

Better Foundations for Secure Software using Trusted Hardware & Verification

Shweta Shinde
ETH Zürich

Security breaches are on the rise

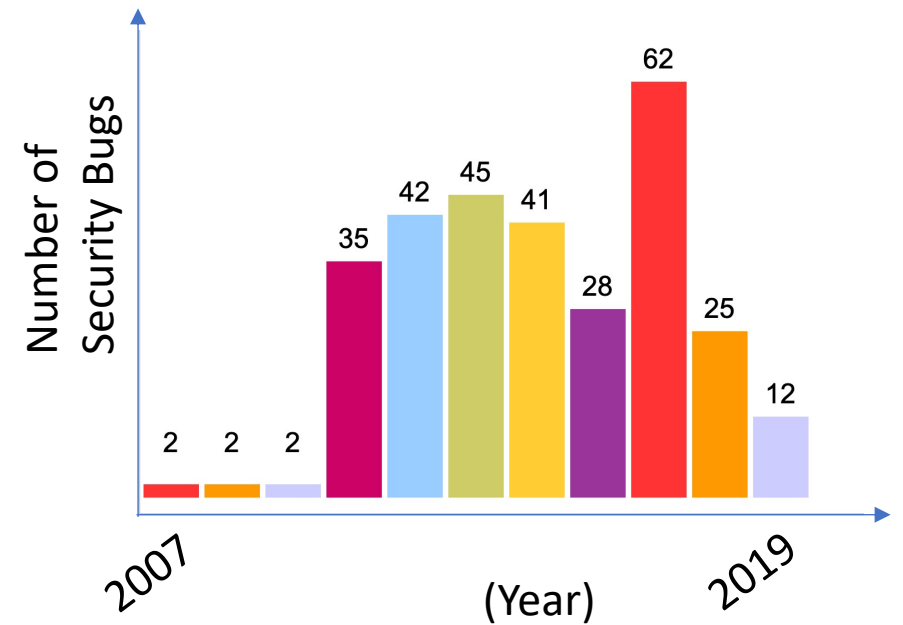

37 billion records exposed
through data breaches
in 2020

Average of 1-25 Bugs per 1000 lines of code


The new normal: Hundreds of bugs a year in Linux

Operating System (Linux Kernel)
27 Million Lines

