



# Cookie Crumbles: Unveiling Web Session Integrity Vulnerabilities

**Marco Squarcina**

TU Wien

 @blueminal

 <https://infosec.exchange/@minimalblue>

 marco.squarcina@tuwien.ac.at



**Pedro Adão**

IST, Universidade de Lisboa

 @pedromigueladao

 <https://infosec.exchange/@pedroadao>

 pedro.adao@tecnico.ulisboa.pt



Joint work with **Lorenzo Veronese** and **Matteo Maffei**

# Who Are We

- **PhD** @ Ca' Foscari, Venice, IT 
- **Senior Scientist** @ TU Wien, Vienna, AT 
- **Web & Mobile (in)Security**
- **CTF player / organizer** since 2009
- Founder of **mhackeroni**   
(5x **DEF CON CTF** finalist)  
Playing with **WE\_OWN\_YOU** 
- IT security education projects with  
**ENISA** , **CSA**, formerly **Cyberchallenge.IT**
- <https://minimalblue.com/>



**Marco Squarcina**

# Who Are We



Pedro Adão

- **PhD** @ Técnico-Lisboa, PT 🇵🇹
- **Associate Prof.** @ Técnico-Lisboa, PT 🇵🇹
- **Programming Lang & Web (in)Security**
- **CTF player** since 2013
- Founder of **STT** and **CyberSecurity ChallengePT**
- **Coach Team PT 🇵🇹** (ECSC 2019-...)
- **Coach Team Europe (ENISA) 🇪🇺**  
(ICC 2022, 2023)





# Have Weak Integrity

2013



# THE DEPUTIES ARE STILL CONFUSED

RICH LUNDEEN

2013

2015

## Cookies Lack Integrity: Real-World Implications

Xiaofeng Zheng<sup>1,2,3</sup>, Jian Jiang<sup>7</sup>, Jinjin Liang<sup>1,2,3</sup>, Haixin Duan<sup>1,3,4</sup>, Shuo Chen<sup>5</sup>, Tao Wan<sup>6</sup>, and Nicholas Weaver<sup>4,7</sup>

<sup>1</sup>Institute for Network Science and Cyberspace, Tsinghua University

<sup>2</sup>Department of Computer Science and Technology, Tsinghua University

<sup>3</sup>Tsinghua National Laboratory for Information Science and Technology

<sup>4</sup>International Computer Science Institute

<sup>5</sup>Microsoft Research Redmond

<sup>6</sup>Huawei Canada

<sup>7</sup>UC Berkeley

### Abstract

A cookie can contain a “secure” flag, indicating that it should be only sent over an HTTPS connection. Yet there is no corresponding flag to indicate how a cookie was set: attackers who act as a man-in-the-middle even temporarily on an HTTP session can inject cookies which will be attached to subsequent HTTPS connections. Similar attacks can also be launched by a web attacker from a related domain. Although an acknowledged threat, it has not yet been studied thoroughly. This paper aims to fill this gap with an in-depth empirical assessment of cookie injection attacks. We find that cookie-related vulnerabilities are present in important sites (such as Google and Bank of America), and can be made worse by the implementation weaknesses we discovered in major web browsers (such as Chrome, Firefox, and Safari). Our successful attacks have included privacy violation, online victimization, and even financial loss and account

man-in-the-middle (MITM). However, there is no similar measure to protect its integrity from the same adversary: an HTTP response is allowed to set a secure cookie for its domain. An adversary controlling related domain is also capable to disrupt a cookie’s integrity by making use of the shared cookie scope. Even worse, there is an asymmetry between cookie’s read and write operations involving pathing, enabling more subtle form of cookie integrity violation.

The lack of cookie integrity is a known problem, noted in the current specification [2]. However, the real-world implications are under-appreciated. Although the problem has been discussed by several previous researchers [4, 5, 30, 32, 24, 23], none provided in-depth and real-world empirical assessment. Attacks enabled by merely injecting malicious cookies could be elusive, and the consequence could be serious. For example, a cautious user might only visit news websites at open wireless





2013

2015

2019

## Cookies Lack Integrity

Xiaofeng Zheng<sup>1,2,3</sup>, Jian Jiang<sup>7</sup>, Jin Wang<sup>4,5,6</sup>

<sup>1</sup>Institute for Network

<sup>2</sup>Department of Computer

<sup>3</sup>Tsinghua National Lab

<sup>4</sup>Internationa

<sup>5</sup>M



### Abstract

A cookie can contain a “secure” flag, indicating that the cookie should be only sent over an HTTPS connection. However, there is no corresponding flag to indicate how a cookie should be handled when it is present in a cookie jar. In this paper, we show that attackers who act as a man-in-the-middle can exploit this lack of semantics. Specifically, an attacker can temporarily poison a user’s cookie jar by injecting cookies with the “secure” flag. These cookies will be attached to subsequent HTTPS connections. This can enable various attacks, such as session hijacking and man-in-the-middle attacks. Similar attacks can also be launched by a web application that communicates with a related domain. Although an acknowledgement of this problem has been made, its impact has not yet been studied thoroughly. This paper aims to close this gap with an in-depth empirical assessment of cookie injection attacks. We find that cookie-related security issues are present in important sites (such as Google, Facebook, and Bank of America), and can be made worse by browser implementation weaknesses we discovered in Chrome, Firefox, and Internet Explorer. Our findings have led to successful attacks have included privacy violations, account takeovers, cookie victimization, and even financial loss.

# The cookie monster in our browsers



@filedescriptor  
HITCON 2019

2013

2015

2019

2023



## Cookies Lack I

Xiaofeng Zheng<sup>1,2,3</sup>, Jian Jiang<sup>7</sup>, Ji

### 8.6. Weak Integrity

Cookies do not provide integrity guarantees for sibling domains (and their subdomains). For example, consider foo.site.example and bar.site.example. The foo.site.example server can set a cookie with a Domain attribute of "site.example" (possibly overwriting an existing "site.example" cookie set by bar.site.example), and the user agent will include that cookie in HTTP requests to bar.site.example. In the worst case, bar.site.example will be unable to distinguish this cookie from a cookie it set itself. The foo.site.example server might be able to leverage this ability to mount an attack against bar.site.example. [...]

An active network attacker can also inject cookies into the Cookie header field sent to https://site.example/ by impersonating a response from http://site.example/ and injecting a Set-Cookie header field. The HTTPS server at site.example will be unable to distinguish these cookies from cookies that it set itself in an HTTPS response. An active network attacker might be able to leverage this ability to mount an attack against site.example even if site.example uses HTTPS exclusively. [...]

Finally, an attacker might be able to force the user agent to delete cookies by storing a large number of cookies. Once the user agent reaches its storage limit, the user agent will be forced to evict some cookies. Servers SHOULD NOT rely upon user agents retaining cookies.

