

The Remote on the Local: Exacerbating Web Attacks Via Service Workers Caches

Dolière Francis Somé
CISPA Helmholtz Center for Information Security
doliere.some@cispa.saarland

Stefano Calzavara
Università Ca' Foscari Venezia
calzavara@dais.unive.it

Marco Squarcina
TU Wien
marco.squarcina@tuwien.ac.at

Matteo Maffei
TU Wien
matteo.maffei@tuwien.ac.at

Abstract—Progressive Web Applications (PWAs) are the new trend in web development, promising several features and similar advantages as native applications. They heavily rely on modern web APIs to offer an engaging user experience. Service Workers are one of the core technologies employed by PWAs. They work as a proxy server for websites, allowing requests and responses to be modified, cached and served to the browser even when the user is offline. In this work we showcase a number of flaws in the Cache API that allow an attacker to void the security policies put in place by web developers, posing serious security and privacy threats. Given that these attacks are enabled by the presence of Service Workers, we demonstrate the impact of our findings by performing a large-scale analysis on the top 110K websites. Finally, we propose a redesign of the Cache API that prevents all the attacks discussed in the paper.

Note to readers: this is a work-in-progress paper for the SecWeb Workshop, co-located with IEEE EuroS&P 2020.

1. Introduction

Progressive Web Applications (PWAs) are the current trend in the tremendous evolution of web applications, from the earlier days of static HTML documents, to rich web applications making extensive use of AJAX (Asynchronous JavaScript and XML) and other advanced APIs in order to provide highly responsive web applications. At the core of PWAs, are service workers. They play the role of an in-browser web application proxy able to intercept HTTP requests and immediately serve HTTP responses present on a cache, or even generate responses on the fly, fetch them from a remote server, cache a copy and return the response to the web application.

Caching HTTP responses can enhance the availability and responsiveness of the application even when the user is offline, since HTTP responses are available right away, directly in the service workers cache. Coupled with the Push API, service workers can also be used to display notifications to users, even when they navigate away of the application, in order to re-engage them back to the application. Last but not least, Chrome recently proposed to users the possibility to install PWAs directly in their devices, providing a more seamless integration with the

Desktop environment. PWAs installed on a user device run in dedicated browser windows, without all the browser menus and address bar, giving them the look of native applications.

So far, service workers gained little consideration from a research perspective. But notably, Lee et al. [15] have demonstrated that the Push API could be abused for phishing purposes, and pointed that privacy attacks could be performed by attackers to learn the set of PWAs a user has visited. They also showed, as well as Papadopoulos et al. [16], that attackers could abuse the persistence of service workers to perform unwanted computations and harmful operations such as mining cryptocurrencies in user's browsers.

In this work, we focus on the Cache API and its current design. Indeed, the cache where service workers store HTTP requests and responses is bound to the entire origin of the application managed by the service worker. In this setting, a malicious script executing in same-origin browsing contexts such as web pages, can tamper with the cache, and the HTTP responses stored there. In the recent years, the Web has witnessed the introduction of a profusion of security policies, such as `Content-Security-Policy`, `Feature-Policy`, `Referrer-Policy`, etc. aimed at mitigating the impact of widespread web attacks such as Cross-Site Scripting, Clickjacking, etc. Worryingly, when a content is cached, so are its security policies meant to protect it. An attacker can simply remove those security policies from the cached content, in order to perform malicious actions that were supposed to be mitigated by these policies. In addition, attackers can also tamper with same-origin and Cross-Origin Resource Sharing (CORS) [4]-compliant content in the cache to bypass defensive programming practices, which rely for instance on the order in which scripts are loaded in web pages. For instance, web applications can use these defensive idioms to ensure that constants and frozen objects [1] cannot be tampered with by potentially malicious code that is executed afterwards in a web page. Unfortunately, by tampering with content in the cache, attackers are able to void those protections.

By manipulating the cache, the attacker lets service workers serve content that they control, in settings where they can perform actions of their choosing without the constraints of security policies or defensive practices. It

is important to stress here that these attacks represent a new class of threats, or more precisely, they exacerbate the power of the standard web attacker by giving them more power to perform malicious actions that were not possible without service workers.

To the best of our knowledge, this work is the first to showcase serious threats around the design of the service worker Cache API and the interplay with potentially malicious scripts running in web pages and other same-origin contexts. Overall, we make the following contributions:

- We demonstrate that the current design of service workers Cache API poses serious security and privacy threats. In particular, we highlight how a web attacker can compromise the cache in order to mount or perform attacks that were not previously possible without service workers.
- We performed a large-scale and depth-scale empirical study on the top 110,000 sites in order to assess the deployment of service workers in the wild. We found real world and highly ranked web applications that are caching security policies that can be bypassed in case of an attack. By doing so, we illustrate that our attack scenarios are not only hypothetical but real.
- We propose a countermeasure to prevent the attacks discussed in this work. In particular we argue that service workers' cache should be made inaccessible to scripts, potentially malicious, running in other same-origin contexts such as web pages. This redesign opens huge opportunities in deploying secure service workers in PWAs.

2. Background

In this section, we provide an overview of progressive web applications (PWAs) with a focus on Service Workers [10], their core component.

2.1. Service Workers

A service worker is an event-driven and browser-managed process triggered by the registration of a JavaScript code hosted by a web application, and registered to manage all or part of an application. It plays the role of a proxy for an application, directly within the browser, for intercepting and managing outgoing HTTP requests issued by the web application to load resources, and the incoming HTTP responses returned by web servers. Coupled with the Cache API, it can be used to store HTTP responses, and then serve those responses to web applications even offline. The Push API [8] helps to engage users even more, by displaying updates in order to catch their attention and make them return to the web application. Figure 1 shows the proxy-like position of a service worker that extends the traditional client-server architecture of web applications, by sitting between the web application and the remote web servers where the application loads resources from.