Hypervisor (XEN)
0.5 Million Lines

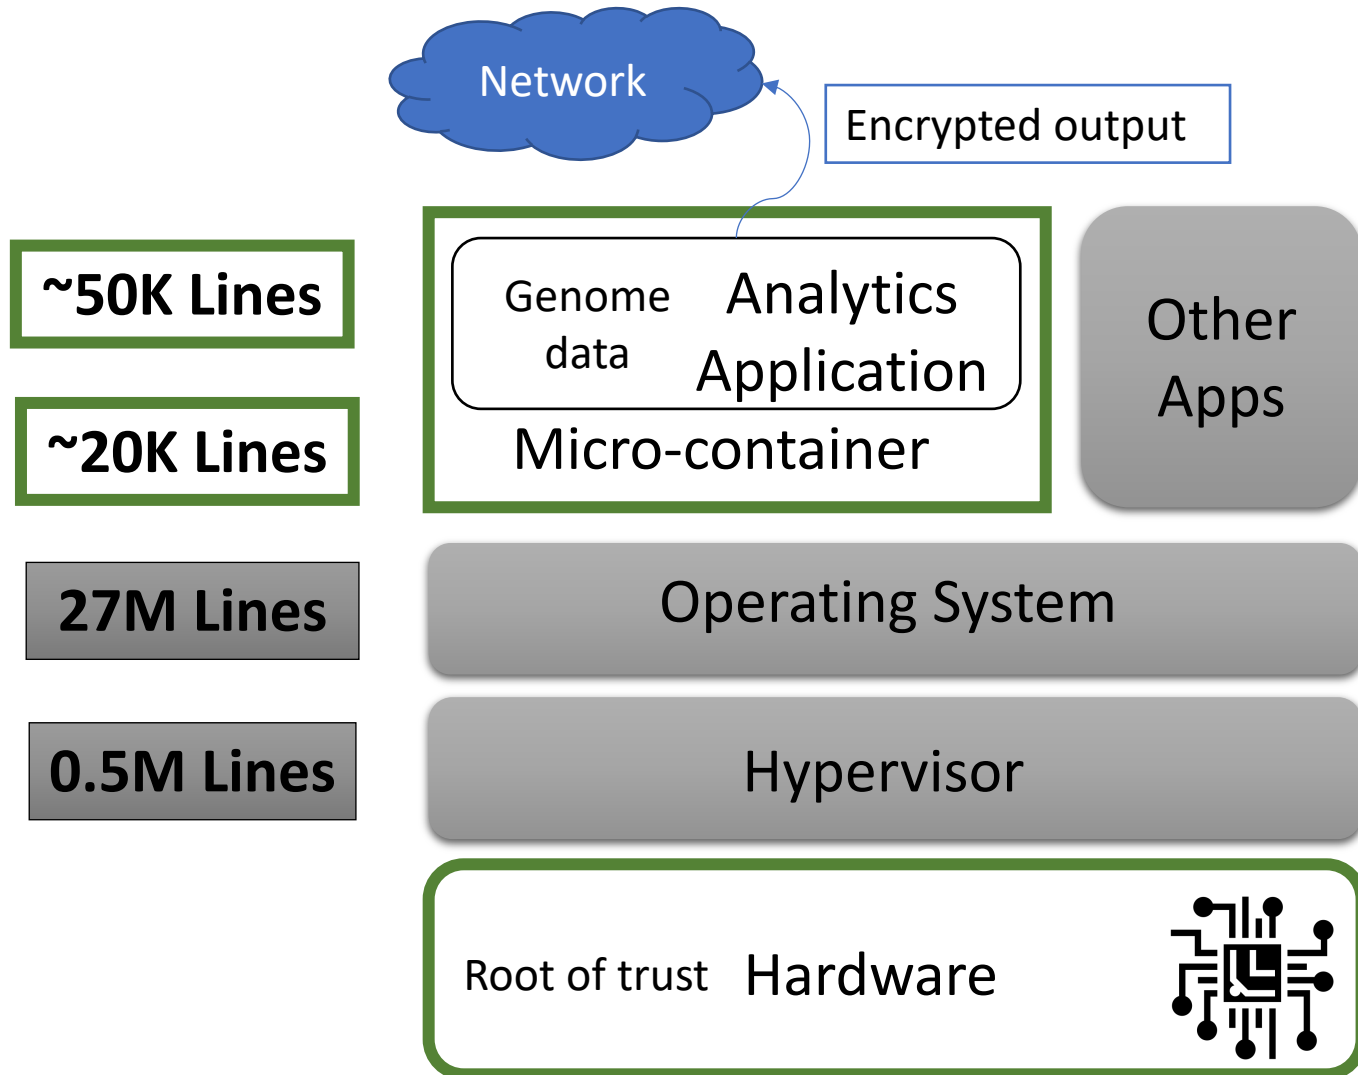


Current computing stack is prone to attacks



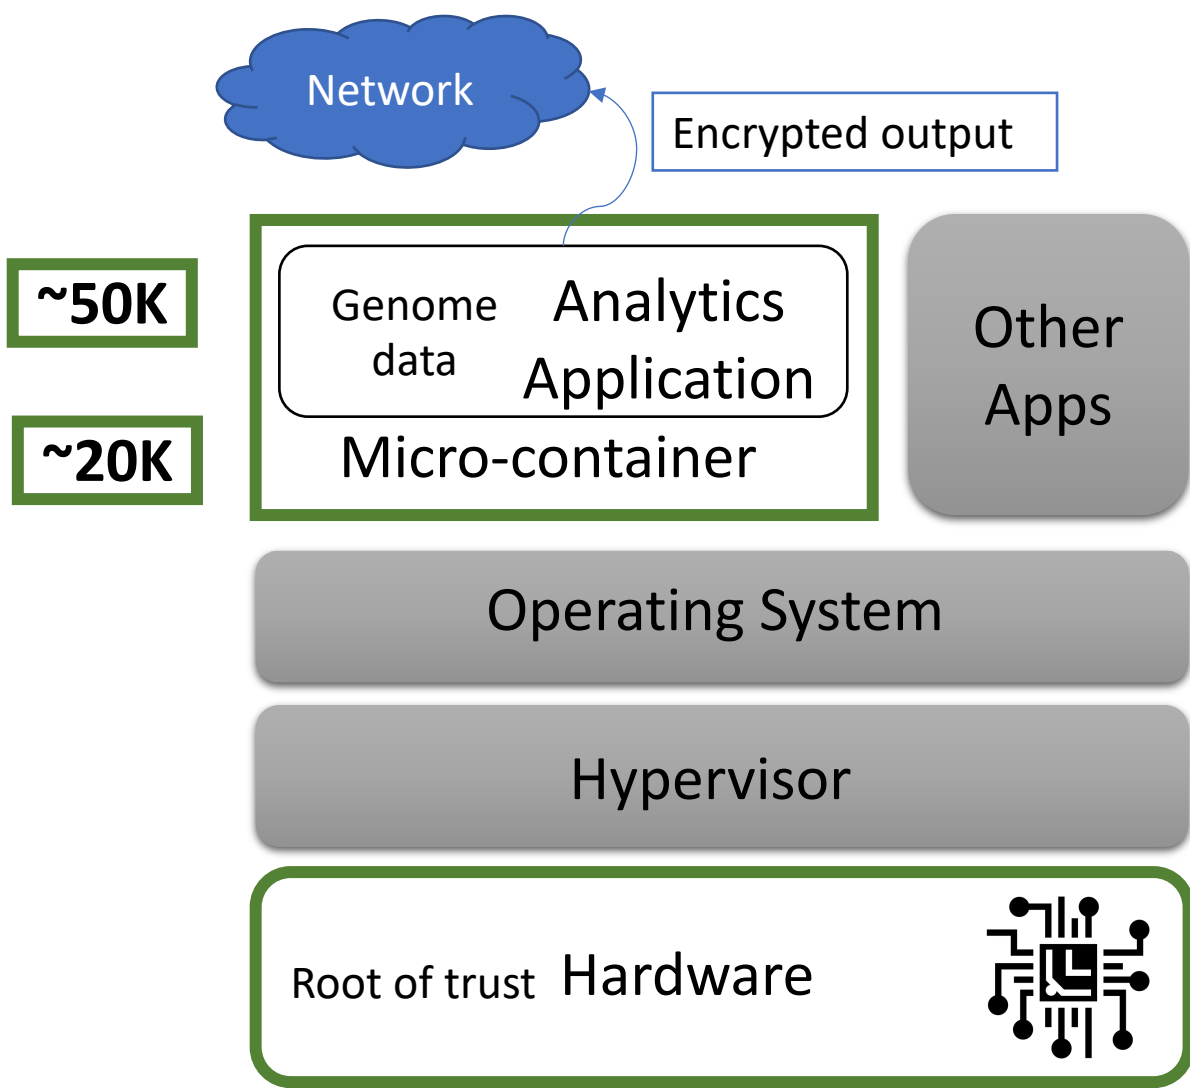
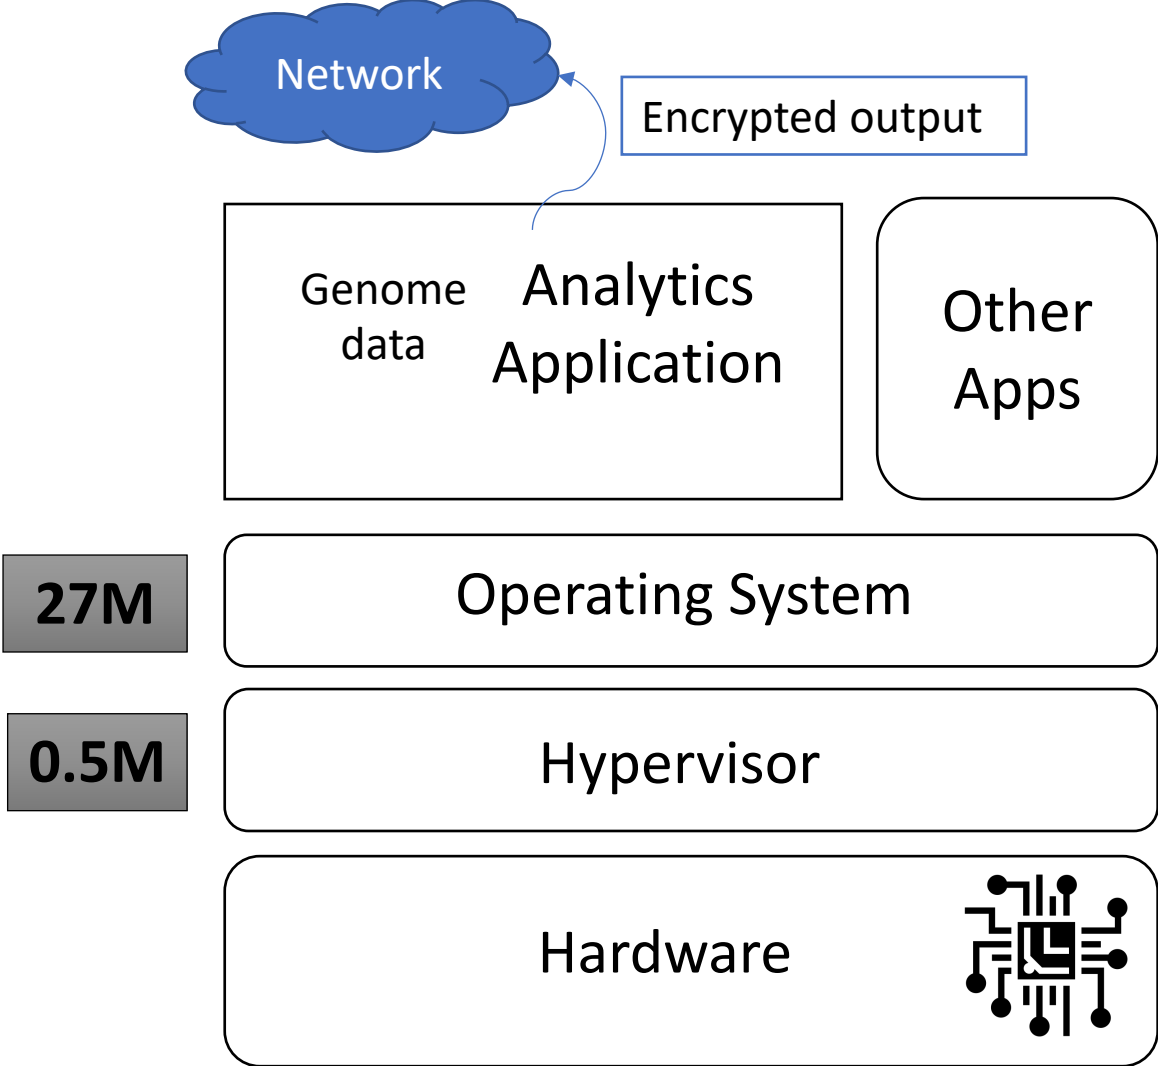
- Encryption or other sophisticated techniques at the application layer
- Bug in lower layers → Compromise the security of the app
- Large size → High probability

