

Persistence of Passwords in Bitwarden’s Browser Extension: Unnecessary Retention and Solutions

Rafael Prates
Instituto Superior Técnico
Universidade de Lisboa

Abstract

Password-based authentication is still the dominant form of authentication on the web, yet users do not adopt password managers for fear of them being insecure, unreliable and other reasons. In this project we modify a password manager to try to comply with certain data security properties as a way to increase adoption of this type of software that has been increasing in importance.

Taking BitWarden’s Google Chrome extension as our chosen password manager, we define password manager states and data security properties regarding the master password that we would like to comply with, perform tests and analyse password retention problems in the application. While the BitWarden extension interacts with many layers, we decided to only change the application layer, as a way to understand how much can be done by the developers of these types of applications.

We then introduce our modified extensions that try to solve the issues presented before and introduce a testing framework that is able to automatically interact with the extension through the graphical user interface to replicate the use case chosen. While our solution does not completely solve the issue, we were able to reduce the problem slightly.

1 Introduction

The need for more passwords has been increasing over the years as password-based authentication is still the dominant form of authentication on the web[8] and as users sign up to more and more services, they require more and more passwords. Unfortunately, because the number of passwords a typical user needs to remember is increasing, users tend to reuse passwords across multiple services or small variations of the same one and reuse previously leaked passwords [38]. Many users tend to follow these practices to avoid the cognitive burden of recalling different passwords [12, 33] especially when passwords that are difficult for an attacker to guess are also hard to memorize for the user. This is compounded by

the fact that users see the high amount of effort required to get just a low percentage of passwords in data breaches and think the ease of use of password reuse outweighs the risk.

PMs come as a solution to this problem by inheriting the responsibility of remembering passwords from the user to the software. This in turn allows the user to associate stronger passwords, that are hard to memorize, but even harder to attack, to the services they use boosting the user’s security and avoiding the unsafe practices of password reuse, all while offering usability features such as automatic strong password generation resilient and auto-completion of log in forms to name a few. PMs also help protect the user against data breaches. By using a strong unique password for each service, even if one service gets its data breached the rest will be safe since there is no password reuse. This also lowers the amount of work a user has to do whenever a data breach occurs, since it is easier to change one password (the one that was breached), than to change all passwords if the user reused the same password or variations of it on all their services. And while having strong unique passwords is not exclusive to using a PM, it certainly makes it easier to do so.

And this is important as the number of data breaches are on the rise [15], leaving many users vulnerable to attacks.

However, PMs are not immune to attacks and a portion of potential users do not feel that using a PM would provide greater security [9] for their secrets, some think they are insecure [1, 2, 4, 34] and some of them outright distrust [2, 34] PMs as a whole. Because of these reasons and others, these users refuse to adopt PMs.

In this project, we decided to focus our attention to a particular type of attack that can be used to steal sensitive data from a computer’s memory. These attacks involve making a copy of the device’s random access memory (RAM) either through a cold boot attack [22] or through vulnerabilities that might exist such as HeartBleed [13], MeltDown[29], Specter[27], etc...

In the context of a PM application, it is expected that secrets, such as passwords or cryptographic material derived from passwords, will have to be in the memory of the process at

some point in time throughout the use of the application. If an attacker is able to extract the memory at the right time, they will be able to extract and steal that information. As such, it is critical that these secrets are deleted from memory as soon as they are no longer necessary. Although there are different types of sensitive data an application can hold, we will only focus on user passwords.

Our study will take BitWarden[5] as our PM, more specifically the Google Chrome (Chrome)'s [19] extension of BitWarden as the case of study. We will be focusing on logging into the PM with a master password (MP) and by dumping the memory of the BitWarden extension process, analyse when the MP is in memory as well as when it should not be according to some desired data security properties. We then implement solutions at the application level to try and eliminate this secret in memory when it is no longer needed.

2 Background work

This chapter presents the relevant background work needed to understand the following chapters. We will be introducing what a password manager (PM) is in general, desired data security properties of a PM, BitWarden's login process, the different layers of abstraction between the application code and the machine and insecure code practices that can retain passwords longer than necessary.

2.1 Password Managers

2.1.1 What are they?

The main feature of PMs consists in storing username information and the correspondent password into what is known as the password vault (PV). The PV is then responsible for storing the user's information in an encrypted fashion to ensure that the user's secrets are kept hidden from attackers.

Since there are various PMs available online, some offer more features such as, password generation, auto-completion of login forms, cloud storage of the PV, cross-device synchronization and the list goes on but in the end the basic concept remains the same. The PV may also store other relevant information like website data, namely URL and icons to name a few, creation date, fill count and other. This information will be referred to as metadata since it is not vital for logging into an account.

2.1.2 How are passwords stored

Password storage is done by creating a password vault, either on the local system or on the cloud. When a PV is stored locally, they are usually an encrypted file residing on the local system's disk. When stored on the cloud, it also encrypted in some fashion, with the details differing from provider to provider. The algorithms used to encrypt a PV file differs from

PM to PM and most require a MP to reverse the encryption and be able to access it.