# cookie monster our browsers



@filedescriptor  
HITCON 2019

rfc6265bis-12

# Recap: Cookie Tossing



https://example.com



https://atk.example.com

Set-Cookie: session=bad; Secure; domain=example.com

Cookie: session=bad

Cookies do not follow the  
**Same Origin Policy**



## Attributes

Expires

Max-Age

Domain

Path

SameSite

## Flags

Secure

HttpOnly

**Path** useful to  
prioritize cookies

**SameSite** does  
not matter here!

# Recap: Cookie Eviction



https://example.com

Cookie: session=good



https://atk.example.com

Name	Value	Domain	Path	E...	S...	11	HttpOnly
session	good	example.com	/	S...	11		✓

# Recap: Cookie Eviction



https://example.com

Cookie: session=good

Name	Value	Domain	Path	E...	S...	HttpOnly
session	good	example.com	/	S...	11	✓



https://atk.example.com

Set-Cookie: x0=\_  
...  
Set-Cookie: x199=\_  
Set-Cookie: session=bad; domain=example.com

```
> for(i=0;i<200;i++) document.cookie='x${i}=_';
< 'x199=_'
> document.cookie = 'session=bad; domain=example.com';
< 'session=bad; domain=example.com'
```

# Recap: Cookie Eviction



https://example.com

Cookie: session=good

Name	Value	Domain	Path	Expiry	SameSite	HttpOnly
session	bad	.example.co...	/	S...	10	



https://atk.example.com

Set-Cookie: x0=\_

...

Set-Cookie: x199=\_

Set-Cookie: session=bad; domain=example.com

Cookie: session=bad

```
> for(i=0;i<200;i++) document.cookie='x${i}=_';
< 'x199=_'
> document.cookie = 'session=bad; domain=example.com';
< 'session=bad; domain=example.com'
```

# Threat Models (Same-site & Network Attacker)

The screenshot shows a web browser window with the title 'Hostile Subdomain Takeover' and the URL 'labs.detectify.com/2014/10/21/hostile-subdomain-takeover-using-herokugithubdesk-more/'. The page content includes a header with the 'detectify' logo, a navigation bar with links for 'App Security /', 'Writeups /', 'How to /', and 'Crowdsolve /'. The main article title is 'Hostile Subdomain Takeover using Heroku/Github/Desktop + more'. Below the title, a date 'October 21, 2014' is shown. The text discusses how hackers can claim subdomains using external services like Heroku, Github, and Desk, and provides a tool for checking DNS entries.

Hostile Subdomain Takeover using Heroku/Github/Desktop + more

October 21, 2014

Hackers can [claim subdomains](#) with the help of external services. This attack is practically non-traceable, and affects at least 17 large service providers and multiple domains are affected. Find out if you are one of them by [using our quick tool](#), or go through your DNS-entries and remove all which are active and unused OR pointing to External Services which you do not use anymore.

## Dangling DNS Records

Discontinued  
Services

# Threat Models (Same-site & Network Attacker)

## Can I Take Your Subdomain? Exploring Same-Site Attacks in the Modern Web

Marco Squarcina<sup>1</sup> Mauro Tempesta<sup>1</sup> Lorenzo Veronese<sup>1</sup> Stefano Calzavara<sup>2</sup> Matteo Maffei<sup>1</sup>  
<sup>1</sup> TU Wien   <sup>2</sup> Università Ca' Foscari Venezia & OWASP



### Abstract

Related-domain attackers control a sibling domain of their target web application, e.g., as the result of a subdomain takeover. Despite their additional power over traditional web attackers, related-domain attackers received only limited attention from the research community. In this paper we define and quantify for the first time the threats that related-domain attackers pose to web application security. In particular, we first clarify the capabilities that related-domain attackers can acquire through different attack vectors, showing that different instances of the related-domain attacker concept are worth attention. We then study how these capabilities can be abused to compromise web application security by focusing on different angles, including cookies, CSP, CORS, postMessage, and domain relaxation. By building on this framework, we report on a large-scale security measurement on the top 50k domains from the Tranco list that led to the discovery of vulnerabilities in 887 sites, where we quantified the threats posed by related-domain attackers to popular web applications.

2021

1520 vulnerable subdomains

attacker is traditionally defined as a web site controlled by the attacker. In this work, we introduce a twist, i.e., its malicious website is hosted under a subdomain of the target web application. For instance, consider the target web application `www.example.com`. If an attacker controls the subdomain `sub.example.com`, then that a related-domain attacker controls the target application. The privileged position of a related-domain attacker endows it, for instance, with the ability to compromise cookie confidentiality and integrity, because cookies can be shared between domains with a common ancestor, reflecting the assumption underlying the original Web design that related domains are under the control of the same entity. Since client authentication on the Web is mostly implemented on top of cookies, this represents a major security threat.

cnn.com, nih.gov, cisco.com,  
f-secure.com, harvard.edu,  
lenovo.com, ...

## Dangling DNS Records

Discontinued Services

Corporate Networks

Expired Domains

Roaming Services

Deprovisioned Cloud Instances

Dynamic DNS Providers

# Threat Models (Same-site & Network Attacker)

The slide features two side-by-side screenshots of web pages. On the left is a paper titled "Can I Take Your Subdomain? Exploring Same-Site" by Marco Squarcina, Mauro Tempesta, Lorenzo Veronese, Stefanos Kotsopoulos, and others. It's published by USENIX. The abstract discusses related-domain attackers who control a sibling domain of their target web application. On the right is "The 2022 Web Almanac" by HTTP Archive, showing a large yellow "2022" and a cartoon illustration of a person climbing a ladder. A pink callout box at the bottom right states: "90% of websites deploy partial HSTS (no **IncludeSubdomain**)". To the right of the slide, there are vertical tabs for "Corporate", "Networking", "Roaming", and "Dynamically".

Can I Take Your Subdomain? Exploring Same-Site

Marco Squarcina<sup>1</sup> Mauro Tempesta<sup>1</sup> Lorenzo Veronese<sup>1</sup> Stefanos Kotsopoulos<sup>1</sup> Giorgio Guglielmi<sup>2</sup> Fabio Volo<sup>2</sup> Luca Paoletti<sup>2</sup> <sup>1</sup>TU Wien <sup>2</sup>Università Ca' Foscari Venezia

THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

### Abstract

Related-domain attackers control a sibling domain of their target web application, e.g., as the result of a subdomain takeover. Despite their additional power over traditional web attackers, related-domain attackers received only limited attention from the research community. In this paper we define and quantify for the first time the threats that related-domain attackers pose to web application security. In particular, we first clarify the capabilities that related-domain attackers can acquire through different attack vectors, showing that different instances of the related-domain attacker concept are worth attention. We then study how these capabilities can be abused to compromise web application security by focusing on different angles, including cookies, CSP, CORS, postMessage, and domain relaxation. By building on this framework, we report on a large-scale security measurement on the top 50k domains from the Tranco list that led to the discovery of vulnerabilities in 887 sites, where we quantified the threats posed by related-domain attackers to popular web applications.

The 2022 Web Almanac

By HTTP Archive

# Web Almanac

## 2022

HTTP Archive's annual **state of the web** report

Our mission is to combine the raw stats and trends of the HTTP Archive with the expertise of the web community. The Web Almanac is a comprehensive report on the state of the web, backed by real data and trusted web experts. The 2022 edition is comprised of 23 chapters spanning aspects of page content, user experience, publishing, and distribution.

cnn.com  
f-secure.com  
lenovo.com, ...