Computation stack for the decades to come



- Thin layer for running applications
- Trusted hardware
- Formal guarantees for defense against
 - Third party attacks
 - Internal bugs in the app

Design Contrast



Building the components of this stack

New Applications [Arxiv'18], [ICDCS'19]

Secure Computation [CCS'13]

Analysis & hardening

[PLDI'14], [FSE'15], [NDSS'19], [CCS'20]

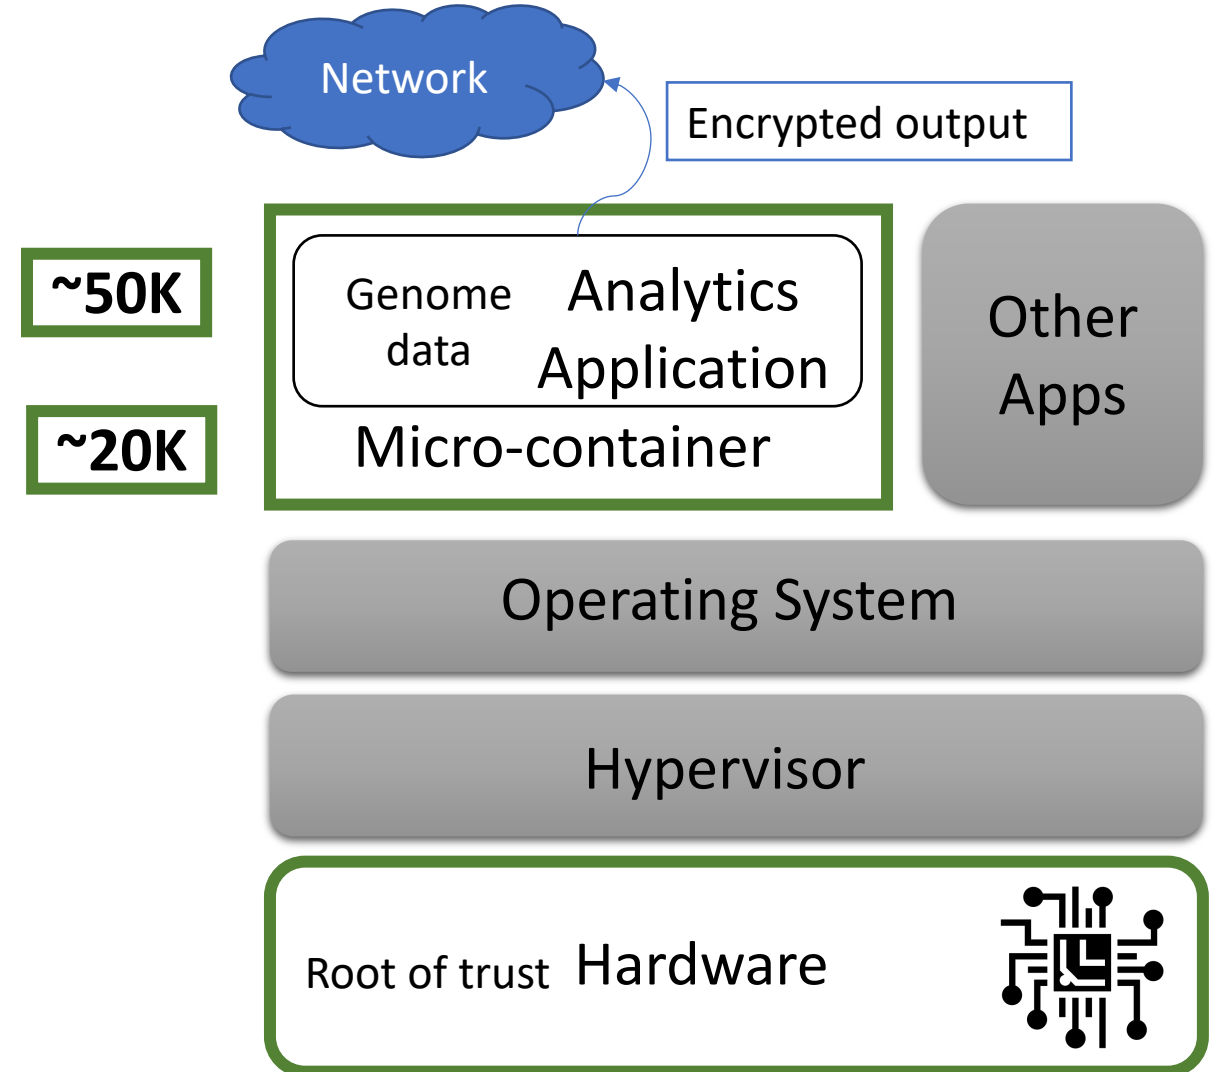
Rich functionality [NDSS'17], [Usenix'22]