The service worker life-cycle is managed by the browser, which can put it to sleep or wake it up depending on whether there are events to be handled or not. Listing 1 shows an example of a service worker registered to manage an entire application.




Figure 1: Service Workers Making Use of the Fetch and Cache API. Upon intercepting a request on the fetch event handler, the service worker first checks whether a response to the request is present on the cache, in which case the response is served. Otherwise, the request is forwarded to the remote server to retrieve the response to the request, a copy of the response is put in the cache, and a copy is returned to the page that made the request.

```
navigator.serviceworker.register('sw.js', {
  scope: '/'
});
```

Listing 1: Registering a Service Worker against an Entire Web Application

The value `/` of the `scope` attribute in Listing 1 stipulates that the service worker will manage the entire application¹. For security reasons, service workers can only be registered on HTTPS websites. Also, the location of the service worker script used during its registration must be from the same-origin as the web application. Nonetheless, sites with insecure JSONP endpoints can be exploited by attackers to install malicious service workers².

The functionality of the service worker itself, as registered in Listing 1, is defined within the file `sw.js`, hosted by the application registering the service worker. There one defines the APIs that are used by the service worker, of which the most notable are described below. Listing 2 shows an excerpt of a service worker code.

As simple as it seems, this service worker is able to provide a full offline experience to users. Intuitively, it intercepts all HTTP requests (Line 1), checks whether a response is present in the cache for the request (Line 3) then serves it (Line 5), otherwise, the resource is fetched from the remote server (Line 6), a copy is cached (Line 10) then served to the web page that issued the request (Line 12). Later on, when the same resource is requested, it is read from the cache and served immediately, without a round-trip request to the remote server.

2.1.1. Fetch API. By using the Fetch API, a service worker intercepts all requests and responses to load web pages of an application (even when they are iframes loaded in other web applications), as well as all other same-origin and cross-origin content that load in the context of those web pages (scripts, images, stylesheets,

1. See https://developer.mozilla.org/en-US/docs/Web/API/Service_Worker_API/Using_Service_Workers for more about the scope of service workers

2. See <https://bugs.chromium.org/p/chromium/issues/detail?id=422966> for a discussion on the attack

```

1  self.addEventListener('fetch', (event) =>
2    {
3      event.respondWith(
4        caches.match(event.request.clone())
5          .then(response => {
6            return response ||
7              fetch(event.request.clone())
8                .then(response => {
9                  caches.open('offline').
10                     then(cache => {
11                       cache.put(event.request,
12                                response.clone())
13                     });
14                     return response.clone()
15                   });
16             });
17    });
18  )
19  });

```

Listing 2: Service worker that caches content

XMLHttpRequest, etc.). This makes a service worker a real proxy for a web application.

- Requests can be canceled, the headers and bodies modified (add/remove/change), or redirected.
- Responses can be generated on the fly from the service worker, read from the cache, or obtained from the network.

Note that HTTPS requests and responses are received in clear by service workers, even when they are third party requests. Notably, the browser imposes Same-Origin Policy (SOP) [9] restrictions when service workers try to access the content of HTTP responses. To be accessible, responses must either be from the same-origin as the service worker, or if they are cross-origins, they must be allowed by CORS.

2.1.2. Cache API. Coupled with the Cache API, a service worker can be used to provide an offline or better user-experience for web applications. Indeed, when a response is obtained from the network, a service worker can write a copy of it in the cache using the Cache API, as shown in Figure 1 and in Listing 2 (Line 10). Therefore, next time a request to the same resource is made, the service worker can serve the cached response, instead of or before fetching it from the network. This is particularly useful when network is not available, or if the developer wants to quickly present the user potentially outdated data that is updated after a response is received from the network.

2.1.3. Other features. Other popular features available to service workers are:

- *Importing scripts.* The SOP imposes that the registered service worker file comes from the same-origin as the application against which it is registered. However, the service worker can load additional scripts that also execute in the context of the service worker. Those scripts can come from third party servers. Many third parties have leveraged this feature to serve scripts to service workers, that provide a specific feature, such as push services found by Lee et al. [15].

- *Push Notification API.* An application asks the user for the permission to display notifications with updates that may be of interest to the user, at any time, even when the user navigates away from the application. When the application calls the Notification API, the browser displays a dialog box that prompts for permission to show notifications. With a granted permission, the server of the application can later on deliver push messages to the user’s browser, which will wake up the service worker to handle the push message.
- *postMessages.* Service workers run in an isolated context, so they cannot directly access the execution context of other same-origin web pages and workers. Nonetheless, they can communicate by using the postMessage API [2].
- *Sync API.* The Sync API is used by service workers to synchronize with remote servers, pending requests once the user got online.
- *IndexedDB.* Another shared storage between same-origin browsing contexts is IndexedDB. If the Cache API is dedicated to storing HTTP requests and responses, IndexedDB is a more general-purpose storage where different contexts can store and share data.
- *SharedArrayBuffer.* SharedArrayBuffer was used for synchronous data sharing. However, due to the famous Spectre attack [14], this API has been removed from many browsers. Nonetheless, Chrome has re-enabled it on browser versions supporting the site isolation mechanism [3].

