



USENIX Security '23

A Verified Confidential Computing as a Service Framework for Privacy Preservation

Hongbo Chen¹, Haobin Hiroki Chen¹, Mingshen Sun²,

Kang Li³, Zhaofeng Chen³, and XiaoFeng Wang¹

¹ Indiana University Bloomington, ² Independent Researcher, ³ CertiK

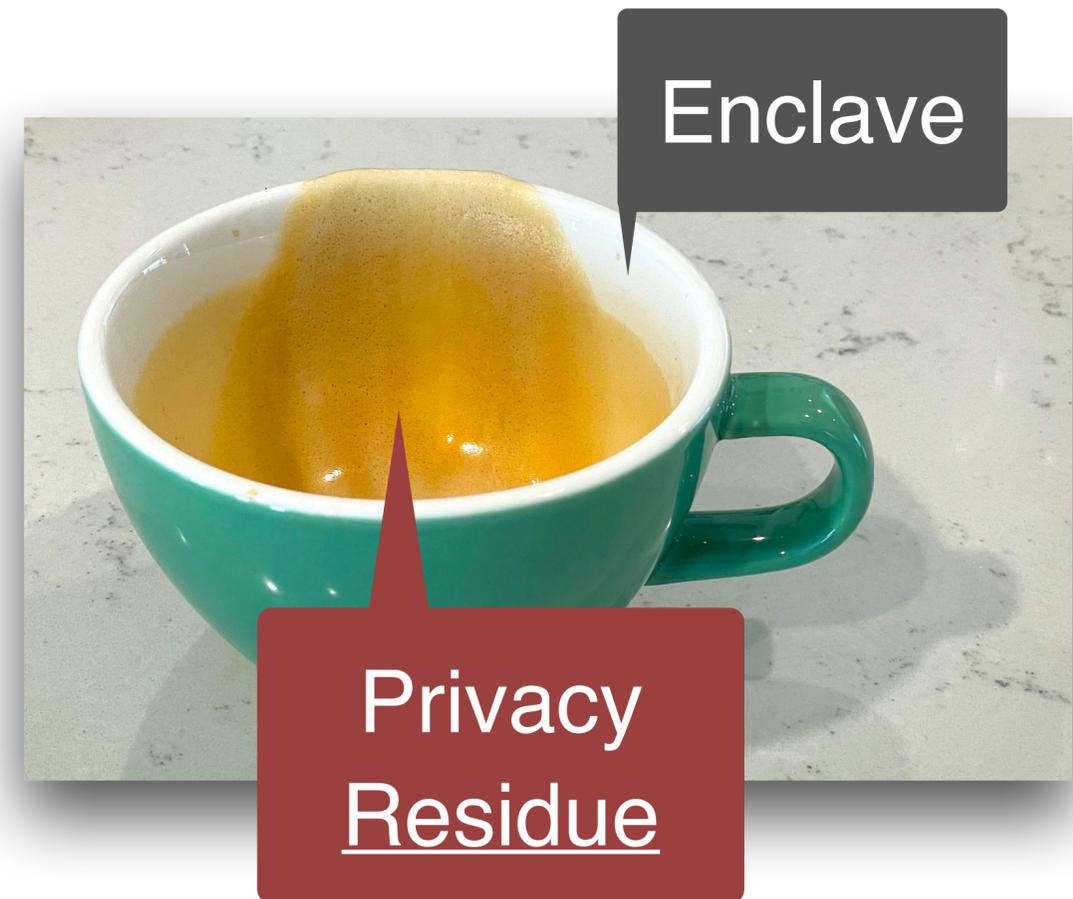
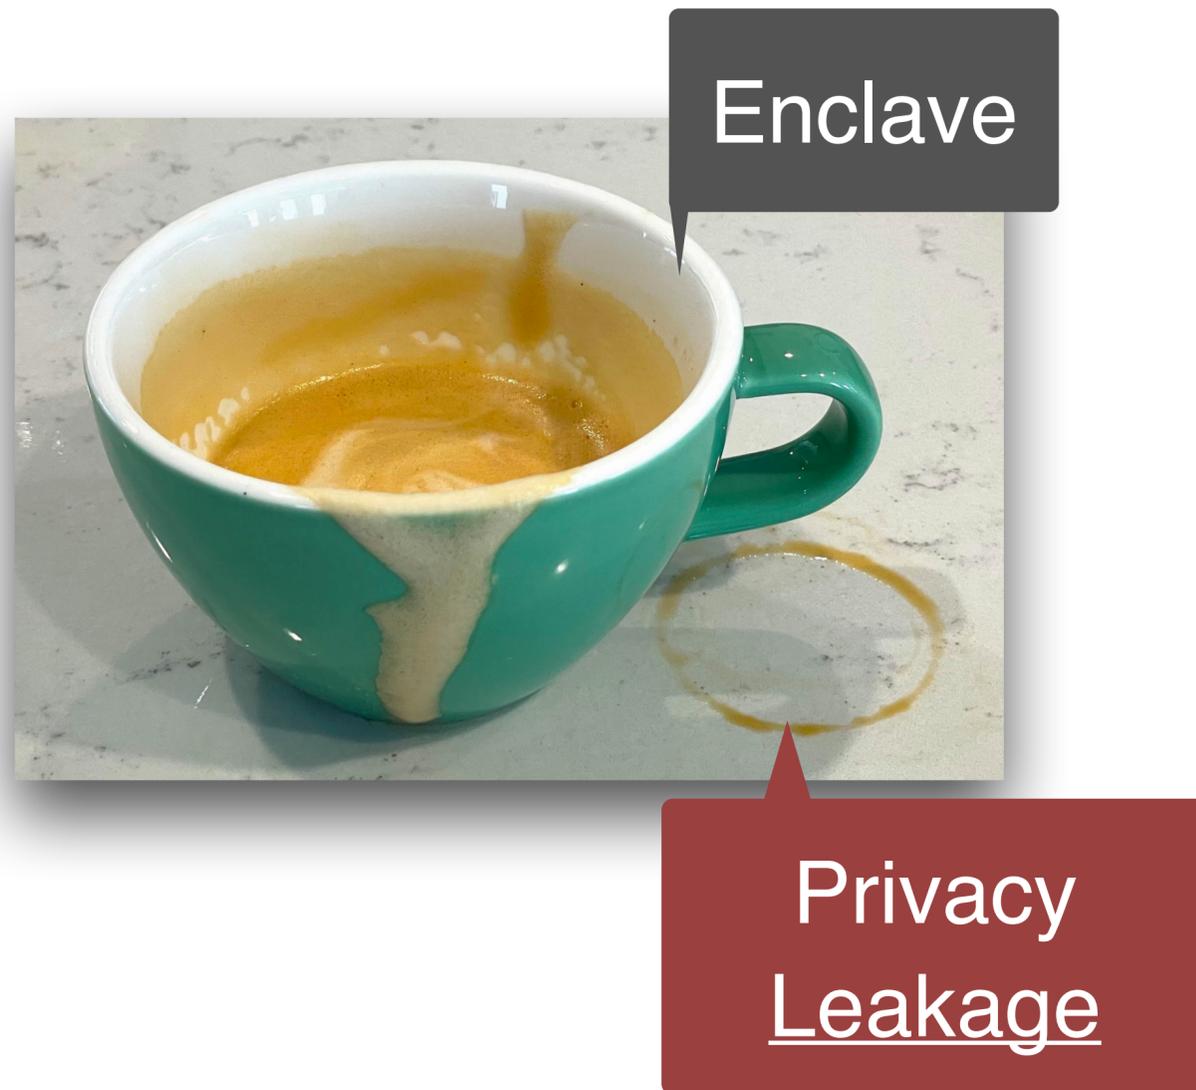


