



(intel) Research Institute for Collaborative Autonomous and Resilient Systems



Trustworthy & Accountable Function-as-a-Service

N. Asokan https://asokan.org/asokan/

@nasokan

Joint work with Fritz Alder, Arseny Kurnikov, Andrew Paverd, Michael Steiner

Function-as-a-Service (FaaS)

Recent instantiation of "serverless computing"

- Customer specifies the function
- Service provider manages runtime, scaling, load-balancing etc.

Differences to Infrastructure-as-a-Service (laaS)

- Relatively short-running function invocations
- Stateless functions (storage provided by separate service)

FaaS is available from established cloud providers

Usual security concerns of cloud computing still apply:

- Confidentiality of data
- Integrity of computation





https://www.theregister.co.uk/2018/07/24/apache_ibm_cloud_vulnerable/

FaaS is available from established cloud providers

Usual security concerns of cloud computing still apply:

- Confidentiality of data
- Integrity of computation

More accurate resource usage measurements required:

• Sub-second compute time measurements

Currently achieved via existing reputational trust, but can we do better?

FaaS can also be provided by non-traditional service providers

- Data centres with spare capacity
- Individuals with powerful PCs (e.g. gamers)

Open source frameworks available

Multiple start-ups in this space



https://golem.network/



https://openwhisk.apache.org/



https://ankr.com

FaaS can also be provided by non-traditional service providers

- Data centres with spare capacity
- Individuals with powerful PCs (e.g. gamers)

Heightened security concerns:

- Service provider identity/location may be unknown
- Service provider may not have security expertise

Very few disincentives for cheating:

• Malicious service provider might inflate resource usage measurements

No reputational trust has been established

System Model & Requirements





Adversary model

Two types of adversaries:

Service provider

- Learn inputs and outputs of function invocations
- Modify inputs and outputs, or execute the function incorrectly
- Overcharge the function provider
 - Falsely inflate resource usage measurements
 - Create fake function invocations

Function provider

• Under-pay the service provider for resources used by the function

Requirements

R1 - Security

- Service provider cannot modify inputs or outputs of a function invocation
- Client assured that output is result of correct execution of intended function on supplied inputs

R2 - Privacy

• Service provider cannot learn inputs or outputs of a function invocation

R3 - Measurement accuracy

Resource measurements must have sufficient accuracy for FaaS billing

R4 - Measurement veracity

• All parties must be able to verify authenticity of resource measurements

Trusted Execution Environments



Hardware support for

- Isolated execution: Isolated Execution Environment
- Protected storage: Sealing
- Ability to convince remote verifiers: (Remote) Attestation

Trusted Execution Environments (TEEs)

Operating in parallel with "rich execution environments" (REEs)

Hardware-assisted TEEs are pervasive



Background: Intel SGX



Trusted Untrusted

CPU enforced TEE (enclave)

Remote attestation

Secure memory

- Confidentiality
- Integrity

https://software.intel.com/sgx

Preliminary design



Design Challenges

Challenge: Sandboxing untrusted functions

Malicious function provider could attempt to reduce in-enclave measurements

• No protection from code in the same enclave



Challenge: Attesting worker enclaves

Default SGX remote attestation involves multiple message round-trips

- Overhead and latency for short-running functions is too high
- Must be repeated for each enclave



Challenge: Encrypting client input



Challenge: Measuring time in enclaves

CPU instructions RDTSC: read timestamp counter AEX: asynchronous enclave exit ERESUME: resume enclave

SGX enclave cannot reliably measure its own running time

- RDTSC value can be manipulated by VMM
- sgx_get_trusted_time() can be arbitrarily delayed

• Enclaves can be transparently interrupted (AEX) and resumed (ERESUME)



Challenge: Measuring time in enclaves



S-FaaS Architecture

Architecture overview

Worker enclave runs function within a sandbox

- e.g. Ryoan
- sandboxing interpreters: e.g. for JavaScript



Challenges

C1: Sandboxing C2: Attesting enclaves C3: Encrypting input C4: Measuring time

Architecture overview

ka: enclave's DH key **kc**: client's DH key

ko: output keykr: resource reporting key



Transitive attestation

Clients and function providers attest worker enclaves indirectly

Measuring Resource Usage in SGX

FaaS is available from established cloud providers

Service	Invocations	Time (GHz-s)	Memory (GB-s)	Network (GB)
AWS Lambda	Х	Ο	Х	
Azure Functions	Х	Ο	Х	
Google Cloud Functions	Х	Х	Х	Х
IBM Cloud functions	Х	0	Х	