Formal verification [Usenix Security'20]

Attacks & Defenses [AsiaCCS'16] [CCS'21]

Trusted Computing Primitives

[TR'15], [Eurosys'20]



Practical Relevance: Initial Adoption

New Applications

Microsoft, Largest Asia-Pacific ISP

Secure Computation

SAP Labs

Analysis & hardening

Dexecure

Rich functionality

Anqlave, Anquan, Community

Formal verification

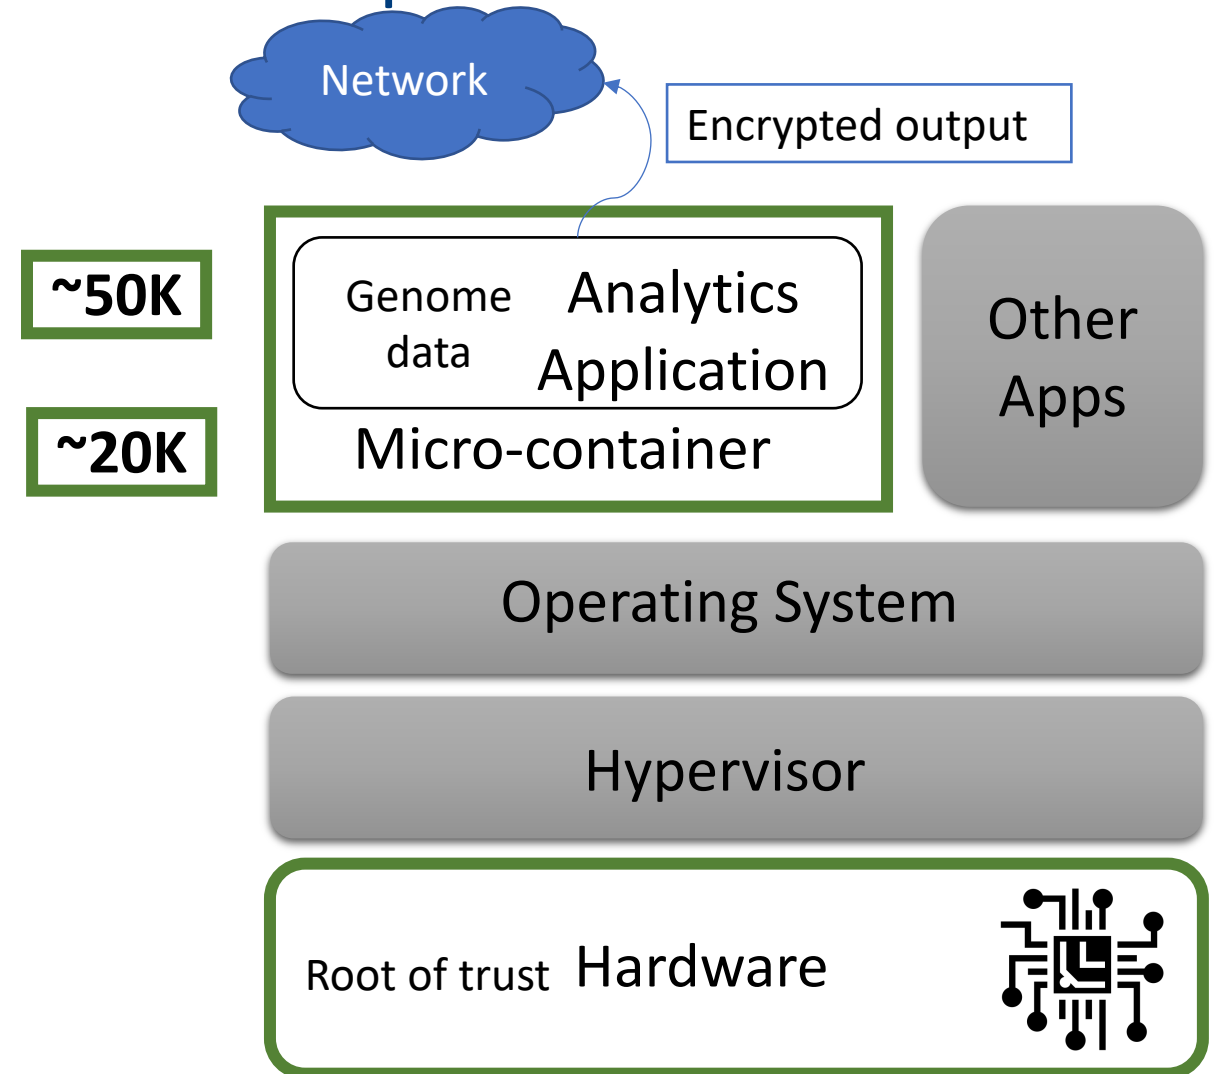
Intel, Google, Microsoft, Anqlave

Attacks & Defenses

Intel, Community

Trusted Computing Primitives

Qualcomm, Seagate, Baidu, Community



1st Component of this stack

New Applications [Arxiv'18], [ICDCS'19]

