

W H I T E P A P E R



Vulnerability of First-Generation Digital Certificates and Potential for Phishing Attacks and Consumer Fraud

Flawed Identity Information in Manually-Vetted Certificates and
the Benefit of Second Generation Automated Vetting

Kirk Hall

GeoTrust, Inc.
April 2005

Contents

1. Introduction.....	1
2. Structure of This Whitepaper.....	1
3. Understanding First Generation Digital Certificates.....	2
- Where did the rules come from? Who made the rules?.....	2
- What data is in a digital certificate?.....	3
- Where do you find this information?.....	4
- Who is in charge of manual vetting? Who makes sure it's done right?.....	4
- How First Generation manual authentication certificates can be faked.....	6
- Examples of spoofed web sites.....	8
- How are misleading certificates obtained?.....	9
- How did this happen?.....	10
- Legal liability issues.....	10
- Bottom line.....	11
4. The Answer: Second Generation Automated Vetting – Cheaper, Faster and the Key to Identity Verification Improvements for the Next Decade.....	12
- This is frustrating! What is the solution?.....	12
- Enhancing trust capabilities.....	13
- Conclusion.....	15

1. Introduction

More and more people are looking for solutions to the massive growth in phishing and other consumer fraud on the Net. Digital certificates, already vital for securing web sites and encrypting communications with consumers so they can't be intercepted and read in transmission, will be a major part of the solution. But how will this happen?


Digital certificates that secure web sites (SSL server certificates) contain certain information submitted by the site that bought the certificate from a public Certification Authority (CA). Some have suggested greater display and reliance on this information to help consumers detect fraud. Can this work, or will it instead give rise to new phishing opportunities and even greater incidents of consumer fraud?

This White Paper takes a look at digital certificates – past, present and future – and their potential for deterring phishing attacks and online fraud. It demonstrates the severe pitfalls from First Generation manual vetting of certificate holders and the inherent unreliability of the identity information they contain (which can easily be faked). These certificates could create serious legal liability for the writers of browser software and for CAs if the identity information they contain is presented to end users as reliable data. Finally, this paper lays out a path for the future, with a description of higher assurance Second Generation automated vetting of Web identities, and discusses ongoing enhancements that will provide a better solution for online identity authentication and reduction of consumer fraud.

GeoTrust has been the leader in developing this Second Generation automated authentication process since 2001, resulting in faster, cheaper, more reliable provisioning of SSL server certificates to businesses in countries around the world. VeriSign and other public Certification Authorities have also implemented automated authentication in recent months, migrating from the First Generation manual authentication process to more reliable automated processes. Prices have fallen as a result, and more consumer sites now are secured by automated technology than ever before. In fact, Second Generation automated authentication is now growing twice as fast as First Generation manual authentication.

2. Structure of This White Paper

This White Paper first looks at the standards and processes for First Generation manual vetting developed over the last 25 years as well as the identity data contained in First Generation digital certificates. The paper then demonstrates the inherent difficulties of strong authentication of entities using manual vetting and the serious potential for phishing attacks and fraud against consumers if the flawed data in First Generation certificates is displayed in browser GUIs and used for identity and trust decisions.



The paper then discusses the advantages of Second Generation automated vetting, pioneered by GeoTrust and now widely adopted across the digital certificate industry, which is displacing the less reliable First Generation manual vetting of the 1990s among web site owners.

Finally, the paper provides a view to the future, and ways that digital certificates and SSL technology can, in fact, provide greater assurance of web site identity to consumers and actually help avoid, rather than enhance, phishing attacks and other online consumer fraud.

3. Understanding First Generation Digital Certificates

The technology surrounding SSL server certificates and Public Key Infrastructures has been well known since the 1970s, and will not be covered here. What's important for our purposes are the two main things that digital certificates can do to secure web sites and help avoid consumer fraud: They encrypt communications between the web site owner and consumer, and they provide certain identity data about the web site owner.

As Internet use expanded during the 1990s from universities and the defense industry (a closed community) to online commerce and broad consumer use, the encryption function has worked brilliantly, but the identity function has not.

Consumers have learned to trust the padlock symbol for sites protected by an SSL server certificate as meaning they can safely transmit their personal and financial data to complete a transaction. For the majority of consumers, they have never clicked on the lock to look "inside" the certificate at the limited identity data stored there by the issuer. A number of browser software makers are considering extracting that identity data and displaying it in the browser toolbar, but that is a flawed approach considering the inherent unreliability of that data.