2.2. Related Work

There have been so far little research around service workers. The first and most notable one is the work of Lee et al. [15]. The authors have been the first to conduct a study of security and privacy issues related to progressive web applications. Their work focused extensively on the Push Notification API, which they have shown could be used to mount phishing attacks, in particular because websites usually use their logos when showing notifications. Nothing prevents attackers from using the same logos when showing their own notifications, thus luring the users into clicking on phishing notifications as if they were from a benign domain. Moreover, they found that browsers such as Firefox would not include the domain name of PWAs in notifications, thus exacerbating the phishing attacks. They also analyzed popular third parties providing push services, and identified flaws that allow a network attacker to tamper with push messages delivered to users. Finally, they abused the persistency of service workers, whose lifespan is beyond that of normal web pages, in order to mine cryptocurrencies (Monero), by using push messages in order to distribute transactions. Likewise, Papadopoulos et al. [16] demonstrated a more resource abuse scenario, where a remote entity makes use of the current powerful web technologies in browsers, including service workers, to perform harmful computations and operations, such as mining cryptocurrencies for instance.

Lee et al. [15] also considered the Cache API, but in a privacy setting, for discovering PWAs installed in a user browser while offline. To do so, while the victim user is

still online, he is tricked into visiting a set of websites. If those websites deploy service workers, they will be installed in the user’s browser. Then, once offline, the attacker tries to load within iframes, resources from the set of websites that the user has been previously tricked into visiting. If they resources successfully load, since the user is offline, it means that the websites have installed service workers and caching content. This allows the attacker to learn the set PWAs of the victim user. To the best of our knowledge, this is the only work that studied the Cache API. As we have mentioned, the goal was to discover the set of PWAs of a victim user.

Our work is the first to consider serious security threats of the current design of the Cache API, and its interplay with potentially malicious scripts running in web pages. In Section 3, we provide more details on our threat model, and illustrate some scenarios with real-world examples.

3. Threat Model and Attacks Scenarios

The motivation for this work is to highlight that an attacker gains additional power on a progressive web application that uses service workers to cache web content in order to improve the application availability offline and/or its responsiveness.

To do so, we consider a standard web attacker who controls a malicious script running in the context of a website deploying service workers and caching web content. The malicious script could be:

- A third party script injected in the web application as the result of a Cross-Site Scripting (XSS) vulnerability exploit.
- A benign script in a first place, that turned wild intentionally or got compromised from the server-side by an attacker.

We consider that the attacker operates a bunch of malicious domains and subdomains under their under control, say for instance `https://attacker.com` and `https://*.attacker.com`. Although already powerful, such attacker could find on the web application various security restrictions that hinder the impact or feasibility of malicious actions on the page, for instance:

- Defense-in-depth security mechanisms such as Content Security Policy (CSP) in the page could highly mitigate the impact of the attacker, such as preventing them from injecting additional content, or making AJAX calls for instance.
- Defensive JavaScript programming practices, such as running first in the application, security critical libraries to freeze sensitive JavaScript objects, can prevent an attacker from tampering with such sensitive objects [13].

These security mechanisms, if properly deployed by the application and enforced by browsers, can highly protect the application even in presence of the normal web attacker described above.

3.1. Web Attacks Enabled by Service Workers

We now show how the web attacker can bypass those restrictive security mechanisms, in order to perform malicious actions, given that the application has




Figure 2: Attacks on Service Workers cache. As the cache is bound to the entire origin of the application managed by the service worker, malicious scripts running in the contexts of web pages can rewrite content stored in the cache. This can be done on (i) (security) HTTP headers or on (ii) HTTP responses content themselves.

deployed service workers which are caching content so as to enhance the application availability and responsiveness. Figure 2 describes the threat model. Since the cache is also accessible by potentially malicious scripts running in the context of web pages, the attacker can leverage their position in a web application to tamper with the content of the cache to perform attacks such as (i) removing security policies from web pages and breaking into even sandboxed contexts, (ii) voiding defensive programming, (iii) persisting attacks even after the attacker eventually got detected and removed, (iv) accessing sensitive user information present in the cache, for instance when a user logs out without the application clearing the cache. We illustrate selected attack scenarios in the following subsections.

3.2. Bypassing Security Policies

In recent years, the web has witnessed the introduction of a profusion of security mechanisms that web applications can deploy so as to mitigate a great number of web attacks. These policies are often shipped as headers of responses to navigation HTTP requests (requests to load HTML documents), where they govern and protect content on entire web pages. When these documents are cached by service workers, so are their security policies, which an attacker can manipulate in order to revert their effects and perform malicious actions, as shown on Figure 2.

3.2.1. Content Security Policy. Content Security Policy (CSP) is a defense-in-depth security mechanism which deployed on a page, defines content that are allowed or disallowed to load [18]. Various directive names are available for different web content types, and directive values intuitively represent the set of whitelisted origins where content of the specific type are allowed to load from. CSP is widely supported by major browsers, and studies have shown that it is deployed by a good number of top ranked web sites. Listing 3 shows an example of (a modified) CSP we found in the cache of Twitter’s service worker, for the page `https://twitter.com/home?precache=1`.

