

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 16, 2013

H. Tschofenig
Nokia Siemens Networks
S. Turner
IECA, Inc.
S. Farrell
Trinity College Dublin
M. Hanson
Mozilla
July 15, 2012

An Inquiry into the Nature and the Causes of Web Insecurity
draft-tschofenig-secure-the-web-02.txt

Abstract

The year 2011 has been quite exciting from a Web security point of view: a number of high-profile security incidents have gotten a lot of press attention but also new initiatives, such as the National Strategy for Trusted Identities in Cyberspace (NSTIC), had been launched to improve the Web identity eco-system. The NSTIC strategy paper, for example, observes problems with Internet security due to the widespread usage of low-entropy passwords and the lack of widely deployed authentication and attribute assurance services.

With this memorandum we try to develop a shared vision for how to deal with the most pressing Web security problems.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 16, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the

document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Terminology	6
3.	Passwords	7
4.	Roadmap	9
5.	From Two-Party to N-Party	12
6.	IANA Considerations	14
7.	Acknowledgments	15
8.	Open Issues	16
9.	References	17
9.1.	Normative References	17
9.2.	Informative References	17
	Authors' Addresses	19

1. Introduction

HTTP is an IETF standard and documented in [RFC 2616](#) [[RFC2616](#)] and provides the core foundation of the browser-based platform but is also widely used for non-browser-based applications in smart phones and Internet tablets. Like any other specification in the IETF HTTP also comes with various security mechanisms. Digest authentication support in HTTP was published in 1997 with [RFC 2069](#) [[RFC2069](#)] and later updated in 1999 by [RFC 2617](#) [[RFC2617](#)]. The HTTP state management mechanism, namely cookies, was initially published in 1997 with [RFC 2109](#) [[RFC2109](#)], and re-written in 2000 by [RFC 2965](#) [[RFC2965](#)].

For client side authentication two different solution tracks have therefore been offered from the IETF, namely TLS client side authentication (at that time the usage of client certificates was envisioned) and also application level authentication via HTTP basic and digest. TLS-based client authentication using certificates was quite complex for end users to configure (and still is complex today). HTTP based authentication on the other hand did not find widespread usage either for a number of reasons. First, the user interface was rendered differently than in regular Web application form making it less attractive for users. At that time HTTP had a semantics that was closer to file system access control and therefore the decision making process was binary, either the user was granted access to the resource or it wasn't. With the HTTP 401 there was no way for a user to, for example, recover from a lost password or other forms of failure cases. The authentication and authorization process was not seen as continuum but rather as a binary decision. For these reasons form-based authentication mechanisms had found widespread acceptance by the Web application developer community. Additionally, many Web sites decided to deploy their own authentication infrastructure and to store cleartext credentials (in most cases a username and a password) into a database. To add to this problem cookies were and still are the most common mechanism for session management, i.e., a non-cryptographic way to bind the initial authentication to the subsequent HTTP protocol exchanges. Cookies introduce various weaknesses into HTTP, including the ability for attackers to perform session hijacking. In the last few years we have seen many press announcements of password leakage due to unauthorized access to these credential databases.

In the last few years a few other standardization efforts were started: [RFC 2965](#) HTTP state management specification was recently revised to capture deployment reality [[RFC6265](#)]. HTTP Strict Transport Layer Security (HSTS) [[I-D.ietf-websec-strict-transport-sec](#)] allows Web sites to declare themselves accessible only via secure connections, and the attempt to

clarify the Web Origin Concept [[I-D.ietf-websec-origin](#)], which covers the principles that underlies the concept of origin as used to scope of authority or privilege by user agents. The HTTPbis Working Group [[I-D.ietf-httpbis-p7-auth](#)] revises [RFC 2616](#) plus those parts from [RFC 2617](#) that describe the authentication schemes.

A lot has changed over the last 10 years in the Web eco-system, as briefly described in the sub-sections below, and various efforts are still ongoing or have recently been started to provide make Web applications even more powerful. Unfortunately, the underlying Web platform had not been able to keep up with these changes and the security weaknesses will only became more apparent. It is time to tackle this problem and to develop a common understanding of the problem and the desired design goals.

From Documents to Mobile Code During the last 10 years the Web has changed quite fundamentally with the widespread usage of JavaScript. While Web pages have for a long time been dynamically generated the ever increasing capabilities of JavaScript, with respect to functionality and performance, have changed the security model. A typical Website collects content from multiple other Web sites and delivers it to the user's browser and by delivering code inside HTML new security challenges have emerged. Also the standardization landscape had been challenged by this new development and [[I-D.tschofenig-post-standardization](#)] documents architectural implications.

