



HAL
open science

Robust Stack Smashing Protection for WebAssembly

Quentin Michaud, Yohan Pipereau, Olivier Levillain, Dhouha Ayed

► **To cite this version:**

Quentin Michaud, Yohan Pipereau, Olivier Levillain, Dhouha Ayed. Robust Stack Smashing Protection for WebAssembly. IEEE Future Networks World Forum 2024, Oct 2024, Dubai, United Arab Emirates. hal-04888542

HAL Id: hal-04888542

<https://hal.science/hal-04888542v1>

Submitted on 15 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Robust Stack Smashing Protection for WebAssembly

Quentin Michaud^{*†}, Yohan Pipereau[†], Olivier Levillain[†], Dhouha Ayed^{*}

^{*}Thales Group, Palaiseau, France

[†]SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France

Abstract—WebAssembly is an instruction set architecture and binary format standard, designed for secure execution by an interpreter. This technology is identified as an alternative for current containerization technologies that is suitable for secure and lightweight orchestration for 5G/6G environments. Previous work has shown that WebAssembly is vulnerable to buffer overflow due to the lack of effective protection mechanisms.

In this paper, we evaluate the implementation of Stack Smashing Protection (SSP) in WebAssembly standalone runtimes, and uncover two weaknesses in their current implementation. The first one is the possibility to overwrite the SSP reference value because of the contiguous memory zones inside a WebAssembly process. The second comes from the reliance of WebAssembly on the runtime to provide randomness in order to initialize the SSP reference value, which impacts the robustness of the solution.

We address these two flaws by hardening the SSP implementation in terms of storage and random generator failure, in a way that is generalizable to all of WebAssembly. We evaluate our new, more robust, solution to prove that the implemented improvements do not reduce the efficiency of SSP.

Index Terms—WebAssembly, Memory bugs, Stack Smashing Protection

I. INTRODUCTION

WebAssembly [16], [12] has been created as a fast and secure-by-design answer to the always increasing need for complex computation in browsers, e.g. 3D workloads.

The success of WebAssembly as a portable Instruction Set Architecture (ISA) and binary format has prompted its adoption in many applications besides browsers, such as its use as a standalone runtime [9]. This has a huge impact on the cloud world and the computing world in general, and is considered as an alternative to Linux-based containers [6], [3], promising to be more portable, lightweight and secure. Some see in the flexibility of WebAssembly a universal binary format that could be distributed seamlessly across operating systems and hardware architectures. As such, WebAssembly is proposed as a novel approach to 5G/6G service orchestration as its properties make it suitable for edge and far edge environments [4], [10].

WebAssembly claims strong security. By default, it provides sandboxing between different WebAssembly instances and

between WebAssembly and its host. It also enforces control-flow integrity, and protection against code reuse attacks. However, the security of WebAssembly has been challenged in several works [17], [14]. First, WebAssembly offers weak protection against memory corruption attacks compared to native binaries. Some vulnerabilities, such as stack-based buffer overflows, have been present in native binaries for a long time, but are mitigated with mechanisms such as Stack Smashing Protection (SSP). This protection was initially absent in WebAssembly [2]. Second, differences in design between WebAssembly and native binaries make the former vulnerable to attacks that are not possible in native binaries. One example is the corruption of heap data using a stack-based buffer overflow.

Stack Smashing Protection has been implemented in WebAssembly after the publication of the papers discussed in the previous paragraph. In this paper, we propose the following contributions: (i) a thorough analysis of SSP in WebAssembly; (ii) some proofs of concept to confirm the weaknesses of the current implementation; (iii) the implementation of a more robust SSP mechanism in LLVM [1] and `wasi-libc`;¹ (iv) an evaluation of our solution.

We open-source all our code contributions: the implementation of SSP in the WebAssembly target of the LLVM compiler;² modifications to `wasi-libc`;³ the adaptation of *CookieCrumblor* [7] (a tool used to assess the robustness of SSP implementations) to WebAssembly; and our proofs of concept.⁴

This paper is structured as follows. First, Section II and III present necessary background and motivation for this work. Then, Section IV contains our security analysis of WebAssembly SSP and our remediation proposals. Finally, Section V provides an evaluation of our work and Section VI concludes and gives some perspective for future work.

II. BACKGROUND

We start by giving a brief introduction to buffer overflows and Stack Smashing Protection. We also provide a quick background on WebAssembly and its inner workings.

This paper has received funding from by the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101139067. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

¹<https://github.com/WebAssembly/wasi-libc/>

²<https://github.com/ThalesGroup/llvm-project/tree/new-wasm-ssp>

³<https://github.com/ThalesGroup/wasi-libc/tree/new-wasm-ssp>

⁴<https://github.com/mh4ck-Thales/Robust-SSP-in-Wasm>

A. Buffer overflow and Stack Smashing Protection

Buffer overflows are an old and well-known vulnerability [15]. They occur when a program stores more data in a buffer than the buffer may hold. Writing to memory out of buffer bounds leads to the corruption of memory adjacent to the buffer. Buffer overflows may also happen during the execution of a WebAssembly program.