Introduction & Background

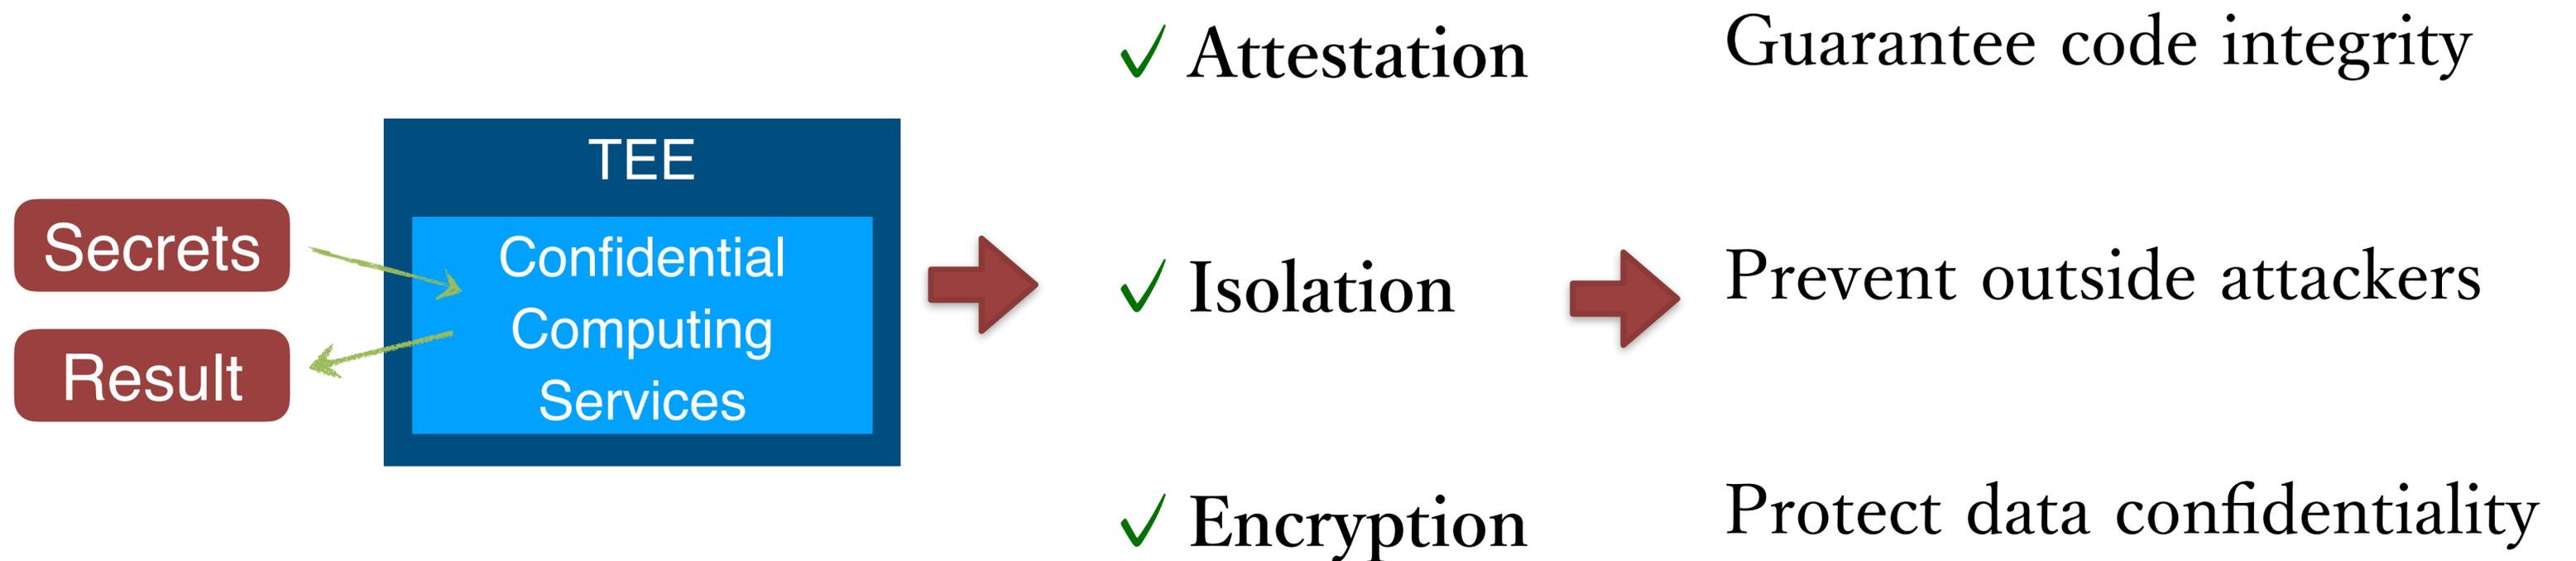
Coffee Incidents



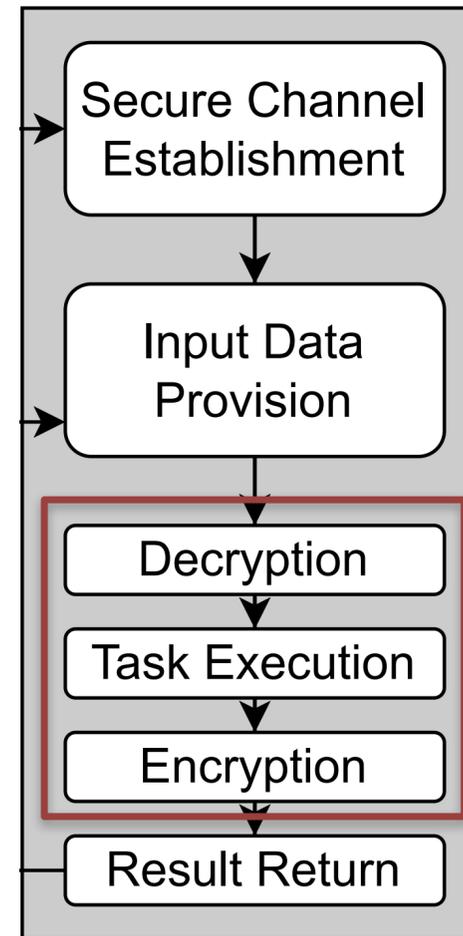
Privacy Incidents



TEE's *Abilities* and Inabilities

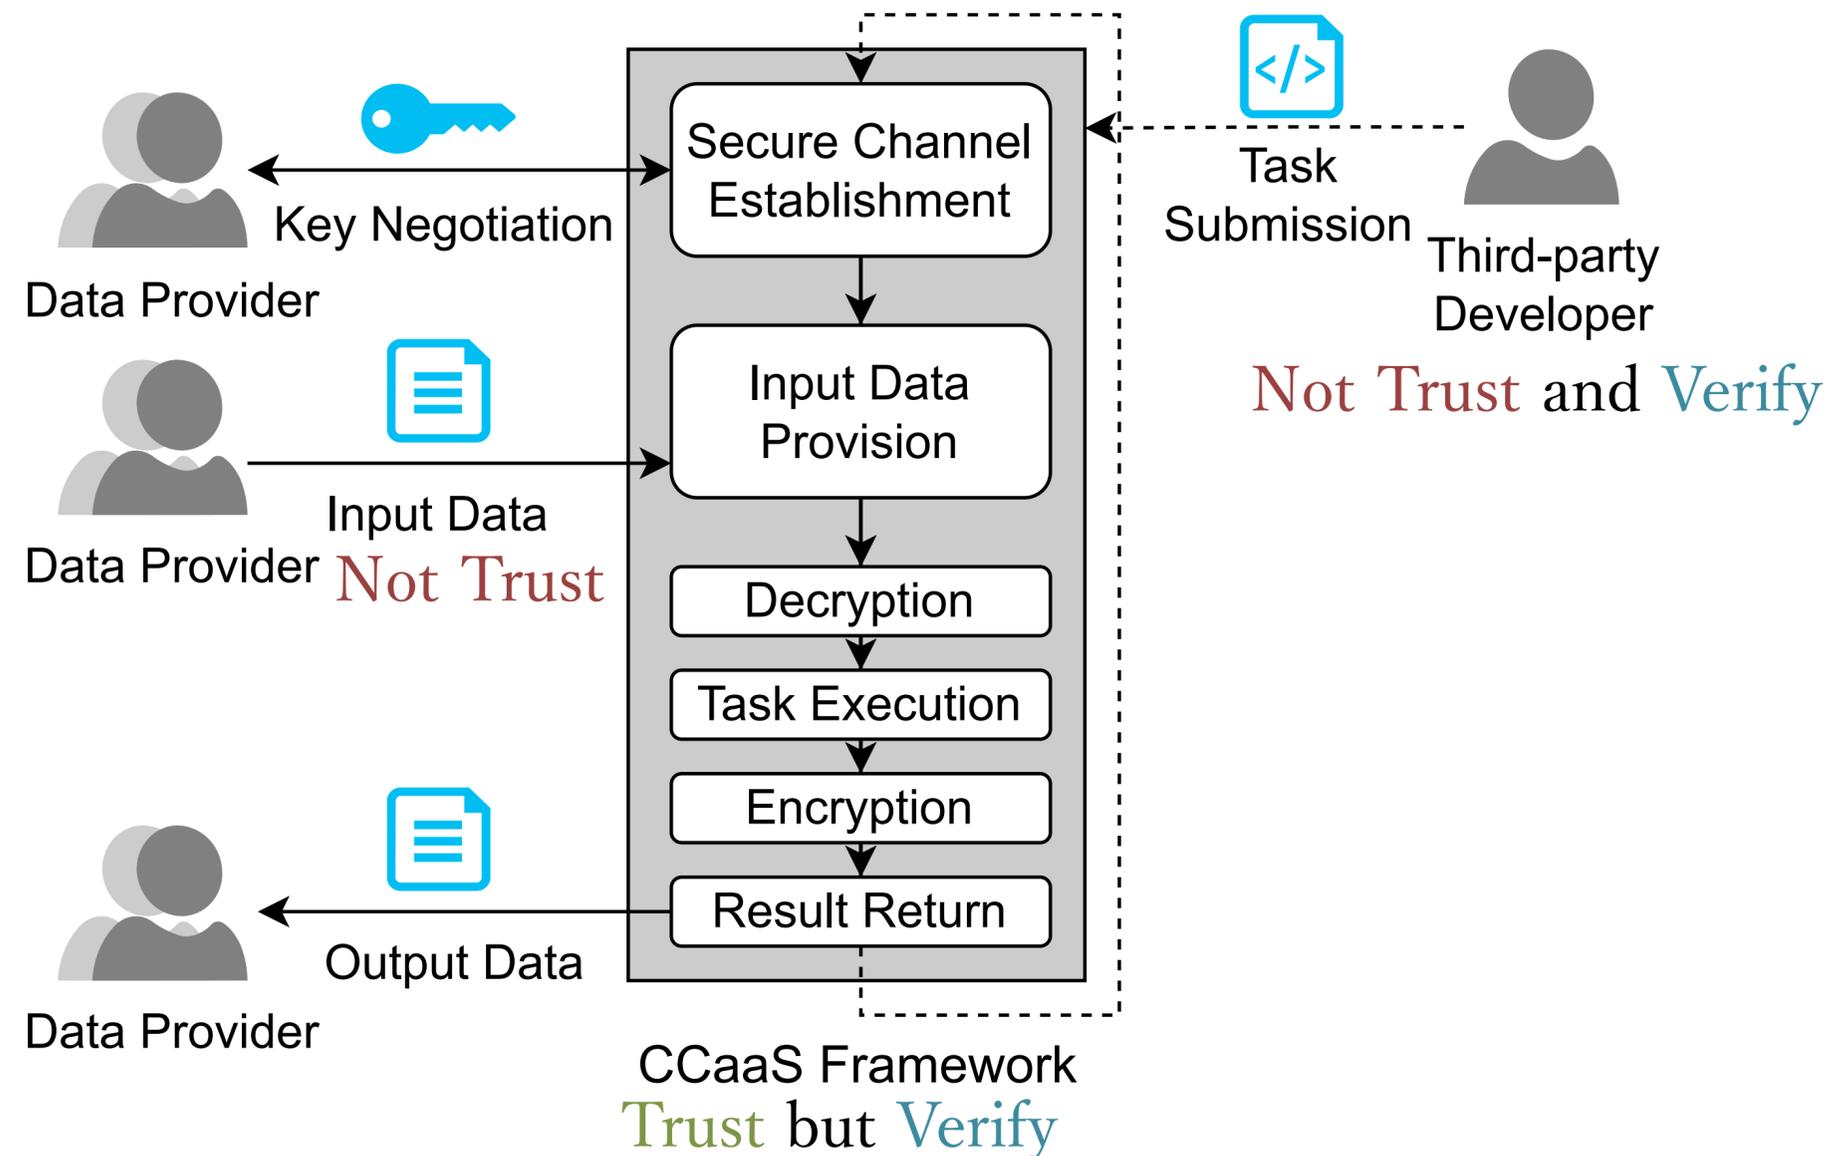


When Confidential Computing Become a **Service**

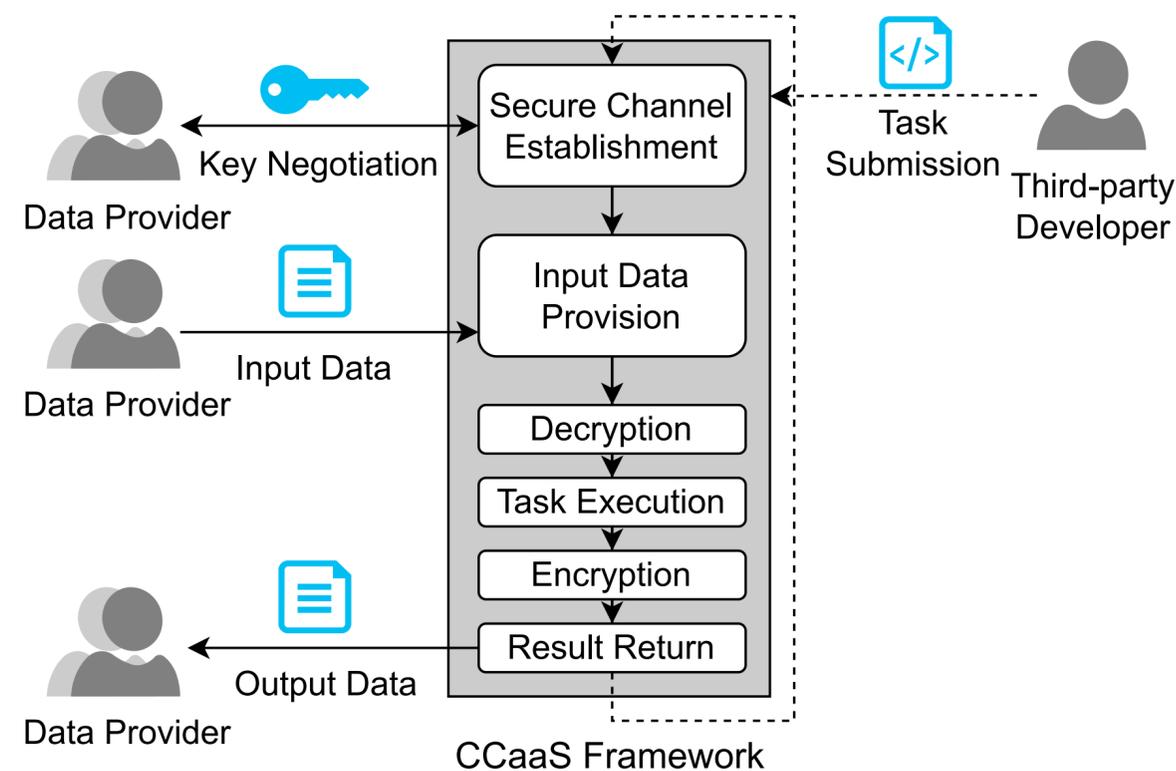


CCaaS Framework

CCaaS for **Multiple** Data Providers



TEE's Abilities and Inabilities



- ◆ **Attestation:** guarantee identity of code
⇒ cannot prove the trustworthiness
- ◆ **Isolation:** prevent outside attackers
⇒ cannot prevent data leakage
- ◆ **Encryption:** protect data safety
⇒ cannot avoid secrets withheld



Our goal: prove to the user that the enclave service cannot threaten their private information.

Proof of Being Forgotten (PoBF)

No Leakage

+

No Residue

All secret and secret-tainted values are within a confined zone during computation.

After the computation (e.g., serving a user), no secret is found in the enclave.

Theoretical Foundation: Enclave Model

Table 1: Generalized model of secure enclaves.

Type	Sym.	Definition
Natural	n	$\in \mathbb{N}$
String	str	$\in \mathbb{S}$
Bool	b	$::= \text{True} \text{False}$
Value	v'	$::= \text{ConcreteN}(n) \text{ConcreteB}(b) \text{Any}$
Sec. Tag	vt	$::= \text{Secret} \text{NotSecret} \text{Nonsense}$