Directives names include `script-src`, `connect-src`, `object-src`, `frame-src`, etc. Among others, this policy mandates that the page can only load scripts (`script-src`)


```

script-src 'self' *.twitter.com;
connect-src 'self' *.twitter.com *.akamaihd
.net *.giphy.com *.twimg.com www.
google-analytics.com app.link;
form-action 'self' *.twitter.com;
frame-src 'self' *.twitter.com;
manifest-src 'self';
object-src 'none';
default-src 'self';

```

Listing 3: CSP adapted from the policy deployed on `https://twitter.com/home?precache=1` cached by Twitter’s Service Worker

from Twitter main domain (*self*) and subdomains (**.twitter.com*). Likewise, to fetch data (*connect-src*) the web page can only connect to Twitter main domain and subdomains, but also to a number of third party domains such as *akamaihd.net*, *giphy.com*, *google-analytics.com*, etc. The other directives read similarly. Notably, *default-src* is a fallback directive which applies to many content types when their related directives are not present in the page. For instance, the *img-src* directive used for images is not present. Therefore, the *default-src* directive in the Listing 3 restricts images to be loaded only from Twitter’s main domain.

Attack Now, we place ourselves in a setting where an attack occurs on Twitter’s homepage, as illustrated in Figure 2. We consider that the goal of the attacker is to load iframes from and connect to a third party domain under the control of the attacker (i.e. *attacker.com*). Thanks to the CSP deployed by Twitter, the attacker fails to load iframes or connect to `https://attacker.com`. In fact, since the attacker domain is not whitelisted on the *connect-src* and *frame-src*³ directives, the attack fails in a first place.

Nonetheless, since the CSP is also present with the content in the service worker’s cache, the attacker proceeds as follows:

- The cached HTTP response is read from the cache. In particular, the attacker removes the CSP header from the response, then writes back to the cache the modified response without the CSP header.
- Now the attacker reloads the page, by executing for instance `window.location.reload()`.
- Upon reloading, the service worker reads and serves the modified HTTP response without the CSP header.
- The attacker is now able to load iframes and connect to domains under his control.

3.2.2. X-Frame-Options. X-Frame-Options is a critical security header, that is used to mitigate clickjacking attacks. The typical values of this policy include:

- `deny` which mandates that the page never gets framed, not even by other same-origin pages.
- `sameorigin` allows framing only by other same-origin pages.

Note that the X-Frame-Options header is deprecated in favor of the *frame-ancestors* directive of CSP. However,

3. Depending on CSP version supported by the browser, the *child-src* directive is preferred to the *frame-src* directive for iframe restrictions.

X-Frame-Options is still widely used to mitigate click-jacking.

As an example, we found the `sameorigin` policy set on `https://www.pinterest.com/offline.html`, a web page cached by the service worker deployed by Pinterest. The policy mandates that `https://www.pinterest.com/offline.html` can only be framed by other pages from *www.pinterest.com*.

Attack Now we consider an attacker, as illustrated on Figure 2 present on `pinterest.com`, willing to frame `https://www.pinterest.com/offline.html` within an attacker’s controlled web page, say `https://attacker.com`. To do so, the attacker navigates the user to `https://attacker.com`, by executing `window.location = 'https://attacker.com'`. Then, from `attacker.com`, he tries to inject `https://www.pinterest.com/offline.html` as an `iframe`. The attack fails in a first place due to the presence of the X-Frame-Options: `sameorigin` header.

To make the attack succeed, the attacker needs to first remove the X-Frame-Options header from the response of `https://www.pinterest.com/offline.html` cached by the service worker of Pinterest.⁴ After removing the X-Frame-Options header, the attacker can proceed to the attack, which this time will succeed.

3.2.3. Referrer Policy. A Referrer Policy controls the value of the `Referer` HTTP request header. Properly configured, it could prevent an application from leaking to third parties, potentially sensitive information that could be contained in a URL. In fact, by default, the `Referer` header included the entire URL, which may contain sensitive information. With a `Referrer-Policy`, one can control more precisely the value of the `Referer` header.

A `Referrer-Policy` typically takes the following common values:

- `no-referrer` (`Referer` header is not sent), `same-origin` (do not send `Referer` header to third parties), `no-referrer-when-downgrade` (send entire URL, but not for insecure requests, i.e., an `http://` request from an `https://` web page);
- `origin` (only send the origin of the webpage), `origin-when-cross-origin` (to third parties, only send the origin of the web page), `unsafe-url` (leak the entire URL);
- `strict-origin` (only send the origin of the web page to secure requests), `strict-origin-when-cross-origin` (to secure third parties, only send the origin of the web page).

Individual HTML elements can be applied a `Referrer-Policy` by setting the `referrerpolicy` attribute. These elements include `<a>`, `<area>`, ``, `<iframe>`, `<script>`, `<link>`.

As an example, `thehackernews.com` deploys a service worker which is caching numerous HTML pages with a `strict-origin-when-cross-origin` `Referrer-Policy`. This policy instructs the browser to

4. If a CSP is also present, with the *frame-ancestors* directive, the attacker can either remove entirely the CSP, or only the *frame-ancestors* directive. In fact, the *frame-ancestors* directive applies to clickjacking when present.

leak to third parties only the origin of web pages, and not the entire URL.

Attack To bypass this policy and leak the entire URL, the attacker only needs to remove the policy from the cached page, then reload it. This time, the browser will leak the entire URL to third parties.

3.2.4. Feature Policy. A Feature Policy can be deployed to restrict an application and its sub-resources in the usage of web APIs and user devices such as camera, microphone, the geolocation, origin relaxation (with `document.domain`), etc. [5] Listing 4 shows an example of `Feature-Policy` cached by the service worker on `computerbase.de`.

```
camera 'none';
document-domain 'none';
geolocation 'none';
microphone 'none';
payment 'none';
sync-xhr 'none';
```

Listing 4: Feature Policy deployed and cached by `computerbase.de`

`Feature-Policy` works similarly to CSP. In fact, it has directives (`camera`, `geolocation`, `microphone`, `sync-xhr`, ...) that are used to name specific APIs and devices, and directive values (such as `'none'`, `'self'`, `https: *.computerbase.de`) which whitelist the origins that can use those APIs.