Where did the rules around digital certificates come from? Who made the rules?

The development of digital certificate and PKI protocols over the past 25 years was focused primarily on technical syntax and overall system structure and design, and paid only scant attention to the specific authentication steps to be followed by CAs prior to certificate issuance. The limited discussion of authentication processes during this period (1) was written mostly by technical PKI experts, and not by commercial users of PKI or by CAs themselves, (2) recognized that different authentication steps will be appropriate for different uses and communities (e.g., closed communities of companies who knew each other versus open worldwide communities), and (3) was very vague in nature, reflecting the limited expertise of the authors in business and commercial practices.

Early industry documents establishing digital certificate and PKI protocols skipped over specific authentication steps to be used prior to issuance of digital certificates, or made only general reference to some sort of authentication process. See the

early standards stated at the IETF's RFC 791 (1981), RFC 822 (1982), and RFC 1422 (1993).¹ The standards for authentication were very weak and delegated. For example, RFC 1422 provides in part:

3.4.1.2 User Registration. Most details of user registration are a local matter, subject to policies established by the user's CA and the PCA [Policy Certification Authority] under which that CA has been certified. *** The CA will employ some means, specified by the CA in accordance with the policy of its PCA, to validate the user's claimed identity and to ensure that the public component provided is associated with the user whose distinguished name is to be bound into the certificate.

RFC 2459, "Internet X.509 Public Key Infrastructure Certificate and CRL Profile" (1999),² which set X.509 v3 standards for certificates in use today, likewise did not prescribe particular authentication standards:

2. Requirements and assumptions *** A certificate user should review the certificate policy generated by the certification authority (CA) before relying on the authentication or non-repudiation services associated with the public key in a particular certificate. To this end, this standard does not prescribe legally binding rules or duties.

The technical experts who set up digital certificates and PKI as we know it today worked only on the technical aspects and architecture; they "punted" on ultimate identity issues, leaving it to closed communities and public CAs to decide what authentication steps they would take before issuing a certificate to a business or individual. They probably thought this was the easiest part of the online identity equation; in fact, it's the hardest.

What data's in a digital certificate? Where do I find it? How did it get there?

Under X.509 v3 standards, a digital certificate is set up with the following data fields about the certificate holder:


CN	Common Name (the web address, such as www.geotrust.com)
O	Organization
OU	Organization Unit (can be a division, department, etc.)
L	Locality
S	State or Province
C	Country

However, public CAs do not all follow the same naming conventions (except for the CN field, which must be uniform across all providers),³ and do not always use all the

¹ See <http://www.faqs.org/rfcs/rfc791.html>; <http://www.faqs.org/rfcs/rfc822.html>; <http://www.faqs.org/rfcs/rfc1422.html>.

² <http://www.faqs.org/rfcs/rfc2459.html>

³ Indeed, the CN field *must* be handled in an identical fashion by all CAs in order to make the basic SSL mechanisms work; any variations from the CN standard will spoil the digital certificate.



fields. They also commonly put other information and phrases in the fields for their own purposes, such as “Subject to CPS [Certification Practice Statement] found at www.address.com” or “Liability Limited” – items that provide no identity information.

Where do you find this information?

It isn't easy. In Internet Explorer, for example, you must double click on the padlock, then choose Details, then Subject. Consumers are generally unaware the data is in there, and have never relied on it for identity purposes. *That's why the data field vulnerabilities discussed below have never made a difference before: No one has ever looked at or relied on the data before.* This will all change if certificate identity data is pulled out and displayed on the browser GUI in the future; browser users will get a false sense of security and trust a displayed identity that is totally false – a perfect phishing environment for the coming decade. The *one* exception is the CN or Common Name field – for reasons explained below, the field *must* be unique, correct and verified, and can be relied upon by consumers to know which web site they are at.

How did the data get inside the certificate? During the early 1990s, each public CA came up with its own First Generation manual vetting process. The process typically requires only the most minimal checking of documents and records, and is prone to mistakes and forgery. Most often you can find an extremely brief and cryptic description of the manual vetting process being used in the CPS (Certification Practice Statement) published by the CA and buried in its online archives. It's hard to make sense of the process – and CAs are quick to state in legal language that they make no promises about the identity of the certificate holder and won't be held legally liable for any fraud or mistakes.

What Are the Phishing and Fraud Vulnerabilities in Manual Vetting?