Start exploring

90% of websites deploy partial HSTS (no **IncludeSubdomain**)

Corporate

Networking

Roaming

Dynamically

# Cross-Origin Request Forgery (CORS)



https://bank.com

POST /action

Cookie:s=x;csrf=y  
- csrf-tok=y

Done!



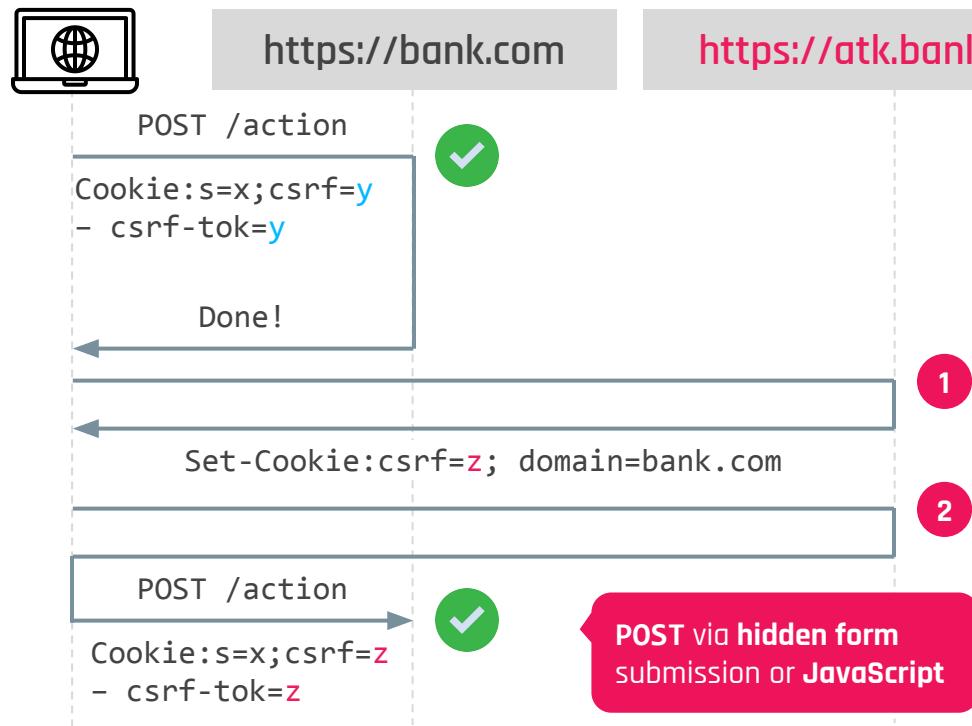
https://atk.bank.com



## Double-Submit Pattern (DSP)

```
if cookie(csrf)==POST(csrf-tok):  
    return True  
return False
```

# Cross-Origin Request Forgery (CORS)



## Double-Submit Pattern (DSP)

```
if cookie(csrf)==POST(csrf-tok):  
    return True  
return False
```

**Wrong assumption:** attacker can only manipulate the token, but not the cookie!

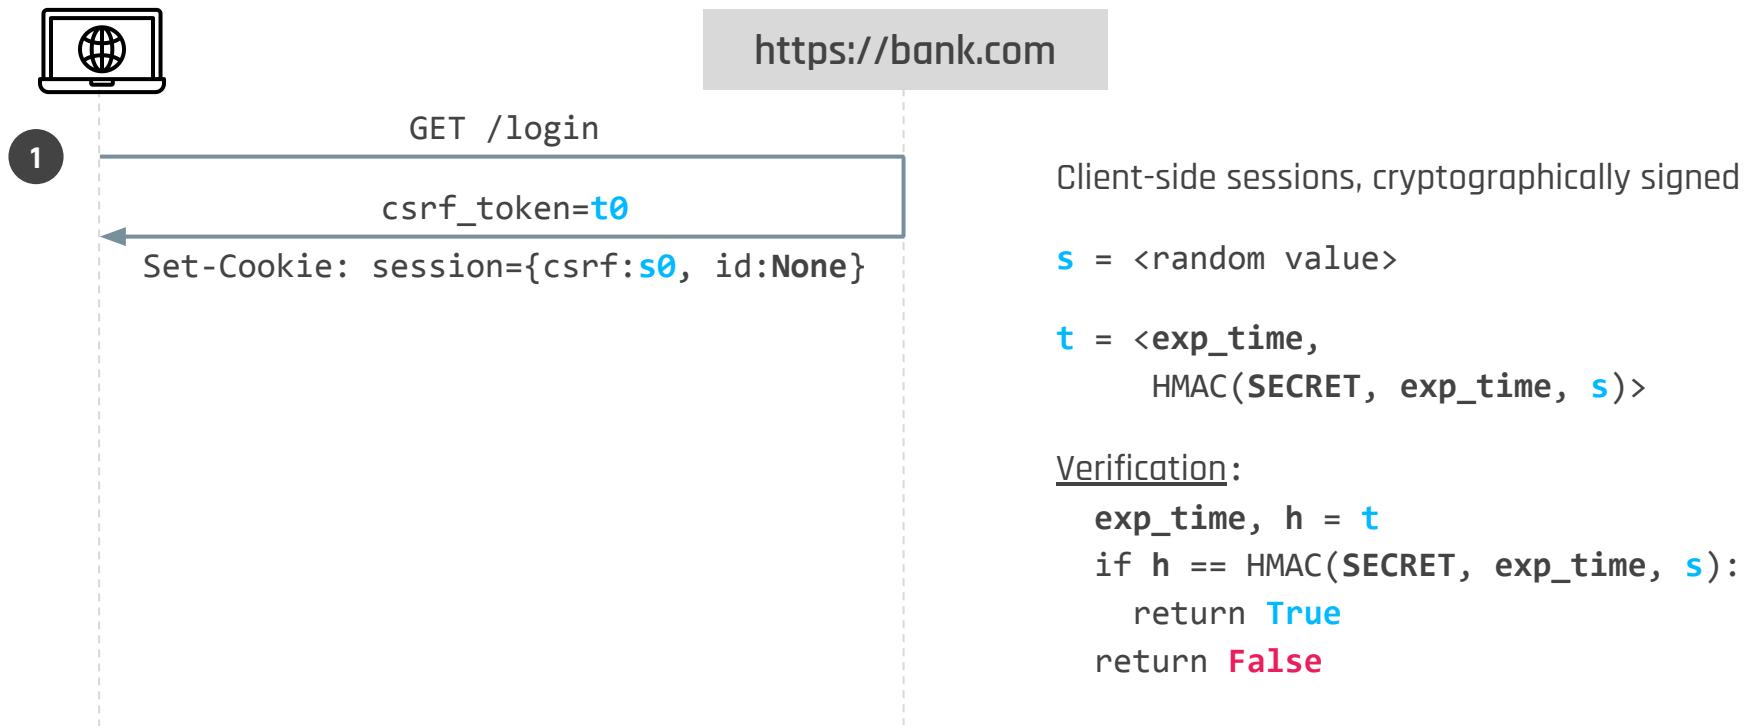
Trivially **vulnerable** against same-site attackers, just **toss** and **submit**!