Attack Listing 4 is an actual feature policy that we found in the cache of the service worker deployed by `computerbase.de`. This policy effectively prevents scripts from using any of the devices and APIs listed in the policy, i.e., the camera or the microphone.

To bypass this policy, the attacker can simply remove it from the content cached by the service worker of the application. The modified policy removes all the restrictions on using the related APIs, allowing the attacker to request and use them, if the user provides the rights to do so.

After removing the `Feature-Policy` and reloading the web page, requesting access to the camera will successfully display a dialog box to the user asking for permission to access the hardware. Even if access to the camera is still subject to a user's decision, a few blogposts have demonstrated different ways a user can be tricked into unpurposedly clicking on the *Allow* button which grants the application access to the camera. This can be done for example by displaying an attacker-controlled dialog box on top of the browser's permission dialog.⁵

3.3. Manipulating Cached Content

In addition to manipulating security policies cached by service workers, the attacker can also manipulate the cached content themselves to perform various attacks. Note that manipulating content in the cache is subject to the Same-Origin Policy restrictions [9].

- Modifying a cached content is only possible if the content is from the same-origin as the service

5. See <https://www.youtube.com/watch?v=i-MiYevFy6Y> for a demo on permission prompt spoofing.

worker, or if the content is the result of a successful CORS request. This is the case for all cached HTML content, which are the target of most of our attacks.

- Adding and deleting or replacing content is allowed for any type of content.

3.4. Bypass Defensive Programming

By defensive JavaScript programming, we consider well-known programming idioms that can be used to restrict the set of actions that can be performed by an attacker controlled script. A typical idiom is shown below where a web page makes a security-critical script `security.js` runs first in a web page,⁶ which defines objects and constants that cannot be altered by an attacker script that runs afterwards.

```
<script src='/security.js'>
<script src='https://attacker.com/atk.js'>
```

3.4.1. Constants. In JavaScript a constant, once defined cannot be reassigned another value or deleted. This idiom can be used by a web page to ensure that a variable is always present and has the same value whenever accessed.

The code below shows the definition of a constant variable, named `token`. Attempts to write to this constant variable will fail, raising a `TypeError` exception.

```
const token = 'c3G07zd20S9qU56Dr5ID';
```

3.4.2. Frozen Objects. A script that runs first can also freeze entire JavaScript objects⁷ so as to prevent them from being modified afterwards.

To achieve such properties on an object, a script that runs first in a page can define and freeze it as shown in Listing 5.

```
let obj = { a: 'hello', b: 'world' };
Object.freeze(obj);
```

Listing 5: Frozen JavaScript object.

Attempting to alter the properties of the Object `obj` afterwards will silently fail. In particular, attempting to delete a property, will return `false`.

3.4.3. Attack. Now we place ourselves in the setting of an attacker as shown in Listing 2. The goal of the attacker is to alter constants or frozen variables. There are two different scenarios to be considered.

The constants and frozen objects are in a same-origin cached content. As we have mentioned earlier, same-origin content in the cache can be altered. The attacker only needs to:

- Read from the cache the original content, and locate the constants and frozen objects that have to be altered.
- For a constant, if the goal of the attacker is to set a single new value, the attacker can directly do so by replacing the old value of the constant with

6. See [13] for a discussion on how to run a script first in a web page.

7. <https://tc39.es/ecma262/#sec-object.freeze>

the new one. Otherwise, if the goal is to make the constant value writable multiple times, he can do so by replacing the `const` keyword, which defines constants, with the `let` keyword instead. By doing so, the variable can be modified later on at will.

- For frozen objects, the attacker simply removes or comments the instruction `Object.freeze` that freezes the object, making it a normal object that can be altered later on at will.
- After performing the desired modifications, the attacker writes back the modified content to the cache.
- Finally, the attacker reloads the page so that all the modifications can be taken into considerations: the constants and frozen objects have lost their protections, and can be freely altered by the attacker.

The constants or frozen objects are in a cross-origin cached content. The attacker cannot modify cross-origin cached content. Nonetheless, the attacker can redefine the order in which scripts are loaded within a page, so as to achieve the modifications on constants and frozen objects. To do so, the attacker:

- Rewrites the HTML content of the web page (which is same-origin therefore alterable), in order to make the attacker script run before the scripts where the constants and frozen scripts are defined.
- On the objects that will be frozen afterwards by the web page, the attacker performs the desired operations: addition, deletion and modifications of properties and their values. Then, and very importantly, the attacker freezes the object. Freezing the object by the attacker is important if the scripts that run afterwards attempt to revert the changes of the attacker, before freezing the object. By doing so, the web page ends up managing objects with the properties and values chosen by the attacker.
- For constants, the intuition is the same. The attacker is the first to define the constant, with the value of his choosing. When other scripts in the web page attempt to define the constant again, the browser will raise an exception.⁸
- For the modifications to be taken into consideration, the attacker just reloads the web page.

3.5. Persisting Attacks

In this scenario, we consider an attacker that has infected an application for a certain time, but eventually got detected and removed from the application. The attacker can write malicious scripts in the cache, as if they were loaded by the application itself. To do so:

- The attacker identifies the web pages where he wants his scripts to execute.
- If the pages are not yet in the cache, the attacker can add them to the cache, by calling `cache.addAll` with an array of the URLs of web pages to store in the cache.
- Finally, the attacker alters the cached web pages, in order to inject the attacker scripts.