Secure Computation [CCS'13]

Analysis & hardening

[PLDI'14], [FSE'15], [NDSS'19], [CCS'20]

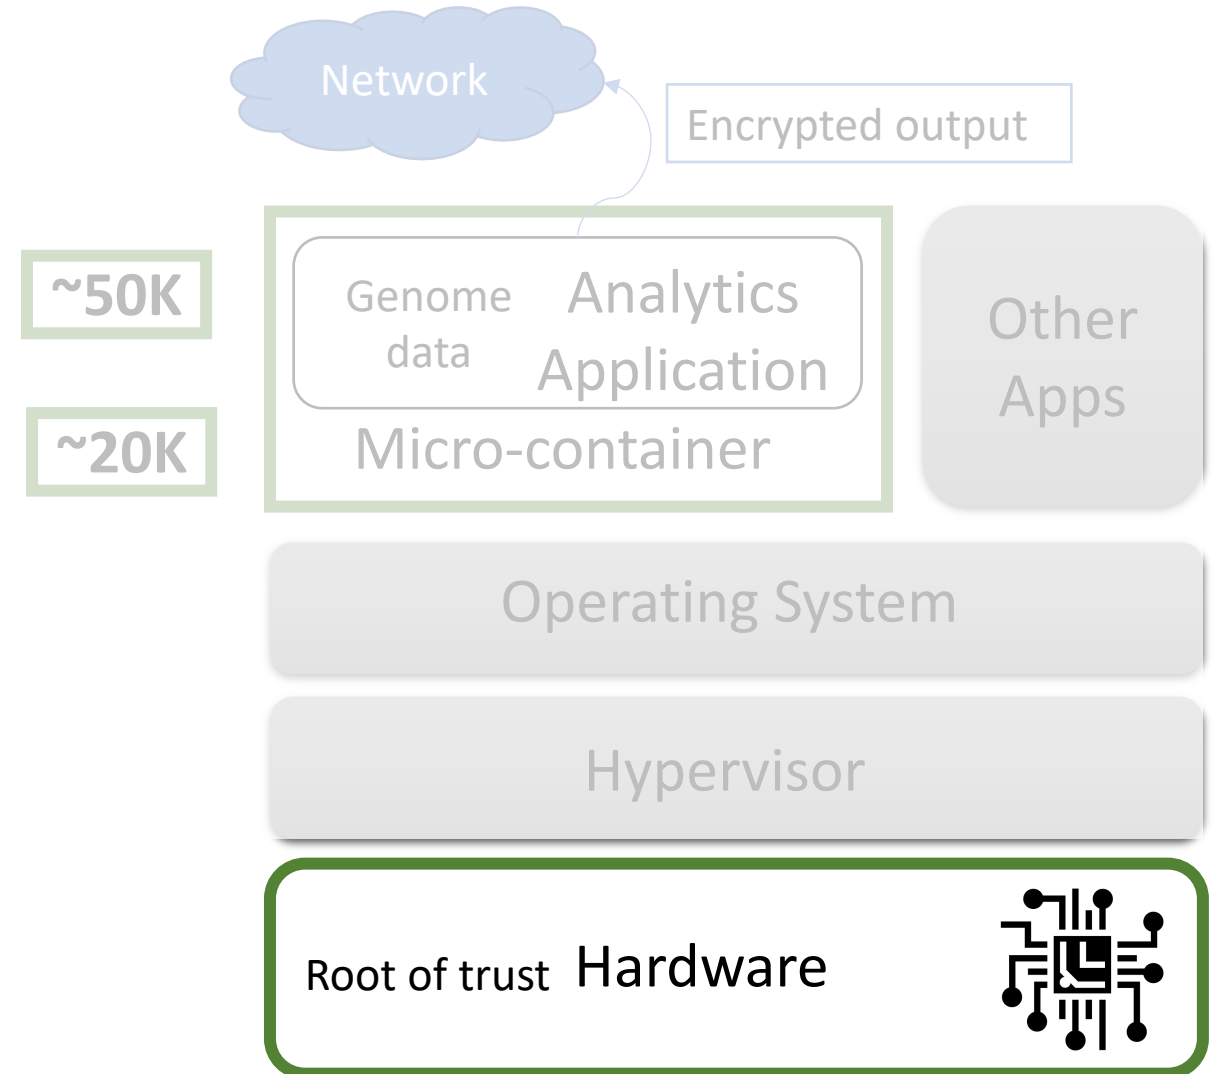
Rich functionality [NDSS'17], [Usenix'22]

Formal verification [Usenix Security'20]