Stack Smashing Protection (SSP), also known as *stack canaries* or *stack cookies* [11] is a defense mechanism available to prevent exploitation of stack-based buffer overflows. SSP provides a detection mechanism for stack-based buffer overflows and terminates the execution of the program after the current function is executed. At program start time, the program initializes a random reference value (named *canary* or *cookie*) and writes it in a memory zone, preferably where overwrite is made impossible, or at least difficult. Each time a function is called, the function prologue is executed which creates a new stack frame and copies the canary reference value in the stack, in a dedicated variable, the *stack canary* located after the buffer. The function epilogue checks this value against the canary reference value stored in safe memory. If the stack canary is different from the reference value, it means that the stack canary has been overwritten and that a stack-based buffer overflow has occurred. In this case, a specific function is called to terminate the process.

Stack Smashing Protection is implemented in two different code bases. The initialization of the reference value and the function called when the stack canary is overwritten is provided in the language standard library (e.g., the GNU C standard library or the musl C standard library). The generation of the specific function prologue and epilogue for setting up and verifying the integrity of canaries is implemented in the compiler.

B. WebAssembly

WebAssembly (commonly abbreviated as Wasm) is a binary format, designed to be compact, easy to parse and fast at execution. A WebAssembly file, containing a WebAssembly program, is named a *module*. An *instance* is a module being executed in a runtime. WebAssembly is also an Instruction Set Architecture (ISA), designed as a stack-based virtual machine. It was designed to be fast and secure by design.

WebAssembly bytecode is executed using a stack-based Virtual Machine (VM). This means that each instruction gets input operands by popping values off a stack, and pushes its eventual results on this stack referred as the *evaluation stack*. There are no registers in the WebAssembly virtual machine. The WebAssembly bytecode is located in a specific memory managed by the virtual machine, that is read to execute an instruction, but that is not directly accessible by the program.

The WebAssembly virtual machine relies on multiple memory regions which are represented in Fig. 1.

The *managed code memory* contains the WebAssembly program code. It is only accessible by the VM, so the WebAssembly code cannot read or modify it.

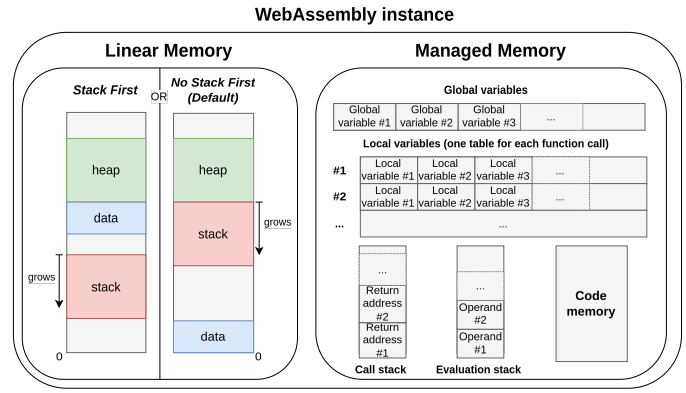


Fig. 1. The memory layout of a WebAssembly virtual machine

The *managed call stack* contains return addresses. These return addresses are of WebAssembly's `i32` type, which is used as the type for memory pointers and addresses.⁵ It is used to keep track of all the ongoing function calls, while preventing control-flow hijacking based on return address overwrite.

The *managed evaluation stack* is used to give parameters to instructions and to store their results. This stack can hold the four WebAssembly basic types, i.e. `i32`, `i64`, `f32` and `f64` that are respectively integers and floating point numbers encoded on 32 or 64 bits.

The *linear memory* is used to store non-scalar types, e.g., strings, arrays, or lists. This linear memory is a single contiguous memory segment with no notion of memory permissions. As such, all memory in the linear memory is readable and writable. The management of this memory is left to the program, but for most programming languages and their respective compilers, the structure used is the same as the one widely used in native binaries, which includes a stack, a heap and a data zone for static or predetermined values. These memory zones contain most of the data used by the program, the data being distributed between the different zones according to the source code and the compiler used.

These three zones can be arranged in several ways in memory, and in practice different WebAssembly compilers made different choices resulting in different layouts. For the purpose of this article, we focus on the two layouts available with the LLVM toolchain, named *stack-first* and *no-stack-first*. These layouts are represented in Fig. 1.

LLVM default memory layout is *no-stack-first*. Later, *stack-first* was introduced to address the problem of silent data corruption by a stack overflow in *no-stack-first*. As of today, *stack-first* has been adopted by default in Rust, in Zig, and LLVM is discussing to make it a default.