Mashups and Data Sharing With the increasing specialization of Web sites developers started to outsource functionality to other sites. Partially this is a user-convenience aspect (e.g., users do not want to create a new address book with every site, publish their latest status on each and every site again and again) but often also driven by business interestes. In any case, the need to access resources hosted on other sites emerged and often these resources were not visible to everyone. Sharing long-term passwords is considered a bad habit and consequently the Web Authorization (OAuth) protocol [[I-D.ietf-oauth-v2](#)] started to become used widely. OAuth avoids the need to share long-term credentials with random Web sites.

The Real-Time Web As HTTP became the protocol of choice for many application developers, also because of it's ability to go through firewalls and NATs, requirements for asynchronous protocol communication had to be addressed as well. HTTP, as a request/response protocol, was initially not designed for pushing data from the server-side to the client as soon as it is available. Long polling requests and other tricks had been used to allow bi-directional communication between the HTTP client and the HTTP

server. More recently the BiDirectional or Server-Initiated HTTP (hybi) working group was created, which only concerns one aspect of real-time communication. To allow one Web browser to communicate directly with another Web browser the same-origin security framework utilized by the browser has to be bypassed and the work on Real-Time Communication in WEB-browsers (rtcweb) was chartered very recently to develop a architecture. More details can be found in [[I-D.ietf-rtcweb-security](#)] and in [[I-D.ietf-rtcweb-overview](#)]. Extending Web clients with real-time communication capabilities opens the doors for a large number of applications that had previously only been available for downloadable applications.

While astonishing progress has been made in getting the Web application eco-system on a par with native applications the foundation of the Web platform is still unable to address many of the most common security vulnerabilities; problems the Internet community had been fighting against for over a decade and had solved for other application protocol frameworks and Internet deployment environments. This document aims to provide an introduction to the problem that will serve as input to an upcoming Web authentication workshop.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

3. Passwords

Passwords have a long history in authentication protocols on the Internet. They appear to be convient for users and are easy to provision to users by many Web site. Still, passwords present a number of challenges, including:

- o Users re-use the same password at multiple sites. This allows a rouge Website provider to attempt to impersonate users on other sites. It also allows a hacker to use stolen passwords obtained from one site to be used at a non-compromised site.
- o Password are stored in cleartext in most cases. In case of a data breach account information, including the password, becomes accessible to an attacker.
- o Users are tricked in typing their password into a Website maintained by phishing attempts. Furthermore, some Websites request username and password for access to protected resources maintained by other Websites. While there are technical ways to avoid the need for such long-term password sharing practice using OAuth some Websites still ask users.
- o Many password based authentication protocols are not secure against eavesdropping, or allow easy ways for offline dictionary attacks.
- o When end systems are compromised as well then a keyboard logger can capture any password sequence a user enters.

So, why do we need passwords at all? It is easy to come up with solutions that use hardware-based mechanisms (e.g., such as OTP tokens), mobile phones, etc. [[Quest](#)] lists some of these mechanisms and makes an attempt to classify them. There are, however, reasons why alternatives have not found widespread deployment on the Internet, such as

- o Passwords are cheap (at least the primary costs) for user's and service providers. Hardware tokens on the other hand have a certain amount of cost associated with them.
- o Provisioning new users with passwords is easy. Tools and processes exist and are widely accepted.
- o Service providers have no external dependency when they manage user accounts themselves (unlike with many third party identity management solutions).

- o Users are familiar with password-based systems and the acceptance is good.
- o Passwords can easily be delegated to others.
- o Users typically feel quite secure when they are using shared secrets and it fits into their mental model of self-securing.
- o Passwords can easily be transferred to multiple devices used by a single user.

Note that the credential type and the actual form of where these credentials are stored (e.g., software, hardware) is orthogonal to the actual identity proofing process. Stronger forms of identity proofing (e.g., requiring in-person passport verification) can be quite expensive. There are also secondary costs in the form of support calls and education if credential provisioning is more complicated, as it is often the case with client certificates.

Regardless how many disadvantages passwords have they will be with us for a long time. As such, our attempt is therefore to start from the current deployment and to look towards a future where fewer of them are used, and if they are used then in a more secure fashion.

4. Roadmap

It is our aim to accomplish three types of goals:

1. Reduce the number of passwords used on the Web
2. Increase security of how passwords are used
3. Broaden the use of other, non-password-based credentials

A non-goal of this document is to evaluate ways for improving identity proofing, which is a requirement for accomplishing higher levels of assurance.