2.1.3 What is stored on the password vault?

A password vault contains entries. There are as many entries as the user wants or needs, but the regular behaviour is having an entry for each different account. Each entry stores the necessary information for logging into the account with the minimum usually being the username and password.

In addition, it is not unusual to store metadata along with the entry on the password vault. This includes but is not limited to:

1. Website metadata such as URL, Icon and name.
2. Password metadata such as creation time, modification time, last use time, fill count and expiration date.
3. User settings namely notes and autofill settings.

Unfortunately not every PM guarantees encryption of all metadata. For example, KeePassX and KeePassXC both encrypt all metadata [30]. Extension-based password managers encrypt most metadata, but all have at least one item they do not. Browser-based managers that rely on the operating system to encrypt their vaults protect the relevant metadata too. Those that do not rely on the operating system have a significant amount of unencrypted metadata.

2.1.4 How are password vaults encrypted?

In general, most app-based and extension-based encrypt their vaults using a MP. Requirements for the MP vary between applications with some not requiring one at all. Browser-based systems (except Firefox) rely on the operating system instead to help them encrypt the password vault and as such implementations vary.

2.2 Desired Data Security Properties of a Password Manager

Since this project is done in the context of the PassCert project, BitWarden was chosen as the base PM. The PassCert project is an effort to create a proof-of-concept PM that through the use of formal verification, guarantees properties on data storage and password generation[21]. which is the BitWarden extension for chrome.

With the basic function of a PM in mind, we will define 2 states in which the application is running: (a) not running; (b) locked (pre-login); (c) unlocked (and running); and (d) locked (session terminated)[31].

We will not analyse the PM in its not running state (when Chrome is not opened or the extension is disabled) as we have decided to focus on the security of the MP in the login process of BitWarden.

2.2.1 Locked (pre-login)

We define BitWarden to be in the "locked (session terminated)" when the vault has not been unlocked in the current session. In this state we concede that the MP is stored in the memory of the program and is visible to potential attackers as the application need the MP to perform the tasks necessary to authenticate the user.

2.2.2 Unlocked (and running)

We define the BitWarden PM to be in the "Unlocked (and running)" state once the BitWarden vault is unlocked. To reach this state, the user must successfully authenticate by entering a valid e-mail and MP.

In this state, the MP is no longer necessary to be in memory as we will discuss in section 2.3.1 and as such, should no longer be in memory.

2.2.3 Locked (and running)

We define the BitWarden PM to be in the "Locked (and running)" state when (a) the vault is locked manually; or (b) when the session is terminated. For this project however, we only considered (b) in our testing. Continuing from the unlocked (and running) state, the MP should not be in memory as well.

2.3 BitWarden

BitWarden is a PM available for different platforms, desktop, mobile, browser extensions and even online. As expected from a PM, it manages account information like usernames and passwords in a vault.

2.3.1 BitWarden's Login Authentication

Figure 1 is an overview of how BitWarden's login authentication works. When the user provides the e-mail address and MP, BitWarden uses Password-Based Key Derivation Function 2 (PBKDF2)[24] with a 100,000 iteration rounds to stretch the MP, using the e-mail address as salt. The generated value is a 256 bit Master Key. This Master Key is used in conjugation with the MP as salt to create the Master Password Hash which is sent to the BitWarden server upon account creation and login, and used to authenticate the user account. Once the Master Key and the Master Password Hash have been generated, the MP is no longer required to be in memory, as the application has everything it needs to authenticate the user.

BitWarden mentions that in their client application they do not store the MP locally or in memory and that they do their best to ensure that any data that may be in the application to function is only held in memory for as long as needed[6]. Ideally, the BitWarden extension for Chrome would follow these principles.

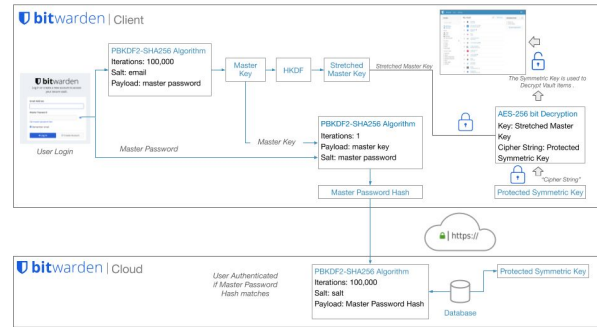


Figure 1: Control flow of logging into BitWarden

2.4 Security issues of keeping secrets in memory for longer than necessary

Keeping secrets longer than necessary is dangerous, as it is quite possible for passwords to remain in memory even after an hour after the program was terminated [25]. An attacker that can obtain a memory dump of the computer could potentially obtain secrets this way, use them to breach the user's vault, steal all the information in it and proceed to breach accounts for other services that were stored in the vault.