The *WebAssembly local and global variables* are another memory mechanism. As for the evaluation stack, they are restricted to the four WebAssembly basic types. The scope of global variables is the entire module, while local variables are only accessible by the function being executed. These variables are manipulated through dedicated instructions and

⁵A proposed extension of WebAssembly is using the `i64` type to address the memory, but this extension is not addressed in this paper.

are stored in a specific table that is not accessible from the linear memory. It is however important to notice that current toolchains do *not* usually map local and global variables found in programming languages onto WebAssembly local and global variables.

By design, WebAssembly does not provide access to the host environment in which the WebAssembly program is executed. It can only be performed using functions provided by the WebAssembly runtime, that will then interact with the host environment as requested and store the results in linear memory or in the evaluation stack, as an internal WebAssembly function would do. It is up to the runtime to implement or not these special functions. In order for a WebAssembly binary to work with a large panel of runtimes and host environments, standardizing such special functions was needed.

For standalone WebAssembly, this led to the creation of the WebAssembly System Interface (WASI). It is composed of a set of modular standards regrouped around different use cases.

WASI [5] is still evolving as a standard, but it is already widely used. Two main versions of WASI exist to this day: WASI preview 1 (WASIp1), released in late 2020 and WASI preview 2 (WASIp2), released in the beginning of 2024. This paper along with its proposed proofs of concept is using WASIp1, as most research was done before the publication of WASIp2. However, we believe that the work conducted in this paper is applicable to WASIp2 as well.

III. MOTIVATION AND RELATED WORK

A. *WebAssembly lack of memory protection*

WebAssembly security has already been studied in several works. Lehmann et al. [14] conduct an in-depth security analysis of the WebAssembly linear memory, and how it is used by programs compiled from various languages. It shows that common memory protections are missing from WebAssembly, and demonstrates how this lack makes code less secure than when compiled to a native binary. It concludes by discussing some mitigations, including the proposition to port protections provided by compilers to WebAssembly. One of these mitigations is Stack Smashing Protection. Our first proof of concept, corresponding to the `-no-ssp` files in our artifact repository, is inspired by their work and proves that buffer overflows in standalone WebAssembly are exploitable in practice.

However, the effectiveness of Stack Smashing Protection in WebAssembly is not guaranteed due to the great differences between WebAssembly and native binaries, and the security of its implementations has not been assessed yet. Other propositions of mitigation mainly require significant work in the WebAssembly specifications and its extensions, and thus have not been adopted yet.

In [17], Stiévenart et al. study a corpus of thousands of C programs vulnerable to stack-based and heap-based buffer overflows. They compare the behavior of these programs when they are compiled as x86 binaries with state-of-the-art protections (including Stack Smashing Protection) and

WebAssembly binaries, that did not have Stack Smashing Protection at the time of the study. They observe that x86 binaries are subject to many crashes, for the most part triggered by SSP. On the contrary, WebAssembly binaries are continuing execution after the buffer overflow and memory corruption most of the time.

The difference is attributed to the absence of SSP in WebAssembly binaries, which allows an attacker to exploit buffer overflows in a stealthier fashion. This means that WebAssembly binaries are more vulnerable to memory corruption due to buffer overflows than native ones. At least, it means that WebAssembly binaries can see their internal memory corrupted and their data integrity violated. In the worst case, it may be the enabler of more complex and dangerous attacks on WebAssembly (such as attacking the WebAssembly VM), as exemplified by [14].

Since Stiévenart et al. work, SSP has been implemented in a subset of WebAssembly using LLVM and `wasi-libc`. This means that there is no SSP available in non-WASI WebAssembly binaries, such as in the browser or depending on Node.js. However, it could still be implemented in toolchains of these other environments using our work as the base for a secure implementation.

Zhang et al. [18] propose *VMCanary*, an alternative implementation of SSP for all of WebAssembly. However, *VMCanary* relies on an extension of the ISA and thus is non-standard, making it incompatible with current WebAssembly runtimes and tooling. On the contrary, our work is based on the existing implementation in LLVM and `wasi-libc`, building on a solution which is fully compliant with the WebAssembly specifications. Our solution has no adherence with any WebAssembly tooling, and its principles can be extended to other toolchains without breaking compatibility.

These papers conclude that WebAssembly is lacking protections that are present in native binaries. Some security features are included in the design of WebAssembly, but there are no guarantees that they fulfill the role of the protections that are missing. The introduction of Stack Smashing Protection on the WebAssembly world can be seen as an improvement, but its effectiveness has not been assessed yet.

B. *Global impact of memory corruption and protections*

In addition to assessing the possibilities of memory corruption in WebAssembly, Lehmann et al. [14] analyze the impacts of such corruption. Their work places buffer overflow vulnerabilities as a way to potentially gain further, more impactful attack primitives. For example, considering a program that reads and writes from and to different files, overwriting the memory contents may allow the attacker to modify a filename and thus trigger an arbitrary file read or write.

Another possibility of exploit is using restricted control flow hijacking, by abusing the `call_indirect` instruction of WebAssembly. This instruction allows WebAssembly to support function pointers, which are required when the compiler cannot statically determine the exact function to call (e.g. callback functions, dynamic methods in object-oriented

programming). This makes the implicitly enforced control-flow integrity in WebAssembly weaker in the case of indirect calls. As a result, the attacker may be able to control the function that will be called, and thus control the code that will be executed.

