovery of trust elements and services offered.

This article will summarize the criticism and challenges that EV certificates currently face – and outline a promising, flexible and even backwards compatible alternative.

Criticism and challenges for EV Certificates

The debate about the value and verification quality of website certificates is certainly not new. But it was dramatically intensified when a large number of companies and institutions backed the Let’s Encrypt initiative to offer free and auto-renewable domain validation certificates. This was seen as an attack on the business model of existing Certificate Authorities that would now rely solely on convincing potential customers of the added value of their organization (OV) and enhanced (EV) certificate offerings [Ref. 1].

Most companies regards themselves to be at the center of the internet and well known by their domain name to all potential users. Also companies are primarily focusing on how they can protect themselves against fraudulent users, but not on how they might protect their users from phishing and related threats. So selling the added value of EV certificates to companies is hard, in particular considering that most of the worlds largest “reference” websites have chosen not to use EV certificates.

Simultaneously several security researchers discovered problems with the issuance of EV Certificates. This related partly to the initial verification process and partly to how the limited information in the certificate could be misinterpreted by end users. It was demonstrated that the challenge is not only whether a certificate points correctly to a real legal entity, but also that legal entities with matching certificates may be created with the sole purpose of having the end user confuse them with a known company name or even a comforting phrase [Ref. 2, 3]

Moreover, Google discovered breaches of trust in the certificate issuing process of several Certificate Authorities, most notably Symantec that reportedly was responsible for at least 30.000 misissuances over a period of a few years [Ref. 4]. Ultimately this led Symantec to sell its CA business to its previous competitor Digicert.

Unfortunately the response from most major certificate authorities to all of this criticism and issues has been defensive and based on dubious statistics rather than proactively seeking a role in solving the real world problems [Ref. 5, 6, 7].

This defensive posture voiced on behalf of the CA's common organization "CA Security Council" has, however, provoked the member Digicert to withdraw from the organization [Ref. 8, 7]. In opposition to the remaining members of the CA Security Council, Digicert claims that they "would prefer, that if CAs are going to engage in website monitoring and information sharing, that it would address the full spectrum of fraud and abuse that exists".

Facing an important crossroad for internet trust

Thanks to the efforts of initiatives like Let's Encrypt and established browser vendors the situation with respect to encrypted communication has changed from being limited to be the new normal. A turning point is reached where warning or blacklisting services with unencrypted communication is currently replacing recommending or whitelisting encrypted communication.

Unfortunately this development does not apply to the intended identity and trust assurance that OV and DV certificates were supposed to offer. Compared to the encryption option which is now typically an integral part of the service offered by hosting providers, the majority of smaller companies will find the acquisition of an EV Certificate both expensive, cumbersome, and without commercial incentives. Also browser vendors appear ready to declare EV certificates dead [Ref. 9].

A joint research paper by Google, University of California, Berkeley, and International Computer Science Institute [Ref. 10] has revealed that phishing is more common and advanced than commonly considered. The report recognizes the need to educate users about password managers and unphishable two-factor authentication as a potential solution.

However, this only applies to phishing in its most narrow sense where users have already created an account with the proper site. Password managers and two-factor authentication does not solve the wider problem of users entering credentials and other information or remitting money to fraudulent parties, whether these are offering seemingly legitimate services or may have hacked a legitimate website. Also the adoption rate of these methods has been very limited so far.

Thus a desirable property of an optimal solution is also to minimize the probability of users getting into initial contact with fraudulent services rather than with proper well reputed services. If verified legal identities could be easily linked to additional identity and reputation attributes, search and comparison services would have better options to guide their users to websites that are both competitive and safe in various respects.

