



Comparing Security in eBPF and WebAssembly

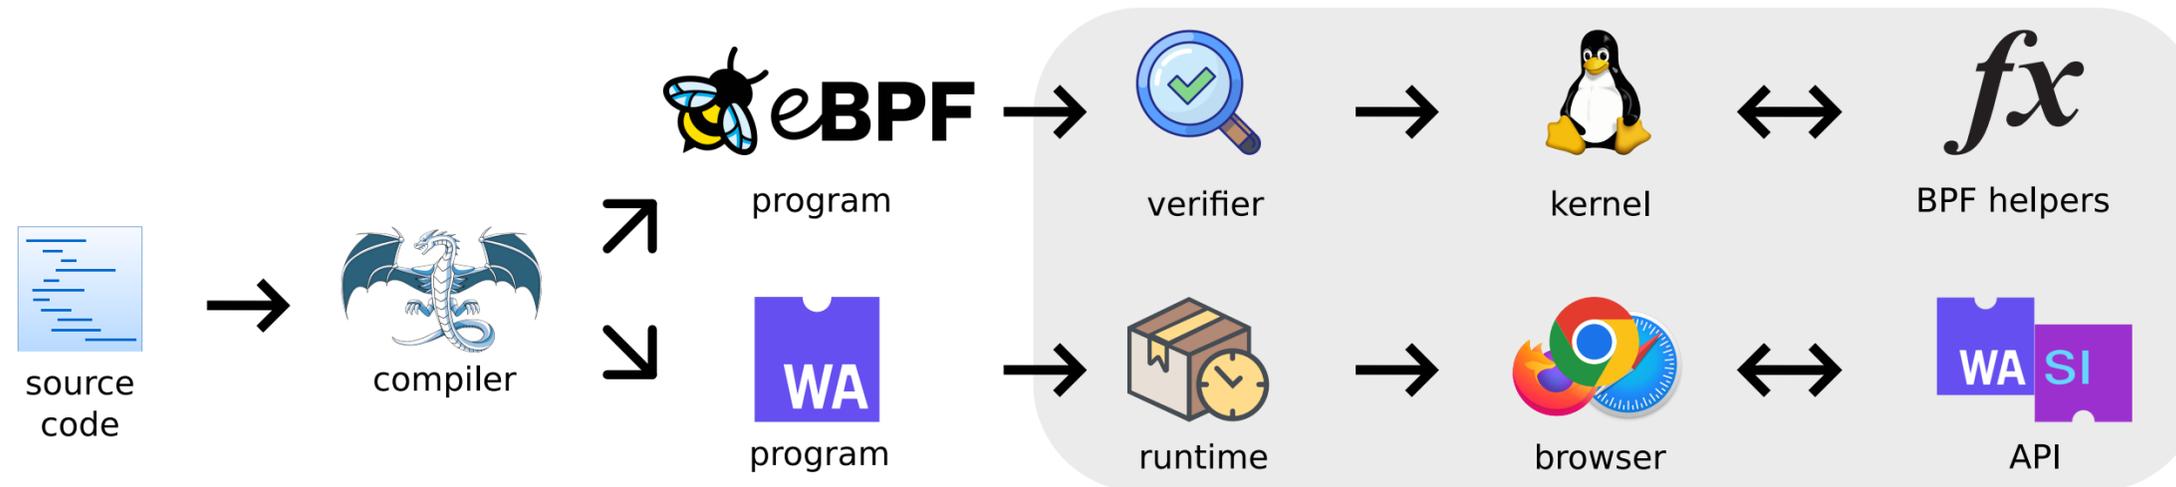
Jules DEJAEGHERE ▸ Bolaji GBADAMOSI ♦ Tobias PULLS ♦ Florentin ROCHET ▸