# Synchronizer Token Pattern (STP)

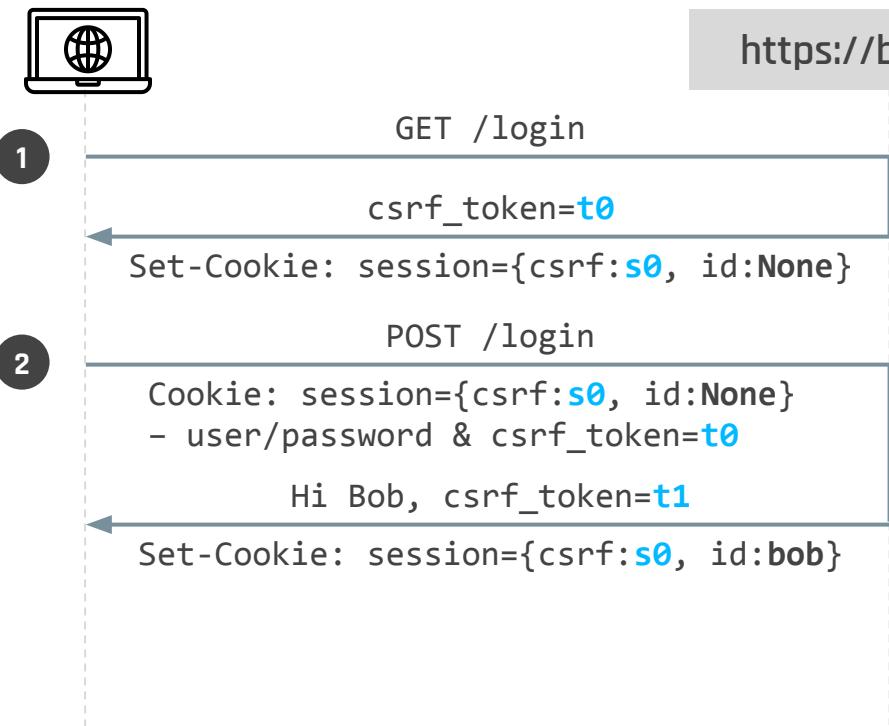
Improves the vulnerable **Double-Submit Pattern (DSP)**

- Session  := <**id**, **CSRF\_secret**> **CSRF secret stored in the session**  
**(can be client or server-side)**
- **CSRF\_token** = generate\_token(**CSRF\_secret**, params...)
- **Verification** (server-side)  
**CSRF\_token == generate\_token(CSRF\_secret, params...)**
- Overwrite the session cookie? Deauth the user, **NO CORF**, attacker sad :/

# Synchronizer Token Pattern (Flask-login + Flask-WTF)



# Synchronizer Token Pattern (Flask-login + Flask-WTF)



Client-side sessions, cryptographically signed

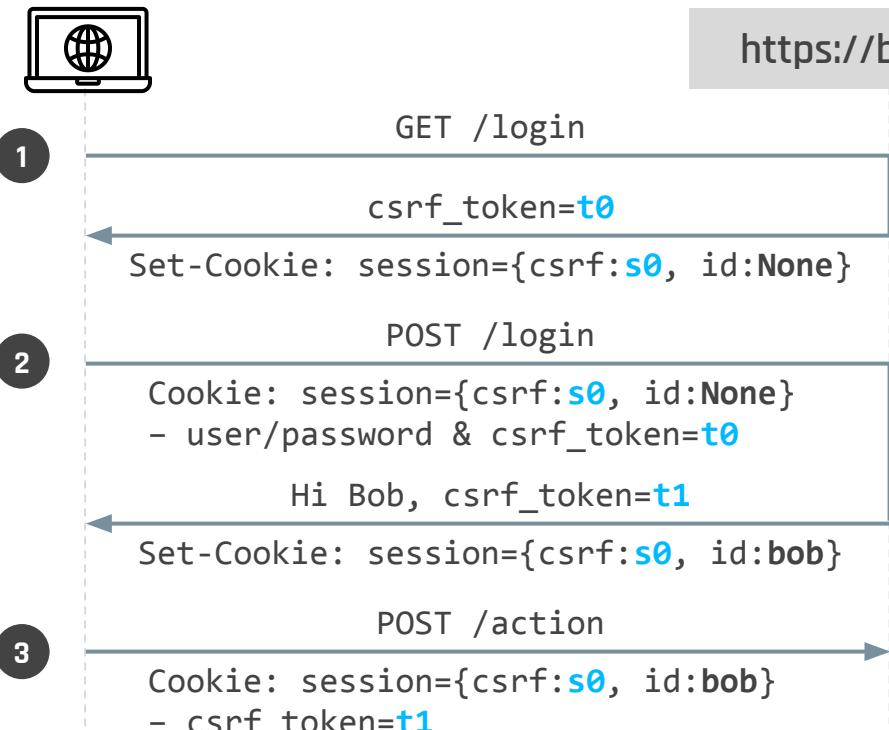
s = <random value>

t = <exp\_time,  
HMAC(SECRET, exp\_time, s)>

Verification:

```
exp_time, h = t
if h == HMAC(SECRET, exp_time, s):
    return True
return False
```

# Synchronizer Token Pattern (Flask-login + Flask-WTF)



Client-side sessions, cryptographically signed

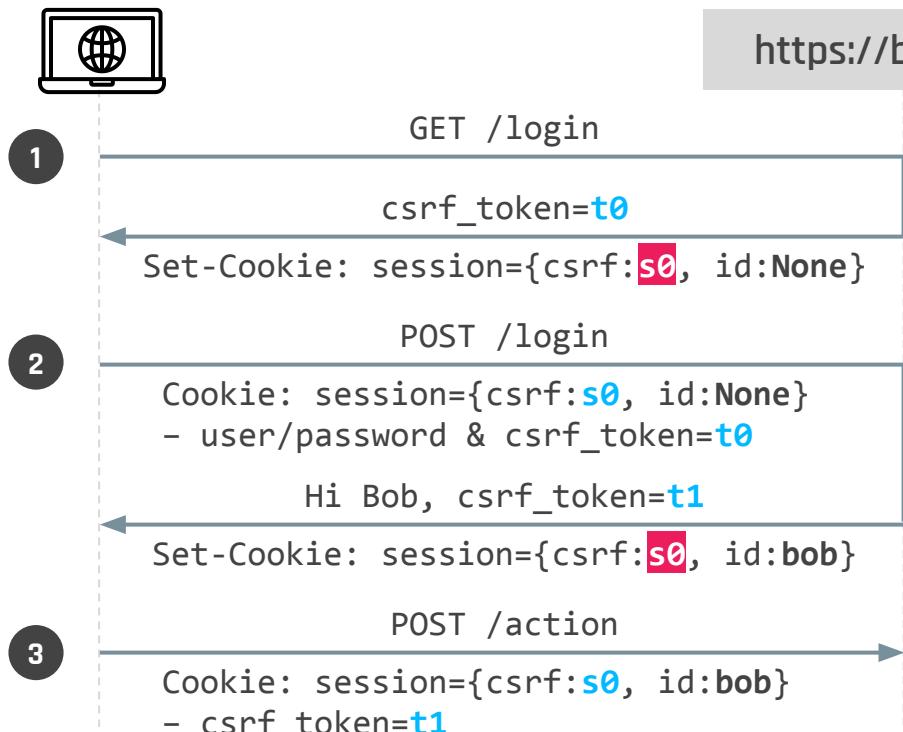
`s` = <random value>

`t` = <exp\_time,  
HMAC(SECRET, exp\_time, s)>

Verification:

```
exp_time, h = t
if h == HMAC(SECRET, exp_time, s):
    return True
return False
```