More specifically, First Generation manual vetting usually involves an online application from the domain holder with the faxing of a few basic business documents (copies of articles of incorporation, local business license, etc.) in an attempt to show identity, which are then briefly checked against light-weight third-party business databases. All this can be forged or faked by a certificate applicant intent on phishing or fraud as described below, meaning the identity data in an issued certificate may be completely wrong and misleading. This vetting process is all the more prone to error considering that each state, province, and country has its own set of business documentation, which no single CA can verify worldwide.

Well, who's in charge of how manual identity vetting is done? Who makes sure identity in certificates is done right?

Basically no one. Some years ago, the major browser makers required public CAs who wanted to maintain their trusted roots in the browser store of trusted roots (see below) to obtain an annual American Institute of Certified Public Accountants (AICPA) WebTrust audit – but those audits *do not* establish any standards for how identity is established in certificates. Instead, the auditors just look at what the CA *says* it does in the vetting process in the CA's CPS, then conducts spot checks of past vetting to see if the CA has been following its own rules. (The basic WebTrust audit

rule is “say what you do, and do what you say.”) If the CA has been fooled by phony documents or the process is flawed, the CA will still receive a clean audit.


Even worse, some public CAs outsource the entire identity verification process to others, without any checking or auditing of the process. The “subcontractors” can – and sometimes do – vouch for the identity of a certificate applicant without doing any checking at all. This flawed identity data is then transmitted back to the public CA, who inserts it in the certificate and signs it with the CA’s key for distribution and use by the applicant. Because WebTrust doesn’t reach these outsourced subcontractors, they are never audited and their shoddy practices are not discovered. Imagine letting consumers rely on that kind of unverified identity data if it is prominently displayed in a browser GUI.

The problem is likely to get much worse in the near future, especially if browser GUIs are changed to display certificate data to consumers for identity reliance and trust purposes. There are literally dozens of “trusted” root certificates already pre-loaded in the browser software currently used by hundreds of millions of consumers around the world. (As one example, open Internet Explorer – go to Tools / Internet Options / Certificates / Trusted Root Certification Authorities – and you’ll find over 100 trusted root certificates listed.) Any certificate issued off of one of these trusted roots – or off a sub-root or a root that has been signed by a trusted root – will be treated as a “trusted” certificate and be automatically accepted by a consumer’s browser software.

Many of these trusted roots have been sold over the years, and many chained subroots have been issued to third parties. If any one of these third parties – who may be in any country of the world – decides to use the root to issue phony or unverified certificates to companies intent on phishing or fraud, and the identity data is prominently exposed to the unwary consumer in the browser GUI, there could be an explosion in phishing attacks.⁴ Perhaps worse, after the first public exposure of this new kind of phishing, consumers might no longer trust secured Web sites *at all*, as the flaws in the SSL technology used to secure the sites would no longer be trusted.⁵

⁴ Note that a digital certificate issued by a criminal CA could not only include false O or organization data, but also false L, S, and C data as well. The Web site could be made to look like a legitimate business located in a US city, when in fact it was entirely offshore and impossible to reach.

⁵ Some may respond that the answer to errors or fraud in issuance of digital certificates is to list the certificate in a Certificate Revocation List (CRL) and/or to delete the trusted root from the browser software, which would make certificates issued off the root untrusted. However, consumers *never* check CRLs to see if a certificate has been revoked (and generally turn off the automatic CRL checking that is available in some applications), so even if a bogus certificate has been revoked after causing harm, *it can still be used successfully against other consumers for the balance of its lifetime, which may be as much as four years.* Second, dropping a trusted root from the browser root store would take years, would effectively revoke all valid certificates previously issued off that root thereby harming thousands of innocent parties, and basically has never been done. Neither is a good solution to dealing with a phishing certificate.



There have been some attempts to come up with strict authentication standards that would apply to all public CAs, but these are years from completion and for the reasons discussed below will likely not scale nor work in a worldwide environment.

Give me some examples of how First Generation manual authentication certificates can contain identity data that is faked or inaccurate.