TagValue	v	$::= (v', vt)$
Mode	mo	$::= \text{EnclaveMode} \text{NormalMode}$
Location	l	$::= \text{Stack}(n) \text{Ident}(str) \text{RV}$
Enc. Tag	et	$::= \text{Zone} \text{NonZone}$
Cell	c	$::= \text{Normal}(v) \text{Enclave}(et, v) \text{Unused}$
Result	r	$::= \text{Ok}(X) \text{Err}(e)$
Error	e	$::= \text{Invalid} \text{NoPrivilege}$
Storable	me	$::= \text{List}(l, c)$

Table 2: Enclave program syntax.

Term	Sym.	Definition
Exp.	e	$::= l v' \text{UnaryOp}(e) \text{BinaryOp}(e1, e2)$
Proc.	p	$::= \text{Nop} \text{Eenter} \text{Eexit} \text{Asgn } l := e$ $ \text{If } e \text{ Then } p1 \text{ Else } p2 \text{While } e \text{ Do } p$ $ p1; p2$

Theoretical Foundation: NoLeakage Theorem

A procedure's execution **does not leak secret**.



- Its initial state is secure;
- All memory writes are within the Zone;
- It aborts when error occurs;

Theoretical Foundation: NoResidue Theorem

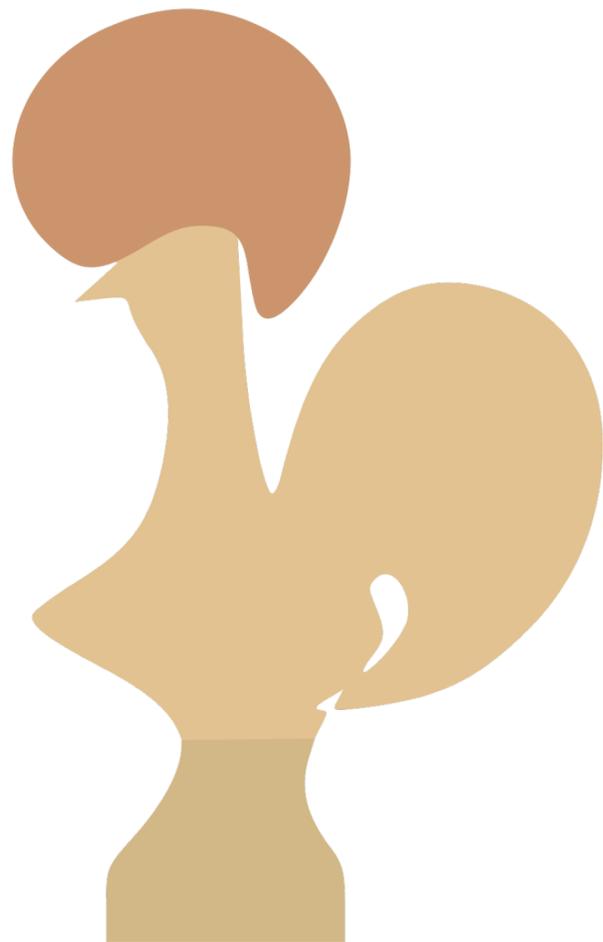
If the Zerorize procedure is executed at the end of a function, then no sensitive data residue is left in the enclave.



`zerorize`

Clears the values stored in the confined zone.

Theoretical Foundation: Checked by Coq



✓ Mechanically Checked by Coq



Realizing the secure enclave service.

Design Goals

Security:

No Leakage

No Residue

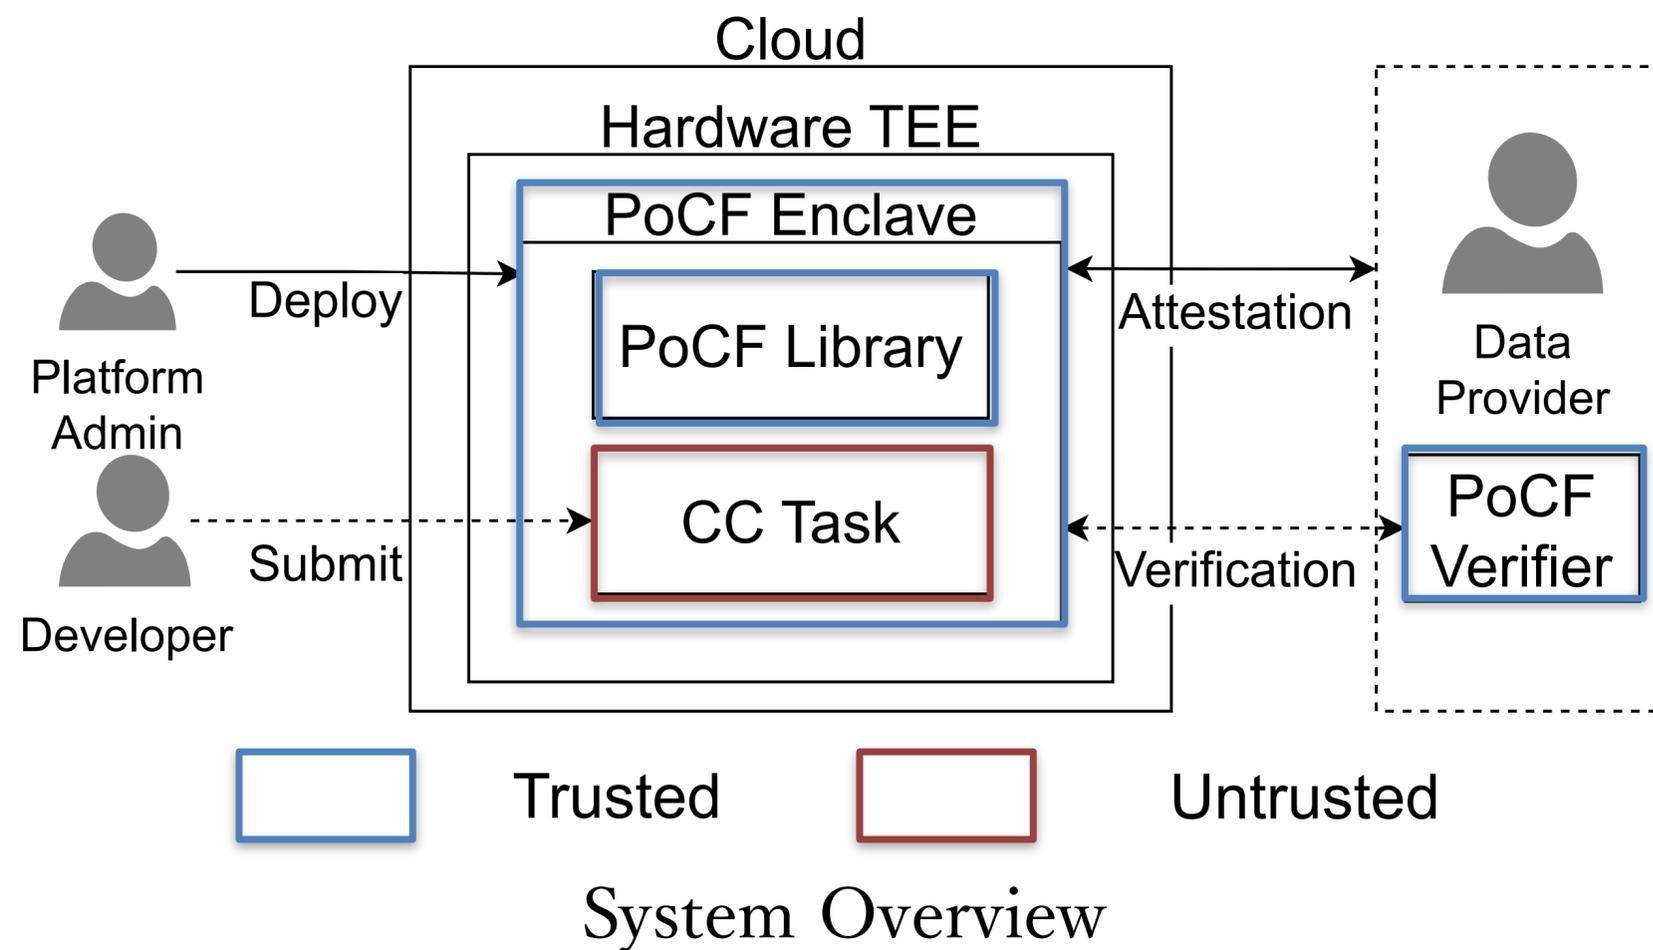
Verifiable

Auxiliary:

- Minimal code modification
- Various hardware TEE support



PoBF-Compliant Framework (PoCF)



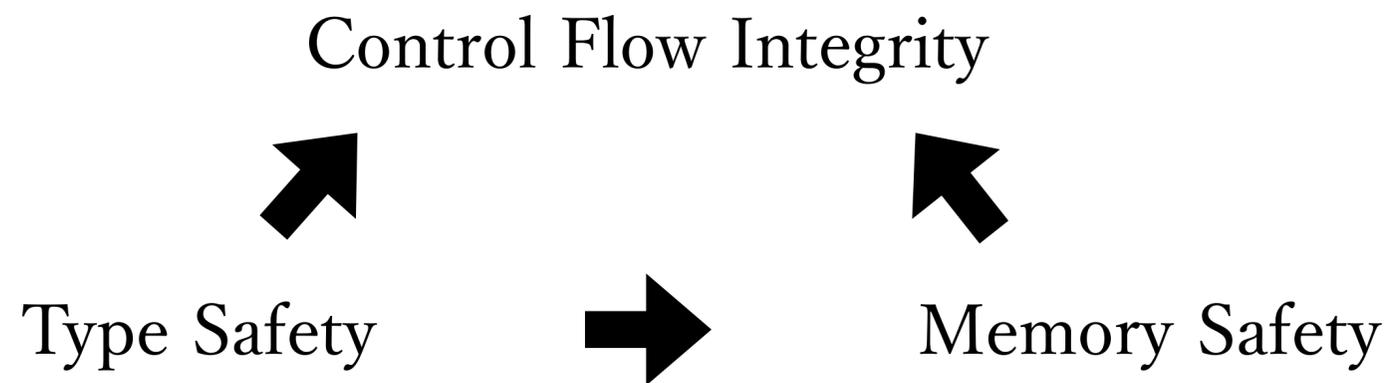
Our Artifacts:

- PoCF Library (TEE-Agnostic)
- PoCF Enclave (TEE-Specific)
- PoCF Verifier

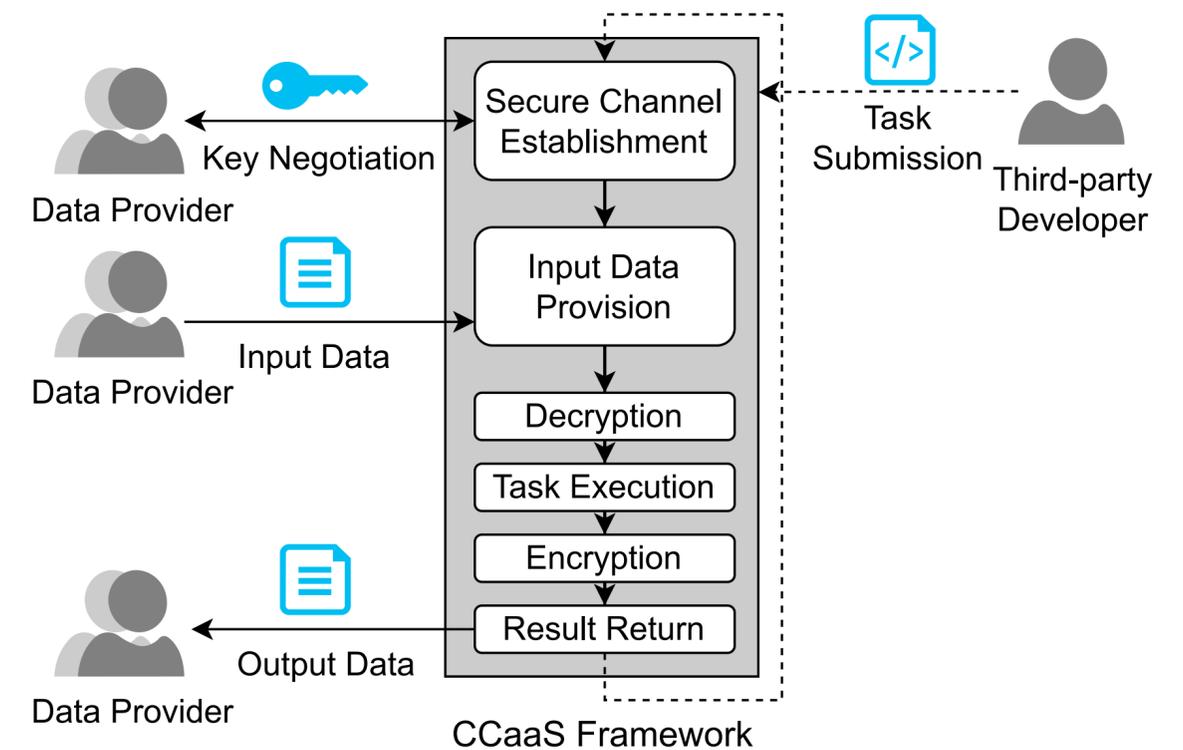
Submitted by 3rd Party Developer:

- CC (Confidential Computing) Task

Pillar of PoCF: Workflow Integrity



Workflow Integrity?