# Synchronizer Token Pattern (Flask-login + Flask-WTF)



Client-side sessions, cryptographically signed

s = <random value>

t = <exp\_time,  
HMAC(SECRET, exp\_time, s)>

Verification:

```
exp_time, h = t
if h == HMAC(SECRET, exp_time, s):
    return True
return False
```

# CORF Token Fixation (Flask-login + Flask-WTF)



https://bank.com



https://atk.bank.com

1

GET /login

csrf\_token=t0

Set-Cookie: session={csrf:s, id:None}

2

Set-Cookie: session={csrf=s, id:None}; domain=bank.com

3

GET /login

Cookie: session={csrf:s, id:None}

csrf\_token=t1

Set-Cookie: session={csrf:s, id:None}

Equivalent to an  
**unauthenticated session fixation**

# CORF Token Fixation (Flask-login + Flask-WTF)



https://bank.com



https://atk.bank.com

4

POST /login

Cookie: session={csrf:**s**, id:None}  
- user/password & csrf\_token=**t1**

Welcome Bob!

Set-Cookie: session={csrf:**s**, id:**bob**}



Bob authenticates

# CORF Token Fixation (Flask-login + Flask-WTF)



https://bank.com

4

POST /login

Cookie: session={csrf:**s**, id:None}  
- user/password & csrf\_token=**t1**

Welcome Bob!



Bob authenticates

Set-Cookie: session={csrf:**s**, id:**bob**}

5

POST /action

Cookie: session={csrf:**s**, id:**bob**}  
- csrf\_token=**t0**



The **CSRF secret s** is not  
refreshed during login!  
The **CSRF token t0** known by the  
attacker is valid for Bob's session!

# CORF Token Fixation

- Bypasses faulty implementations of the **Synchronizer Token Pattern**
- Caused by the **CSRF secret** in the session **not being renewed** upon login
- The attacker does not need to know the CSRF secret, but only an **unauthenticated session id** and a **valid CSRF token** for that session
- Works against **server-side** and **client-side** session handling implementations
- User already logged-in? No problem, **force a deauth** and toss the attacker's pre-session, either via eviction or request to /logout endpoint

# CORF Token Fixation (CodeIgniter4)



https://bank.com



1

GET /login

csrf\_token=t0

Set-Cookie: session=sess0

`_ci_last_regenerate|i:1690849755;  
csrf_test_name|s:32:"47be9758fe558  
98f1958bd201764a0be";`

CSRF secret s0

# CORF Token Fixation (CodeIgniter4)



https://bank.com



1

GET /login

csrf\_token=t0

Set-Cookie: session=sess0

2

POST /login

Cookie: session=sess0

- user/password & csrf\_token=t0

Welcome Bob!

Set-Cookie: session=sess1



\_\_ci\_last\_regenerate|i:1690849755;  
csrf\_test\_name|s:32:"1f5b0c83a29e9  
f9725d219e53a6d2be1";user|a:1:{s:2  
:"id";s:1:"1";}CSRF secret s1

# CORF Token Fixation (CodeIgniter4)



https://bank.com



1

GET /login

csrf\_token=t0

Set-Cookie: session=sess0

2

POST /login

Cookie: session=sess0

- user/password & csrf\_token=t0

Welcome Bob!

Set-Cookie: session=sess1

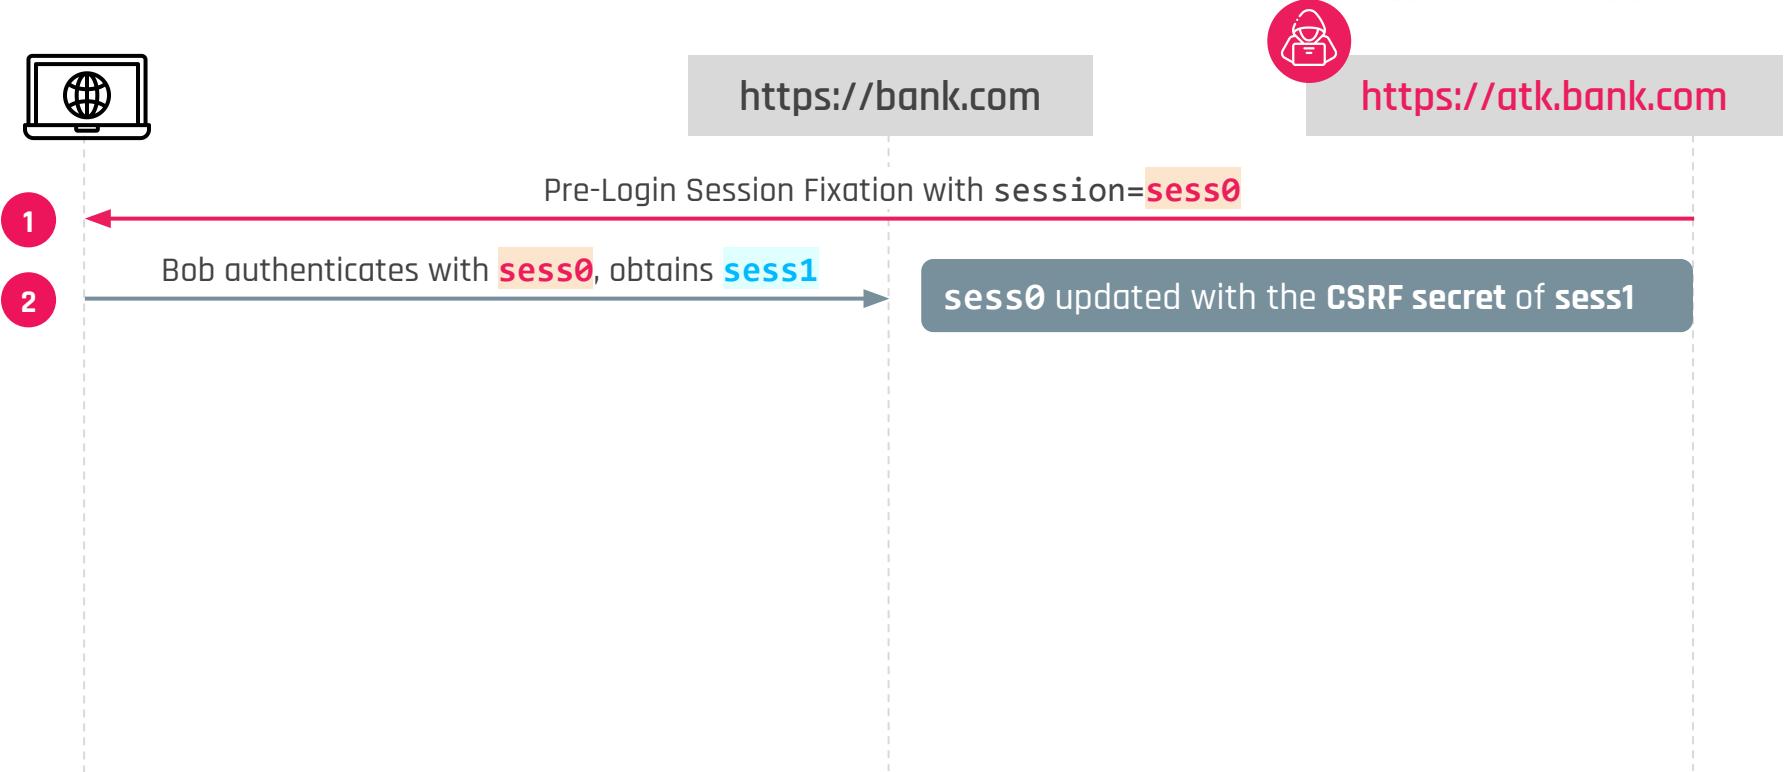
\_\_ci\_last\_regenerate|i:1690849755;  
csrf\_test\_name|s:32:"1f5b0c83a29e9  
f9725d219e53a6d2be1";