2.5 Layers of abstraction

The application we are studying runs on top of several other layers of abstraction. Namely the application is written in Typescript with the Angular framework. This Typescript code is then transpiled to JavaScript (JS) using Angular's transpiler. The JS resulted from this transpilation is then interpreted by browser (in our case, Chrome) using the browser's built-in JS interpreter. In Chrome's case, it uses the V8 JS engine[36] and the Blink[7] rendering engine. These two communicate with each other to interpret and display the page. Since Chrome is written in C++, it interfaces with the standard C++ runtime libraries and any other libraries necessary by Chrome's code. Lastly, Chrome also has to interface with the operating system (OS).

Our project focuses only the application layer as it would be unfeasible to modify everything in the pipeline.

2.6 Password retention risks

2.6.1 Application sent to background

Modern OSes implement Virtual Memory which, amongst other things, can allow for secondary memory to act as main memory. This allows applications to swap memory in RAM to the disk when system resources are low. This poses a significant risk as a secret could be the memory of a process when its memory gets written into the disk, prolonging the amount of time a secret is exposed. An attacker could then

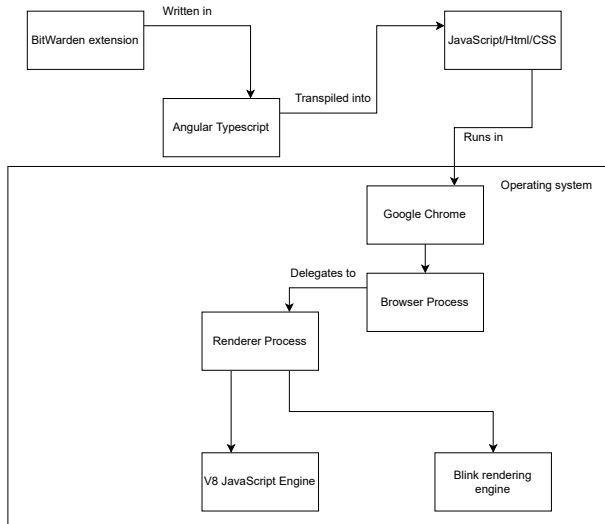


Figure 2: The different layers of abstraction

steal the secondary memory storage medium (a solid-state drive, hard-disk drive, etc...) and analyse the paging files to extract secrets that were written to the medium through virtual memory.

2.6.2 Crash dumps

If the computer crashes while the application contains a secret in its memory those secrets could be written to the crash dump file, exposing them. For example, Windows[18] and Ubuntu [26] (if the kernel crash dump utility is installed) dump the contents of RAM at the point of crash into a file.

2.6.3 Delayed garbage collection

The memory of JS applications are managed by a garbage collector (GC). When a piece of memory is no longer referenced by any variable it will remain in memory until the GC decides to reuse it. The amount of time that it takes until that piece is reused by the GC is undefined. It can be minutes or even hours [25], depending on the algorithm used and the system's resources and workload. This comes into play when we take immutable types in JS into consideration. The primitive type String in JS is immutable and thus can not be overwritten manually by the developers, leaving the deletion up to the GC when it eventually reuses the piece of memory that contained the String object and overwrites it with some other value. This is a known problem when using immutable data types for sensitive information, such that even Java's Cryptographic Architecture recommends using mutable data structure types [23] for passwords and secret data.

2.6.4 Different layers of abstraction and lack of secure API

Since the BitWarden extension runs on top of other layers, any communication between these layers has the possibility of creating buffer copies of sensitive data and having that data retained for a period of time, outside of the control of the application itself. Without any options of a secure API between layers, the application developers have no control over how the layers that the application interacts with treat the data.

2.6.5 Function Arguments Copies

Whenever a primitive type is passed to a JS function, a new copy of it is made[17]. This presents a problem because if a secret is passed to a function using a String data type, then a new immutable copy of the secret is created hence increasing the amount of copies of it in memory. Numerous copies of the secrets in memory means that it is harder for the GC to delete and reuse those memory pieces in a timely manner, as discussed in subsection 2.6.3.

3 Password retention on the BitWarden extension

3.1 Threat model

Our work assumes an attacker has access to a snapshot of the memory of the computer through a memory disclosure attack or through a memory dump of the system. This could be from a crash dump file, like we discussed in section 2.6.2 or through a cold boot attack[22].

3.2 Memory Dump Analysis

We performed several tests on the BitWarden extension for the Chrome browser by dumping the memory of the Chrome process responsible for running the BitWarden Extension. These memory dumps were done several times at different points in time, throughout the use of the extension.

3.2.1 Test steps

The points of time in which we perform a memory dump will be referred to as a "test step" in the testing procedure. As such, the following is an overview when the process' memory is dumped:

- **Step 0** - After typing the e-mail in the e-mail field, but before typing the MP into the password box;
- **Step 1** - After half the MP was typed into the password box;
- **Step 2** - After the MP was fully typed;

- **Step 3** - After logging in and unlocking the BitWarden Vault;
- **Step 4** - After simulating a task;
- **Step 5** - After terminating the BitWarden session.