Unfortunately, there are all too many examples. Here are just some:

- *Fraudulent documents*. With current design software and scanners, applicants can create convincing – but fake – corporate and other identity documents.
- *Real documents, fraudulently presented*. Fraudsters can easily get copies of real enterprise identity documents from public records and present them to public CAs.
- *Shell corporations and empty business data accounts*. In many states and provinces, a corporation can be formed in 10 minutes for less than \$100 using pre-printed forms. Creating an account with a business data agency – such as Dun & Bradstreet – can be done online in minutes. For many public CAs, a newly-created shell corporation with an “empty” business data account would be sufficient for certificate authentication purposes.
- *Duplicate names and naming conventions (XYZ Corp, XYZ Ltd., XYZ LLC)* – There can be 50 distinct “Acme Corporations” in the fifty states, with more in other countries. In addition, consumers are generally unfamiliar with the exact corporate name for a trusted company – Co., Corp., Ltd. – opening other opportunities for fraudulent imitation. Certificates might be issued in all these names. Which is real, and which is used by an imposter?
- *Worldwide scalability* – These business identity problems multiply exponentially for other countries. Is a Société Anonyme in France the same as a Fideicomiso in Mexico or a Joint Stock Company of the Closed Type in Russia? Who would set the standards for authentication for these international entities? (No authoritative body does it at present.) The process doesn’t scale. And what if an apparently valid certificate is issued to a shell US subsidiary of a foreign company, and all the consumer data provided to the US subsidiary is instantaneously transmitted to the foreign owner for phishing purposes?
- *No common O field vetting among public CAs* – There are no common First Generation manual vetting procedures among the top public CAs, and will not be for many years to come, if ever.
- *Non-corporate businesses (partnerships, sole proprietorships)* – There are many legitimate business entities that don’t have to be officially registered anywhere, such as partnerships and sole proprietorships. Public CAs don’t have the ability to determine their validity.
- *Corporations may be dissolved, sold, merged, etc.*, but the data in their old certificate remains unchanged and therefore be misleading.

- *The owner of the domain shown in O field of a certificate may have nothing to do with owner of the business using the site* – Many Web hosting companies and Internet service providers will secure multiple independent sites using a single digital certificate in the top level domain name, which can't contain any confirmed identity information about the individual sites being secured.

In addition, many sites outsource certain functions (e.g., web-based purchasing and billing) to other sites that are branded to look like the original site. The identity data in the outsourced site will not relate to or be the same as the original site. Consider the needs of small businesses – Yahoo's merchant sites, for example – or the needs of large business – such as United Airlines as discussed at the end of Section 4 below – where outsourcing is used. This creates a serious potential for confusing the consumer and opens tremendous phishing holes.

Many more examples could be cited. All illustrate the intrinsic unreliability of the identity data contained in the O, OU, L, S, and C fields of digital certificates obtained through First Generation manual vetting, and demonstrate why the data can't be used for identity and trust purposes. The one consistent exception is data in the CN field, which can be trusted and is at the center of modern Second Generation automated vetting, as discussed in Section 4 below.

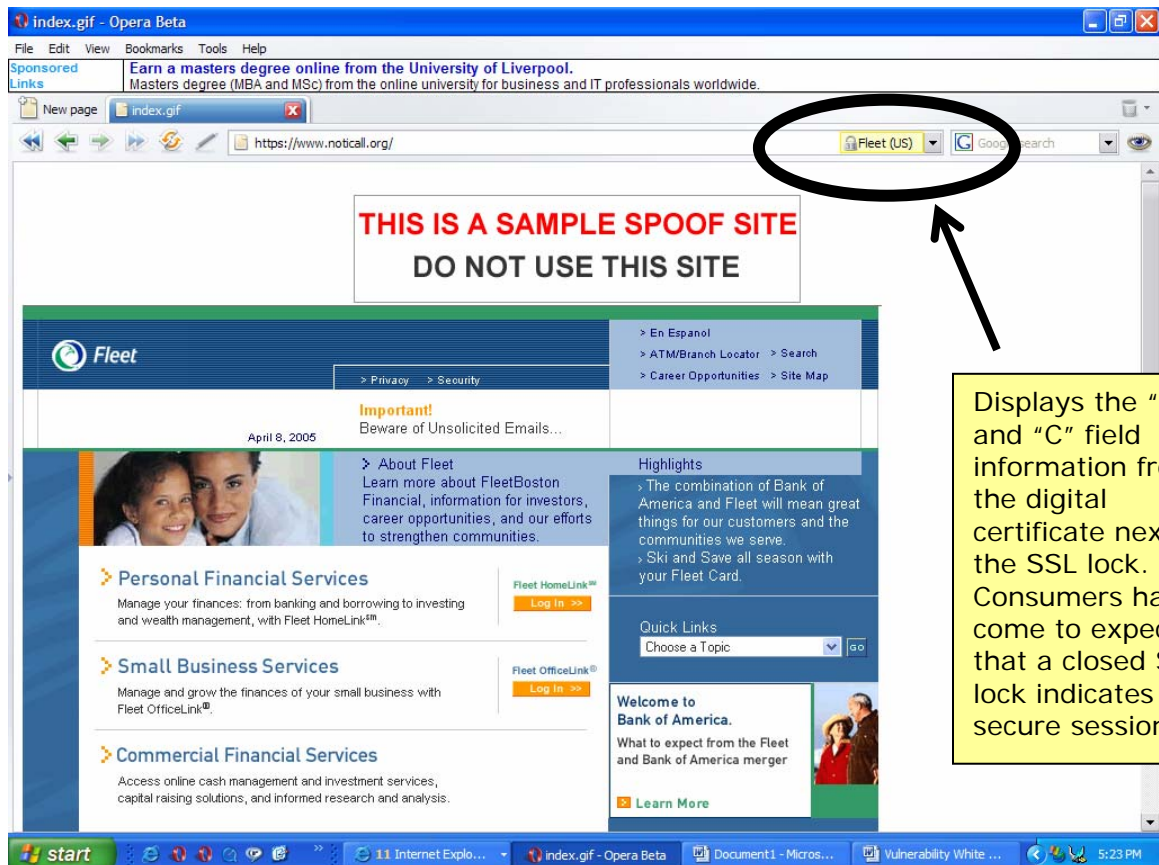
Are there any examples of fake or inaccurate data inside a certificate from a public CA?