CSRF secret s1

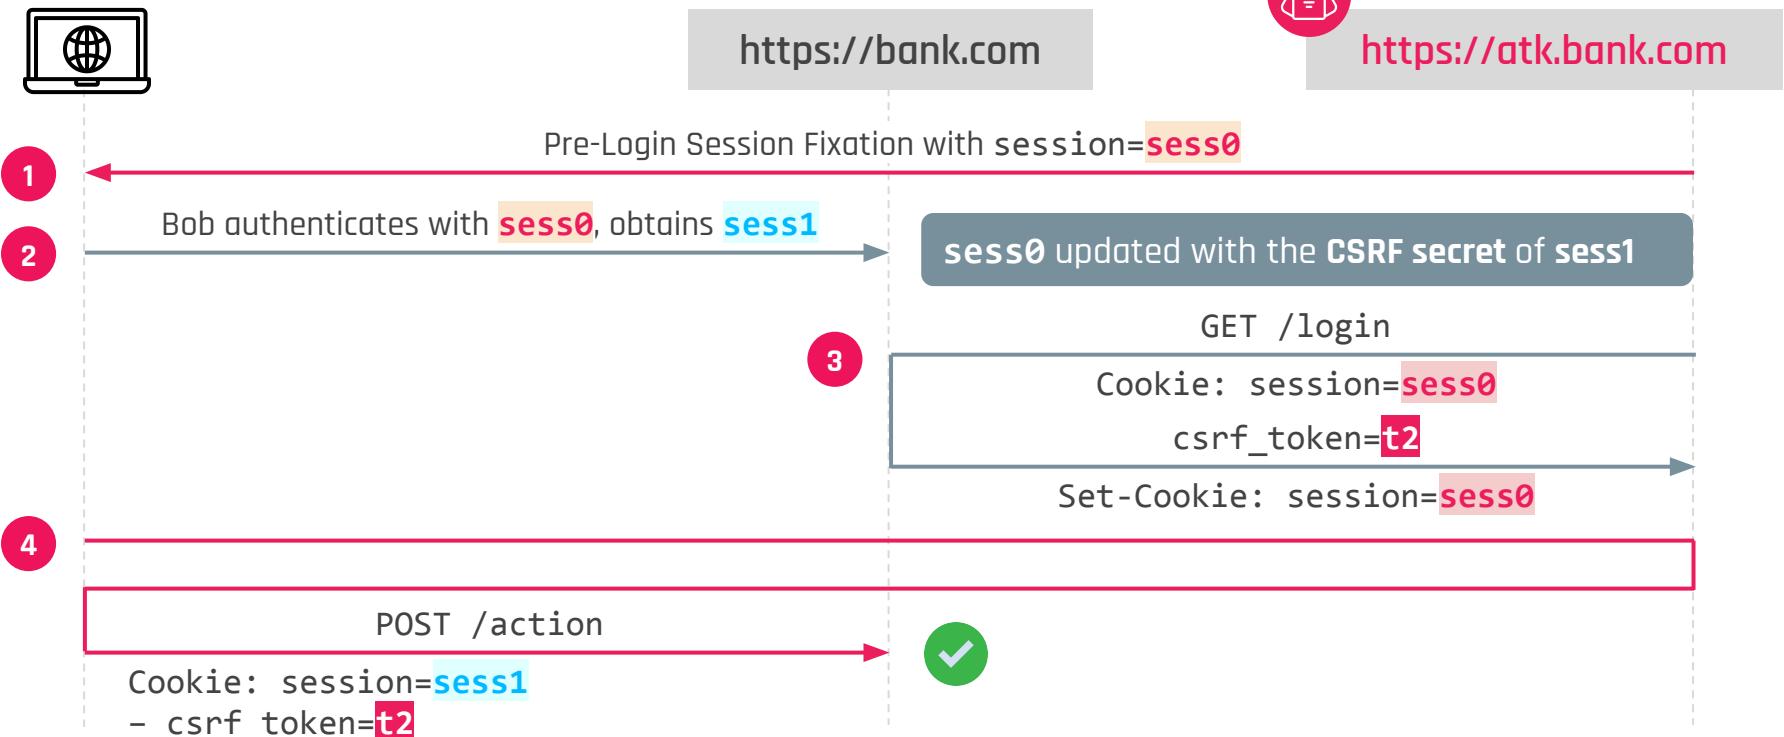
\_\_ci\_last\_regenerate|i:1690849755;  
csrf\_test\_name|s:32:"1f5b0c83a29e9  
f9725d219e53a6d2be1";user|a:1:{s:2  
:"id";s:1:"1";}:

CSRF secret s1

# CORF Token Fixation (CodeIgniter4)

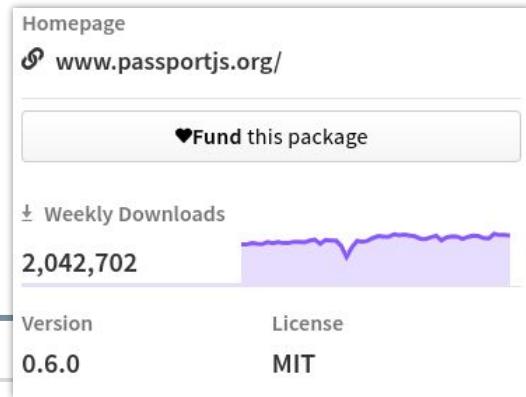


# CORF Token Fixation (CodeIgniter4)



# Web Frameworks Analysis

Framework (9/13 vulnerable)	Broken STP	Default DSP	Session Fixation	CVE-2022-25896
<b>Express</b> (passport + csurf)	●		●	CVE-2022-25896
<b>Koa</b> (koa-passport + csrf)	●			
<b>Fastify</b> (fastify/passport + csrf-protection)	●	●	●	CVE-2023-29020 CVE-2023-27495 CVE-2023-29019
<b>Sails*</b> (csurf)	●		●	
<b>Flask</b> (flask-login+flask-wtf)	●			
<b>Tornado</b>		●		
<b>Symfony</b> (security-bundle)	●			CVE-2022-24895
<b>CodeIgniter4</b> (shield)	●	●		CVE-2022-35943
<b>Yii2</b>		●		

