

X.509 User Certificate-based Two-Factor Authentication for Web Applications

Thomas Zink, Marcel Waldvogel¹

Abstract: An appealing property to researchers, educators, and students is the openness of the physical environment and IT infrastructure of their organizations. However, to the IT administration, this creates challenges way beyond those of a single-purpose business or administration. Especially the personally identifiable information or the power of the critical functions behind these logins, such as financial transactions or manipulating user accounts, require extra protection in the heterogeneous educational environment with single-sign-on. However, most web-based environments still lack a reasonable second-factor protection or at least the enforcement of it for privileged operations without hindering normal usage.

In this paper we introduce a novel and surprisingly simple yet extremely flexible way to implement two-factor authentication based on X.509 user certificates in web applications. Our solution requires only a few lines of code in web server configuration and none in the application source code for basic protection. Furthermore, since it is based on X.509 certificates, it can be easily combined with smartcards or USB cryptotokens to further enhance security.

Keywords: multi-factor-authentication, authentication, crypto token, S/MIME, certificate, X.509

1 Introduction

At the time of this writing, the most commonly used user authentication scheme for login to digital services of all kinds is still a combination of username and password. If an attacker guesses or steals the users password he can effectively steal the users digital identity, access the user's personal information, and perform actions in the user's name. The password is often the weakest link[FH07].

For many applications this level of security might be enough. But sensitive information or the ability to perform critical actions require additional layers of security. Multi-factor authentication (MFA) is a methodology that introduces additional, independent authentication factors that all need to be validated when authenticating a user. In some application domains, like e-banking, multi-factor authentication is already an established mechanism. Many well-known service providers have also adopted 2FA (two-factor authentication, often called '2-step verification'), including Apple, Google, Microsoft, and even Steam.

The factors differ in nature and include:

- Something you know (knowledge of a secret, like a password or PIN)

¹ University of Konstanz, 78457 Konstanz, <first.last>@uni.kn




Figure 1: The different stages of securing a web application with certificate-based second factors. Adding a software certificate provides an additional factor without linking it to the users identity. Application support is required to link users with their certificate. Instead of storing the certificate in software, it can be stored on a hardware token, to guarantee it stays private.

- Something you have (possession of an object, like a smart card or cryptographic token, aka cryptoken)
- Something you are (physical features, like fingerprint or retinal patterns)

Independence of the authentication factors is critical to ensure that a compromise of a single factor does not affect the others. Although the authentication factors are independent, their affiliation to the authenticating identity must be apparent and verified by the service. In fact, many authentication schemes, that claim to offer MFA, actually fail to associate all authentication factors with the same single identity, as discussed in [Section 3](#).

MFA introduces additional complexity in applications as well as identity management [Si16]. Many standard services like LDAP or AD do not support MFA out of the box. As a result, additional authentication services and user management is usually required.

In this paper, we suggest another unique form of two-factor authentication for web applications that requires little changes in the application and builds on existing identity management infrastructures. We utilize standard X.509 user certificates (aka S/MIME certificate) as proof for the user’s identity and as a first factor for authentication. After the user has provided his certificate and thus proven his identity he is prompted for his password. This scheme can easily be combined with standard hardware solutions like smartcards or USB cryptokens ([Figure 1](#)). It is surprisingly easy to implement, yet affective and user friendly, and can be adapted quickly and flexibly for any type of web application without the need for another online service.

2 Related Work

For many decades now, passwords have been considered a necessary evil. Already in 1984, Thompson[Th84] noted that password checking can easily be circumvented. A few years

later, Muffett[Mu91] published *Crack* which allowed brute-force cracking of passwords (and enabled system administrators to identify passwords vulnerable to this attack). The primary target of network sniffing, keyloggers and several trojans[MSK05] also were passwords. Herley and Florêncio[HF06] present a non-technical way to fool keyloggers by interspersing the password with random character sequences which are entered when the password field does not have the input focus.

Given the success of break-ins to steal passwords by the million and even billion in recent years[Th16]. The requirement of some organizations to force their users to regularly change their password without any exposure has led users to increasingly chose weak passwords, write them on Post-Its, or share them between multiple accounts instead of using password managers[FH07, Cr16]. Even passwords stored in encrypted form are worthwhile, as the users often use low-entropy passwords and share them between several accounts, reinforcing the importance of the weakest link.