A more in-depth look at the testing procedure will be presented in section 5.

3.3 Results and observations

In the memory dumps, we refer to both a partial MP and a full MP. A full MP is the set of characters that compose the entirety of the MP, whether they be encoded in 8-bit (UTF-8/ASCII) or 16-bit (UTF-16) encoding. Given a full MP p of length i , a partial MP is defined by the subset $p_j, j \in \{\frac{i}{2}, \dots, i-1\}$. This applies to both UTF-8/ASCII encodings and UTF-16 encoding.

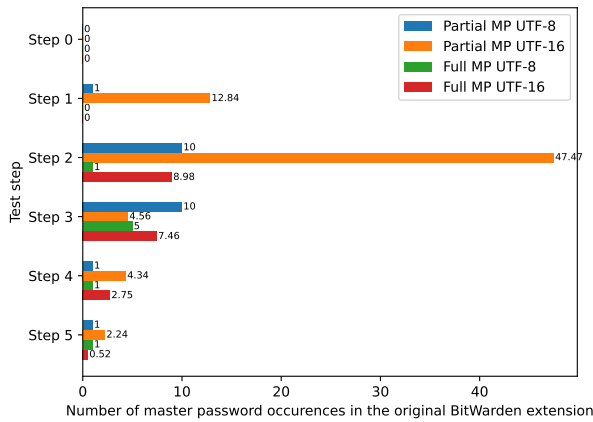


Figure 3: Number of partial and full master password occurrences in memory in the original BitWarden extension by test step

In fig. 3, we can see the results of the testing procedure performed on the original BitWarden extension. In Step 0, there are no occurrences of the MP in memory. This is to be expected, since nothing of MP has been typed yet. Also, the memory given to the process is zeroed out by Linux’s virtual memory manager before making it available[20] to the process. This means that even if the process was given an address of memory that previously stored the MP from a previous test, it would not interfere with the results of the current test, since the memory given would be wiped beforehand. In step 1 and 2, as the user is typing the MP we see the number of occurrences rise. We mainly see occurrences of the partial and full MP encoded in UTF-16. This is because Chrome’s JS engine, the V8 JS engine, implements the ECMAScript standard which states that the primitive string type is to be encoded in UTF-16[14].

As mentioned in section 2.2, it would be ideal if after the vault is unlocked (step 3), the MP would no longer be accessible in memory. However, our findings show that is not the case. There are still occurrences of the full MP and occurrences of the partial MP after logging in and unlocking the vault. Even after simulating a simple task, with the intention of increasing the system’s load and resource usage to promote memory clean up by the GC, not all occurrences were cleared. Terminating the session and logging out decreases the occurrences but does not completely erase everything. This means the MP can be obtained from memory, for an indefinite amount of time as we mentioned in section 2.6.3, even after the the user closes the BitWarden vault.

3.3.1 Observation #1: Prefixes of the Master Password in Memory

The underlying data structure responsible for storing the MP as it is being typed, is an immutable string. Since we can not modify immutable data types, a new object has to be created. So given a password p with length i , when a new character c is added a new string object is created with the result of $p_i + c$. Likewise, when a character is removed, a new object is created with the value of p_{i-1} . The previous memory addresses are no longer referred to and cleanup is left up to the GC system. As shown in fig. 3, we were able to identify memory addresses that contained prefixes of the full MP.

3.3.2 Observation #2: Prefixes stay longer in memory if the input is left untouched

As discussed in subsection 3.3.1, new string objects are created whenever the user types or deletes characters while typing the MP. In our testing, we discovered that if the MP is typed relatively quickly without much delay between keystrokes (around a second or so), the amount of prefixes in memory would decrease compared to when the user types a portion of the MP, leaves it untouched for more than a second or two and then continues typing the rest of the MP.

3.3.3 Observation #3: Unlocking the vault does not clear the master password

After unlocking the vault, the MP is no longer required to be in memory as mentioned in section 2.3.1, but continues to be, even after the user locks the vault and terminates the session.

3.4 List of problems

While analysing the BitWarden’s extension source code of the login process we compiled a list of problems within it. Lee et al.[28] found similar issues in Android applications and were able to eliminate leftover passwords from memory by solving them.

3.4.1 Problem #1: Use of immutable data types to hold the master password

In the login component of BitWarden, the MP is stored in an immutable String. Like we mentioned in section 2.6.3, application has no way to erase the content of that String, leaving the deletion of the secret up to the GC.

3.4.2 Problem #2: No zeroization of the master password

Due to the problem mentioned above, BitWarden does not go through the effort to try and zero out the content of the variable that holds the MP.

3.4.3 Problem #3: Use of String to communicate master password from interface to component

In an Angular application, the application can access the information in a native element of the DOM through a Control Value Accessor[10]. In our case, BitWarden interacts with the input field responsible for the MP in the login page to receive the MP written by the user into the login component to perform the login process. Unfortunately, Angular only provides the DefaultValueAccessor[11] in input fields of type text. This accessor provides whatever is written into the input field to the application, as a String type.