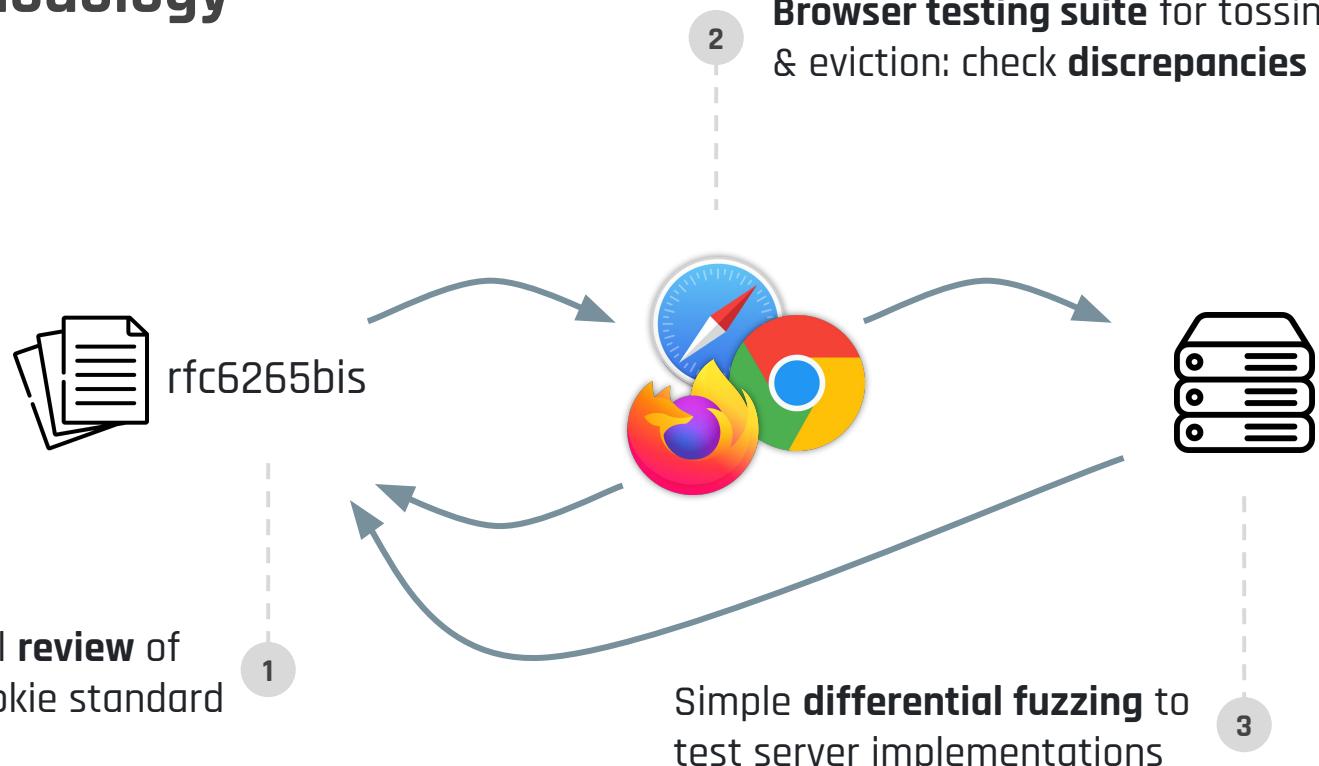


\*affects the bootstrap template app



# Are Getting Better?

# Methodology



# Strict Secure & Prefixes (Host-)

HTTP Working Group  
Internet-Draft  
Updates: [6265](#) (if approved)  
Intended status: Standards Track  
Expires: March 9, 2017

M. West  
Google, Inc  
September 5, 2016

**Deprecate modification of 'secure' cookies from non-secure origins**  
**`draft-ietf-httpbis-cookie-alone-01`**

block setting cookie without  
the **Secure** flag if the cookie  
jar contains Secure cookie  
with the same name

HTTP Working Group  
Internet-Draft  
Updates: [6265](#) (if approved)  
Intended status: Standards Track  
Expires: August 26, 2016

M. West  
Google, Inc  
February 23, 2016

**Cookie Prefixes**  
**`draft-ietf-httpbis-cookie-prefixes-00`**

# Strict Secure & Prefixes (Host-)

HTTP Working Group  
Internet-Draft  
Updates: [6265](#) (if approved)  
Intended status: Standards Track  
Expires: March 9, 2017

M. West  
Google, Inc  
September 5, 2016

**Deprecate modification of 'secure' cookies from non-secure origins**  
**`draft-ietf-httpbis-cookie-alone-01`**

block setting cookie without  
the **Secure** flag if the cookie  
jar contains Secure cookie  
with the same name

HTTP Working Group  
Internet-Draft  
Updates: [6265](#) (if approved)  
Intended status: Standards  
Expires: August 26, 2016

```
> document.cookie = '__Host-sess=bar; Path=/; Secure; Domain=example.com'  
< '__Host-sess=bar; Path=/; Secure; Domain=example.com'  
> document.cookie  
< ''
```

**Cookie Prefixes**  
**`draft-ietf-httpbis-cookie-prefixes-00`**

High-integrity cookies,  
cannot be set from a  
sibling domain!



Set-Cookie:

foo=

=foo

=foo=

--foo

foo

Valid? Or Invalid?



If valid, think of the corresponding Cookie: header...



Set-Cookie:

foo=

=foo

=foo=

--foo

foo

Valid? Or Invalid?



If valid, think of the corresponding **Cookie:** header...



Set-Cookie:

foo=

=foo

=foo=

==foo

foo

Valid? Or Invalid?



If valid, think of the corresponding **Cookie:** header...



Set-Cookie:

foo=

=foo

=foo=

--foo

foo

Valid? Or Invalid?



If valid, think of the corresponding Cookie: header...



Set-Cookie:

foo=

=foo

=foo=

==foo

foo

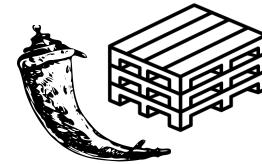
Valid? Or Invalid?



If valid, think of the corresponding Cookie: header...



# Trivia

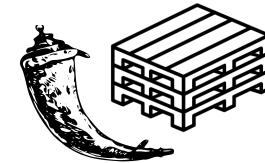


Werkzeug <2.2.3

Set-Cookie:	Cookie:	Key	Value	Server <key, value>
foo=	foo=	foo		<foo, >
=foo				
=foo=				
==foo				
foo				



# Trivia

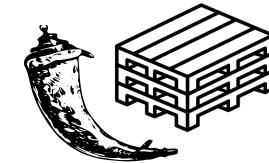


Werkzeug <2.2.3

Set-Cookie:	Cookie:	Key	Value	Server <key, value>
foo=	foo=	foo		<foo, >
=foo	foo		foo	
=foo=	foo=		foo=	
==foo	=foo		=foo	
foo	foo		foo	



# Trivia



Werkzeug <2.2.3

Set-Cookie:	Cookie:	Key	Value	Server <key, value>
foo=	foo=	foo		<foo, >
=foo	foo		foo	<foo, >
=foo=	foo=		foo=	<foo, >
==foo	=foo		=foo	<foo, >
foo	foo		foo	<foo, >



# Trivia



## [RFC6265bis] Accept nameless cookies. (#1018)

[Browse files](#)

Set

foo

=fo

=fo

==f

foo

This patch alters the cookie parsing algorithm to treat  
`Set-Cookie: token` as creating a cookie with an empty name and a value  
of "token". It also rejects cookies with neither names nor values (e.g.  
`Set-Cookie: ` and `Set-Cookie: =`.