Countermeasures including one-time passwords in software[Ha95] or hardware tokens [BLP03] have been introduced to overcome the password weakness. However, they have only been deployed in selected environments, as widespread compatibility to existing software has been limited, especially in open and heterogeneous environments such as those seen in research and higher education.

There, various forms of authentication over several protocols (local, LDAP, AD, web-based, ...) are used, often in combination to ensure access even if some systems have failed. For services offered among institutions of higher education, Shibboleth is now becoming the de-facto standard, even beyond web applications[Si12]. Work is under way to include two-factor authentication into Shibboleth[Si16]. This currently requires use of Shibboleth, which is not common for in-house applications yet and requires configuration at the Identity Provider, not (only) at the Service Provider.

3 State of the Art

Support for Multi-factor authentication in daily applications is steadily increasing. Amazon, like many banks, sends SMS text messages to the mobile phone. Apple relies on the ARM processor's TrustZone to keep keys from prying hackers. Github uses TOTP (temporal OTP) using a smartphone authenticator application, it also allows U2F using YubiKey. Similarly, Google also makes use of TOTP and U2F.

TOTP solutions are recommended to be coupled with a central server to avoid reuse of the same token for different services. This creates a dependence on additional servers and the network. Given that this solution is aimed at (distributed) administration in a heterogeneous network, a partial network or service downtime might prevent other systems from running.

LinOTP (and it's fork PrivacyIdea) is a generic OTP solution for Linux and can also integrate into the web server. However, it requires a separate user management and modifications to the application to link authenticating users to tokens. Without these changes, the

assumed two-factor authentication degrades to two single-factor authentications of possibly distinct users. Moreover, LinOTP provides a self provisioning service for users. While this enables users to individually configure a second authentication factor, the service itself is only secured with the users password, undermining the whole 2FA.

X.509 Client certificates for authentication do not see widespread distribution, at least in consumer products. They are mostly found in areas with restricted access, usually in the form of smartcards. Web server support for client certificates has been around for years, however, the common use case is what we call a ‘certificate wall’. That is, access to a resource on the server is only granted with a valid certificate signed by a specific authority. Again, this is not a valid form of multi-factor authentication, since the identity of the user is not crosschecked with the identity provided by the client certificate.

4 Threat Model

Our solution was triggered by the requirements of the following two high-risk groups:

Help-desk support personnel. To underline the openness of higher education institution, face-to-face support help desks are increasingly used, often residing in a public space such as the library. The opening times often extend beyond the staffing period of the desk.

Therefore, the machines will be left unattended overnight and could be tampered with, including installing software or hardware keyloggers. Physically locking the machines away daily would require significant effort.

Furthermore, the personnel operating these machines often has significant privileges, such as creating a new user or resetting a user’s password.

Roaming system administrators. In a user-friendly setting with diverse client setups throughout an education and research organization, support staff and system administrators will have to visit users at their desk. Often, some configuration steps have to be performed from these computers. However, it remains unclear whether the machine is clean from trojans or keyloggers.

It can generally be assumed by the user in front of the machine that the machine has not been tampered with. However, a small risk remains.

Similar thinking generalizes this to personal machines (laptops or desktops) in office rooms. An intruder might have stealthily broken into the room and modified the computer, i.e., performing an *Evil Maid* attack[Sc09]. Even a machine of a vigilant system administrator can be infected by malware, especially in a targeted attack. The risk there is even smaller, but remains¹.

Our threat model therefore includes

¹ The analysis of one local incident suggests that this has happened at least once.

- potentially untrusted terminals (public, customer, possibly hacked, keylogger),
- operated by a small, controlled user group with elevated privileges,
- logging in to (one or several) web services with elevated privileges.

The goals therefore are to

1. avoid additional effort on this group for unprivileged operations,
2. minimize the impact for privileged operations, and
3. prevent stealing of credentials with long-term validity required to perform these privileged operations, while
4. requiring minimal intervention to those web services.

5 Solution

Our solution is based on the observation, that the vast majority of web applications use the user's email address – or the local part of the email address – as the user's login name. This holds for most public web services, like Amazon or eBay, but also for many public or private organizations and companies that operate their own email infrastructure. Usually, the user's email user name/address is also his unique identifier (UID) used for organization-wide authentication. If the email address is not used for authentication in the organization, another UID scheme usually is in place. X.509 certificates must also be unique for a user and often demand the usage of a 'distinguished name' as subject. This subject can also include the UID, which can then be used for authentication instead of the user's email address.