FaaS billing policies of established cloud providers (X = explicit; O = implicit)

Types of measurements

Symbol	Description	Units
t	Total compute time of the function	multiples of T
Т	Duration of each tick in CPU cycles	GHz-s
m _{int}	Time-integral of memory usage	GB-s
m _{max}	Maximum memory used by the function	GB
net	Total number of network bytes sent and received	GB

Measuring compute time

High level idea: two concurrent threads in the enclave (timer & worker)

Measuring compute time

High level idea: two concurrent threads in the enclave (timer & worker)

Intel SGX internals

Enclave data structures TCS: Thread Control Structure (C)SSA: (Current) Save State Area

Intel Transactional Synchronization Extensions (TSX)

Special instructions enabling Hardware Lock Elision (HLE)

Read set

- Memory addresses read by the transaction (added upon access)
- Transaction will abort if address is concurrently written

Write set

- Memory addresses written by the transaction
- Transaction will abort if address is concurrently read

Roll-back

• All operations since the beginning of the transaction are reverted

Starting a function

SSA stack Marker 0x12...

Timer thread algorithm

```
while(processing == true) {
          // begin TSX txn
  XBEGIN
  if(worker.ssa == marker) // add worker.ssa to txn read set
     for(i=0; i<LOOP_COUNT; i++) // LOOP_COUNT depends on T
        nop;
     t_internal++;
  XEND // end TSX txn
  t_external = t_internal // update external counter
```

Worker thread interrupted

Worker thread resumed

SSA stack Marker 0x12..

Custom ERESUME handler

```
.text
.globl custom eresume handler
.type custom eresume handler,@function
custom eresume handler:
  push %rax
                                     # Save registers
  push %rbx
  lea g worker ssa gpr(%rip),%rax
                                   # Load pointer
  mov (%rax),%rbx
                                     # Dereference pointer
  movl $12345,(%rbx)
                                     # Write SSA marker value
  pop %rbx
                                     # Restore registers
  pop %rax
  jmp *g_original_ssa_rip(%rip) # Resume execution
```

Completing a function

Evaluation: Pre-function latency

1. Create Docker container

Warm-start

- - 4. Perform key-agreement 5. Return empty response

Baseline: $3179 \text{ ms} (\sigma = 40 \text{ ms})$ S-FaaS: 3249 ms (σ = 38 ms) Latency increase: ~2%

5. Return empty response

Cold-start

2. Create enclave 3. Provision function 4. Perform key-agreement

> **Baseline:** 204 ms (σ = 106 ms) S-FaaS: 210 ms (σ = 149 ms) Latency increase: ~3%

Measuring Memory and Networking

Memory

- Instrumented allocators used by interpreter
- Measurements updated on every allocation/free

m _{int}	Time-integral of memory usage
m _{max}	Maximum memory used by the function

Network

• Payloads measured inside enclave

Integration with OpenWhisk

Integration with OpenWhisk

Docker containers

Evaluation

Evaluation: Accuracy

Synthetic function with well-defined compute and memory requirements

• fibonacci(k) calculates the first k numbers in the Fibonacci sequence

Compute time

- Expected to be linear in k
- Can be compared with measurement outside the enclave

Memory time-integral

- Expected to be quadratic in k (k-element list pre-allocated at start of function)
- Harder to measure outside enclave

Evaluation: Accuracy

Evaluation: Accuracy

Evaluation: Performance

Pre-function latency

- Measure cold-start and warm-start latency
- Tested using an empty function to isolate pre-function latency
- Baseline: equivalent operation (same interpreter) without SGX

Resource measurement overhead

- Measure overhead of S-FaaS resource measurement mechanisms
- Octane JavaScript benchmarks (excluding graphical tests)
- Baseline: equivalent operation without resource measurement