Since we decided for this project to focus only on changes that could be done at the application layer, we have to use what Angular provides. However, in section 4.2 we present how we tried to mitigate this problem.

4 Our solution

In this section, we will describe the fixes we implemented at an application-level in the BitWarden Chrome Extension to try and reduce the amount of MP copies found after logging into the BitWarden Vault.

Two versions of the program were created with one fundamental difference: one uses a separate Angular component that deals with the input of the MP and communicates with the login parent component, while the other version has the login component deal with the MP input directly, without using a separate Angular component. This way, we can compare if the Angular communication between components is a factor that could lead to password leakage in our program. Each of these versions also has a variation where we inlined a function in a critical area to see how it would affect the amount of occurrences of the MP in the memory of the process.

This gives us four different versions in total: (a) Child component; (b) Child component - inlined; (c) No child component; and (d) No child component - inlined.

4.1 Common changes between all versions

The login component performs a few steps with the MP to complete the login process. First, it checks if the MP is not empty and then delegates the login to the authentication service, passing the e-mail and MP of the user. The authentication service then passes the MP and other information to the crypto service for (a) making the Master Key (makeKey); and (b) making the Master Password Hash (hashPassword). Once the Master Password Hash is complete, it uses it and the e-mail to perform the login and the login component receives a form response and proceeds with the rest of the login process.




Figure 4: Flowchart of how the master password is used in the login component

The first change we implemented was to change the immutable variable type of the MP (string) in the LoginComponent to a mutable data structure.

In our case, we decided to use a JS ArrayBuffer [3] for a few reasons. First of all, it is a mutable data type fixing the problems specified in Problem #1 (section 3.4.1) and giving us an opportunity to address Problem #2 (section 3.4.2) as well. Secondly, BitWarden’s login code flow makes a few calls to services of the application with the MP, those being the authentication and cryptographic service as discussed above. All of these function calls already have support for the ArrayBuffer data structure. This allowed us to change the function parameters of the functions calls to accept an ArrayBuffer data type instead of the previously used string. Since we are no longer passing a primitive JS type (string) to several functions, but instead a reference type (ArrayBuffer), JS will instead pass the argument (in this case the MP) by sharing (call by sharing). This avoids making extra copies of the MP when it is passed down to a function. This addresses Problem #3 discussed in section 3.4.3. We replicated these changes in the crypto service as well.

After BitWarden performs the login process, Problem #2 is addressed as the MP is no longer required to be in memory. To do so, the ArrayBuffer containing the MP is manually overwritten before the variable is set to null to help it being marked for garbage collection. Because JS’s memory management is automatic it will either be reused throughout the execution of the program or freed up and used elsewhere, like another process running on the computer. This process might take a while, but since we cleared the buffer previously, the MP contained is no longer there.

The next change we had to implement was due to how

Angular/JavaScript reads the MP from the Document Object Model (DOM). The two different approaches (both the one using a child component and one without using one) are explained below.

4.2 Implementation with a login child component - Child component

To try and receive the MP in a more secure manner, we devised a new component responsible solely on bridging the communication between the native input element in the DOM and the login Angular Form. To do this, we implemented Angular's Control Value Accessor interface on the component to change the default behaviour. Later on, the login component of Angular receives the ArrayBuffer with the full MP from this child component and proceeds to perform the login.

4.2.1 MasterPasswordCustomInputComponent

This component is responsible for listening to input changes from the password input field coming from the DOM. Our approach consists of receiving the input through a function, and storing it in an ArrayBuffer. We also take the care to zero out the data of the previous ArrayBuffer, which contained the previous password input.

Unfortunately, the application is forced to receive the input from the DOM as a string, as it is the only text format that Angular supports. Like we said previously, this creates a copy of the input and it is also immutable, leaving the deletion up to the GC. To mitigate further damage, we do not store the input in the function and only use it when we absolutely must, to transform the input information to the ArrayBuffer.

After the user types the entire MP and clicks the login button, the child component sends the ArrayBuffer with the full MP to the login component. The login component then performs the login process as usual.

4.3 Implementation with a login child component - inlined

We use a function (*Utils.fromUtf8ToArray*) to convert the string input to an ArrayBuffer. This means we are passing a string into the function which creates an additional copy of the argument as discussed in section 2.6.5. Since this function is called every time the input field changes (whenever a new character is added or deleted), this could lead to potentially even more copies of the input being created.

We made an extension, called the child component inlined, where we inlined the *Utils.fromUtf8ToArray* function to see if it would make a significant difference in the amount of password leakage.

4.4 Implementation without a login child component

As a way to understand if the communication between Angular components could produce extra copies of the MP in memory, we also devised a version without an extra component, where the login component deals with the input of the MP. In this version, the Control Value Accessor interface was directly implemented in the login component.

4.5 Implementation without a login child component - inlined

Similarly to what we did in the child component inlined extension, we simply inlined the *Utils.fromUtf8ToArray* function.