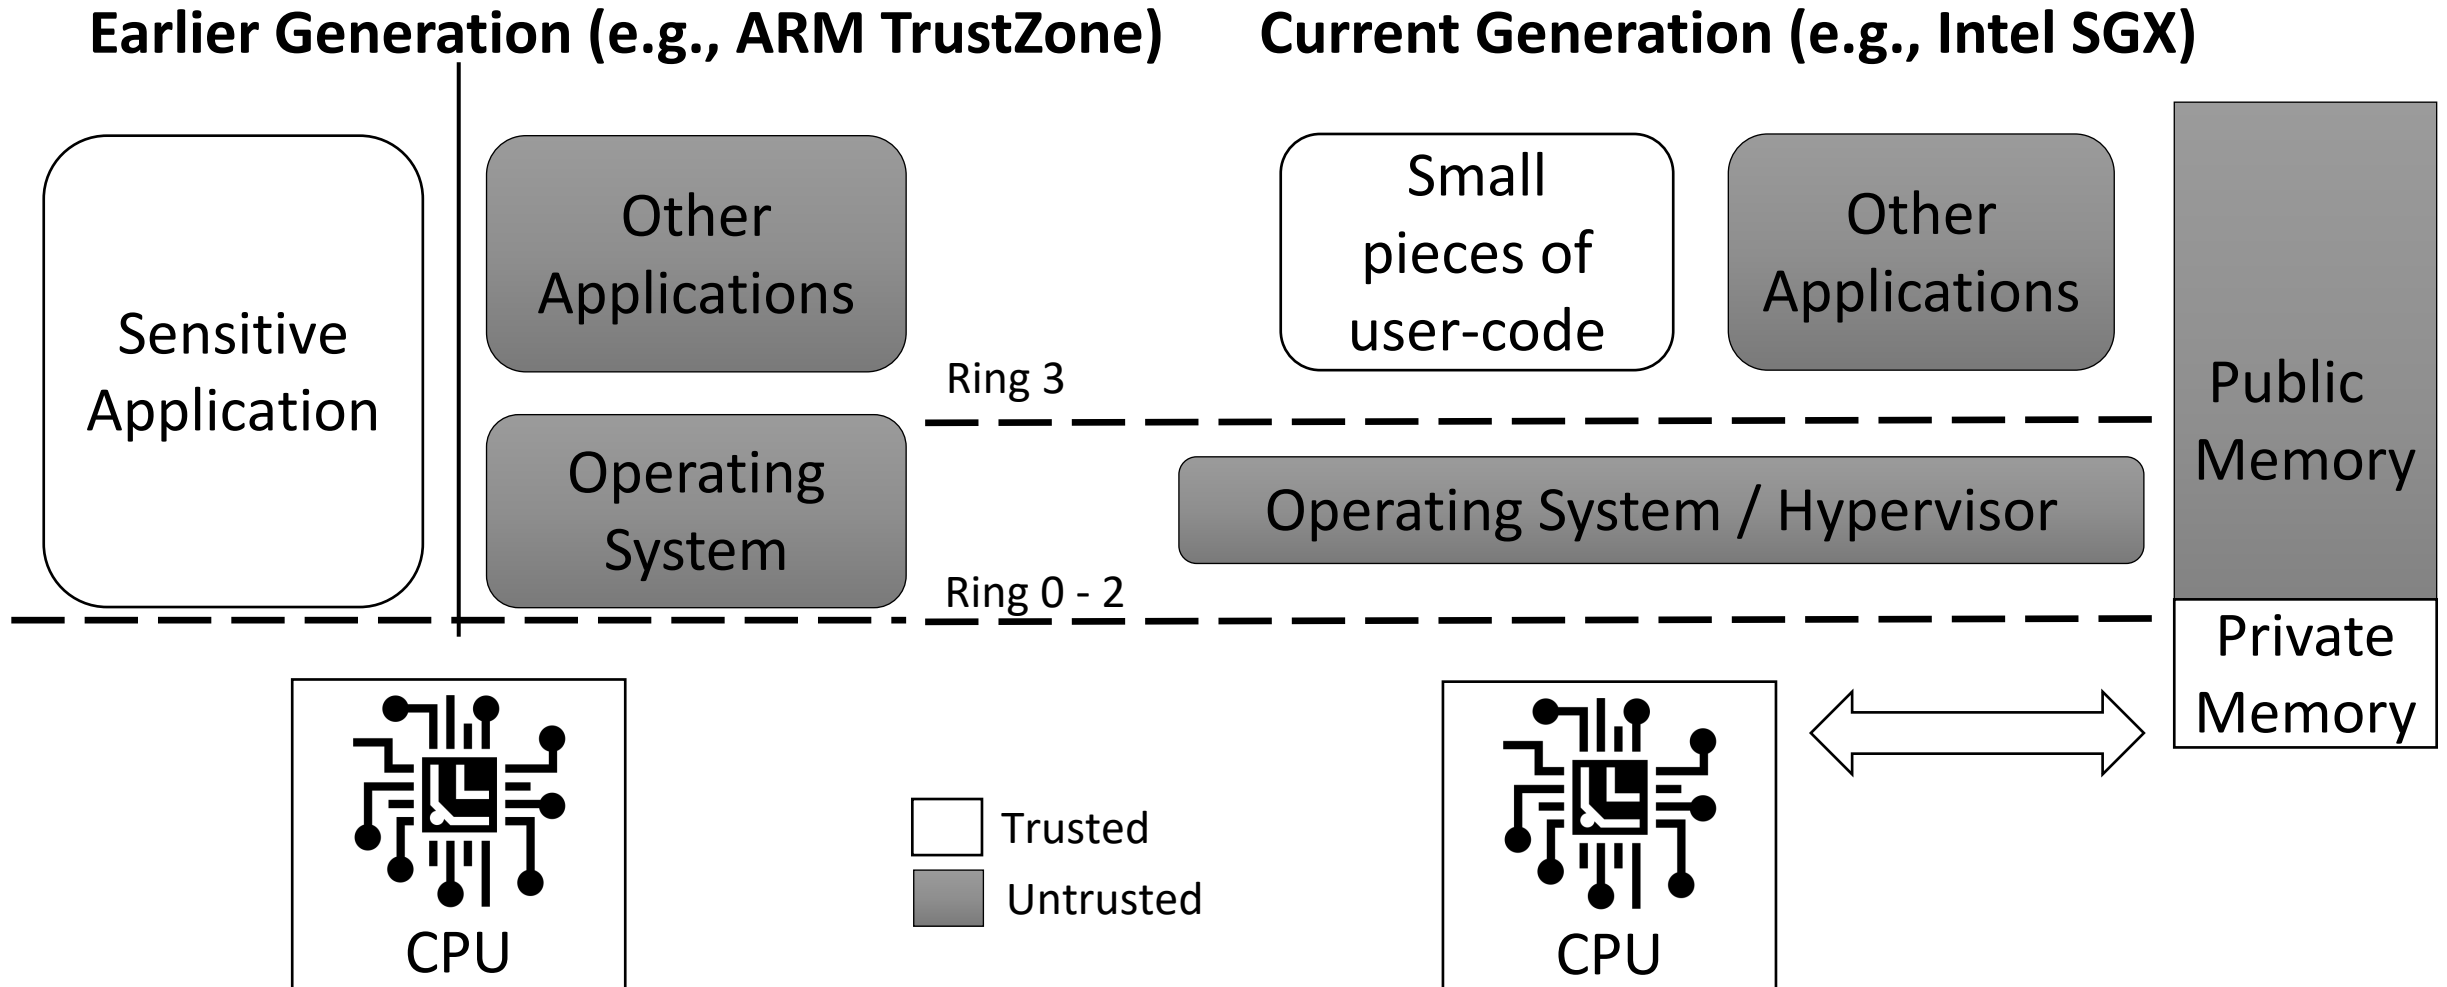
Attacks & Defenses [AsiaCCS'16] [CCS'21]