8. If the exception is not anticipated and caught by the web page, this may cause the page to malfunction.

In this category, we also include (malicious) browser extensions, that the user has installed but eventually got uninstalled at some point. The extension can also store scripts in the cache, that will persist even after the removal of the extension from the user's browser. Moreover, Firefox extensions offer capabilities that outmatch those endowed to the threat model we consider. Indeed, extensions have the ability to manipulate HTTP responses. Therefore, they can install a service worker on web applications that have not deployed any service worker. They can even further prevent the application from registering its own service workers, and therefore can persist the extension malicious presence in the user browser. This powerful extension attacker is out of the scope of this work.

3.6. Accessing Sensitive User Information

In web applications, an origin remains the same no matter whether a user is authenticated, or not. Hence, if the cache is not cleared by the developer, content written into the cache when a user is authenticated persist even when the user is logged out. An attacker who is present only in logged out pages, can still access sensitive information previously written into the cache while the user was logged in.

4. Methodology

We performed an empirical study to assess the deployment of service workers and to identify websites that fall into our threat model. We actually performed 2 data collections. The first one, as summarized in Figure 3, was fully automated. Moreover, we also performed a second, semi-automated collection, where we manually logged into a set of websites using Google and Facebook single-sign on. The goal of the second collection was to assess whether more interesting service worker use cases occur after authenticating to websites.

4.1. Fully Automated Data Collection

Our data collection methodology is summarized in Figure 3

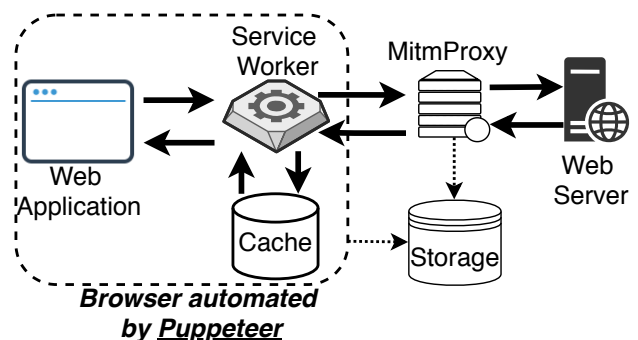


Figure 3: Data Collection Methodology. Service workers deployed by web applications are hooked and their APIs monitored at runtime.

4.1.1. Browser Configuration. We used Chrome for collecting the data, because it is the leading browser in the implementation of service workers APIs, and in PWAs. For instance, so far only Chrome proposes users the ability to install their favorite PWAs. Moreover, Lee et al. [15] noticed that some websites required the user to grant the Notification permission before service workers were registered. Since they were using Firefox in their study, they had to manually go through 2,911 sites in order to grant the push notifications. We do not have such limitation in Chrome, where the browser can be configured to grant the push notifications to any website requesting it. For our work, we created a base user profile in Chrome which we configured to grant the push notifications to all sites. This can be done in `chrome://settings/content/notifications`, and adding `https://*` to the set of allowed sites. We then used a copy of this base profile in order to handle each site.

4.1.2. Mitmproxy Setup. Mitmproxy is an interactive HTTPS proxy for intercepting and manipulating HTTP responses [12]. We set up Mitmproxy to intercept requests issued by the automated Chrome browser. The main goal is to identify websites that are loading service workers. In fact, Chrome adds the informative `service-worker` header to requests issued by web applications while registering service workers. Lee et al. [15] did a pattern-matching in their work in order to identify websites using the Push API. We decided instead to perform a runtime monitoring of service workers to identify the APIs that they call, and in particular their accesses to the cache.

When Mitmproxy intercepts a request to load a service, it waits for the response from the website, and prepends a JavaScript code that we have defined. This code runs first within the service worker. In this code, we extensively make use of JavaScript Proxy API [6] in order to intercept all API calls made by the original service worker. The arguments of each API calls are then saved into a local database for further analysis. It is worth noting that our hooks fully preserve the original functionality of service workers, by only passively monitoring the service workers APIs.

4.1.3. Automation with Puppeteer. To automate the collection process, we used Puppeteer [7], which among others, allows us to run multiple Chrome instances in parallel, each for one topsite. To support all the multiple Chrome instances, we also run multiple Mitmproxy instances, on different ports. In summary, here is how we handled a top site:

- For each top site, we look for popular and related (sub)domains on the first 5 results pages of Google and Bing search engines.
- We set up listener that registers all HTTP requests and responses. Indeed Puppeteer provides useful information on when a response is served by an installed service worker or not.
- We instruct Puppeteer to inject a script in all web pages in order to monitor potential accesses to the cache from outside of the service worker. This will serve as a ground truth to motivate some of our countermeasures (See Section 6).
- We visit the homepage of the site, and wait for a couple of seconds and then extract related links

(same domain and subdomains). We merge those links with those from the search engines.

- We group the links per origin, randomly select at most 50 links of them from different origins, with at least 2 links per origin. The intuition of selecting 2 links per origin is motivated by the design of service workers. In fact, when a service worker is registered, it manages only requests of webpages that are navigated after its installation. With 2 links per origin, if a service worker is registered on the first link, then by navigating the second link, we have more chance to trigger more APIs calls.

5. Cache Attacks via Service Workers

We considered 110,000 websites from the top Tranco List [17]. The results are presented per number of origins, and number of sites. The motivation being that service workers operate on a per-origin basis, but we also show the numbers of sites in order to give a fair comparison with previous results.

5.1. Sites Deploying Service Workers

Figure 4 shows the origins deploying service workers that we found, and the different APIs used by those service workers.