Typestate Specification

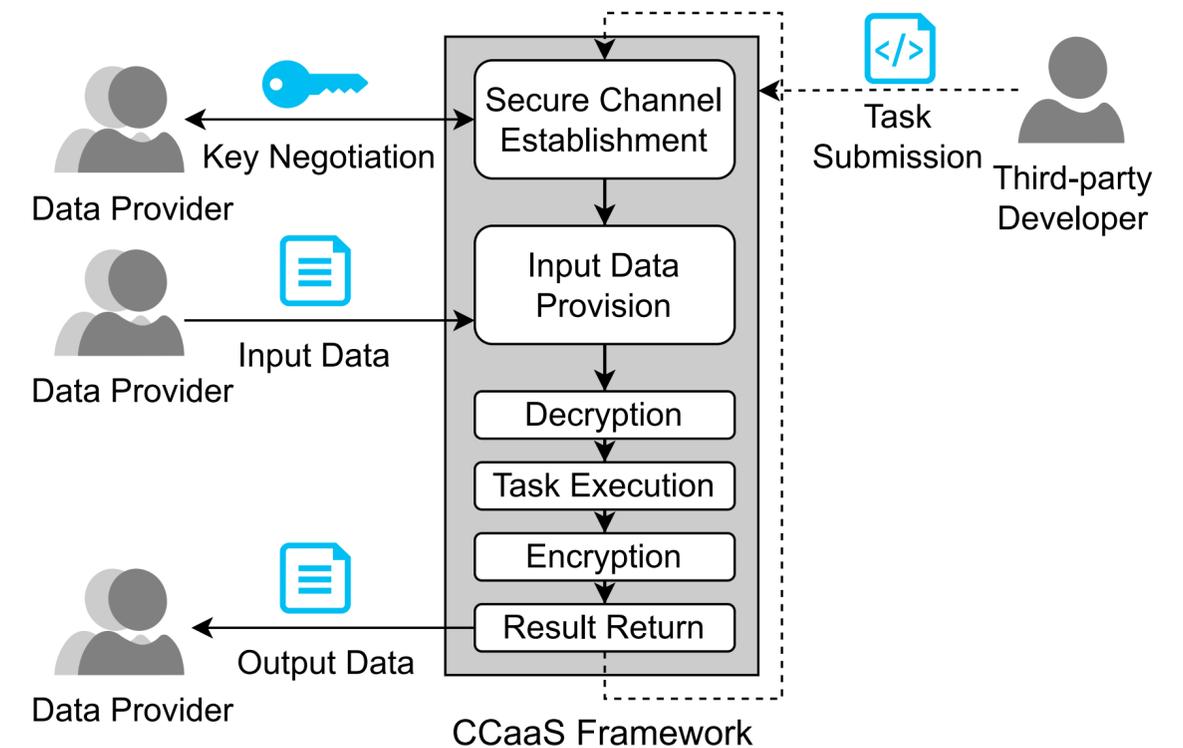
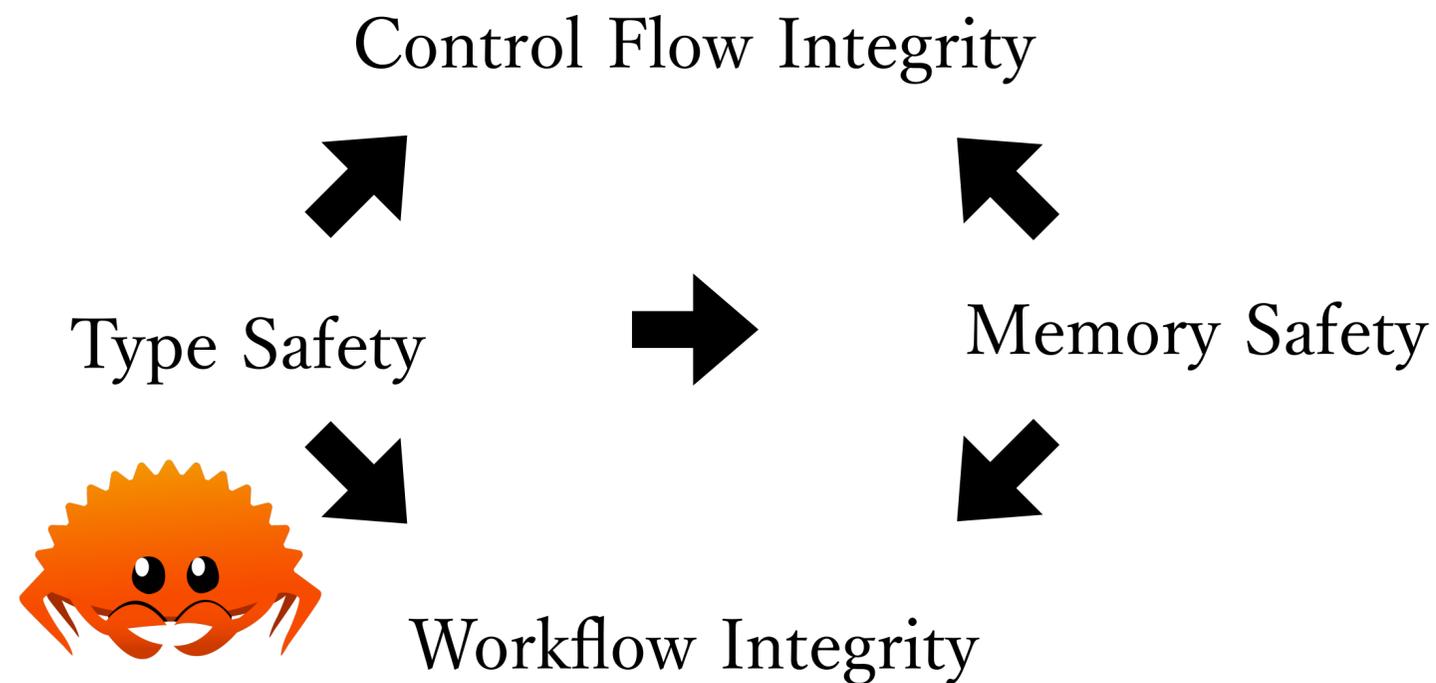
- ✓ Specified.
- ✓ Enforced by Rust.
- ✓ Verified by Prusti.
- ✓ Statically checked.
- ✓ Finally, workflow integrity guaranteed with minor runtime cost!



Listing 1: Typestate abstraction and specification.

```
1 pub struct Task<S, K, D> where
2     S: TaskState + DataState + KeyState,
3     K: Zeroize + Default, D: EncDec<K>,
4     <S as DataState>::State: DState,
5     <S as KeyState>::State: KState,
6 {
7     data: Data<<S as DataState>::State, D, K>,
8     key: Key<K, <S as KeyState>::State>,
9     _state: S,
10 }
11
12 pub trait TaskState {
13     #[pure]
14     fn is_initialized(&self) -> bool {false}
15     #[pure]
16     fn is_finished(&self) -> bool {false}
17     // Other similar functions are omitted.
18 }
19
20 pub struct Initialized;
21 #[refine_trait_spec]
22 impl TaskState for Initialized {
23     #[pure]
24     #[ensures(result == true)]
25     fn is_initialized(&self) -> bool {true}
26 }
27
28 #[ensures((&result)._state.is_allowed_once())]
29 // Other similar specifications are omitted
30 pub fn cc_compute(self) ->
31     Task<ResultEncrypted, Invalid, EncryptedOutput>;
```

Workflow Integrity by Rust & Typestate



NoResidue Instrumentation

- ✓ Heap: modified Memory Allocator
- ✓ Global: not mutable
- ✓ Stack and Registers: Instrumentation

No Residue



Listing 1: Typestate abstraction and specification.

```
1 pub struct Task<S, K, D> where
2   S: TaskState + DataState + KeyState,
3   K: Zeroize + Default, D: EncDec<K>,
4   <S as DataState>::State: DState,
5   <S as KeyState>::State: KState,
6 {
7   data: Data<<S as DataState>::State, D, K>,
8   key: Key<K, <S as KeyState>::State>,
9   _state: S,
10 }
11
12 pub trait TaskState {
13   #[pure]
14   fn is_initialized(&self) -> bool {false}
15   #[pure]
16   fn is_finished(&self) -> bool {false}
17   // Other similar functions are omitted.
18 }
19
20 pub struct Initialized;
21 #[refine_trait_spec]
22 impl TaskState for Initialized {
23   #[pure]
24   #[ensures(result == true)]
25   fn is_initialized(&self) -> bool {true}
26 }
27
28 #[ensures((&result)._state.is_allowed_once())]
29 // Other similar specifications are omitted
30 pub fn cc_compute(self) ->
31   Task<ResultEncrypted, Invalid, EncryptedOutput>;
```

NoLeakage Verification

✓ **Edge function calls:** does not leak secret.

- E.g., OCALL in SGX and call to the hypervisor in SEV

✓ **Static** taint analysis

- Key's tracking: tpestate
- Data tracking: MIRAI's taint analysis

No Leakage



PoCF Verifier



PoCF: Publicly Available

Verifiable

- Once CC Task Submitted: the deployer verifies it.
 1. Pass Verification: PoCF Enclave compiled.
- Data providers:
 1. Obtain the source code.
 2. Conduct verification.
 3. Calculate measurement.
 4. Feed data.
- Trusted builder: proprietary code.



Evaluation

Summary of Evaluation Results

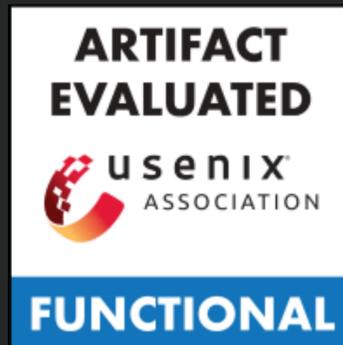
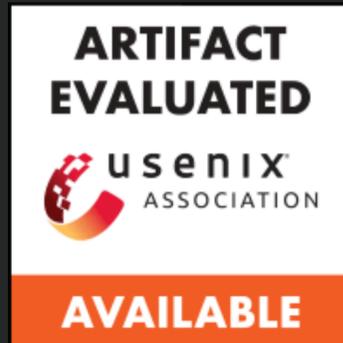
1. PoCF reaches its design goals.
2. PoCF incurs negligible overhead in CPU-bound tasks.
3. PoCF exhibits degradation in IO-bound tasks (lack of IO optimizations).
4. The data flow tracking tool is not very accurate.





Questions?

You're welcome to try and star our artifact!



Github: [ya0guang/PoBF](https://github.com/ya0guang/PoBF)