A major problem with EV Certicates is that trust is not necessarily mutual or associative. Currently users need to trust their browser vendor, that again needs to trust various CA's that are in turn offered financial compensation for a positive validation of a company. The user has no way to check this validation, that is more often than not performed by a CA residing in a completely different part of the world than where the user and the company in question is located. Just as registration of real

estate, cars, and drivers licences is typically done by entities having awareness of local challenges, there are reasons to apply the principle of “Think Global, Act Local” and replace unnecessary remote entities with native or local trust anchors.

This problem of indirection becomes even more substantiated when EV certificates are required to contain multiple claims for which different parties are authoritative. An example of this problem is mentioned in [Ref. 11] where Certificate Authorities are requested to add a specific financial authority claim to the basic identity claim of certain service providers operating under the new EU Payment Services Directive. The roots of trust for these two claims can be two different national governments within the EU, so it would seem advantageous to separate the handling of the claims.

So the question is how we can most effectively deliver identity validation:

a) as a free service that may be implemented automatically by hosting providers

b) as a basis for association with additional trust elements addressing various types of fraud

c) aggregated and provided by an entity that the user trusts

d) in a way that simultaneously provides commercial incentives for legitimate companies

e) and functions whether the company trades from its own domain or via platforms

f) is adapted towards and favourizing nations that are frontrunners with respect to digitization, to incentivize nations and states lagging behind, while providing suitable workarounds

g) will function reliably when end users are replaced by initially less clever AI services acting on the users behalf

Companies are now getting natively digital

In many countries national or local authorities have been registering legal entities for more than 100 years. While the internet was developed and grew these registries were still fully based on paper filed in physical lockers. Therefore, trustworthy intermediaries were strictly needed to transform these paper-based identities and make them suitable for online purposes. That was the task fulfilled by EV certificates.

By now, most jurisdictions have replaced the paper-based archives with digital systems and many local registries across all continents have followed the requests from governments to provide open access to their data. Also many registries has increased their security measures, e.g. to require national eID’s supporting two-factor login for owners to update registrations.

Traditionally these registries have been listing the physical address of the companies – and in some cases even more addresses to distinguish e.g. between the production address and a legal contact point. This was to enable authorities as well as potential customers or investors to contact or visit the company for further information about the operations and services of the companies.

Amazingly it has taken more than 50 years of internet and 25 years of WWW before the business registers have noticed that many companies are now completely digital and only have formal physical addresses. Recently 5 countries have taken steps to register a web site address as part of a company’s basic official record. The first implementers were Norway (150.000 companies) and Finland (110.000) followed by Greece (50.000), Denmark (48.000), and New Zealand (36.000). The

number of companies pr. country/state that have registered a website address can generally be discovered via the OpenCorporates service [Ref. 12]. Note that this index currently does not include the 48.700 Danish companies and possibly companies from other countries for which Open-Corporates has not yet updated the fetching algoritms and subsequently refetched all companies.

The 48.000 Danish registrations have been recorded over little more than a year and must be seen in relation to a complete base of approximately 300.000 active companies in Denmark. This is despite a message next to the registry's website entry field saying "Currently this entry is not being used for anything and will not be publicly shown" (NOTE: It is available in datadumps and via API access [Ref. 13]) It demonstrates that nearly all users registering a new company or updating an existing record will complete a website entry even without understanding the potentially derived benefits.

While adding an extra field to its register is a relatively simple issue for a business registry – and completing the entry has been shown to be trivial for users, this is the decisive prerequisite for automating the currently manual process of verifying the mutual relationship between a company and its primary website. Consequently it makes sense to pursue a strategy that incentivizes this natural development in order to eventually reach a point where identity and reputation verification becomes an integral part of search and selection criteria rather than an individual blacklisting exercise.

Proposal for a new website trust model

The aim of this proposal is to establish a framework that fulfils the requirements a) through g) as mentioned above. The description is kept as non-technical as possible while recognizing existing technologies that may be (re)used. It is assuming that all connections are properly encrypted, and for simplicity it disregards a number of complicating issues, e.g. privacy concerns, the indirect company registration process in the USA, and caching models for optimization of speed.