Trusted Computing Primitives

[TR'15], [Eurosys'20]



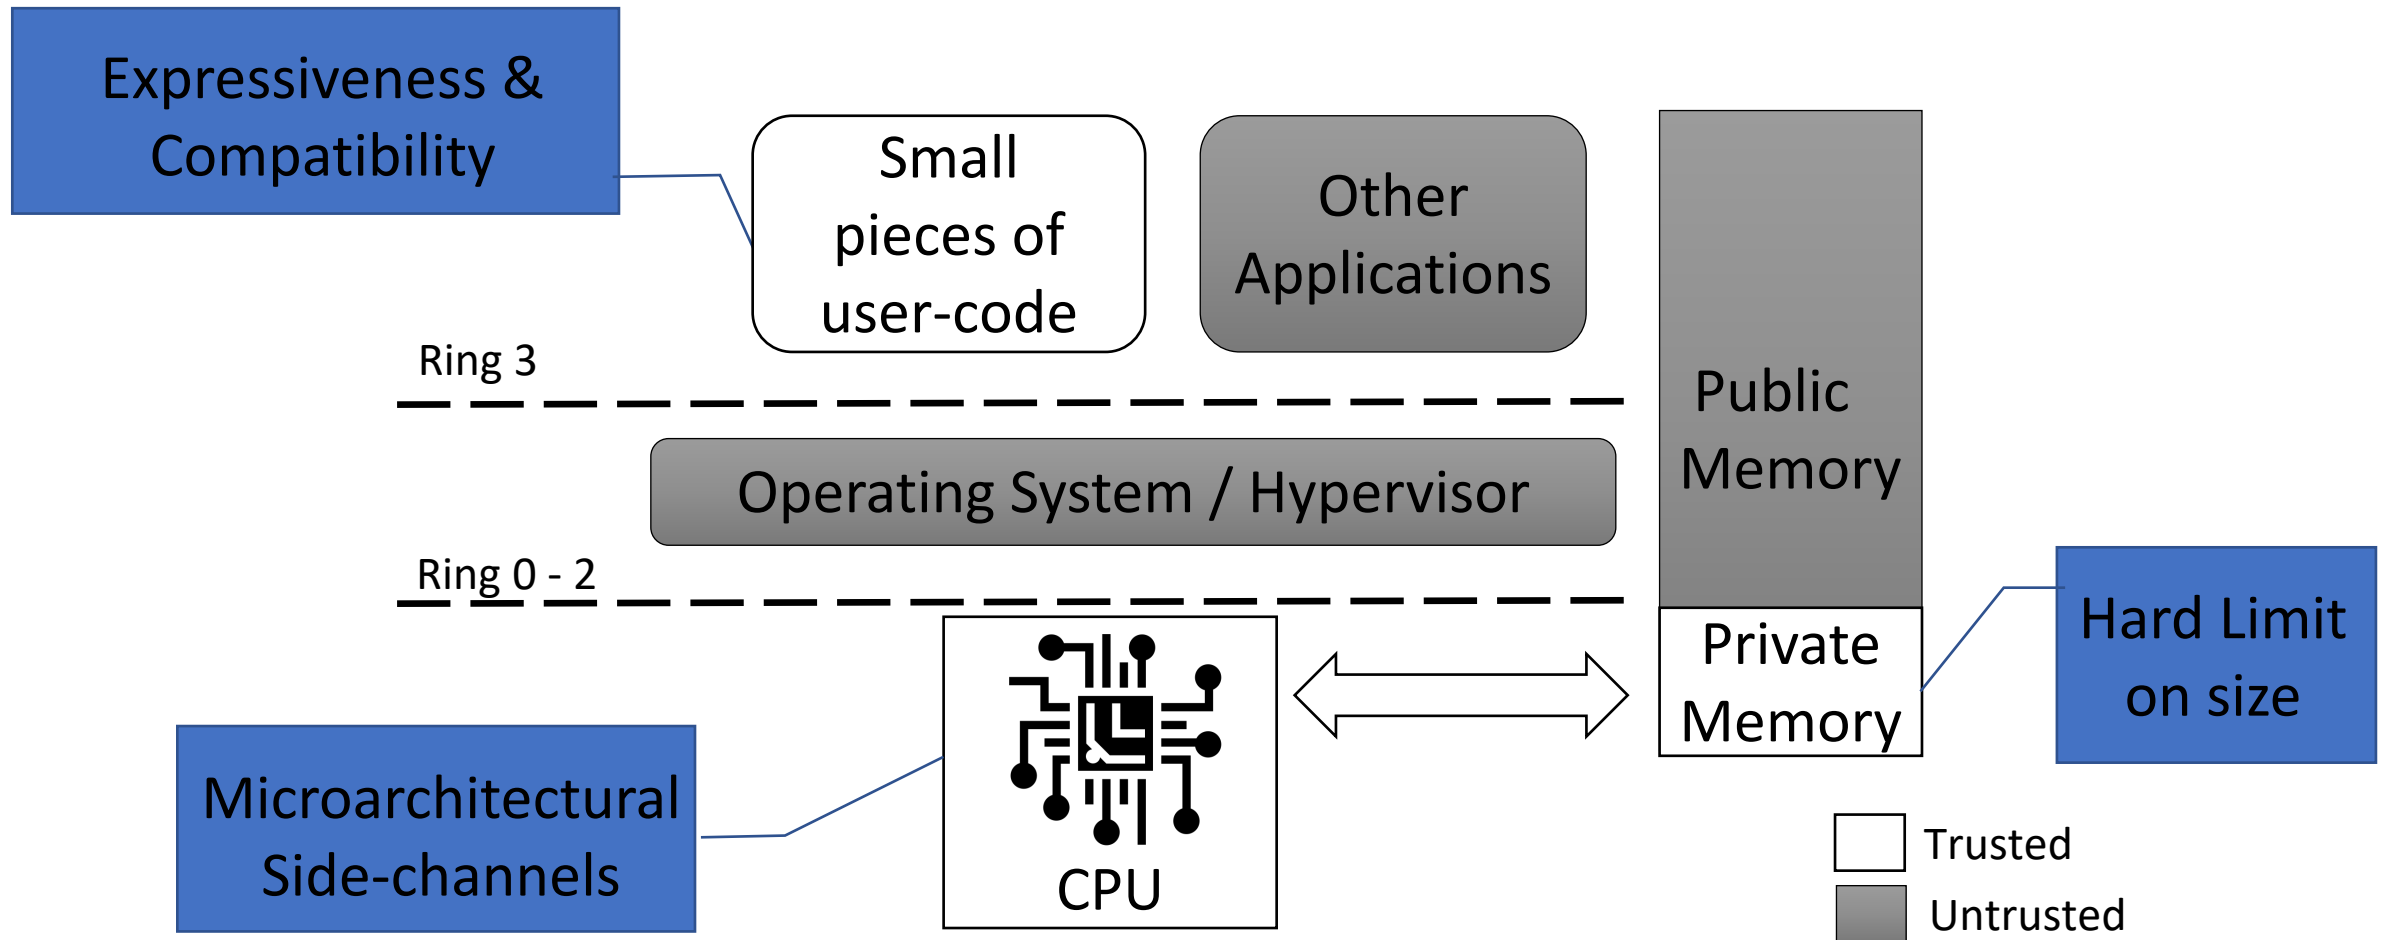
Trusted Execution Environments (TEEs)