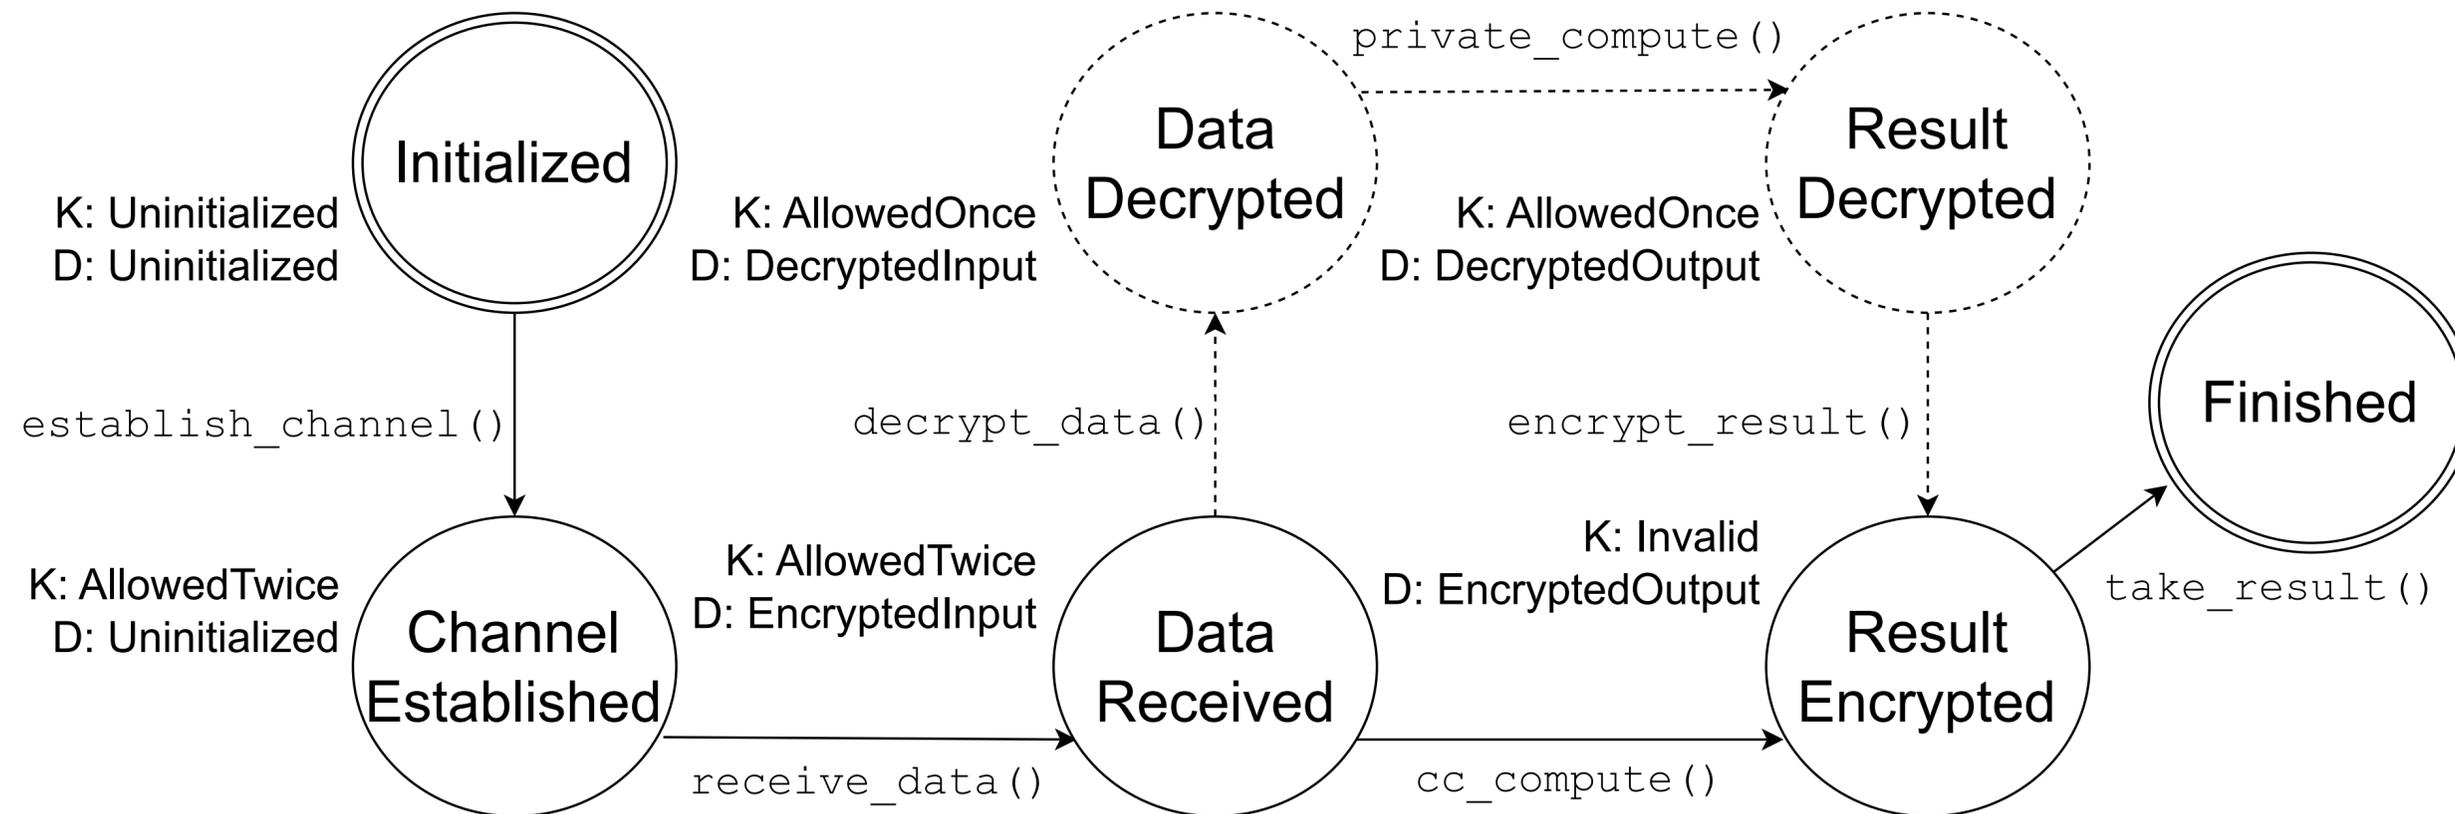


Thanks!



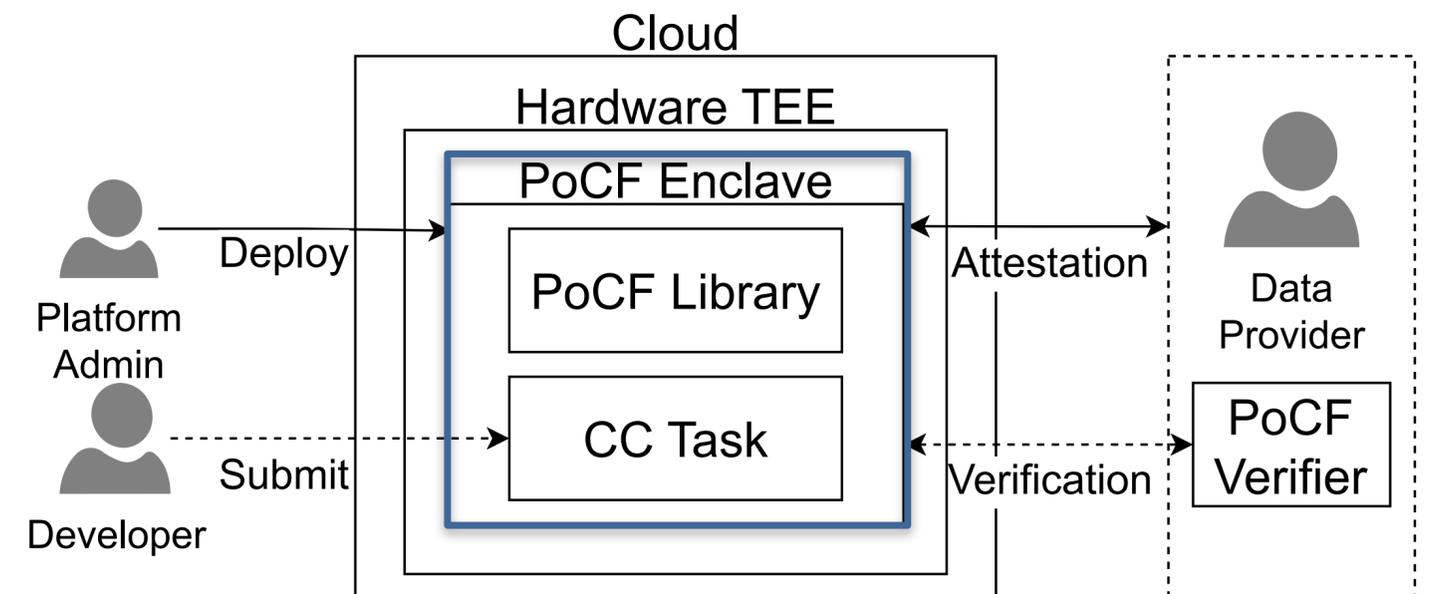
Backup Slides

PoCF Library: TEE-Agnostic State Machine



PoCF Enclave: TEE-Specific Enclave Service

- Intel SGX
 - DCAP & EPID Attestation
 - Teaclave (Rust) SGX SDK
 - ECALL & OCALL
- AMD SEV on Azure
 - Azure Attestation Service
 - Standard Library



Effortless Porting

- Verifier invocations wrapped.
- Seamless use of standard library

```
1 #[macro_export]
2 macro_rules! ocall_log {
3     ($str: expr) => {
4         let s = alloc::format!($str);
5         log(s)
6     };
7     ($formator:expr, $($arg:expr),+ $(,)? ) => {
8
9         let s = alloc::format!($formator, $($arg),+);
10        log(s)
11    };
12 }
13
14 #[macro_export]
15 macro_rules! println {
16     () => {
17         ocall_log!("\n")
18     };
19     $($arg:expr),+ $(,)? ) => {
20         $(
21             #[cfg(mirai)]
22             verify!(does_not_have_tag!($arg, SecretTaint));
23         )*
24         ocall_log!($($arg),+);
25     }
26 }
```



Taint Analysis: Accuracy of MIRAI

Table 4: The precision test of MIRAI categorized by Rust features.

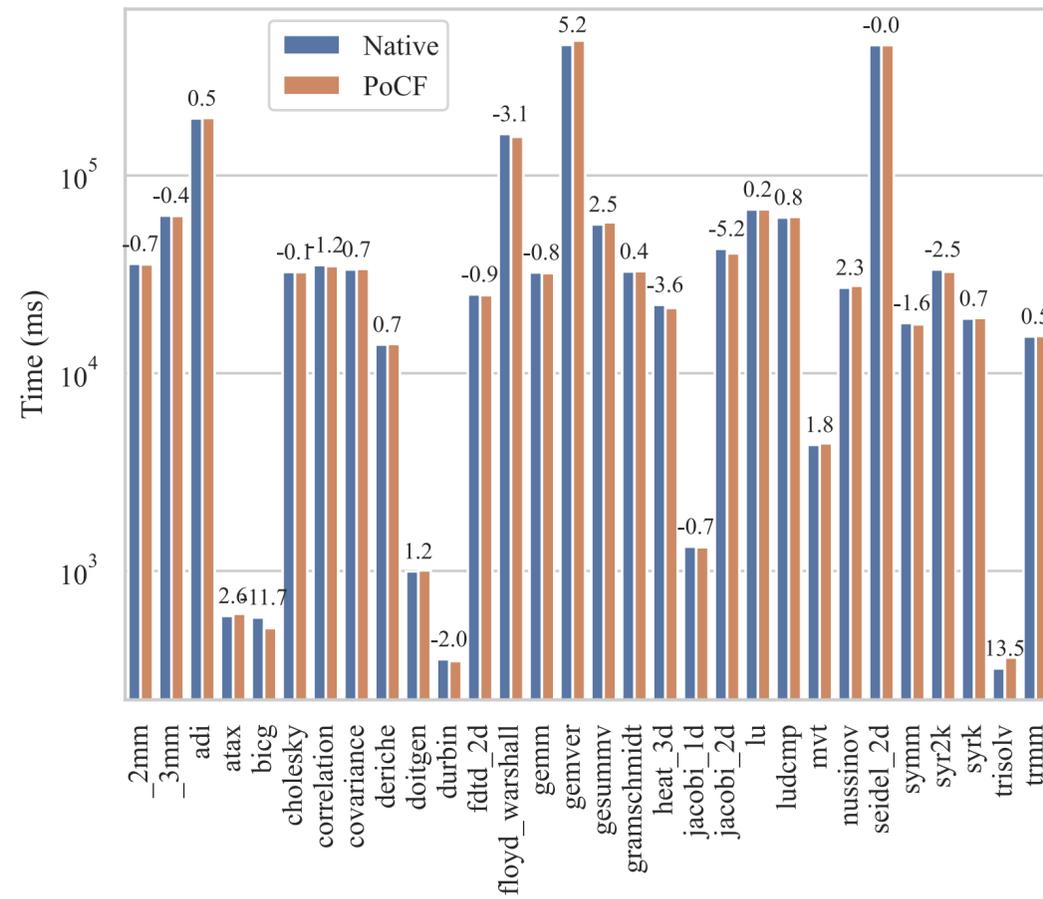
Test Name	Covered Rust Features	Expected	Actual	Missed Feature(s)
untrusted_input	Traits, generics, and arrays	✓	✓	/
control_flows	Loops, branches, and pattern matches	✗: 1; ○: 5	○: 6	/
ownership_transfer	Ownership and borrow check	✗: 2	✗: 2	/
pointers	Smart and raw pointers	✗: 4	✗: 1	Leakage by Rc<T>, Box<T>, and *const T.
complex_structs	Collections and structs	✗: 4	✗: 1	Tag propagation from field to the whole struct

All the tests were analyzed by MIRAI using its strictest analysis level, i.e., MIRAI_FLAG=diag=paranoid.