We do not believe that the technical community should be attempting to come up with the single and best solution to satisfy these three goals. We would like to leave room for innovation and allow many different solutions to co-exist to best suite their deployment.

Subsequently, we try to highlight a few guiding principles in an attempt to come with a way forward.

Move Authentication down into the Platform:

Exposing authentication protocol functionality to the user and requiring Web application developers to write security related code has proven to lead to problems. Avoid user interaction related to security whenever possible but keep in mind that authorization decisions, particularly with regard to data sharing, require a consent. Ensure that library support is available for Web developers to allow them easy integration of security functionality into their applications. Unfortunately a protocol design also needs to consider the transition scenario where the Web endpoints are not yet upgraded to support the new functionality and that the authentication functionality is not yet available.

Design for Growth:

No single authentication mechanism nor credential type is able to fulfill all use cases. Design for later extensions and develop the protocol architecture in such a way that components are interchangeable. In particular, there are a number of authentication mechanisms already in use in other deployment environments.

Context Matters:

Users require context for all disclosures and the sequence of interactions matters. A monolithic authentication protocol that provides mutual authentication is less likely going to capture the context related disclosures. Server-side authentication is the first interaction that will have to be provided to guarantee genuine content as well as the prerequisite of an early setup for a confidentiality protected channel. Client side authentication may, however, come at a much later stage of the application interaction. It is often bundled with an authorization decision where different application execution paths depend on the level of authorization.

Discussion: Is it indeed given that client authentication will have to happen at a later stage given that platform-level authentication proliferates and "authenticated by default" becomes the norm? If so, then strong signals in UIs of authenticated status, identity selection, and anonymous/pseudonymous modes become more important. One could compare this to the evolution in the telephony communication where caller ID information was initially not provided but became the norm later and blocking the caller ID instead became the expectation.

Transform Long-Term Passwords to Short-Term Credentials:

One of the function of authentication protocols is to transform long-term credentials into short term secrets. Long-term credentials, such as passwords, require substantial protection in a protocol exchange and therefore this interaction often leads to a computationally expensive, multi-roundtrip protocol exchange. We do, however, encourage protocol designers to make heavier use of this transformation step into short term credentials. Furthermore, the initial step of entity authentication cannot be seen in isolation of the ultimate purpose of application protocol interaction that requires session management to take place. While this session management today happens in most cases in a non-cryptographic way (i.e., without data origin authentication) we believe it is time revisit this practice.

Keep the User Experience in Mind:

Design your protocol stack in such a way that developers up the stack can give good advice to users. The use case analysis should

include common failure scenarios since error paths need as much expressiveness as success paths, whereby expressiveness refers to the ability to communicate with the user about failure cases.

5. From Two-Party to N-Party

It would be short sighted to write about a topic like this without touching a commonly desired way to reduce the number of long term credentials: federated logins

Federated login allows a user to utilize his credential obtained from one organization, acting as the Identity Provider, for accessing a resource at another entity, who acts as a Relying Party. While this approach addresses some of our design goals it causes secondary problems to appear - particularly related to privacy.

The following issues in this transition from a two-party to a three-party model are to observe:

Introduction:

How do the three parties find each other? In particular, how does the user (via his user agent) inform the relying party about the identity provider it wants to use? How does the relying party inform the user agent (and user) about the identity providers it is able and willing to interact with? How does the relying party find the identity provider for a given user?

Mutual Authentication:

How do we ensure that each party is authenticated to each other?

Authorization and Trust:

What information should the user share with the relying party and how can he be reassured that the information is used in the way he permitted? What information is needed by the Relying Party for the application specific functionality? How is the identity provider able to protect its users against misbehaving relying parties?

Collusion:

How should a user be protected against identity providers and relying parties conspiring?

Security:

How can it be ensured that the interactions between the three parties are not manipulated or attacked?

Note: While this text talks about three parties there may well be more parties involved in the exchange. The role of the identity consists of a credential provider and an attribute provider that may be provided by different parties. Furthermore, attributes associated with personal data may be contributed by multiple attribute providers, not just by a single entity. There may also be additional parties involved in the communication between the identity provider and the relying party the trust path from the identity provider to the relying party.

6. IANA Considerations

This document does not require actions by IANA.

7. Acknowledgments

The content of this document has been created based on discussions with a number of persons, including

- o Jeff Hodges
- o Michael Garcia
- o Adam Barth
- o Brad Hill
- o Dan Mills
- o Ed Felton
- o Tara Whalen
- o Andy Steingruebl
- o Tim Polk
- o Dirk Balfanz
- o Nico Williams
- o Tobias Gondrom
- o Julian Reschke

We would like to thank them for their input. We would also like to thank the participants of the May 2011 W3C Identity in the Browser workshop for their discussion feedback.