▸ University of Namur ♦ Karlstad University

1st Workshop on eBPF and Kernel Extensions

September 10, 2023, New York

Overview: lifecycle of eBPF and Wasm programs

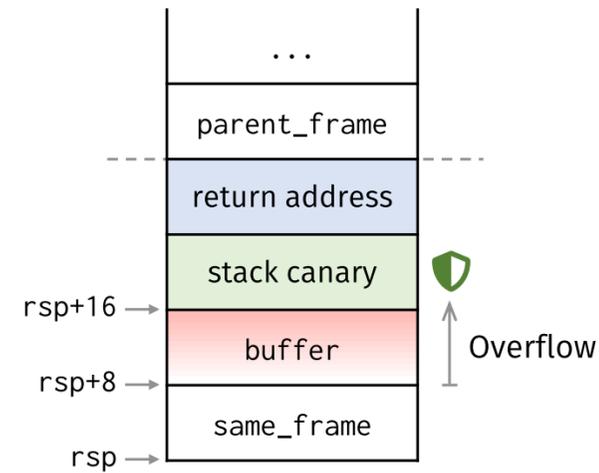


Possible lifecycle for eBPF and Wasm programs

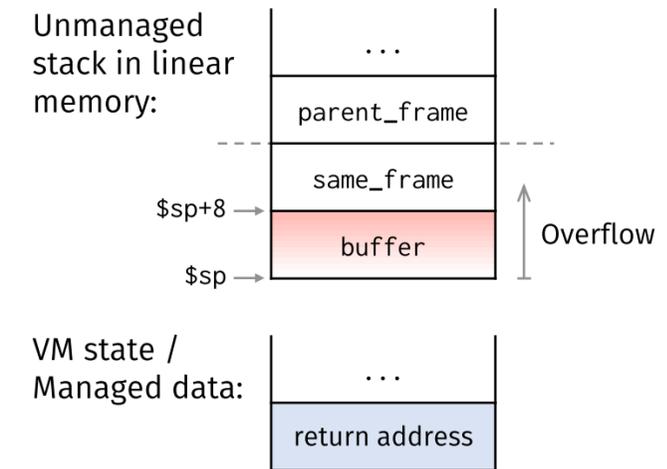
WebAssembly: selected key points

- Binary instruction format
- **Managed stack & linear memory**
- Bounded memory
- Web first but supports non-web embeddings
- **Checked indirect function calls**
- Default to no host access
- **1:1 mapping between binary ⇔ text format**

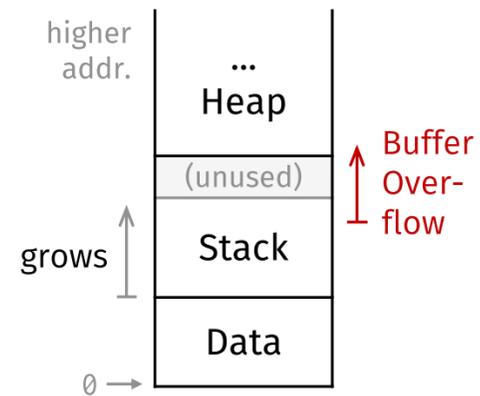
WebAssembly: managed stack & linear memory



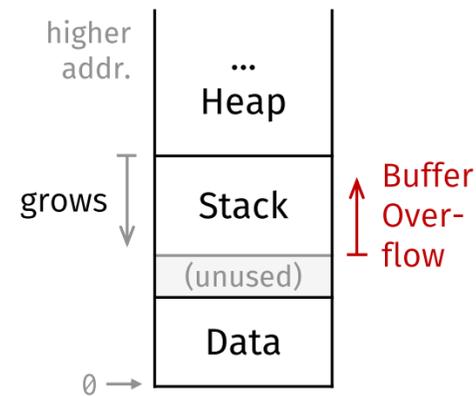
Stack layout on x86-64 with canaries and reordering



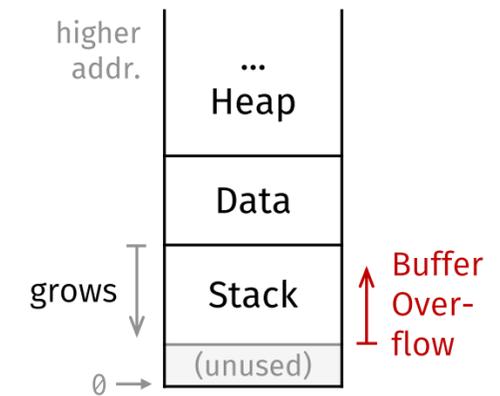
Linear memory and VM state in WebAssembly



emcc 1.39.7 (*fastcomp* backend, deprecated)