Base level ad hoc verification using OpenDiscovery

The foundation for base level verification is the mutual reference between a company's official legal record and its primary domain or chosen trust service provider. This is referred to as OpenDiscovery, and described further in [Ref. 14], which also provides a simple Proof of Concept.

The OpenDiscovery resolution process resembles the DNS resolution process used successfully since the early days of the internet to crossreference between internet domains and IP-addresses.

The basic scenario for OpenDiscovery is when a user visits a specific website. The resolver initially discovers the claimed official company ID from the visited website using one of several comple-mentary techniques, including .well-known; page header tag; Country ID + Serial Number from EV certificate; a proprietary whitelist; or DNS.

It then completes the 3-step discovery process shown in the Figure below. The resolver (and Root Provider) in the figure must be operated by a party trusted by the end user. Below are a few examples of the flexibility enabled by this OpenDiscovery approach.

Specific methods to process claims are out of scope for this article. It is envisioned, however, that the methodology and syntax currently being developed for discovery of 3. party claims between existing and (potentially) new members of a federation "OpenID Connect Federation 1.0" [Ref. 15] with some amendments could serve as one out of possibly several options.

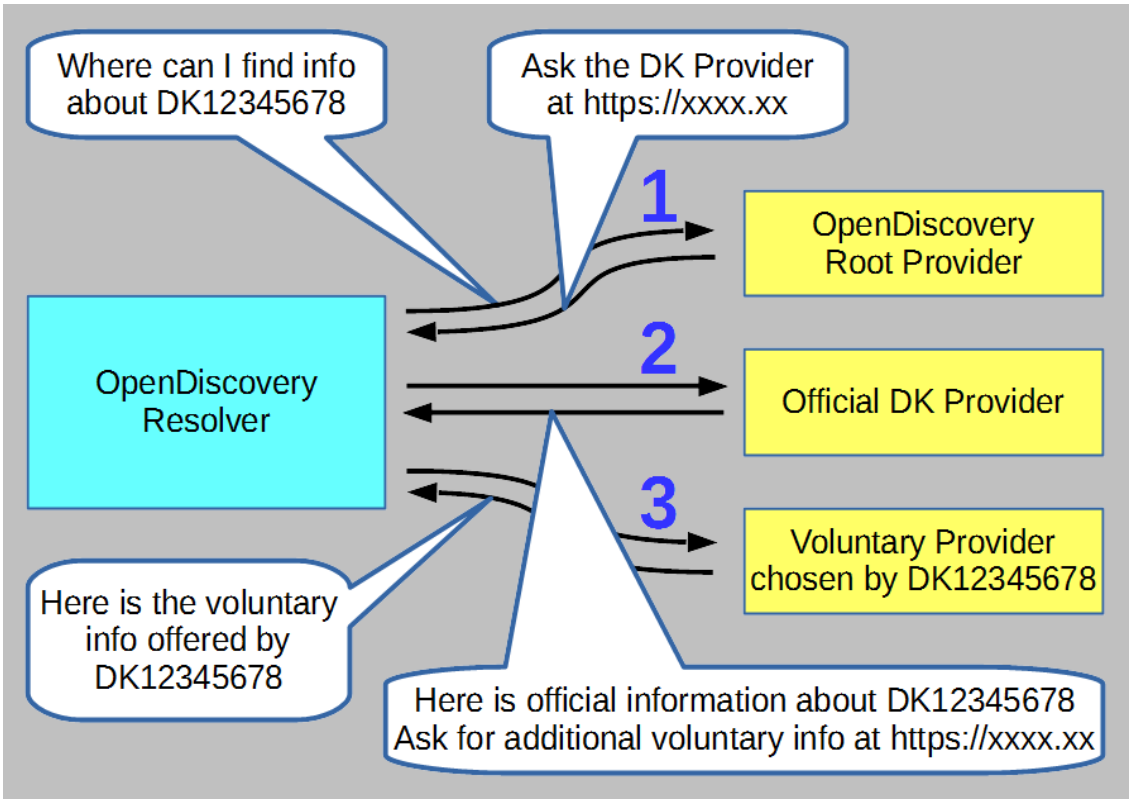

CASE 1: The local business registry has NOT implemented support for website addresses