Hilbig et al. [13] study a dataset of more than 8,000 WebAssembly binaries collected from various sources in late 2020. Among other research questions, they investigate which tools and source languages are used to produce WebAssembly binaries. This question is more and more relevant as the popularity of WebAssembly grows and WebAssembly binaries are increasingly used in new domains. More specifically, as the tools and use cases for WebAssembly diversify, the work needed to spread the new security propositions for WebAssembly becomes longer and longer.

One of the findings of Hilbig et al. is that 64.2% of WebAssembly binaries are written in C or C++, which are memory-unsafe languages. This strongly suggests that the work on assessing memory safety in WebAssembly is relevant. Furthermore, it underlines the importance of Stack Smashing Protection for the global security of WebAssembly binaries and the WebAssembly ecosystem.

Another finding is that nearly 80% of all collected binaries are compiled with the help of the LLVM toolchain. Thus, implementing a security mechanism in LLVM, such as Stack Smashing Protection, would allow introducing increased protection in most WebAssembly programs without additional engineering efforts.

C. Potential weaknesses in SSP implementations

Implementing SSP does not mean that a binary is fully protected against stack-based buffer overflows. SSP can be bypassed when the underlying assumptions are not met. Indeed, weak implementations of SSP allow an attacker to target the SSP mechanism in order to exploit a stack-based buffer overflow undetected.

Bierbaumer et al. [7] conduct an analysis of the implementation of SSP across various platforms (OS, architectures and libraries) to identify potential implementation weaknesses. They propose a list of security properties that robust SSP implementations should satisfy, and a framework named *CookieCrumbler* to automatically assess the implementations.

The authors assume a buffer overflow that is contiguous and located from a buffer in the stack. This means that the overflow does not allow the attacker to skip some bytes in memory. The security properties that robust SSP implementations should satisfy are as follows:

- P1** The canary value placed behind user-controlled buffers must be unknown to the attacker.
- P2** The reference value is placed at a location in memory that is distinct from the location of canaries and ideally mapped read-only.
- P3** If a canary is corrupted, the program execution terminates immediately (or as soon as possible) without accessing any attacker controlled data.

The main goal of Bierbaumer et al. was to prove these properties wrong due to implementation weaknesses. Their findings show that the robustness of SSP implementations is heterogeneous, and that some implementations are indeed vulnerable, allowing an attacker to completely bypass the protection. Making the same analysis for the implementation of SSP in WebAssembly is interesting, as no such work has been done to the best of our knowledge.

In addition, the work of Bierbaumer et al. was mainly targeting x86 binaries, alongside a few other results on other platforms such as ARM or PowerPC. The inner workings of these native platforms are very far from the one of WebAssembly. Therefore, the implementation of Stack Smashing Protection may differ a lot from the ones of native platforms, and the evaluation of the security of such an implementation is even more relevant.

IV. SECURITY ANALYSIS OF WEBASSEMBLY SSP

A. Description of existing WebAssembly SSP and methodology

The implementation of SSP cannot be uniform across the whole ecosystem of WebAssembly. More precisely, an SSP implementation in WebAssembly relies on three elements, that are dependent on the target use. The *compiler*, that will provide the code for loading and checking the canaries; a *library*, that will provide the code for initializing the canary reference value and the abort function that is called if a canary is overwritten; and the *host environment*, as by nature, SSP needs randomness, that WebAssembly cannot provide by itself, so it is reliant on the host and on the way it can access or request resources from the host.

To the best of our knowledge, there is only one existing implementation of SSP in WebAssembly. This implementation relies on both LLVM (providing the compiler) and *wasi-libc* (a C standard library targeting WASI). As such, it is restricted to standalone WebAssembly.

In order to assess the robustness of the Stack Smashing Protection implementation, we use the properties introduced in Section III-C. These criteria can be evaluated independently. We use several methods to assess each of the properties, including source code analysis, disassembly of compiled binaries, and the *CookieCrumbler* tool provided by the authors.

B. Evaluating the generation of canaries

We first assess whether the reference value is unknown to the attacker (property **P1**). Reference values are generated using (pseudo-)randomness. However, not all randomness guarantees a complete unpredictability. Furthermore, one may wonder if the attacker can alter the generation of randomness, and thus compromise the generation of the reference value.

In standalone WebAssembly, the randomness is provided using WASI. At the time of writing, the *wasi-libc* only supports WASIp1. In this version, randomness can be acquired from the host using the `random_get` function. `random_get` is able to return an error code if it is not able to provide randomness. In the following paragraphs, we detail

how this method changes across the different underlying host platforms.