Unfortunately, yes. False certificates have been obtained on a number of occasions from multiple public CAs *without ever submitting any false information in the process*. Go to <http://www.geotrust.com/resource/advisory/sslorg/index.htm> to see live examples. For this white paper, we have included two screen shots to illustrate how certificates with misleading organizational names could be used to enhance phishing schemes:

Example 1: An example of a spoofed Chase site with a misleading organizational name of "Chase" in O Field. Note the display of O (Organization) and C (Country) field data ("Chase (US)") in the yellow field at top right of the GUI. This is not a genuine Chase Bank site, but an example of how the organization field could be used by a phisher to more convincingly perpetrate fraud.

Displays the "O" and "C" field information from the digital certificate next to the SSL lock. A closed SSL lock indicates a secure session, so a consumer might logically think this was a safe site. In fact, it is a spoofed site.

Example 2: Another site with a misleading organizational name of “Fleet” in the O Field. Note the display of O and C field data (“Fleet (US)”) in the yellow field at top right of the GUI. This is not a genuine Fleet Bank site.



How are misleading certificates obtained?

On several occasions we have seen examples where a certificate was easily obtained that contained inaccurate or potentially fraudulent identity information. These were obtained from multiple public CAs *without ever submitting any false information in the process*. Again, phishers and fraudsters haven’t bothered to obtain fake digital certificates to date because consumers have never looked at or relied upon the identity information in the certificate (they don’t even know it’s there) – *but this will all change if next-generation browser GUIs extract and display certificate data in an attempt to provide users with site identity information for trust decision purposes*.

There is a better way – Second Generation automated vetting – as discussed in Section 4 below.

This is frustrating – how did this happen?

It happened because of a fundamental misconception by the early designers of Internet and digital certificate protocols – that authentication of business entities was easy and that universal credentials (not just credentials for closed environments) could be created using manual identity vetting as currently practiced. The designers of the system were engineers without experience or interest in the legal and policy issues surrounding verified business identity issues. And phishing and other identity fraud on unsuspecting consumers was not a problem.

Why didn't anyone notice before? Because *no one has ever looked at or relied upon the identity data found inside digital certificates*. That could all change if the data is given prominent browser display in the future.⁶

First Generation manually vetted certificates can cost as much as \$1,495.00 from some CAs, and take days to obtain. Why does anyone put up with it?

We can only guess that certificate buyers put up with old-style, expensive manual vetting because their technical people grew up with it – it was something you just had to do – and didn't realize the serious flaws with the vetting process and the data contained in a certificate. More and more companies have migrated to the newer, Second Generation automated vetting process for their digital certificates, saving money and time. In fact, recent Netcraft survey data from April 2005 shows that Second Generation automated certificates outsold First Generation manual certificates better than 2-to-1 over the past six months.

What about legal liability from faked or inaccurate identity data in manually vetted certificates?

A company can face major legal liability today in surprising ways. Consider the following example. Earlier this year, a woman sued the local Yellow Pages directory publisher in Portland, Oregon and won an astounding jury award of \$1,600,000.⁷ The woman had gone to a physician (a dermatologist) who was listed in a Yellow Pages ad under the heading "Plastic and Reconstructive Surgery." His ad said he was "Board Certified" – but he was only board certified as a physician and dermatologist, and *not* as a plastic surgeon.