Basic verification is achieved by using an EV Certificate. Globally the OpenDiscovery process will provide the user with added information about a company by skipping the two first discovery steps to discover any self-asserted claims as well as additional 3. party validated claims issued to the Company ID (or API endpoints to obtain such claims) via .well-known discovery. If the company uses several domains for its services, it will only be required to use an EV Certificate for the chosen primary domain, as this may contain the company's claims regarding the auxiliary domains. Also, already in many countries executing step 2 of the discovery process will provide additional authoritative information about the company, e.g. the date of incorporation, whether it is operational or under bankruptcy proceedings, the number of employees, and even financial information.

CASE 2: The local business registry has implemented support for website addresses

If the domain visited is the primary domain of the company, basic verification is based on the first two steps of the depicted discovery process. If the domain is not the primary domain, the third step must be included to discover the visited domain as an auxiliary domain controlled by the company. The third step must also be completed to discover additional 3. party validated claims issued to the Company ID (or API endpoints to obtain such claims) via .well-known discovery.

CASE 3: The website is owned by a large company that does not worry about user trust

One of the primary issues with EV Certificates has been the missing use of EV certificates by most internet giants. On this basis, many lower ranking companies have decided to similarly avoid the use of EV certificates. This proposal allows providers of the resolver functionality that are trusted by their users to whitelist the most commonly used domains in the effort to facilitate a consistent

user experience and simultaneously strengthen the motivation among other companies to implement OpenDiscovery.

CASE 4: The website (domain or subdomain) represents a multitude of different companies

Increasingly companies are replacing or complementing individual websites with presence on a joint website, malls, market places, and social media. Examples include:

- Large corporations with many individially registered daughter companies
- Ticket resellers that are authorized by the ticket issuers (e.g. airlines)
- Companies having profiles on social media or market place platforms

In this case the overall website owner will use .wellknown for general discovery of information related to his own identity. While using individual header tags to signal the respective ownership by individual guest companies that are in turn authorizing their specific location on the overall website in their own primary discovery location.

CASE 5: Isolation of basic company identity claims from supplementary claims

OpenDiscovery lets any company authoritatively specify the location of certificates used to validate any third party claims they provide to others with reference to their own legal entity ID. This lessens the argument to use third parties to aggregate third party validated claims from different entities in one certificate.

Advanced prefetch application of Opendiscovery

One of the deficiencies of the current EV certificate approach is that the certificates do not by themselves provide any information as to whether a specific website is trustworthy or not. Hence it is not justified to prioritize potential vendors based on their possession of an EV certificate.

OpenDiscovery radically changes this situation by allowing company’s to publish any combination of third party validated claims as part of their identity and reputation. This might be claims from an insurance company that an artisan is properly insured against failed deliveries or damages. A well reputed trust mark offering money back guarantees for a webshop. Or that a website has been rated accessible for certain types of disabled persons.

Consumers are normally not looking for a particular service provider, but directly for purchasing goods or services under safe conditions. Hence, the most effective approach is to avoid fraudulent encounters by applying security and privacy related prioritization criteria as an integral part of users general search criteria. The combined discovery, prefetching and use of ID and reputation related claims along with company’s self-asserted claims regarding their services becomes even more important in relation to user-assistive AI tools for the upcoming data-driven economy.

To efficiently accomplish such prefetching of business data, business registries need to provide open data services allowing third parties to bulk download business registry data wholly or partially in order to discover any relevant company ID’s. Fortunately this functionality is already available in a large number of countries or local jurisdictions globally. For countries and jurisdictions, where this is not the case, there would be a need for a simple Company ID discovery service. Such a basic discovery service, which in itself does not involve any trust issues, could be operated by a single company or via a joint permissioned ledger.