Many organizations also either have their own certificate authority, or use the services of a public or commercial certificate authority. Especially the public authorities are likely to operate their own public certificate authority for use in government agencies or educational institutions.

While client certificates slowly take off, particularly in areas where trust or confidentiality is crucial, x.509 certificates are still usually found on servers and used for host authentication and SSL/TLS encryption. Because of this, support for x.509 client certificates has long been neglected. However, interest in stronger authentication schemes and x.509 client certificates as well as cybersecurity in general has risen in recent years. This is due to a combination of events, including Snowden's leak of thousands of confidential NSA documents and well-known attacks on cloud providers (e.g. Apple's iCloud "celebrity breaches"). As a result application developers and service providers have started to adopt a variety of multi-factor authentication schemes, from relying on trusted devices (Apple) to hardware tokens (Amazon, Github).

Still, aside from specialized applications that use smartcard authentication, x.509 client certificates are usually only used on web servers to protect locations or directories. For

example, anyone providing a valid certificate signed by a specific authority is allowed to access the login screen of an application (figure 2). While this adds another authentication factor this cannot be considered 2FA, since there is no process that actually validates that the owner of the certificate is the same entity that tries to login to the application. The reason for this is because the certificate request is issued by the web server while the login is requested by the application.




Figure 2: A typical certificate “wall”. The server requests a valid certificate from the client to allow access to a location.

5.1 2FA with client certificates

In order to enable true multi-factor authentication the application needs to check the identity of the client certificate and somehow match that to the identity of the user trying to login. This requires additional effort in identity management to provide this missing link.

We observed, that for most web-based applications, users login with their email address or the local-part of their email address. Since x.509 user certificates are issued for a specific email address, it is surprisingly easy to extract the user name directly from the certificate on the server and use this as a validated remote user for the application. If the web server supports client certificate parsing, this can even be done directly on the web server. If not, the web server needs to expose the client certificate to the application which in turn parses the certificate and extracts the user name.

Apache2 supports the necessary options since version 2.4.12. Prior to that, the client certificate must be exposed to the application. Listing 1 shows the configuration for apache2. The important part is the export of the certificate and the environment variables.

List. 1: Apache config prior to version 2.4

```

<VirtualHost *:443>
...
SSLOptions +ExportCertData +StdEnvVars
...
SSLVerifyClient require
SSLVerifyDepth 3
SSLRequire ( %{SSL_CIPHER} !~ m/^(EXP|NULL)-/ )
...
</VirtualHost>
  
```

On the application side, parsing the client certificate and extracting the email address can be achieved in few lines of code ([listing 2](#)).

List. 2: Email extraction with apache < 2.4.12

```
$clientcert = openssl_x.509_parse($SERVER[ 'SSL_CLIENT_CERT' ] );
$subjectaltname = $clientcert[ 'extensions' ][ 'subjectAltName' ];
$email = substr( $subjectaltname , strpos( $subjectaltname , ":" ) + 1 );
```

With apache2 version 2.4.12 or later, extraction of the user's email address can be done with a single line of code ([listing 3](#)).

List. 3: Email extraction with apache > 2.12.24

```
$email = $SERVER[ 'SSL_CLIENT_SAN_Email_0' ] ;
```

In any case, [listing 2](#) shows the appropriate login form for the application. The form uses the formerly extracted email address as readonly value for the user name. The user is thus not able to actively change the login name. Since the web server verifies the authenticity of the client certificate, the user's email address is also verified.

List. 4: Login form with readonly user name

```
<form>
<label><b>Username: </b></label>
<input type='text' name='user' value='<?php echo $email; ?>' readonly /> <br />

<label><b>Password: </b></label>
<input type="password" placeholder="Enter Password" name="psw" required />

<button type="submit">Login</button>
<button type="button">Cancel</button>
</form>
```

5.2 Securing web applications

The proposed scheme can be used to efficiently secure any kind of web application with multiple authentication factors. A remaining issue is the storage location of the client certificate itself. Saving the certificate in the browsers store exposes the private key to a number of attacks that might lead to its loss.