After problems with the surgery, the woman discovered the doctor was not certified as a plastic surgeon and sued the Yellow Pages saying she was misled by the ad. The Yellow Pages said they publish thousands of ads, and couldn't possibly know

⁶ There may be situations where detailed manual vetting of an organization will be worthwhile in the future, such as in a closed environment the mortgage banking industry where the parties know each other and want to use commonly-issued digital certificates for advanced purposes such as legal digital signatures on original mortgage documents. See, for example, the authentication processes put forward by Secure Identity Services Accreditation Corporation of the Mortgage Bankers Association. See www.sisac.org. However, closed system manual vetting will never work to create a universal certificate that can be trusted and used anywhere in the world wide web for identity purposes because of the reasons outlined in this paper.

⁷ See

<http://www.aberdeennews.com/mld/aberdeennews/news/11001236.htm?template=contentModules/printstory.jsp>

whether individual ads are true or not. (A representative of the publisher was quoted as saying "We publish 260 directories in 14 states. We don't validate every claim.") The jury didn't agree, and hit the Yellow Pages for \$1,600,000 in damages. The case will likely be appealed to a higher court, but in the meantime the Yellow Pages must post a bond for the full judgment amount.

There could be a similar legal liability for browser makers, and perhaps others, if digital certificate data is displayed in a browser GUI and relied upon by a consumer for identity and trust purposes. As noted above, most identity data inside First Generation manually vetted certificates is inherently unreliable, a fact that will be instantly attractive to phishers. A consumer may say, "I thought this was my bank's *real* web site because I saw my bank's name displayed by the browser as confirmed identity information. A phisher then stole and used my personal data as a result."

There are, in fact, known examples of certificates wrongly issued by public CAs in the name of a bank or other legitimate business. If consumers are phished by means of one of these false certificates displayed on the browser GUI, they may argue that the browser maker is responsible for misleading them as to the web site's identity, and major legal liability could result. Courts may choose to ignore the language in browser User Agreements (the "fine print") that says the browser maker can't be sued for false information, saying that's not fair to consumers and that the browser maker encouraged the consumer to rely on the data by its prominent display.⁸

Bottom line:

First Generation manual vetting is inherently flawed, and digital certificate identity data (except for the CN common name field showing the Web address, which can't be faked and is unique) is too unreliable to be used for identity purposes.

⁸ Consider also the recent case of [Lopez v. Bank of America \(BOA\)](#) in Florida. Lopez alleges that his bank, BOA, convinced him to go from telephone-based commercial banking to online banking and that a hacker obtained his password through by means of a virus attack against his computer. The cybercriminals then transferred \$90,000 from his BOA account to Latvia. Lopez's legal theory is that the Bank induced him to go to online banking, knew of the risk from the virus but failed to warn its customers, and was negligent in allowing a large transfer to a known center of cybercrime. Most legal observers think the case will go to trial, and banks are very concerned. See http://www.theregister.co.uk/2005/02/08/e-banking_trojan_lawsuit/ and also <http://www.financialcryptography.com/mt/archives/000433.html>. Similar legal claims might be asserted if incorrect or fraudulent certificate data is prominently displayed to consumers in a browser GUI.

4. The Answer: Second Generation Automated Vetting – Cheaper, Faster and the Key to Identity Verification Improvements for the Next Decade

Consumers are desperate for better identity solutions to avoid phishing and online fraud. They once trusted email messages that claimed to be from familiar senders, until they learned the senders could be phishers in disguise who were only interested in stealing their data. Later, they stopped trusting apparently familiar Internet addresses (long, complicated addresses that started with a familiar and trusted web address, but ended with a different phisher's address), and after that stopped trusting web pages set out in the spitting image of a trusted web page belonging to their bank, favorite auction site or credit card provider. Now, many people will not trust *any* unsolicited email, even when validly sent by their actual business partner.


This is really frustrating – the Internet is growing exponentially, and needs a better identity trust framework. What is to be done?

The answer to the growing online identity problem starts with an increased use of SSL and digital certificates generally to secure web sites and verify their *actual* Internet address in a way that can't be faked or entered in error. This will be accomplished as public CAs gradually abandon their 1990s First Generation manual vetting processes (which cost a lot in time and money without adding any reliable and trusted identity data) and move to Second Generation automated processes. This has already started to happen, with GeoTrust taking the lead in 2001 and several other public CAs now following suit.