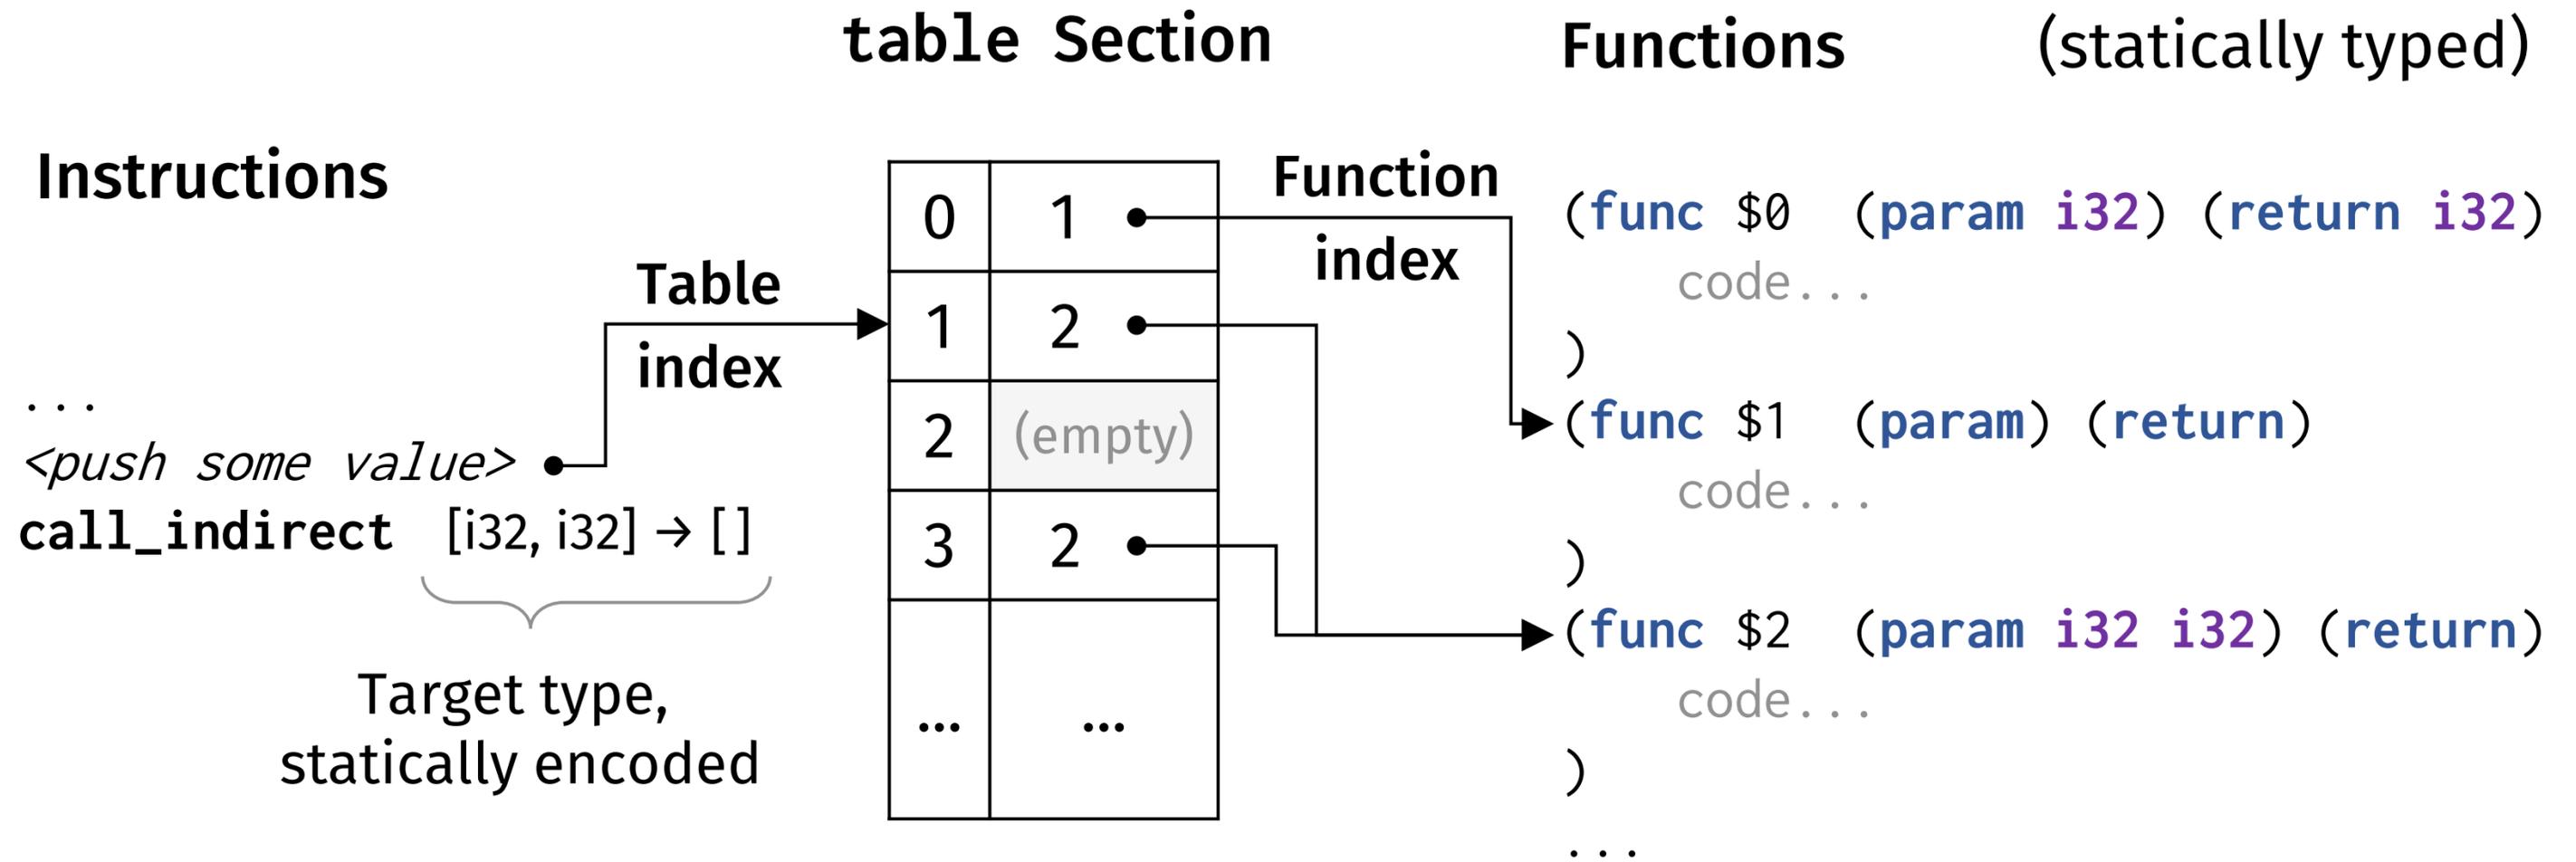
emcc 1.39.7 (*upstream* backend), clang 9 (WASI)