5 Evaluation

5.1 Preparing the testing environment

A Vagrantbox was created that runs Linux Ubuntu containing Chrome, the official BitWarden extension for Chrome, our modified extensions, all the required packages for running the Python scripts and all the required packages to create a local BitWarden server. Additionally, it also sets up the local BitWarden server and starts it in preparation for the testing procedure. We originally performed tests using BitWarden's servers, however, once we automated the testing procedure we started performing the tests on our own locally hosted BitWarden server for several reasons: (a) the traffic and usage patterns performed in the testing procedure violate BitWarden's Terms of Service; (b) logging into BitWarden does not require captchas to be solved when the server is locally hosted which simplifies our testing procedure; (c) no need to create a BitWarden account in BitWarden's hosted servers; and (d) the testing procedure is no longer dependent on the availability of BitWarden's services.

Table 1 shows the relevant software and versions used in our testing.

Name	Version
Linux Ubuntu	20.04.4 LTS
Oracle VM VirtualBox	6.1.32r149290
Google Chrome	90.0.4444.51-1
BitWarden Google Chrome Extension	1.55.0, 8 Dec 2021
PyAutoGui	0.9.53

Table 1: Versions of the software used in the testing procedure

5.2 Automatic testing using Python and PyAutoGUI

To check the differences in password leakage between the normal BitWarden extension and our modified extensions, we devised a script using Python and PyAutoGui[32] to automatically perform a testing procedure to replicate a use case of the application. Automating the testing procedure made testing the different extensions easier and made the tests across the different extensions more consistent with each other, giving us more confidence over the results. We also have the advantage of being able to perform a large number of tests per extension this way, giving us more statistical relevance. Lastly, replicating how a user interacts with the graphics user interface (GUI) of the extension made the test closer to in behaviour to how a user perform a task in a use case in the BitWarden extension, as well as replicating any type of memory pattern that might happen from such behaviour, leading us to more accurate results.

The script uses PyAutoGui to perform several GUI interactions with the OS, Chrome and the BitWarden extensions (both the official one and our modified ones). The following is a summary of what the program does:

- Resets Chrome’s settings to a default known state
- Opens Chrome and points the BitWarden extension to use the locally hosted server
- Closes Chrome (this is to avoid having multiple processes with the BitWarden name which breaks our script)
- Opens Chrome again
- Clicks on the extensions icon and then the BitWarden extension
- Clicks on the log in button and types the e-mail in the e-mail field
- **Test step 0:** Performs a memory dump before writing the MP in the login page
- **Test step 1:** Types half of the MP and performs a memory dump
- **Test step 2:** Types the second half of the MP (now complete) and performs a memory dump
- **Test step 3:** Unlocks the vault and performs a memory dump once it is open
- To simulate a task, the script opens up a new tab, and plays a video on Vimeo[37] for around a minute
- **Test step 4:** After a minute has passed, it performs a memory dump again

- Clicks on the extensions icon and then the BitWarden extension
- Goes to the settings tab
- **Test step 5:** Finally, it terminates the session on the BitWarden extension, performs a memory dump and proceeds to close Chrome

5.3 Analysing the created memory dumps

To facilitate the analysis of the memory dumps, a Python script was also made, responsible for going through all the memory dumps created at the different steps of the testing procedure. It opens the memory dumps in binary mode and simply reads the memory into the program and scans the memory dump for four different things:

1. The first half of the MP in 8-bit encoding (UTF-8/ASCII) and 16-bit encoding (UTF-16)
2. The full MP in 8-bit encoding and 16-bit encoding as well

To calculate the number of partial MP occurrences, we count the amount of times the first half of the MP has appeared in memory and subtract it with the number of times the full MP was in memory as well, since the partial MP is a prefix of the full MP.

It then creates a CSV file with the results, ordered by test number. The file contains the number of occurrences of the partial and full MP, in both encodings, at the different steps of the testing procedure.

5.4 Results

What test step corresponds to which phase of the testing procedure is in section 5.2. The total number of occurrences of the full MP was calculated by summing the occurrences of the full MP in both UTF-8/ASCII encodings and UTF-16 encoding. Likewise, the same process was done for calculating the total number of partial MP occurrences, by summing the partial MP occurrences in both encodings described previously.

5.4.1 Full master password

The results of our extensions against the original BitWarden extension are shown in fig. 5.

In test step 0 and 1, the full MP is not in memory, as it has not been fully typed yet. In test step 2, when we finish typing the full MP, our extensions have more copies of the full MP in memory. However as we said in section 2.2, we concede that the MP can be in memory while the user has not performed the login process. In test step 3, when we perform the login and unlock the vault, our extensions are able to reduce the copies present in memory when compared to the

original extension. We see this trend in later steps for the other extensions as well, except in test step 4 and 5, where the no child component extension is slightly worse when compared to the original one.

The child component inlined extension is the one that performed the best right after step 3, having less full MP occurrences than the rest, but performed slightly worse than others in step 4 and 5.