GeoTrust's Second Generation automated vetting process works because it confirms domain (web address) control by the applicant in real time, and is backed up by real time email and telephone validation, combined with sophisticated fraud-detection algorithms.

The second major part of the solution is greater use and browser display of the web site's *confirmed Internet address*⁹. This is important because all public CAs treat verification of the CN field in the same way, and any discrepancy between the web site's actual Internet address and the *confirmed* Internet address as stated in the CN common name field of a certificate *will automatically generate a significant warning to the consumer via a pop-up notice*. This avoids such dangerous phishing methods as the use of long complicated address strings that start with a valid address but end with a direction to a phony site used for phishing. *The prominent display of the CN*

⁹ Some may warn that consumers can still be tricked by confusingly similar domain names that could actually registered to phishers or are otherwise the wrong address (for example, www.ebay123.com presumably is not related to the real www.ebay.com). There are two answers to this. First, the automated vetting process can include sophisticated algorithms that block issuance of certificates to variations of the popular top level domain names most commonly copied by phishers, and the list of blocked sites can be modified hourly based on phishing reports. Second, the other identity enhancements discussed below will help alert a consumer to at least look more closely at the site before assuming it's valid.



*field in the browser GUI will help users know where they are and learn which sites to trust. Put another way, the CN field (web address) is the *only* piece of data in a digital certificate that's confirmed, guaranteed to be unique, and is registered with an official public domain registry.¹⁰*

The third major part of the solution is the automatic use and display of other trust information that is useful to the consumer, including the following:

- The name and logo of the CA who issued the certificate. Consumers will soon learn from news reports which CAs to trust and which CAs use sloppy procedures and should not be trusted.
- Automated checking against "black lists" for reported bad Web sites.
- The display of trusted and verified seals such as TRUSTe and Better Business Bureau Online.
- The display of user rating data on sites like the business data compiled by BizRate and others.
- Active use of CRLs (Certificate Revocation Lists) and OCSP (Online Certificate Status Protocol) responders to check the revocation status of all certificates. This could be accomplished most easily if browser makers set the default for the CRL checking function to "on." False certificates would then no longer be accepted once revoked.
- Cooperation by Public CAs in working together to find common solutions. This could include creation of black lists available to all.
- Most important, the application of sophisticated algorithms and progressive heuristic techniques by the issuing CA for ongoing fraud detection, as presently used by GeoTrust. These can be changed hourly as needed based on new reports of phishing and identity fraud. Suspicious certificate issuance can be blocked, and already issued certificates can be revoked *and* entered on a black list for immediate warning display in a browser GUI.

The results of all these tests can be displayed via the browser GUI to help consumers decide which web sites to trust, which to be cautious of, and which to avoid entirely.

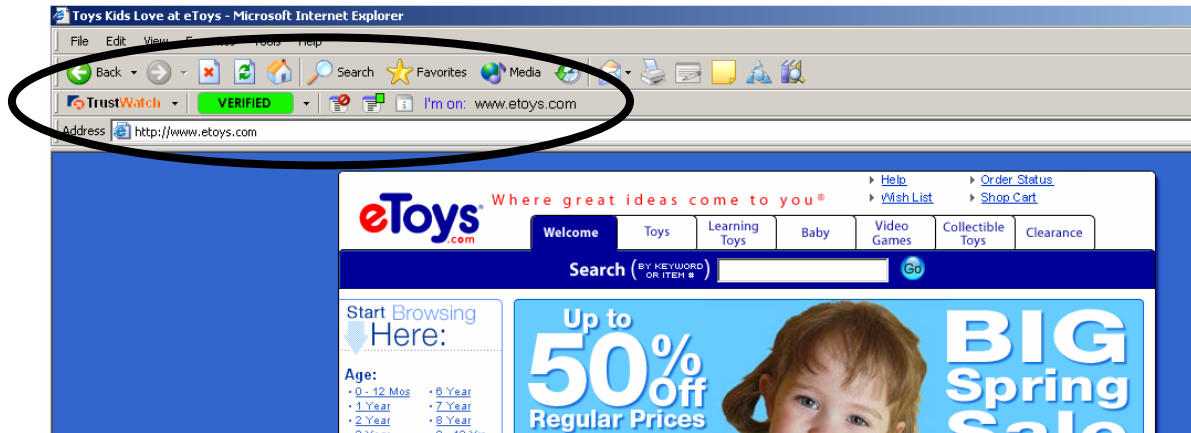
Enhancing Trust Capabilities

Various browser plug-ins and toolkits are already available for this purpose, including GeoTrust's TrustWatch.¹¹ TrustWatch applies its rules and algorithms in real-time (including checking on certificate CN field data for certificates located *anywhere on the visited web site*, even when the consumer has only started with the site's

¹⁰ As an example, all .org domain names must be officially registered with the Public Interest Registry according to ICANN rules, and official registration data can be viewed through the registry's Whois page. See http://www.pir.org/whois_search/.