Fortunately, our solution can be easily combined with hardware tokens like smartcards or USB cryptotokens to ensure that the private key stays private, even if the user's machine is compromised. Not all browsers support hardware-based certificate stores, though. Depending on the required trust level, the application can be secured accordingly.

5.2.1 Example use cases

The ease with which our authentication scheme can be implemented allows many applications.

Web Server or Proxy Using http auth and authentication modules, the whole web server can be secured. To quickly add 2FA support to multiple applications or servers, an authentication proxy could even be used.

Locations/Specific Applications Secure only specific locations or applications served by the web server, using true two-factor authentication.

Specific Actions Using the option `QUERY_STRING` to match post parameters or keywords in the URL, additional authentication factors can be requested for these specific actions.

6 Analysis

Our solution perfectly fulfills the goals stated in [Section 4](#), as follows:

1. Unprivileged users or operations incur no extra burden.
2. Privileged operations are performed with zero effort for the basic (*software certificate*) case of [Figure 1](#) and only a little bit more effort for the advanced protection modes. This applies to both setup and actual use.
3. The credentials that can be stolen do not provide any long-term access to privileged operations. If the computer is infected and under carefully targeted control, the current session may be hijacked; but there is no known general protection against this risk.
4. The web application hidden behind the 2FA does not need to be modified. Minimal changes are required to the web server (or a proxy inserted) for a first strong protection step.

As administrative commands can be identified by URL path or parameters, environment-specific definitions of what is considered a privileged operation can be easily applied. As access to these URLs is only granted to specific certificates, our 2FA solution even protects against several implementation or verification errors inside the application itself. As no online verification is needed with a central service (unlike other services like U2F or linOTP), this system does not increase the damage caused by a partial network or service outage at the university. The user certificates distributed for this purpose can be used to secure other services as well, including email, or vice versa, if policy allows. Then two birds are caught with one stone.

7 Evaluation

We successfully implemented our client certificate-based authentication in our support front-end application used by supporters in our public help desk center. The workstations used to login are publicly exposed and leakage of passwords is a real threat. We already use the DFN PKI for x.509 certificates which allowed us an easy transition towards the certificate authentication.

Each support staff member receives a USB cryptoken to store his personal certificate. The workstations' browsers are configured to access the certificate on the token, which is secured by a person identification number (PIN). To successfully log into the support front-end the supporter has to provide the token, his PIN and then his personal password.

We conducted a user study among the support staff to evaluate the quality of the solution. Unfortunately, only six staff members reacted. However, according to a study by Nielsen [NL93] a small sample is enough to identify a majority of usability problems.

Table 1 shows questions about the awareness of the user regarding the security of his password and the offered solution.

Table 1: Questions regarding user awareness

Question	Yes	No
Do you have to enter your password on untrusted devices	83%	17%
Do you fear that your password is in danger	83%	17%
Do you think the solution is justified	66.7%	33.3%

We also wanted to know the impact the solution has on the users' daily work and how well it is perceived (table 2).

Table 2: Questions regarding user experience

Question	not	slightly	moderate	very
How much does the solution disrupt your work	16.7%	16.7%	16.7%	50%
How satisfied are you with the solution	16.7%	16.7%	16.7%	50%

Finally, we asked if the users know other methods for multi-factor authentication, and if they perceive our solution as better, worse oder equally good.

Table 3: Comparison to other solutions

solution	don't know	better	worse	equal
Smartcard	16.7%	16.7%	0%	66.7%
SMS-TAN	0%	16.7%	83.3%	0%
U2F	16.7%	16.7%	0%	66.7%
HOTP/TOTP	16.7%	16.7%	33.3%	33.3%

In summary, most users are aware that their credentials are in danger and that additional security measures are justified. However, users feel that the additional security impedes

their daily work. Compared to other solutions our proposed scheme is always perceived as at least equally well or even better.

7.1 Comparison to other methods

We shortly compare our proposed method with the previously mentioned schemes with respect to the changes required to applications; dependence on additional services, hardware or software; and administrative overhead (e.g. user management).

Apart from smartcards, all other schemes require at least additional user management, in case of U2F or some OTP solutions even additional services that need to be maintained, patched, and require high availability. In addition, applications that want to utilize these additional authentication factors need to be changed significantly.