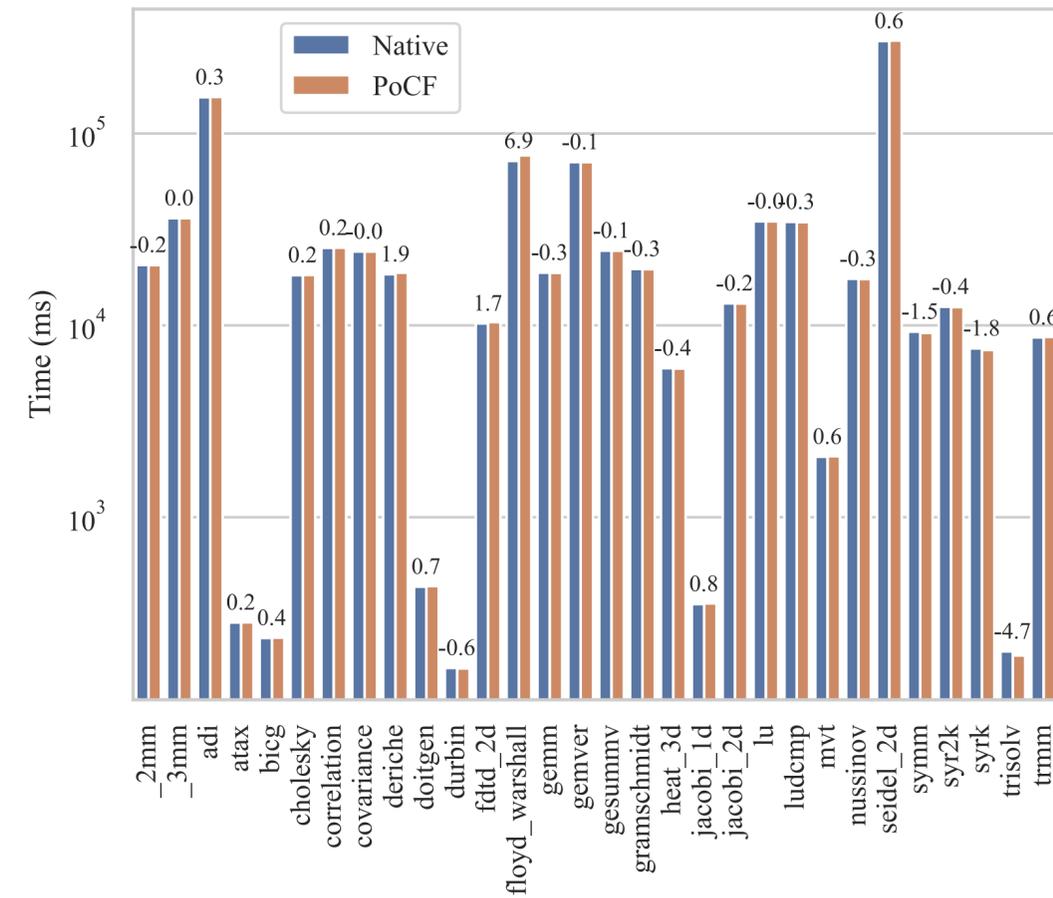
✓: No data leakage; ✗: Has data leakage; ○: Possible data leakage. The number behind “✗” or “○” denotes the number of data leakages.



Microbenchmark: Polybench

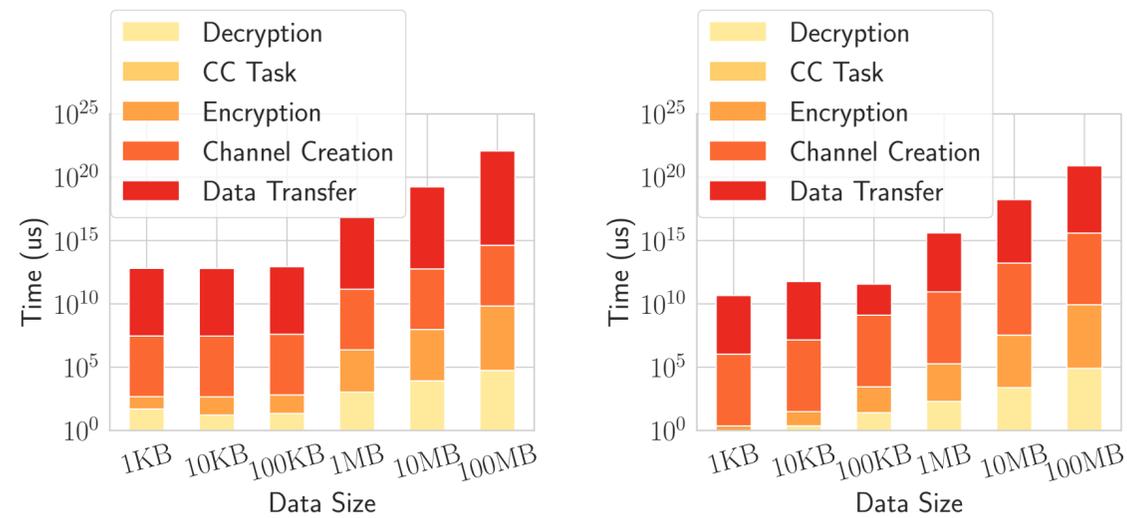


(a) Polybench: Performance of POCF and NATIVE on SGX.



(b) Polybench: Performance of POCF and NATIVE on SEV.

Microbenchmark: Overhead Analysis



(a) Cost breakup of PoCF on SGX. (b) Cost breakup of PoCF on SEV.

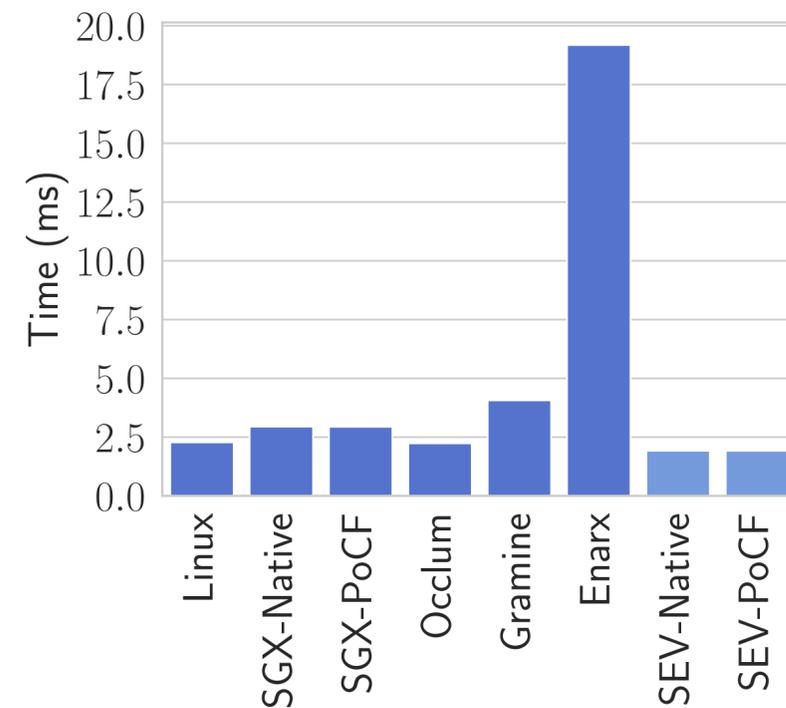
Figure 5: Identity task: Performance breakup of PoCF.

Table 5: Identity Task: Time (ms) under Different Data Sizes.