Closes #159.

↳ main (#1018)

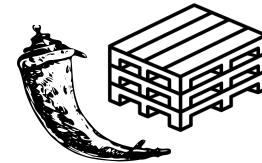
↳ draft-ietf-httpbis-unprompted-auth-02 ... b68e4ff

committed on Jan 10, 2020

1 parent c43cdæ commit 0178223

100

100



Werkzeug <2.2.3

Server <key, value>

<foo, >

<foo, >

<foo, >

<foo, >

<foo, >

# Bypassing \_\_Host-



<http://atk.bank.com>



<https://bank.com>

Set-Cookie: \_\_Host-sess=good;  
Secure; Path=/

# Bypassing \_\_Host-



http://atk.bank.com



https://bank.com

Set-Cookie: \_\_Host-sess=good;  
Secure; Path=/

Set-Cookie: \_\_Host-sess=bad; Path=/app;  
domain=bank.com



# Bypassing \_\_Host-



http://atk.bank.com



https://bank.com

Set-Cookie: \_\_Host-sess=good;  
Secure; Path=/

Set-Cookie: \_\_Host-sess=bad; Path=/app;  
domain=bank.com

Cookie: \_\_Host-sess=bad;  
\_\_Host-sess=good;

# Bypassing \_\_Host-



http://atk.bank.com



https://bank.com

Set-Cookie: \_\_Host-sess=good;  
Secure; Path=/

Set-Cookie: \_\_Host-sess=bad; Path=/app;  
domain=bank.com

Cookie: \_\_Host-sess=bad;  
\_\_Host-sess=good;

CVE-2022-2860\*

CVE-2022-40958\*

**Fixed in browsers and rfc6265bis** by  
blocking nameless cookies with value  
starting for \_\_Host- or \_\_Secure-

\* Reported almost simultaneously with **Axel Chong**, our  
issues were merged to jointly discuss mitigations and  
additional security implications. See also

<https://github.com/httpwg/http-extensions/issues/2229>



# Bypassing \_\_Host- (after the fix)



## Amazon API Gateway

CVE-2022-2860\*

CVE-2022-40958\*

- **Serialization collisions** could still be used to bypass \_\_Host- against chains of 
- Fixed in **AWS Lambda proxy integration for HTTP APIs** after our report

**Fixed in browsers and rfc6265bis** by blocking nameless cookies with value starting for \_\_Host- or \_\_Secure-

\* Reported almost simultaneously with **Axel Chong**, our issues were merged to jointly discuss mitigations and additional security implications. See also

<https://github.com/httpwg/http-extensions/issues/2229>

# Bypassing Strict Secure



http://atk.bank.com



https://bank.com

Set-Cookie: session=good; Secure

Set-Cookie: =session=bad; Path=/app;  
domain=bank.com

Cookie: session=bad; session=good;

Set-Cookie: =session=bad

Name	Value	Domain	Path	E...	S...	H...	Secure
session	good	bank.com	/	S...	11		
	session=bad	.bank.com	/app	S...	11		

# Bypassing \_\_Host- (with the help of the **server**)

- Popular programming languages / Web frameworks **diverge from the spec**
- Client / server inconsistencies. Security implications?



PHP <8.1.11

CVE-2022-31629

Cookie: \_\_Host-sess=bad  
Cookie: \_Host-sess=bad  
Cookie: ..Host-sess=bad

Parsed as the  
**same cookie**

**register\_globals** heritage:  
' ' . [ are replaced by \_ in the  
\$\_COOKIE superglobal array



# Desynchronization Issues

- 1 https://bank.com set a secure   
Set-Cookie: sess=good; Secure

- 2 http://bank.com sets a non-secure  via JS  
document.cookie = 'sess=bad'

EXPECTATION

sess=bad is not set (Strict Secure )

REALITY



# Desynchronization Issues

CVE-2023-29547

- 1 https://bank.com set a secure Set-Cookie: sess=good; Secure

- 2 http://bank.com sets a non-secure via JS  
document.cookie = 'sess=bad'

Fixed in Firefox 112

Caused by restrictions imposed by the FF implementation of **Site Isolation (Project Fission)**

EXPECTATION

sess=bad is not set (Strict Secure )

REALITY

Cookie not set, but document.cookie at http://bank.com returns sess=bad



# Desynchronization Issues

- 1 <https://atk.bank.com>

Fixed in Firefox 115

```
» for(let i=0; i<400; i++) document.cookie = `a${i}=_; domain=bank.com`
⚠ Some cookies are misusing the recommended "SameSite" attribute 400
← "a399=_; domain=bank.com"
» document.cookie.split('; ').length
← 400
» window.open("https://bank.com")
```

- 2 Delete via Set-Cookie (exp. date), Clear-Site-Data header, or manually
- 3 The first 240 are still in Document.cookie in the original and opened window (survives reloads and schemeful navigations)



# Desynchronization Issues

- 1 <https://atk.bank.com>

Fixed in Firefox 115

```
» for(let i=0; i<400; i++) document.cookie = `a${i}=_; domain=bank.com`  
⚠ Some cookies are misusing the recommended "SameSite" attribute 400  
← "a399=_; domain=bank.com"  
» document.cookie.split('; ').length  
← 400  
» window.open("https://bank.com")
```

Could introduce vulnerabilities in frontends trusting `document.cookie` to set **custom HTTP headers** like **ASP.NET** and **Angular**

- 2 Delete via Set-Cookie (exp. date), Clear-Site-Data header, or manually
- 3 The first 240 are still in Document.cookie in the original and opened window (survives reloads and schemeful navigations)



... and that's the way the cookie crumbles!

# Takeaways

- Battle-tested Web frameworks and libraries had **concerning session integrity vulnerabilities**.  
Causes & consequences?
- **Legacy design** is still cursing standards and modern applications: can we **move on without breaking the Web?**
- Developers are falling behind in **keeping track of changes to Web standards**.

Composition issues or lack of understanding of the threat models? Apps in the wild?

Backward compatibility issues?  
How to make deployment easier without trading on security?

Lack of cohesiveness between browser vendors, developers, and authors of Web standards?

# Paper & Artifacts

<https://github.com/SecPriv/cookiecrumbles>

- + **Details on the Web framework vulnerabilities (including PoCs)**
- + **Outcome of the responsible disclosure**
- + **Cookie measurement (nameless and \_\_Host-), dataset & code**
- + **Formal modeling of (patched) Web frameworks**
- + **Browser test suite & server-side testing code**

# Thank You! Questions?



**Marco Squarcina** (TU Wien)

@blueminal

<https://infosec.exchange/@minimalblue>  
 marco.squarcina@tuwien.ac.at

<https://github.com/SecPriv/cookiecrumbles>

**Pedro Adão** (IST, Universidade de Lisboa)

@pedromigueladao

<https://infosec.exchange/@pedroadao>  
 pedro.adao@tecnico.ulisboa.pt