Our solution can easily be used with smartcards or with USB cryptotokens that have x.509 support. In case of smartcards, card readers must be provided. USB tokens work without any additional hardware requirements except for the standard USB port. In practice, however, we realized that driver and application support for both smartcards and cryptotokens is still in its infancy. But that also holds for other second factor methods.

8 Conclusion and Future Work

Sections 6 to 7 clearly show how well the system performs and how easy and flexible it is to set up and use. This applies to both service provider and users.

According to our experience, security mechanisms are best deployed, when they can be installed with just a few keystrokes, no questions asked. The instructions in this paper, while easy to follow, are already too complex for wide deployment. We therefore plan to create finished packages including Docker containers for proxies that can be integrated into existing applications or wrapped around these applications as a proxy.

Acknowledgments

This work was supported in part by the Ministry of Science, Research and the Arts (MWK) of the State of Baden-Württemberg through the funding of project bwITsec.

References

- [BLP03] Biryukov, Alex; Lano, Joseph; Preneel, Bart: , Cryptanalysis of the Alleged SecurID Hash Function. Cryptology ePrint Archive, Report 2003/162, 2003. <http://eprint.iacr.org/2003/162>. 2

- [Cr16] Cranor, Lorrie: , Time to rethink mandatory password changes. <https://www.ftc.gov/news-events/blogs/techftc/2016/03/time-rethink-mandatory-password-changes>, Marz 2016. Retrieved 2017-02-01. 2
- [FH07] Florêncio, Dinei; Herley, Cormac: A Large-scale Study of Web Password Habits. In: Proceedings of the 16th International Conference on World Wide Web. WWW '07, ACM, New York, NY, USA, S. 657–666, 2007. 1, 2
- [Ha95] Haller, Neil: The S/KEY One-Time Password System. RFC 1760, IETF, Februar 1995. <https://tools.ietf.org/html/rfc1760>. 2
- [HF06] Herley, Cormac; Florêncio, Dinei: How to Login from an Internet Café Without Worrying about Keyloggers. In: Proceedings of SOUPS 2006. Juli 2006. Retrieved 2017-02-01. 2
- [MSK05] McClure, Stuart; Scambray, Joel; Kurtz, George: Hacking Exposed: Network Security Secrets and Solutions. Osborne/McGraw-Hill, 3. Auflage, 2005. 2
- [Mu91] Muffett, Alec David: , CRACK: A Sensible Unix Password Cracker. Message-ID <1991Jul15.183637.6511@aber.ac.uk> in newsgroups alt.sources, alt.security, Juli 1991. Retrieved 2017-02-01 from <https://groups.google.com/d/msg/alt.sources/NA1VIInmbvpk/GLdI4Dgv95MJ>. 2
- [NL93] Nielsen, Jakob; Landauer, Thomas K.: A mathematical model of the finding of usability problems. In: Proceedings of ACM INTERCHI'93 Conference. S. 206–213, April 1993. 7
- [Sc09] Schneier, Bruce: , “Evil Maid” Attacks on Encrypted Hard Drives. https://www.schneier.com/blog/archives/2009/10/evil_maid_attac.html, Oktober 2009. Retrieved 2017-02-01. 4
- [Si12] Simon, Michael; Waldvogel, Marcel; Schober, Sven; Semaan, Saher; Nussbaumer, Martin: bwIDM: Föderieren auch nicht-webbasierter Dienste auf Basis von SAML. In: 5. DFN-Forum Kommunikationstechnologien: Verteilte Systeme im Wissenschaftsbereich. S. 119–128, 2012. 2
- [Si16] Simon, Michael: , 2FA mit Shibboleth mit LinOTP. Presentation at 65. DFN-Betriebstagung, Berlin, 2016. Retrieved 2017-02-01 from https://www.dfn.de/fileadmin/3Beratung/Betriebstagungen/bt65/BT65_AAI_Shib-LinOTP_Simon.pdf. 1, 2
- [Th84] Thompson, Ken: Reflections on Trusting Trust. Commun. ACM, 27(8):761–763, August 1984. 2
- [Th16] Thielman, Sam: Yahoo hack: 1bn accounts compromised by biggest data breach in history. The Guardian, Dezember 2016. Retrieved 2017-02-01 from <https://www.theguardian.com/technology/2016/dec/14/yahoo-hack-security-of-one-billion-accounts-breached>. 2