clang 9 (WASI with *stack-first*), rustc 1.41 (WASI)

Illustrations from Lehmann, D. et al. [1]

WebAssembly: checked indirect function calls



Indirect function calls via the table section

Illustration from Lehmann, D. et al. [1]

WebAssembly: binary ↔ text

```
1 #[no_mangle]
2 pub extern "C" fn add(left: i32, right: i32) -> i32 {
3     left + right
4 }
```

```
$ rustc lib.rs --target wasm32-wasi --crate-type cdylib -C opt-level=3
$ wasm2wat lib.wasm
```

```
1 (module
2   (type (;0;) (func (param i32 i32) (result i32)))
3   (func $add (type 0) (param i32 i32) (result i32)
4     local.get 1
5     local.get 0
6     i32.add)
7   (table (;0;) 1 1 funcref)
8   (memory (;0;) 16)
9   (global $__stack_pointer (mut i32) (i32.const 1048576))
10  (global (;1;) i32 (i32.const 1048576))
11  (global (;2;) i32 (i32.const 1048576))
12  (export "memory" (memory 0))
13  (export "add" (func $add))
14  (export "__data_end" (global 1))
15  (export "__heap_base" (global 2)))
```

Rust function compiled to a WebAssembly module in textual format

Comparing eBPF and WebAssembly



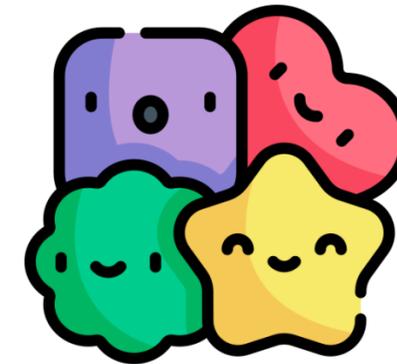
VS



Threat model



The verifier acts as the gatekeeper to ensure kernel safety



Untrusted code can run without compromising the host



Memory safety



- Few limitations on what programmers can write
- Verifier ensures safety
- No proof, then no execution

```
1 #include <linux/bpf.h>
2 #include <bpf/bpf_helpers.h>
3 SEC("xdp")
4 int buffer(void *ctx) {
5     int a[3];
6     int i;
7     for (i = 0; i < 100; i++) {
8         bpf_printk("%d ", a[i]);
9     }
10    return 0;
11 }
12 char LICENSE[] SEC("license") = 'l'
```

■ Code will not run



- Limited set of constructs
- Grammatically correct, then execution allowed
- Runtime checks

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 int main() {
4     int a[3];
5     int i;
6     for(i = 0; i < 100; i++) {
7         printf("%d ", a[i]);
8     }
9     return 0;
10 }
```

▶ Code will run

Control flow integrity



- CFI enforced by the verifier
- Flagging of programs violating CFI
- Verifier ensures termination

```
1 #include <linux/bpf.h>
2 #include <bpf/bpf_helpers.h>
3 char _license[] SEC("license") =
4 int a;
5 SEC("socket")
6 int prog(void *ctx){
7     while (1) {
8         a++;
9     }
10    return 0;
11 }
```

■ Code will not run



- CFI achieved via semantics
- Jump only to the beginning of valid constructs
- Indirect function calls prevent call redirection

```
1 int a;
2 int main() {
3     while (1) {
4         a++;
5     }
6     return 0;
7 }
```

▶ Code will run (forever)

API access



- Many helper functions available by default
- Each program type can only call a subset of the helper functions
- Access to helper functions is restricted if unprivileged BPF is enabled



- Default to no host access
- API implementation is provided by the host
- Standardized: e.g. WebAssembly System Interface

Side-channels



- Constant blinding to avoid code as constant and JIT spraying
- Retpoline when tail calls cannot be converted to direct calls
- Impossible path verification



- Out of scope for the language, in scope for the runtime
- Bound checking when accessing function table (e.g. `call_indirect`)
- No bound verification for linear memory by default (relying on page fault), can be enabled in some settings



Conclusion



- Checks ahead of the execution
- Does not execute if policy violation is found
- Code is trusted but the code is not trustworthy
- Access to many kernel-provided helpers, by default

- Checks at runtime
- Traps when policy violation occurs
- Code is untrusted
- No access to host resources, unless explicitly granted

Takeaways

What are the performance impacts of eBPF and WebAssembly?

Is one approach more efficient than the other?

What can we learn from both technologies?

How could we measure and captures the differences?



Comparing Security in eBPF and WebAssembly

Jules DEJAEGHERE ▸ Bolaji GBADAMOSI ♦ Tobias PULLS ♦ Florentin ROCHET ▸

▸ University of Namur ♦ Karlstad University

1st Workshop on eBPF and Kernel Extensions

September 10, 2023, New York

References

- [1] Lehmann, D., Kinder, J. and Pradel, M. 2020. [Everything old is new again: Binary security of WebAssembly](#). *29th USENIX security symposium (USENIX security 20)* (Aug. 2020), 217–234.

This presentation has been designed using images from [Freepik - Flaticon.com](#).

Meltdown and Spectre icons are from [Meltdown and Spectre website](#).