References:

1. **On The (Perceived) Value of EV Certs, Commercial CAs, Phishing and Let's Encrypt**
   Troy Hunt [2017-07-19]
   https://www.troyhunt.com/on-the-perceived-value-ev-certs-cas-phishing-lets-encrypt/
2. **Are EV certificates worth the paper they're written on?**
   Scott Helme [2017-12-04]
   https://scotthelme.co.uk/are-ev-certificates-worth-the-paper-theyre-written-on/
3. **Phishing with EV** (2 parts: First and Final)
   James Burton [multiple edits 2017 - 2018]
   https://www.typewritten.net/writer/
4. **Intent to Deprecate and Remove: Trust in existing Symantec-issued Certificates**
   Ryan Sleevi (Google) [2017-03-23]
   https://groups.google.com/a/chromium.org/forum/#!topic/blink-dev/eUAKwjihhBs%5B1-25%5D
5. **Relative incidence of phishing among DV, OV, and EV encrypted websites**
   Chris Bailey et al. (Entrust Datacard / Comodo) [2 versions: 2017-09-13 / 2018-04-16]
   https://casecurity.org/wp-content/uploads/2017/09/Incidence-of-Phishing-Among-DV-OV-and-EV-Websites-9-13-2017-short-ve....pdf
   https://casecurity.org/wp-content/uploads/2018/06/Summary-Report-Incidence-of-Phishing-04-16-2018.pdf
6. **Twitter Comments to Entrust/Comodo Report**
   Ryan Sleevi [2018-05-15]
   https://twitter.com/sleevi_/status/996581564099244032
7. **New CA Focus on EV Certs Won’t Stop Phishing**
   Fahmida Y. Rashid (Decipher / Duo Security) [2018-07-03]
   https://duo.com/decipher/new-ca-focus-ev-certs-wont-stop-phishing
8. **Notice of Withdrawal from the CA Security Council**
   Jeremy Rowley (Digicert) [2018-06-15]
   https://www.digicert.com/blog/notice-of-withdrawal-from-the-ca-security-council/
9. **Extended Validation Certificates are Dead**
   Troy Hunt [2018-09-18]
   https://www.troyhunt.com/extended-validation-certificates-are-dead/
10. **New Research: Phishing Is Worse Than You Thought**
    Curtis Franklin (Security Now – article based on Google Research Paper) [2017-11-10]
    https://www.securitynow.com/author.asp?section_id=610&doc_id=738125
11. **CABForum Discussion on Subject information in EV certificates**
    Nick Pope (Thales / ETSI TC ESI) [2018-06-07]
    https://cabforum.org/2018/06/06/minutes-for-ca-browser-forum-f2f-meeting-44-london-6-7-june-2018/#Subject-information-in-EV-certificates-specified-in-clause-92-of-the-EV-Guidelines-and-whether-this-allows-for-the-inclusion-of-X520-organizationIdentifier
12. **Number of Legal Entities pr Country with registered Website Addresses**
    OpenCorporates (Query) [2018-10-18]
    https://opencorporates.com/companies?types_of_data_held=Website
13. **Number of legal Enties in Denmark with registered Website Addresses**
    Danish Business Registry (API Query) [2018-10-18]
    https://datacvr.virk.dk/data/cvr-hj%C3%A6lp/indgange-til-cvr/system-til-system-adgang
14. **OpenDiscovery**
    Henrik Biering (Peercraft) [2017 – 2018]
    https://www.opendiscovery.biz/
15. **Federations - trust between entities**
    Andreas Åkre Solberg and Roland Hedberg et al. (UNINETT /) [2018-08-02]
    https://storage.googleapis.com/openid-connect/oidcfed-05.html