We can first assess the behavior of `wasi-libc` if `random_get` returns an error code. The code initializing the reference value is present in the `init_ssp` function, whose relevant extracts of the source code is available in Fig. 2.

```
void __init_ssp(void *entropy)
{
    if (entropy) memcpy(&__stack_chk_guard,
        ↪ entropy, sizeof(uintptr_t));
    else __stack_chk_guard =
        ↪ (uintptr_t)&__stack_chk_guard * 1103515245;
```

Fig. 2. Extract of the `init_ssp` C function

In this listing, the `entropy` variable contains either 0 if the return code of `random_get` is different of zero, or a pointer to the generated randomness otherwise. We can see in the code that if `random_get` returns an error code, the reference value is set to a deterministic value. Indeed, dereferencing the `__stack_chk_guard` variable will always return the same value, as in WebAssembly there is no randomization of the memory addresses. Each variable is thus stored at the exact same memory location at each execution. This location can be extracted directly from the WebAssembly binary before execution. If the attacker does not have access to the WebAssembly binary, it can also be easily bruteforced. Along with the code, this value can be multiplied with the 1103515245 constant to obtain the reference value. This means SSP in WebAssembly is fragile, as a failure to get randomness through the `random_get` function will systematically result in a predictable reference value.

However, we do not know in which situations the `random_get` function may return an error code. This does not depend on the `wasi-libc` source code, and as such we need to consider the software used to provide randomness to WebAssembly as an indirect part of the Stack Smashing Protection. The implementation of how the `random_get` WASI function is providing randomness depends on the WebAssembly runtime, and subsequently the host. As such, the Stack Smashing Protection in WebAssembly is inherently dependent on the runtime implementation.

In order to further assess the robustness of the SSP implementation in WebAssembly, we need to evaluate the implementation of runtimes. Evaluating thoroughly runtimes and hosts is impractical due to the great amount of possibilities. In order to get a glimpse of the attacking possibilities, we choose to evaluate the most common WebAssembly standalone runtimes on a classic Linux machine. We evaluate the robustness of the implementations using two methods:

- M1** We block the `getrandom` Linux syscall that is commonly used to acquire randomness on Linux.
- M2** In addition to **M1**, we block all access to the `/dev` folder on Linux, which contains the other common source of randomness, the `/dev/urandom` block device.

To assess whether the implementations of various Linux runtimes are correctly providing randomness, we use a simple C program compiled to WebAssembly, that displays the value of the reference value. This value is supposed to change at each execution of the program. If the value repeats itself throughout several executions, it means that the implementation is not able to provide randomness and is returning an error with `random_get`.

We describe here the methodology used to assess the robustness of runtimes regarding their implementation of `random_get`. The experimentation was made on the latest version available of the most popular standalone WebAssembly runtimes at the time of the experiment. The machine used was running Arch Linux with a Linux kernel of version 6.8.5, but the experiment is not dependent on the operating system nor the Linux version up to a point, and should be reproducible in any recent Linux distribution.

The mentioned files (`poc.c` and `seccomp.c`), along with the detailed commands used for the experiment, are made available in our GitHub repository.⁶

These testing methods are simulating potential attacks, misconfigurations, or other cases. For example, a WebAssembly runtime in a hardened container may have restricted access to some Linux resources available in the `/dev` folder.

For each configuration, we execute our test program twice. If the reference value holds the same value, it means that the implementation is not able to provide randomness, and thus returns an error with `random_get`. This situation is marked with **X**. If the reference value holds a different value, it means that `random_get` returned randomness. This does not mean that the provided randomness is secure, merely that the runtime chose to provide randomness and not return an error. This situation is marked with **✓**. The results of this evaluation are presented in Table I.

TABLE I
SUMMARY OF THE DIFFERENT CONFIGURATIONS W.R.T. THE ACCESS TO RANDOM SOURCES

Test configuration	Baseline	M1	M2
wasmtime (v19.0.1)	✓	✓	crash
wasmedge (v0.13.5)	✓	✓	✓
wasmer (v4.2.8)	✓	✓	crash
iwasm (v1.3.2)	✓	X	X
wasm3 (v0.5.0)	✓	X	X
wasmi (v0.31.2)	✓	✓	crash

Two runtimes out of the six tested are failing to provide randomness with the situation **M1**. In situation **M2**, the same runtimes are failing to provide randomness, along with three more runtimes that are crashing when trying to provide randomness in this situation. The remaining runtime is seemingly able to provide randomness. However, further inspection of the source code is required to ensure the quality of the returned randomness.

Reconsidering the global problem again, we find that the shifting of randomness acquisition from the host (through the `libc`) to the runtimes may be a problem for the robustness

⁶<https://github.com/mh4ck-Thales/Robust-SSP-in-Wasm/>

of the SSP implementation. Most tested runtimes are either unable to provide randomness, triggering `wasi-libc` to use a predictable value, or are crashing when trying to provide randomness. One may argue that crashing, at least, does impeach the potential exploitation of weak SSP. However, a runtime crash is not desirable, especially as `random_get` has the possibility to return an error, letting the `wasi-libc`, and as such the program, handle such a case.