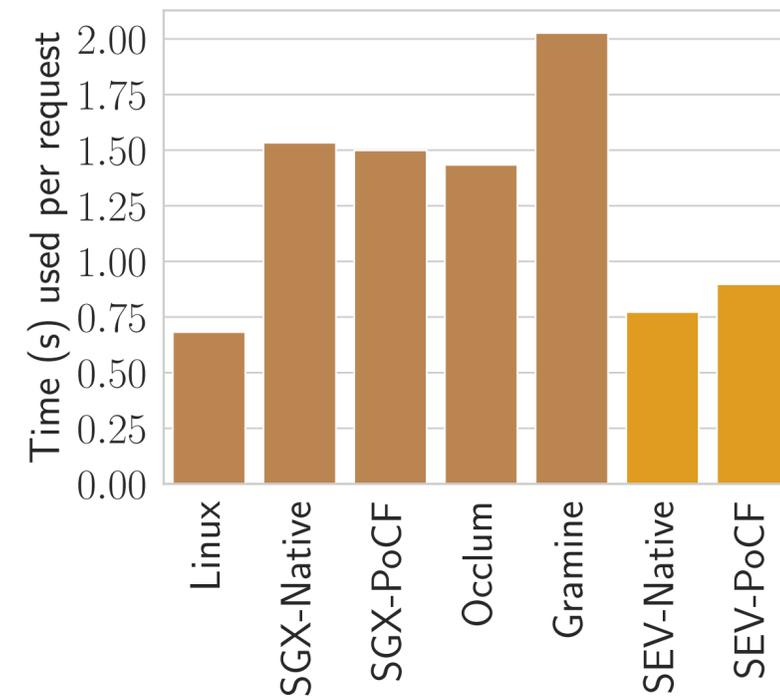
Config	1KB	10KB	100KB	1MB	10MB	100MB
NATIVE X	275.8	281.1	296.3	536.7	3026.5	28018.3
P w/o T X	278.3	280.4	298.6	541.1	3033.9	28022.9
P w/ T X	277.3	287.4	301.7	545.0	3043.7	28215.0
NATIVE V	489.1	487.3	449.7	495.6	502.0	923.3
PoCF V	489.5	492.3	454.4	499.8	506.5	934.8

P: PoCF without data flow tracking; T data flow tracking; X: SGX; V: SEV

Macrobenchmark: AI Inference



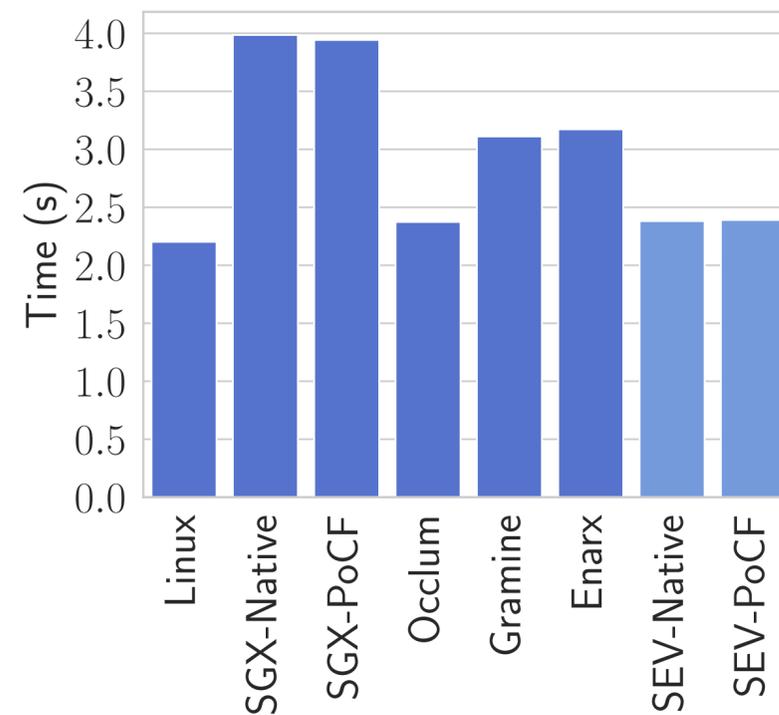
(a) Single-threaded.



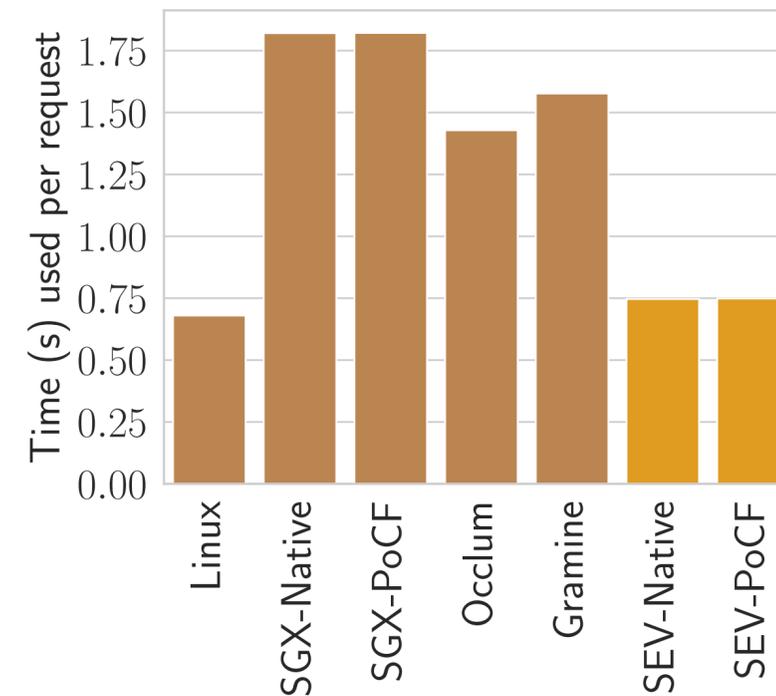
(b) Multi-threaded.

Figure 7: Macrobenchmark: AI inference execution time.

Macrobenchmark: FASTA



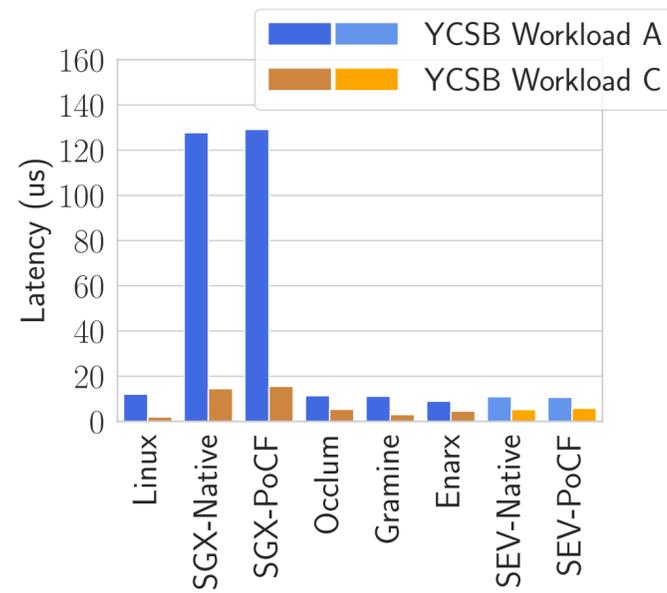
(a) Single-threaded.



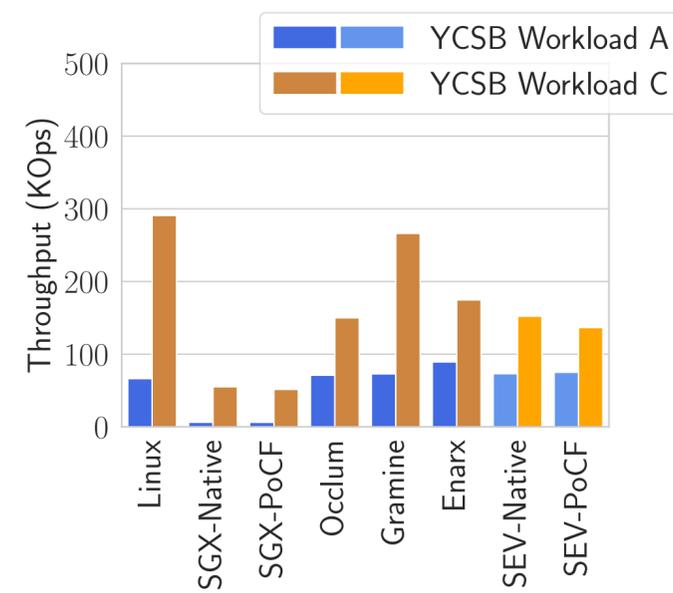
(b) Multi-threaded.

Figure 8: Macrobenchmark: FASTA execution time.

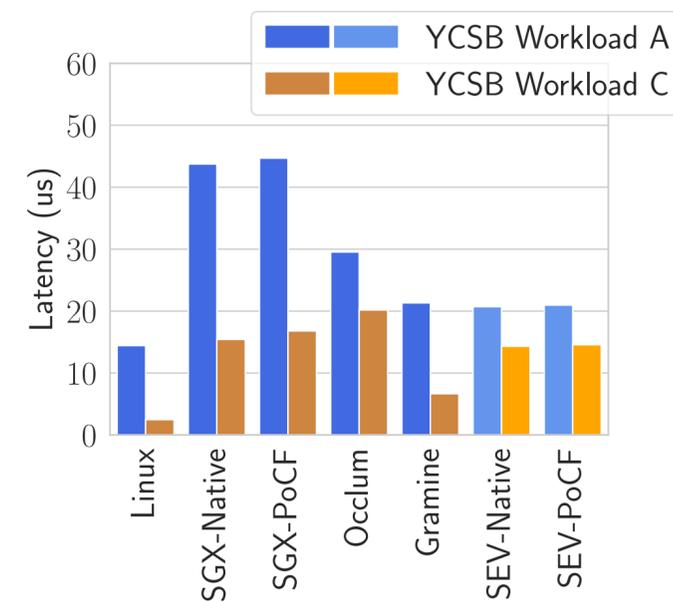
Macrobenchmark: In-memory KVDB



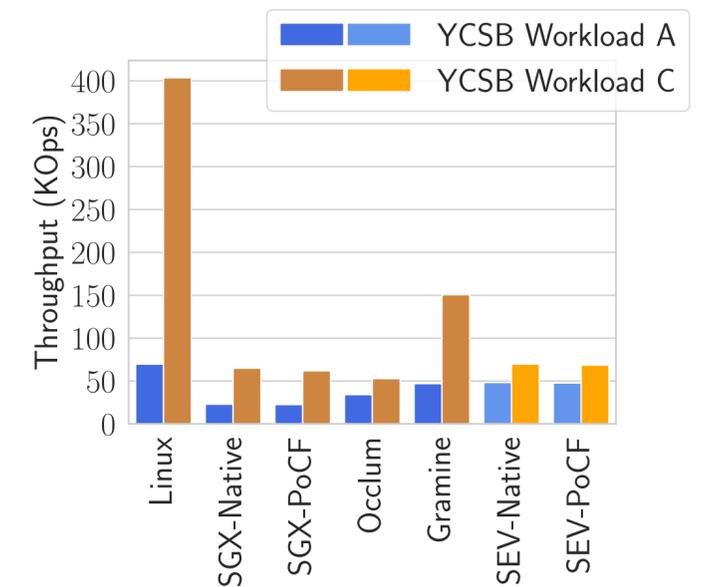
(a) Single-Thread Latency.



(b) Single-Thread Throughput.



(c) Multi-Thread Latency.



(d) Multi-Thread Throughput.