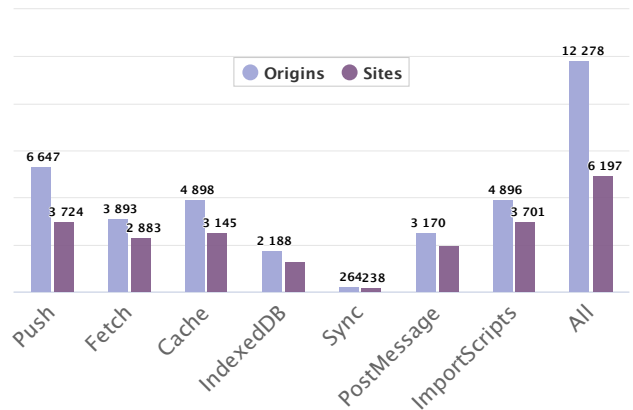


Figure 4: Top sites deploying service workers, and the functionality mostly used by those service workers, in number of sites, and number of origins.

Notably, Push Notification is the most used API by service workers, following by the import of additional scripts within service workers context. Le Pochat et al [17] observed a similar trend in their results. But interestingly, we observed that a half of the websites deploying service workers, also make use of the Cache API for storing HTTP requests and responses. Comparatively, Le Pochat et al [17] found only 12.3% of cache usage among websites deploying service workers. We think we achieve better results for two main reasons. First our methodology of monitoring service workers is definitely more precise than pattern matching. But the main reason why we found many sites using the Cache API is because we used Chrome for data collection, where in order to propose users to install PWAs, websites must first register a service worker with the Fetch API, which certainly leads to more

sites making extensive use of the Cache API to provide offline experience to users.

Other popular APIs include the Fetch API, which is used to intercept and respond to HTTP requests and responses, either directly from the cache when the responses are present, or by fetching the resource from web servers otherwise. The IndexedDB storage is also widely used, while the Sync API is not very widespread. As service workers run in an isolated environment, many websites implement `postMessages` in order to let other contexts such as web pages, communicate with them.

5.2. Cache Analysis

Figure 5 shows the main operations performed on the cache by sites using the Cache API.

The most used operation is the opening of the cache, which is required for other operations. After opening, writing a response in the cache is the most used operation. Service workers either use the `Put` operation to save an HTTP response that is received from a server as the result of an HTTP request issued by web pages under the scope of service workers. Or they use the `addAll` operation in order to prefill the cache with a set of content that would be available right away when web pages request them. Later on, when a web page requests a resource, the service worker would check whether the content is already present in the cache. This is done with the `Match` operation. Cached content help improve the responsiveness and availability of websites, even when the user is offline. The less popular operations are deletion from the cache, or the addition of individual content (`add`). Content in the Cache can be organized by cache names. The `has` operation is used to check whether a cache name is present.

Figure 6 presents the types of content we found cached by service workers. Interestingly, the content types of interest to us are the most widely cached. This includes HTML, JavaScript and the result of AJAX responses. Served by service workers when they are requested by web pages instead of issuing a request to a remote web server, those content help to improve the responsiveness of web applications. Nonetheless, as we have described in our threat model, the cache and its content are also accessible to attackers, who can mount attacks that we have described in Section 3.3.

5.3. Security Policies

In this section, we illustrate our threat model regarding the bypass of security policies. Figure 7 presents the security policies found only on cached web pages (HTML documents). For the needs of our attacks, we consider only the HTML documents from those sites, since the security policies considered here only make sense when deployed on HTML documents. Appendix A further breaks down the cached security policies according to the most used directives.

`X-Frame-Options` is the most widely used policy among the cached web pages (Figure 7). An attacker present on those sites can remove such policies on the cached pages, thus opening the application to clickjacking attacks, as we have described in Section 3.2.

A good number of websites also cache web pages with `Content-Security-Policy`. Those policies are used to restrict script or plugin injection, clickjacking (*frame-ancestors*) or for TLS enforcement. A few websites also use CSP to restrict connections endpoints for AJAX, frame injection or form submissions. The effect of those policies can be voided by attackers, as they can simply remove the CSPs from the cached web pages, in order to perform actions that were not permitted with the policies in place.

For `Referrer-Policy`, it appears that it is used on cached web pages to allow third parties to learn via the `Referer` header, only the origin of web pages from which third party requests are issued, and not the entire URL of web pages. One can observe that the other `Referrer-Policy` values are rather interesting policies, whose effects can be reverted by an attacker which can remove the policies from the cached web pages.

`Feature-Policy` is less used on cached web pages. Nonetheless, we observe that the websites using it rather deploy very restrictive policies. Unfortunately, an attacker can void the effect of such restrictive policies and request access to the related APIs.

6. Countermeasures

The attacks that we have discussed in this work have all to do with the fact that the service worker cache is bound to the entire origin, and thus can be accessed by potentially malicious scripts running in same-origin web pages.

To mitigate this issue, we think that browser vendors should make the cache of service workers inaccessible to other same-origin browsing contexts. This prevents potentially malicious scripts running in web pages from tampering with the cache, as depicted by Figure 8, effectively addressing all of the issues discussed in this work.