8. Open Issues

This document version serves as a starting point for a discussion. As such, there are several things not yet mentioned, such as

- o The introduction section should also point to SPDY as recent development in the are of HTTP evolution.
- o The browser security model should be briefly described. As a comparison, in the browser, we use cross-site communication techniques (redirects, JavaScript) and SSL. In the OS/platform, we use trusted APIs, e.g., signed code, OS-level APIs. In the hardware, we use trusted computing bases (e.g., SIM cards or locked-down platforms).
- o The current document does not discuss how the relying party can trust the information it receives from the identity provider nor how the identity provider makes sure that the relying party adhere to any privacy requirements it has. Various models for accomplishing this trust have been mentioned in the past, including trust frameworks as used by the General Services Administration (GSA) Identity, Credential and Access Management (ICAM), or the work envisioned by Application Bridging for Federated Access Beyond web (abfab) [[I-D.ietf-abfab-arch](#)].
- o The document should also discuss the problems related to the PKI as used by web browsers, the procedures for how trust anchors are provisioned, and the lack of liability in the PKI.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.
- [RFC2617] Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A., and L. Stewart, "HTTP Authentication: Basic and Digest Access Authentication", [RFC 2617](#), June 1999.
- [RFC2109] Kristol, D. and L. Montulli, "HTTP State Management Mechanism", [RFC 2109](#), February 1997.
- [RFC6265] Barth, A., "HTTP State Management Mechanism", [RFC 6265](#), April 2011.
- [RFC2965] Kristol, D. and L. Montulli, "HTTP State Management Mechanism", [RFC 2965](#), October 2000.

9.2. Informative References

- [I-D.ietf-oauth-v2]
Hardt, D. and D. Recordon, "The OAuth 2.0 Authorization Framework", [draft-ietf-oauth-v2-29](#) (work in progress), July 2012.
- [RFC5849] Hammer-Lahav, E., "The OAuth 1.0 Protocol", [RFC 5849](#), April 2010.
- [I-D.ietf-websec-origin]
Barth, A., "The Web Origin Concept", [draft-ietf-websec-origin-06](#) (work in progress), October 2011.
- [I-D.ietf-websec-strict-transport-sec]
Hodges, J., Jackson, C., and A. Barth, "HTTP Strict Transport Security (HSTS)", [draft-ietf-websec-strict-transport-sec-11](#) (work in progress), July 2012.
- [RFC2069] Franks, J., Hallam-Baker, P., Hostetler, J., Leach, P., Luotonen, A., Sink, E., and L. Stewart, "An Extension to

HTTP : Digest Access Authentication", [RFC 2069](#), January 1997.

[I-D.ietf-abfab-arch]

Howlett, J., Hartman, S., Tschafenig, H., Lear, E., and J. Schaad, "Application Bridging for Federated Access Beyond Web (ABFAB) Architecture", [draft-ietf-abfab-arch-03](#) (work in progress), July 2012.

[I-D.ietf-httpbis-p7-auth]

Fielding, R., Lafon, Y., and J. Reschke, "HTTP/1.1, part 7: Authentication", [draft-ietf-httpbis-p7-auth-19](#) (work in progress), March 2012.

[I-D.tschafenig-post-standardization]

Tschafenig, H., Aboba, B., Peterson, J., and D. McPherson, "Trends in Web Applications and the Implications on Standardization", [draft-tschafenig-post-standardization-02](#) (work in progress), May 2012.

[I-D.ietf-rtcweb-overview]

Alvestrand, H., "Overview: Real Time Protocols for Browser-based Applications", [draft-ietf-rtcweb-overview-04](#) (work in progress), June 2012.

[I-D.ietf-rtcweb-security]

Rescorla, E., "Security Considerations for RTC-Web", [draft-ietf-rtcweb-security-03](#) (work in progress), June 2012.

[Quest]

"The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes, In Proc. IEEE Symp. on Security and Privacy 2012 (Oakland 2012)", July 2012.

Authors' Addresses

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

Phone: +358 (50) 4871445
Email: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

Sean Turner
IECA, Inc.
3057 Nutley Street, Suite 106
Fairfax, VA 22031
USA

Phone:
Email: turners@ieca.com

Stephen Farrell
Trinity College Dublin
Dublin, 2
Ireland

Phone: +353-1-896-2354
Email: stephen.farrell@cs.tcd.ie

Mike Hanson
Mozilla

Phone:
Email: mhanson@mozilla.com