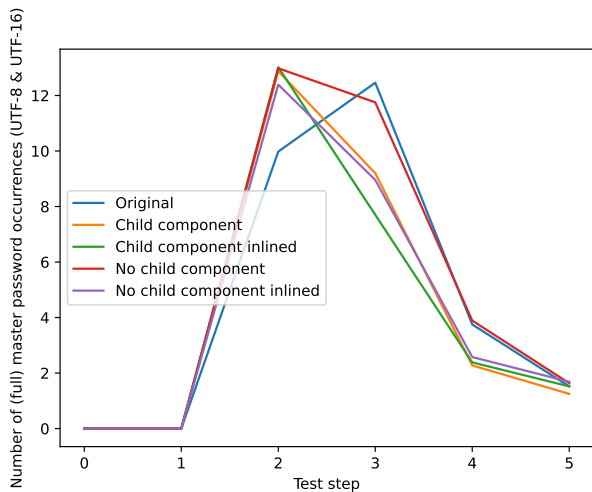


Figure 5: The number of occurrences of the full MP in memory per test step

5.4.2 Partial master password

The results of our extensions against the original BitWarden extension are shown in fig. 6

In test step 0, we see no references as the MP has not been typed yet. In test step 1, when the first half of the MP has been typed, we see similar results to the original BitWarden extensions. The rest of the test steps almost mimic the results obtained in section 5.4.1. We have an increase in occurrences in step 2 but after logging in and unlocking the vault in step 3, we see a reduction of copies in memory over all extensions, except the no child component extension. It stands to note that the no child component extension performs worse when compared to the original BitWarden extensions across the board, when it comes to the partial MP.

The child component and child component inlined extensions performed very similarly, with the child component inlined extension being slightly better in steps 4 and 5 compared to the child component extension.

5.4.3 Child component and no child component

Our no child component extension compares worse than the rest of our extensions and even BitWarden’s original extension in some cases. We find this behaviour odd, as the results




Figure 6: The number of occurrences of the partial MP in memory per test step

between the no child inlined component (where the difference between it and the no child component is simply a function being inlined) and the child component extensions (both the normal and inlined version) are much similar between each other. Without access to external tools, it is hard to say what causes this discrepancy.

6 Conclusion and Next Steps

In conclusion, even though our modified extensions were able to reduce the occurrences of the full MP in memory after logging in and unlocking the vault, they were unable to completely remove every trace of the MP, even after performing tasks on the system. As such, an attacker that gains a snapshot of memory after the vault was unlocked, is likely to be able to successfully retrieve the MP of the user.

This shows that changes in the application layer are not enough to completely eliminate leftover MP’s references in the memory of the process, and efforts must be made in every step of the stack to ensure that sensitive data is dealt with and properly disposed of to ensure the desired data security properties.

6.1 Future work

Modifying the Angular framework to create a new, more secure way to bridge communication from the native DOM to application would be something to work on. On the browser side, while Chrome is closed-source, there are open-source alternatives (like Chromium[35] and Firefox[16]) that could be changed to introduce a secure API to deal with sensitive data on the DOM so communication with web applications

could be more secure. Finally, one could change the OS in which the application is running, to ensure that communication between the OS and the other layers would follow the same secure API protocols that we desire.

Acknowledgments. This work was partially funded by the PassCert project, a CMU Portugal Exploratory Project funded by Fundação para a Ciência e Tecnologia (FCT), with reference CMU/TIC/0006/2019 and supported by national funds through FCT under project UIDB/50021/2020.