P1 is depending both on `wasi-libc` and on the runtime. We conclude that the `wasi-libc` implementation is weak if runtimes are failing to provide randomness, and that several runtimes do in fact fail to provide randomness in some situations. **P1** is thus not verified in several of the tested runtimes.

C. Evaluating the SSP reference value location

In this part, we assess the property **P2** which states that the location of the reference value must not allow for a bypass of the SSP. Indeed, if the reference value can be overwritten by a buffer overflow, this can be used to bypass the canary protection. The attacker just needs to overwrite both the canary and the reference value to the same value. Two properties can be used to protect against such an attack:

P2a The reference value is not located in a position that is accessible with the overflow of the target buffer.

P2b The memory in which the reference value is located is not writable, or some memory between the buffer and the reference value is not writable.

In order to assess **P2a**, we modified the *CookieCrumbler* tool from Bierbaumer et al. [7] for WebAssembly. The functionalities allowing to check if the range between the buffer and the reference value is writable and the code testing the threads were removed, as they are not relevant to WebAssembly.

We execute *CookieCrumbler* compiled with the `clang` LLVM compiler in both the *stack-first* and *no-stack-first* layouts. The results are presented in Fig. 3.




Fig. 3. The results of *CookieCrumbler* in the *stack-first* and *no-stack-first* layouts

We draw the following conclusions:

- A buffer overflow in the *no-stack-first* situation cannot access the reference value, but it is important to note that a stack overflow could. **P2a** is thus verified in the *no-stack-first* layout.

- With the *stack-first* layout, a buffer overflow from any memory zone, as soon as the overflow is long enough, can overwrite the reference value and bypass the canary protection. **P2a** is thus not verified in the *stack-first* layout.

Regarding **P2b**, the very design of the WebAssembly linear memory makes it impossible to verify this condition. Indeed, with the lack of memory permissions in WebAssembly, all addresses in the linear memory are writable. This makes the mapping of the memory containing the reference value as read-only impossible. Likewise, all the addresses located between the buffer and the reference value are guaranteed to be writable.

P2b is thus not verified in both the *stack-first* and *no-stack-first* layouts. Consequently, **P2** is not verified in both layouts. However, the two layouts are not equal in terms of robustness. While the *stack-first* layout does not verify **P2** at all, the *no-stack-first* layout does not allow an overwrite of the reference value with a stack-based buffer overflow. This layout may still be exploited using another attack primitive alongside the stack-based buffer overflow, but this is a more complex attack.

D. Evaluating quick termination on canary corruption

This part is assessing if the Stack Smashing Protection mechanism is aborting quickly in case of a canary corruption, i.e. **P3**. If the canary value is corrupted, data in the linear memory is probably corrupted as well. This means that the program must abort as soon as possible in order to prevent the use of corrupted data. In all SSP implementations, the detection of canary corruption is made at the end of each function. Thus, the detection of a memory corruption is bounded by the duration of the execution of the current function.

The abort procedure is implemented in `wasi-libc`, more precisely in the `__stack_chk_fail` function. Its source code is shown in Fig. 4.

```
void __stack_chk_fail(void)
{
    a_crash();
}
```

Fig. 4. The `__stack_chk_fail` C function

To verify that the `a_crash` function is indeed aborting as soon as possible, we disassemble the compiled `__stack_chk_fail` WebAssembly function to get its assembly code in the WebAssembly Text (WAT) format, shown in Fig. 5.

```
(func $__stack_chk_fail (type 7)
  unreachable
  unreachable
)
```

Fig. 5. Disassembly of the `__stack_chk_fail` WebAssembly function

This function is called directly as soon as the corruption is detected. By inspecting the code, we can see that the

function seems to abort the program directly, by executing a WebAssembly `unreachable` instruction. Thus, this SSP implementation aborts immediately once the canary value is detected as corrupted. We conclude that **P3** is verified.

E. Main findings and remediation proposals

Among the three criteria given to assess the robustness of an SSP implementation, only the quick termination criteria **P3** is verified by the SSP implementation in standalone WebAssembly. The criteria on the unpredictability of the canary value **P1** can be violated in some WebAssembly runtimes which do not crash when access to the host random number generator is impossible. This lack of randomness can be used to guess the canary value.

Sadly, there is no reliable way to prevent against a weak randomness if it is coming from the host or the runtime. However, if the runtime is correctly implemented, it should return an error with `random_get` if it detects that the host or itself is not able to provide strong enough randomness. The library is then in charge of dealing with the error.

To deal with an error from the `random_get` function, the library may try to call the function later. However, this is not generally a relevant approach since it often comes from a permanent failure situation.

Developers might be tempted to generate a random value themselves from the library, but they would need to find another source of randomness using WASI, which seems improbable. Falling back on using the current time, despite being a popular idea, is *not* a robust solution.