To further motivate the need for a private cache for service workers, a question that remains is to assess how many web applications will malfunction if service worker's cache is no more accessible to web pages. In other words, we want to assess whether there are already websites manipulating the cache from the context of web pages. The numbers are shown in Figure 9. As one can see, the main operation consists in opening the cache, in order to write, search or delete content. Preventing cache accesses from web pages will certainly impact the few number of sites that operate on the cache from webpages. Nonetheless, all those operations could be directly performed by service workers. In fact, if those operations are really necessary, and since the developer has full control over the service worker and web pages, `postMessages` can be used to communicate appropriate instructions to service workers, so that the proper operations can be performed on the Cache. As shown by Figure 4, the usage of `postMessage` is already widespread among service workers.

7. Conclusion

In Progressive Web Applications, service workers make use of the Cache API in order to improve the responsiveness of web applications and provide users an offline browsing experience even when network is unavailable. In




Figure 5: Common Operations on Cache




Figure 7: Security Policies




Figure 8: Make the service worker cache inaccessible to web pages




Figure 9: Cache Accesses from Web Pages

this work, we have shown that since in its current design, the cache is bound to the entire origin of web applications deploying service workers, malicious scripts running in web pages for instance can tamper with the cache in order to defeat security protections such as removing security policies from web pages and bypassing defensive JavaScript programming idioms. These represent a new class of attacks that are only possible with service workers caching content. We performed an empirical study and found a good number of topsites caching content with security policies, making them vulnerable to the threats considered in this work. As a countermeasure, we suggest browsers vendors make the cache of service workers inaccessible to web pages and other same-origin contexts.




Figure 6: Type of Content Cached By Service Workers

References

- [1] Object.freeze() . https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Global_Objects/Object/freeze.
- [2] PostMessage API . <https://developer.mozilla.org/en-US/docs/Web/API/Window/postMessage>.
- [3] Site Isolation . <https://www.chromium.org/Home/chromium-security/site-isolation>.
- [4] CORS protocol - Fetch Specification. <https://fetch.spec.whatwg.org/#http-cors-protocol>.
- [5] Feature Policy. <https://w3c.github.io/webappsec-feature-policy/>.
- [6] JavaScript Proxy API. https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Global_Objects/Proxy.
- [7] Puppeteer chrome automation tool. <https://pptr.dev/>.
- [8] Push API. <https://w3c.github.io/push-api/>.
- [9] Same-Origin Policy. https://developer.mozilla.org/en-US/docs/Web/Security/Same-origin_policy.
- [10] Service Worker API. https://developer.mozilla.org/en-US/docs/Web/API/Service_Worker_API.
- [11] 26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019. The Internet Society, 2019. [Online]. Available: <https://www.ndss-symposium.org/ndss2019/>
- [12] r. Aldo Cortesi, Maximilian Hils. Mitmproxy - an interactive HTTPS Proxy. [Online]. Available: <https://mitmproxy.org/>
- [13] K. Bhargavan, A. Delignat-Lavaud, and S. Maffei, "Language-based defenses against untrusted browser origins," in *Proceedings of the 22th USENIX Security Symposium, Washington, DC, USA, August 14-16, 2013*, S. T. King, Ed. USENIX Association, 2013, pp. 653–670. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/presentation/bhargavan>
- [14] P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom, "Spectre attacks: Exploiting speculative execution," in *2019 IEEE Symposium on Security and Privacy, SP 2019, San Francisco, CA, USA, May 19-23, 2019*. IEEE, 2019, pp. 1–19. [Online]. Available: <https://doi.org/10.1109/SP.2019.00002>
- [15] J. Lee, H. Kim, J. Park, I. Shin, and S. Son, "Pride and prejudice in progressive web apps: Abusing native app-like features in web applications," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018*, D. Lie, M. Mannan, M. Backes, and X. Wang, Eds. ACM, 2018, pp. 1731–1746. [Online]. Available: <https://doi.org/10.1145/3243734.3243867>
- [16] P. Papadopoulos, P. Ilia, M. Polychronakis, E. P. Markatos, S. Ioannidis, and G. Vasiliadis, "Master of web puppets: Abusing web browsers for persistent and stealthy computation," in *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society, 2019. [Online]. Available: <https://www.ndss-symposium.org/ndss-paper/master-of-web-puppets>

- [17] V. L. Pochat, T. V. Goethem, S. Tajalizadehkhoob, M. Korczynski, and W. Joosen, "Tranco: A research-oriented top sites ranking hardened against manipulation," in *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society, 2019. [Online]. Available: <https://www.ndss-symposium.org/ndss-paper/tranco>
- [18] M. West. (2016) Content Security Policy Level 3. W3C Working Draft. [Online]. Available: <http://www.w3.org/TR/CSP3/>

Appendix A. Cached Security Policies

Directive	#Origins	#Sites
Content-Security-Policy		
script-src	189	185
object-src	146	145
base-uri	132	131
frame-ancestors	129	117
upgrade-insecure-requests	114	89
worker-src	111	110
default-src	76	74
connect-src	70	69
form-action	48	48
frame-src	45	42
font-src	43	41
block-all-mixed-content	20	17
child-src	20	17
manifest-src	13	12
X-Frame-Options		
SAMEORIGIN	701	641
DENY	107	94
allow-from	14	13
Referrer-Policy		
strict-origin-when-cross-origin	166	160
no-referrer-when-downgrade	68	60
origin-when-cross-origin	25	25
no-referrer	16	16
same-origin	16	15
strict-origin	10	10
origin	6	6
unsafe-url	3	3
Feature-Policy		
geolocation	17	16
microphone	16	15
sync-xhr	12	10
camera	16	15
magnetometer	14	13
gyroscope	14	13
payment	13	12
midi	10	9
speaker	8	7
fullscreen	8	7
accelerometer	6	6
usb	6	6

TABLE 1: Directives and Values of Cached Security Policies