References

- [1] *65% of people don't trust password managers despite 60% experiencing a data breach*. July 2020. URL: <https://www.passwordmanager.com/password-manager-trust-survey/>.
- [2] Nora Alkaldi and Karen Renaud. "Why do people adopt, or reject, smartphone password managers?" In: (2016).
- [3] *ArrayBuffer - JavaScript*. URL: https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Global_Objects/ArrayBuffer (visited on May 10, 2022).
- [4] Salvatore Aurigemma, Thomas Mattson, and Lori Leonard. "So much promise, so little use: What is stopping home end-users from using password manager applications?" In: (2017).
- [5] *BitWarden Open Source Password Manager*. URL: <https://bitwarden.com/> (visited on May 10, 2022).
- [6] *BitWarden Security Whitepaper*. URL: <https://bitwarden.com/help/bitwarden-security-white-paper/> (visited on May 10, 2022).
- [7] *Blink (Rendering Engine)*. URL: <https://www.chromium.org/blink/> (visited on May 10, 2022).
- [8] Joseph Bonneau, Cormac Herley, Paul C Van Oorschot, and Frank Stajano. "The quest to replace passwords: A framework for comparative evaluation of web authentication schemes". In: *2012 IEEE Symposium on Security and Privacy*. IEEE. 2012, pp. 553–567.
- [9] Sonia Chiasson, Paul C van Oorschot, and Robert Biddle. "A Usability Study and Critique of Two Password Managers." In: *USENIX Security Symposium*. Vol. 15. 2006, pp. 1–16.
- [10] *Control Value Accessor Angular*. URL: <https://angular.io/api/forms/ControlValueAccessor> (visited on May 10, 2022).
- [11] *DefaultValueAccessor Angular*. URL: <https://angular.io/api/forms/DefaultValueAccessor> (visited on May 10, 2022).
- [12] Matteo Dell'Amico, Pietro Michiardi, and Yves Roudier. "Password strength: An empirical analysis". In: *2010 Proceedings IEEE INFOCOM*. IEEE. 2010, pp. 1–9.
- [13] Zakir Durumeric, Frank Li, James Kasten, Johanna Amann, Jethro Beekman, Mathias Payer, Nicolas Weaver, David Adrian, Vern Paxson, Michael Bailey, et al. "The matter of heartbleed". In: *Proceedings of the 2014 conference on internet measurement conference*. 2014, pp. 475–488.
- [14] *ECMAScript 2023 Language Specification*. URL: <https://tc39.es/ecma262/#sec-literals-string-literals> (visited on May 10, 2022).
- [15] Phoebe Fasulo. Aug. 2018. URL: <https://securityscorecard.com/blog/cybersecurity-data-breaches-statistics-on-the-rise>.
- [16] *Firefox Browser*. URL: <https://www.mozilla.org/> (visited on May 10, 2022).
- [17] *Functions - JavaScript*. URL: <https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Functions> (visited on May 10, 2022).
- [18] *Generate a kernel or complete crash dump - Windows*. URL: <https://docs.microsoft.com/en-us/windows/client-management/generate-kernel-or-complete-crash-dump> (visited on May 10, 2022).
- [19] *Google Chrome*. URL: <https://www.google.com/chrome/> (visited on May 10, 2022).
- [20] Mel Gorman. *Understanding the Linux virtual memory manager*. Prentice Hall Upper Saddle River, 2004.
- [21] Miguel Grilo, João F. Ferreira, and José Bacelar Almeida. "Towards Formal Verification of Password Generation Algorithms used in Password Managers". In: *arXiv preprint arXiv:2106.03626* (2021).
- [22] J Alex Halderman, Seth D Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A Calandrino, Ariel J Feldman, Jacob Appelbaum, and Edward W Felten. "Lest we remember: cold-boot attacks on encryption keys". In: *Communications of the ACM* 52.5 (2009), pp. 91–98.
- [23] *Java™ Cryptography Architecture (JCA) reference guide*. URL: <https://docs.oracle.com/javase/6/docs/technotes/guides/security/crypto/CryptoSpec.html#PBEEEx>.
- [24] B. Kaliski. 2000. URL: <https://tools.ietf.org/html/rfc2898>.
- [25] Stavroula Karayianni, Vasilios Katos, and Christos K Georgiadis. "A framework for password harvesting from volatile memory". In: *International Journal of Electronic Security and Digital Forensics* 7 4.2-3 (2012), pp. 154–163.

- [26] *Kernel Crash Dump*. URL: <https://ubuntu.com/server/docs/kernel-crash-dump> (visited on May 10, 2022).
- [27] Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, et al. “Spectre attacks: Exploiting speculative execution”. In: *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 1–19.
- [28] Jaeho Lee, Ang Chen, and Dan S Wallach. “Total Recall: Persistence of Passwords in Android.” In: *NDSS*. 2019.
- [29] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, and Mike Hamburg. “Meltdown”. In: *arXiv preprint arXiv:1801.01207* (2018).
- [30] Sean Oesch and Scott Ruoti. “That Was Then, This Is Now: A Security Evaluation of Password Generation, Storage, and Autofill in Browser-Based Password Managers”. In: *Proc. of USENIX Security Symp.* 2020.
- [31] *Password Managers’ Secrets Management: ISE*. Oct. 2020. URL: <https://www.ise.io/casestudies/password-manager-hacking/>.
- [32] *PyAutoGui’s documentation*. URL: <https://pyautogui.readthedocs.io/en/latest/> (visited on May 10, 2022).
- [33] Shannon Riley. “Password security: What users know and what they actually do”. In: *Usability News* 8.1 (2006), pp. 2833–2836.
- [34] Elizabeth Stobert and Robert Biddle. “A password manager that doesn’t remember passwords”. In: *Proceedings of the 2014 New Security Paradigms Workshop*. 2014, pp. 39–52.
- [35] *The Chromium project*. URL: <https://www.chromium.org/chromium-projects/> (visited on May 10, 2022).
- [36] *V8 JavaScript engine*. URL: <https://v8.dev/> (visited on May 10, 2022).
- [37] *Vimeo*. URL: <https://vimeo.com/> (visited on May 10, 2022).
- [38] Chun Wang, Steve TK Jan, Hang Hu, Douglas Bossart, and Gang Wang. “The next domino to fall: Empirical analysis of user passwords across online services”. In: *Proceedings of the Eighth ACM Conference on Data and Application Security and Privacy*. 2018, pp. 196–203.