This is why we believe the only acceptable course of action when `random_get` fails is to abort the program during its preamble, thus avoiding running a program with a weak SSP. While this stance may be controversial on availability and practical considerations, it is the only safe way to enforce security against a weak randomness coming from the host or the runtime.

The criteria on safe location of canary reference value **P2** is also violated since the WebAssembly SSP reference value is stored in linear memory without protection against a vulnerable stack buffer. The WebAssembly linear memory does not allow to store the canary reference value safely, as it may always be overwritten no matter where it is stored. As a result, it is necessary to store the canary reference value in another WebAssembly memory region. Moreover, we need to be able to access to this value from the whole WebAssembly module.

Global variables are the only memory mechanism that meet these requirements. They can only be accessed using WebAssembly instructions, and they are stored in a safe, VM-managed memory. Thanks to the WebAssembly protections, an attacker cannot execute arbitrary code to try and access the canary reference value.

The weaknesses found in this analysis are exploitable in practice, as our second proof of concept, corresponding to the files ending with `-ssp` in the artifact repository⁷, illustrates.

⁷<https://github.com/mh4ck-Thales/Robust-SSP-in-Wasm>

To protect against such attacks, we implemented our remediation proposals in the LLVM and `wasi-libc` projects. This modified toolchain is the one evaluated in the following section.

V. EVALUATION

In this section, we propose to evaluate the efficiency of our implementation of SSP in WebAssembly. We use an approach similar to Stiévenart et al. [17] which compares the execution of programs of the Juliet test suite v1.3 [8]. The Juliet test suite is a large collection of vulnerability scenarios written in C and organized by MITRE CWE numbers. In our experiment, we only analyze CWE121 and CWE122 tests which respectively correspond to stack-based and heap-based buffer overflows. We observe the root cause of crashes in the test and classify them in four categories: *silent execution*, *memory fault*, *SSP fault*, *timeout*. A *silent execution* is an execution which terminates without a crash. Since all executions lead to an out-of-bound write operation, a silent execution corresponds to a failure to detect a buffer overflow. A *timeout* occurs as some programs never terminate, which forces us to use a timeout value of 20 seconds. A *memory fault* is an execution aborted by a memory fault such as `SEGV` or `SIGBUS`. An *SSP fault* is a crash triggered by the SSP mechanism.

In our experiment, we consider five configurations selected according to two parameters. The first parameter is whether the binary is a native x86 binary or a WebAssembly binary. The second parameter is the presence or absence of SSP. In all configurations, we use LLVM with `clang` and `clang++` compilers in version 17. WebAssembly configurations use the `wasmtime` runtime and `wasi-sdk` in version 21. We focus exclusively on the *stack-first* memory layout after observing that using the default memory layout of LLVM or *stack-first* yields similar results.

1) *Observations*: The results of our experiment are presented in Fig. 6. For CWE 121, we observe that 24% of crashes are caused by memory faults for WebAssembly with SSP disabled. In x86 binaries using SSP, we observe that 53% of crashes are caused by an SSP fault. Both the existing implementation and our proposal are able to detect 60% of buffer overflows. This proves that our solution is as performant as the original one.

For CWE 122, we observe 22% of memory faults for WebAssembly with SSP disabled. x86 with SSP results in 21% of SSP faults. Both the existing implementation of SSP in WebAssembly and our proposal are able to detect 20% of buffer overflows.

The results presented here are consistent with figures reported by Stiévenart et al. [17].

2) *Interpretation*: First, native and WebAssembly configurations using SSP mitigate more than half stack-based buffer overflows (CWE 121). This confirms that SSP in WebAssembly is efficient at mitigating stack-based buffer overflows, compared to the situation without protection. Surprisingly, we observe that some heap-based buffer overflow (CWE 122) of the Juliet test suite crash because of an SSP fault. This




Fig. 6. Execution outcome of each binary in the Juliet test suite

behavior is not expected since a heap overflow grows farther from stack memory, i.e., from the canary. We found that all CWE 122 SSP faults occur because the corresponding Juliet tests have been mistakenly tagged as CWE 122, while they are effectively stack-based buffer overflow (CWE 121). This confirms the expected result that SSP cannot detect heap-based buffer overflows.

Second, our implementation of SSP has the same coverage as the existing implementation. However, as pointed out in Section IV, the existing SSP implementation can easily be bypassed.

Third, our implementation is not able to cover the entirety of buffer overflows, in particular a buffer overflow is not detected when the overflow does not reach the canary. This can happen with small overflows, when e.g. other variables are allocated between the vulnerable buffer and the top of the stack frame. However, this defect is common to all SSP implementations.

These results validate the effectiveness of SSP in WebAssembly, and prove that our proposed implementation is as safe and efficient as the existing one.

VI. CONCLUSION

In this paper, we focused on the mitigation of stack-based buffer overflows in WebAssembly with the Stack Smashing Protection mechanism. SSP is particularly interesting as it is one of the few binary protections that does not require to modify the WebAssembly specification.