¹¹ See <http://www.trustwatch.com/about.html>

unsecured home page) and displays the result in a color-coded (green-yellow-red) GUI that is easily seen and understood by consumers. Additional enhancements are coming in the near future.



There are other components of future solutions as well. For example, consumers can “vote on” the trustworthiness and credibility of a given SSL-secured site and report phishing and fraud attempts related to the site. Browser interfaces can then display in real time the cumulative result of that consumer feedback. Sites that have too few responses or too many negative results will deserve caution and little trust, while other sites with many positive results will deserve strong trust from the consumer. This is how eBay, Amazon and Yahoo establish trust relationships across millions of unknown buyers, sellers, and other users every day -- *not* by means of any First Generation manual vetting process. Such a process will democratize trust decisions and create a distributed trust environment.

Another potential solution is for trusted SSL secured sites to “vouch” for their critical companion sites by means of a displayed and verifiable hyperlink. Consider the example of United Airlines. When a consumer visits the main United page, www.united.com and clicks on “Flights”, the consumer is immediately redirected to a secured page at the address <https://www.itn.net> that has the same look-and-feel of a United page. Who is itn.net? A search of the relevant Whois registry for the domain and related domains¹² shows a confusing array of names in various cities (itn.net, Travelocity.com Internet Services, alston.com; GetThere.com, etc.). Who are consumers really dealing with? The system works today because the consumer is relying on the look-and-feel of the site (branded as United Airlines) and previous successful transactions initiated through a trusted initial Internet address, www.united.com. But this won't work for unknown web sites that are new to a viewer. Consumer confidence in redirected web pages might be increased if the original site includes a legend on the redirected page stating as follows:

¹² See footnote 9. There is an ICANN-approved official registry for each domain (such as .net and .com), and the registry must maintain an online Whois searchable data base to allow the public to determine who has registered a particular domain name.

“You are now at a trusted partner’s site to complete your transaction with us. We vouch for the trustworthiness of this site. To confirm, please click here: https://www.originaldomainname.com/trusted_partners,”

At that point, the consumer could confirm that the site is included in a list of trusted partners by clicking on the hyperlink and seeing that the original site has listed the partner and its address on a secured page in the original domain. This would create transitive trust environment and enhance the user experience.

The addition of other identity security techniques, sophisticated algorithms, and trust practices is possible – it’s happening now with many browser makers, public CAs, and application providers – so long as the Internet industry moves away from the First Generation manual vetting processes initially designed in the early 1990s, and does *not* try to rely in the inherently inaccurate identity information contained inside the data fields (O, OU, L, S, and C) of those certificates as a means of creating a universal credential. That old process is generally a waste of time and money, and increased reliance on the flawed data it produces will create vast new opportunities for phishing and other online identity fraud. Instead, with the use of Second Generation automated vetting of identities tied to the CN common name field that can’t be faked, additional identity and anti-fraud techniques can be layered in depth to help consumers make more sophisticated trust decisions that simply are not possible today.

Conclusion:

Be wary of supposed “confirmed” identity information contained inside First Generation manually vetted digital certificates – a process developed with little thought in the early 1990s. The manual vetting process is inherently vulnerable to mistakes and fraud, and no public CA takes responsibility for the actual identity of a certificate holder. (Some public CAs don’t even bother to vet the business entity information before issuing a certificate.) This creates the perfect environment for new and potentially devastating phishing attacks if the data is now displayed in browser GUIs for identity verification purposes.

Second Generation automated vetting is gaining rapid acceptance among certificate buyers and many public CAs since GeoTrust’s introduction of the concept in 2001, outselling First Generation manual certificates better than 2-to-1 over the past six months. This lowers the cost and time for certificate issuance, and helps expand the use of SSL to secure the web. By relying on the CN or common name for web sites that can’t be faked, Second Generation automated vetting establishes the groundwork for layering of other new, sophisticated algorithms and trust techniques and browser displays to help consumers make *better informed trust decisions*. *It’s the growing trend for the next decade.*