Benchmark environment

• Core i5-6500, 8GB RAM, Ubuntu 16.04, Intel SGX SDK 2.2.1

Evaluation: Pre-function latency

Cold-start

- 1. Create Docker container
- 2. Create enclave
- 3. Provision function
- 4. Perform key-agreement
- 5. Return empty response

 Baseline:
 3179 ms (σ = 40 ms)

 S-FaaS:
 3249 ms (σ = 38 ms)

 Latency increase:
 ~2%

Warm-start

- 1. Create Docker container
- 2. Create enclave
- 3. Provision function
- 4. Perform key-agreement
- 5. Return empty response

 Baseline:
 204 ms (σ = 106 ms)

 S-FaaS:
 210 ms (σ = 149 ms)

 Latency increase:
 \sim 3%

Evaluation: Resource measurement overhead

Function	Baseline	S-FaaS						
		No encryption		Encryption		Encryption & receipt		
Box2D	3.019	3.118	3.3%	3.121	3.4%	3.135	3.8%	
DeltaBlue	1.446	1.524	5.4%	1.529	5.7%	1.537	6.3%	
NavierStokes	4.155	4.418	6.3%	4.447	7.0%	4.473	7.7%	
RayTrace	0.779	0.848	8.9%	0.850	9.1%	0.852	9.4%	
Richards	1.719	1.767	2.8%	1.767	2.8%	1.799	4.7%	
Overall	-		5.3%		5.6%		6.3%	

Trade-offs and limitations

Need for an additional thread

• State-of-the-art SGX side-channel defences^(*) require control of both sibling hyperthreads

Timing granularity

- Choice of T affects extent of under- or over-reporting
- S-FaaS service providers can specify T for each function

Architecture-specific calibration

Timing loop must be calibrated for different CPU architectures

(*) SGX side-channel defenses:

Cloak: Gruss et al., "Strong and Efficient Cache Side-Channel Protection using Hardware Transactional Memory", Usenix SEC 2017 HyperRace: Chen et al., "Racing in Hyperspace: Closing Hyper-Threading Side Channels on SGX with Contrived Data Races", IEEE S&P 2018 Varys: Oleksenko et al., "Varys: Protecting SGX enclaves from practical side-channel attacks", Usenix ATC 2018

Suggested SGX enhancements

Secure tick counter

• Provide a trustworthy tick counter that can be accessed without leaving the enclave

Custom ERESUME handlers

- Allow enclaves to specify an in-enclave handler to be called on each ERESUME
- Could also be used to detect frequent AEX events indicative of side-channel attacks

Integration with distributed systems

Smart contracts to pay for outsourced computation

- S-FaaS function receipts and resource measurements can be verified in smart contracts
- Straight-forward integration with payment networks
 - Particularly beneficial to non-traditional service providers

Leader election based on useful work

- Similar to Resource-Efficient Mining for Blockchains (Zhang et al.)
- Uses "useful computation" to determine who mines next block

Deployment considerations

Incremental deployment

- Initially, S-FaaS requires no changes on client-side (no client attestation or encryption)
- Clients can individually start to verify attestation and/or encrypt inputs

Implementations with other TEEs

- S-FaaS could be ported to e.g. ARM TrustZone
- TrustZone secure world still requires functions to run in a suitable sandbox, but timing would be simpler because secure world cannot be arbitrarily paused

Conclusions

FaaS increasingly popular with cloud providers and non-traditional service providers

- Requires strong security: data confidentiality and integrity of computation
- Requires accurate and trustworthy resource consumption measurement

S-FaaS demonstrates how to secure current FaaS architectures using Intel SGX

ACM CCS Cloud Computing Workshop 2019 https://ccsw.io/

Code available on GitHub

What if SGX is broken?

Back to current state of FaaS security and resource measurement

- TEEs useful in two kinds of settings:
 - 1. improving security
 - 2. improving other attributes while preserving security

S-FaaS is Type 1. TEE compromise is a bigger concern in Type 2

- Application-specific ways of detecting / mitigating effects of TEE compromise, e.g.,
 - post-mortem auditing of signed receipts
 - statistical mechanisms like in PoET and Zhang et. al.