We evaluated the existing implementation of SSP in WebAssembly. Two weaknesses were identified: the possibility to overwrite the canary reference value and a fragile fallback in case of a random generator failure.

An SSP solution for WebAssembly that mitigates these weaknesses was specified and implemented. The solution improves the robustness of the existing SSP implementation by proposing secure storage of the canary reference value and a hardened fallback in case of a random generator failure, without any loss of efficiency in detection.

We evaluated our solution and demonstrated that it mitigates a significant portion of stack-based buffer overflows, while be-

ing more robust than the already existing one. This proves the positive impact of this protection on WebAssembly security, leading us to believe that SSP should become a default in all WebAssembly binaries in the future.

The theoretical analysis detailed in this paper is generalizable to all WebAssembly toolchain implementations. We publish as open-source software the tools used for our analysis, as well as our implementation of SSP. We hope our work and the related code will be useful to help the community to build safe and secure WebAssembly applications and tooling.

REFERENCES

- [1] "The llvm compiler infrastructure," accessed on 2024-03-21. [Online]. Available: <https://llvm.org/>
- [2] "Security - WebAssembly (memory safety)," accessed on 2024-04-17. [Online]. Available: <https://webassembly.org/docs/security/#memory-safety>
- [3] "wasmCloud," accessed on 2024-04-17. [Online]. Available: <https://wasmcloud.com/>
- [4] "Distributed Compute and Communications in 5G," *5G Americas*, Nov. 2022.
- [5] "Wasi, the webassembly system interface," 2024, accessed on 2024-04-19. [Online]. Available: <https://wasi.dev/>
- [6] "WebAssembly on Kubernetes: from containers to Wasm (part 01)," Mar. 2024, accessed on 2024-04-17. [Online]. Available: <https://www.cncf.io/blog/2024/03/12/webassembly-on-kubernetes-from-containers-to-wasm-part-01/>
- [7] B. Bierbaumer, J. Kirsch, T. Kittel, A. Francillon, and A. Zarras, "Smashing the stack protector for fun and profit," in *ICT Systems Security and Privacy Protection*, L. J. Janczewski and M. Kutylowski, Eds. Cham: Springer International Publishing, 2018, pp. 293–306.
- [8] T. Boland and P. E. Black, "Juliet 1.1 c/c++ and java test suite," *Computer*, vol. 45, no. 10, pp. 88–90, 2012.
- [9] L. Clark, "Standardizing WASI: A system interface to run WebAssembly outside the web – Mozilla Hacks - the web developer blog," accessed on 2024-04-17. [Online]. Available: <https://hacks.mozilla.org/2019/03/standardizing-wasi-a-webassembly-system-interface>
- [10] CNCF, "Bringing WebAssembly to telecoms with CNCF wasmCloud," Jan. 2024. [Online]. Available: <https://www.cncf.io/blog/2024/01/05/bringing-webassembly-to-telecoms-with-cncf-wasmcloud/>
- [11] C. Cowan, C. Pu, D. Maier, J. Walpole, P. Bakke, S. Beattie, A. Grier, P. Wagle, Q. Zhang, and H. Hinton, "StackGuard: Automatic adaptive detection and prevention of Buffer-Overflow attacks," in *7th USENIX Security Symposium (USENIX Security 98)*. San Antonio, TX: USENIX Association, Jan. 1998.
- [12] A. Haas, A. Rossberg, D. L. Schuff, B. L. Titzer, M. Holman, D. Gohman, L. Wagner, A. Zakai, and J. Bastien, "Bringing the web up to speed with WebAssembly," in *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2017. New York, NY, USA: Association for Computing Machinery, Jun. 2017, pp. 185–200.
- [13] A. Hilbig, D. Lehmann, and M. Pradel, "An Empirical Study of Real-World WebAssembly Binaries: Security, Languages, Use Cases," in *Proceedings of the Web Conference 2021*. Ljubljana Slovenia: ACM, Apr. 2021, pp. 2696–2708.
- [14] D. Lehmann, J. Kinder, and M. Pradel, "Everything Old is New Again: Binary Security of WebAssembly," in *Proceedings of the 20th USENIX Security Symposium*, 2020, pp. 217–234.
- [15] A. One, "Smashing the stack for fun and profit," *Phrack*, vol. 7, no. 49, November 1996.
- [16] A. Rossberg, "WebAssembly Core Specification," W3C, Tech. Rep., Dec. 2019. [Online]. Available: <https://www.w3.org/TR/wasm-core-1/>
- [17] Q. Stiévenart, C. De Roover, and M. Ghafari, "The Security Risk of Lacking Compiler Protection in WebAssembly," Nov. 2021, arXiv:2111.01421 [cs].
- [18] Z. Zhang, W. Zheng, B. Hua, Q. Fan, and Z. Pan, "VMCanary: Effective Memory Protection for WebAssembly via Virtual Machine-assisted Approach," in *2023 IEEE 23rd International Conference on Software Quality, Reliability, and Security (QRS)*, Oct. 2023, pp. 662–671.