# 6.5810: FAASM

#### Adam Belay <abelay@mit.edu>



# Logistics

• Sign up to meet + discuss your final project (happening tomorrow)

#### Recap

- Serverless isolation mechanisms
- So far: GVisor (library OS), Firecracker (virtualization), and UKL (Unikernels)
- Today: FAASM (software fault isolation)
  - Some similarities to CloudFlare Workers (recently open sourced)
  - <u>https://github.com/cloudflare/workers.cloudflare.com</u>

# FAASM's plan

- Rely on Web Assembly for isolation
  - Key advantage: Functions can share memory directly
- Everything else done the standard way:
  - Cgroups: memory and CPU quotas
  - Network namespaces + traffic shaping: Bandwidth limits
  - Etc.

# Quick aside: Software fault isolation

- Restricts a function to accesses of its own memory
- Works by inserting checks before memory and control transfer instructions
- Provides a logical isolation domain inside a process (despite sharing the same page table and memory)
- In the past: Binary translation has been used to achieve SFI
- This work targets a special instruction set architecture instead

# Quick aside: Web Assembly (WASM)

- History: Highly optimized Javascript JITs already existed (e.g., V8)
- Can we leverage this existing infrastructure to run native code?
- Plan: Define a "web assembly" language for portability
  - This means you must compile to web assembly
  - But can support C, C++, Rust, etc. (any LLVM front end)
- JIT translates web assembly to native assembly
- Checks inserted to preserve memory safety

#### FAASM Design



# Faaslet

- Isolation
  - Access constrained to a contiguous memory region
  - By default, all memory is private, but shared mappings are possible
  - Control flow integrity ensure function can't jump out of its code
- Faaslets can be snapshotted to reduce startup times
- Each Faaslet runs in a single thread
- Standard mechanism (e.g., cgroups) enforce resource limits

#### Host interface

Class	Function	Action	Standard	
Calls	<pre>byte* read_call_input() void write_call_output(out_data) int chain_call(name, args) int await_call(call_id) byte* get_call_output(call_id)</pre>	Read input data to function as byte array Write output data for function Call function and return the call_id Await the completion of call_id Load the output data of call_id		
State	<pre>byte* get_state(key, flags) byte* get_state_offset(key, off, flags) void set_state(key, val) void set_state_offset(key, val, len, off) void push/pull_state(key) void push/pull_state_offset(key, off) void append_state(key, val) void lock_state_read/write(key) void lock_state_global_read/write(key)</pre>	Get pointer to state value for key Get pointer to state value for key at offset Set state value for key Set len bytes of state value at offset for key Push/pull global state value for key Push/pull global state value for key Append data to state value for key Lock local copy of state value for key Lock state value for key globally	none	
Dynlink	<pre>void* dlopen/dlsym() int dlclose()</pre>	Dynamic linking of libraries As above	DOGIN	
Memory	<pre>void* mmap(), int munmap() int brk(), void* sbrk()</pre>	Memory grow/shrink only Memory grow/shrink	POSIX	
Network	<pre>int socket/connect/bind() size_t send/recv()</pre>	Client-side networking only Send/recv via virtual interface		
File I/O	<pre>int open/close/dup/stat() size_t read/write()</pre>	Per-user virtual filesystem access As above	WASI	
Misc	int <b>gettime</b> () size_t <b>getrandom</b> ()	Per-user monotonic clock only Uses underlying host /dev/urandom		

Table 2: Faaslet host interface (The final column indicates whether functions are defined as part of POSIX or WASI [57].)

# Host interface enables faster communication

- Normally: All intermediate results placed in a key value store over the network (e.g., S3)
- Here, functions can pass byte arrays directly to one another inside the same process using shared memory
- But some benefit lost if functions don't fit on same machine



Figure 2: Faaslet shared memory region mapping

# State programming model

- Build on distributed data objects (DDOs)
  - Key -> value abstraction
- Two tiers; local tier provides shared memory access, global tier transfers over the network
- Push() and Pull() operations explicitly transfer data from local to global tier

# Faasm scheduling

- Distributed scheduling
- Goal: Place calls on instances with warm functions and shared data



Figure 5: FAASM system architecture

# Faasm snapshotting

- Web assembly represents memory as a simple array of bytes
- FAASM's plan: Fully load a function, then copy the byte array and store it for future execution in a shared object store
- Eliminates code generation cost and other initialization
- Copy-on-write mappings can be used to restore a clean copy efficiently

#### Evaluation: SGD (machine learning)



Figure 6: Machine learning training with SGD with Faaslets (FAASM) and containers (Knative)

#### Evaluation: Python compute overhead



(b) Python Performance Benchmark

## Startup times and initial memory use

	Docker	Faaslets	<b>Proto-Faaslets</b>	vs. Docker
Initialisation	2.8 s	5.2 ms	0.5 ms	<b>5.6K</b> ×
CPU cycles	251M	1.4K	650	<b>385K</b> ×
PSS memory	1.3 MB	200 KB	90 KB	$15 \times$
RSS memory	5.0 MB	200 KB	90 KB	$57 \times$
Capacity	~8 K	~70 K	>100 K	12  imes

Table 3: Comparison of Faaslets vs. container cold starts (no-op function)

#### Another perspective on compute overhead



Not So Fast: Analyzing the Performance of WebAssembly vs. Native Code. Jangda et. Al. ATC'19

#### Discuss: Does FAASM achieve Amazon's goals?

- 1. Isolation
- 2. Density
- 3. Performance
- 4. Compatibility
- 5. Fast switching
- 6. Soft allocation

#### Conclusion

- Web assembly + FAASM enables very fast context switching, startup times, snapshotting, and data sharing
- But CPU throughput remains an issue
  - Can it be optimized further?