Inflexible Design & Closed Implementation

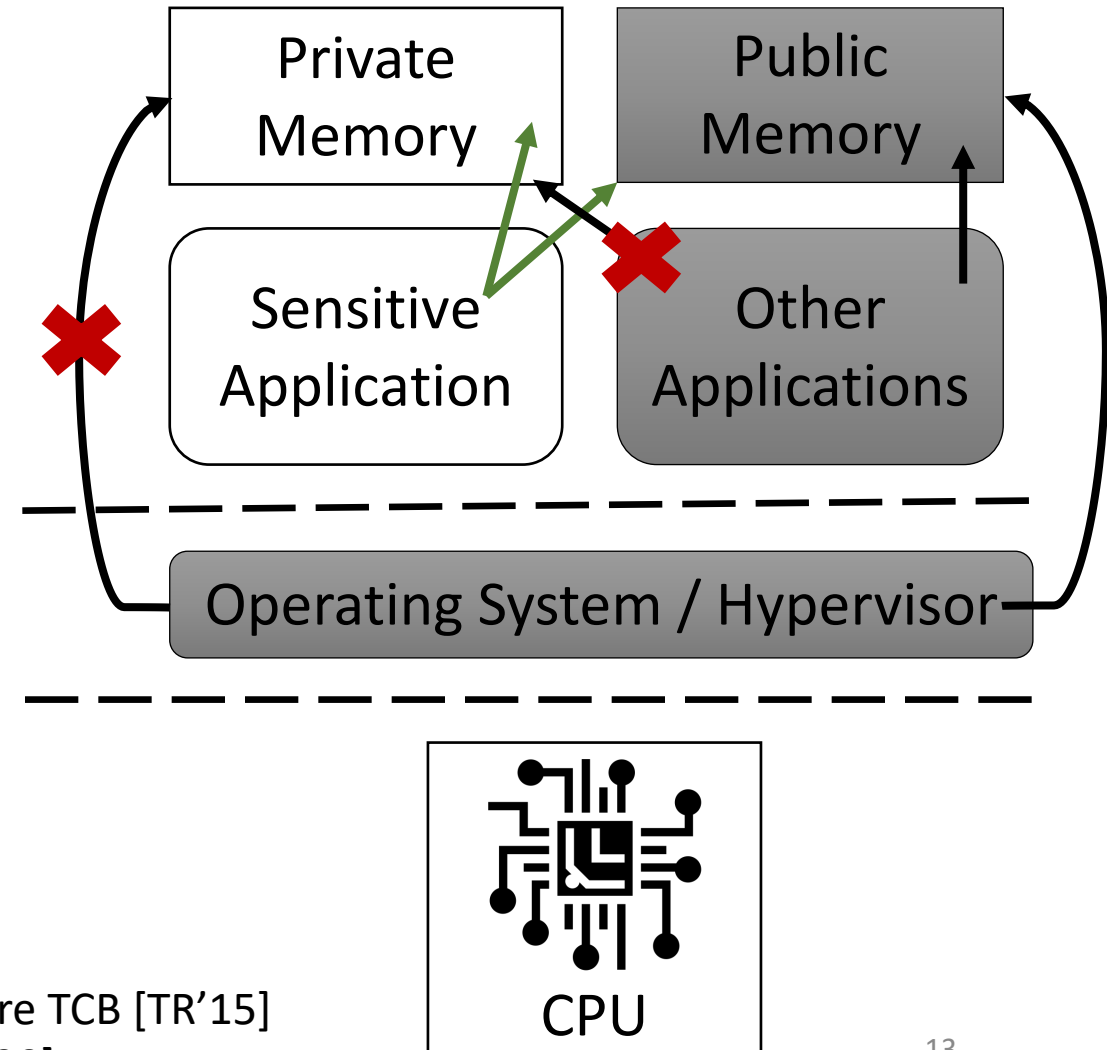
- TEEs in commercial hardware: Intel SGX, ARM TrustZone, AMD SEV
- One particular design point in the space
 - Intel SGX – small server/desktop apps (e.g., DRM, cryptography)
 - ARM TZ – vendor-provisioned mobile apps (e.g., fingerprint, ledger)
 - AMD SEV – full VM isolation only (e.g., cloud computing)
- Implemented on closed-source hardware
 - Slow iteration dictated by a company
 - Adding new features/defenses is cumbersome

Limitations of Commercial TEEs



Better TEEs

- Main Observation:
 - Physical memory isolation
 - Simpler ways to achieve
- Similar abstraction to Intel's TEE
- Novelty: Designed to maintain
 - Compatibility
 - Performance



Focus on commercial TEEs (e.g., Intel SGX),
since they are widely available

2nd component of this stack

New Applications [Arxiv'18], [ICDCS'19]
Secure Computation [CCS'13]

Analysis & hardening
[PLDI'14], [FSE'15], [NDSS'19], [CCS'20]

Rich functionality [NDSS'17], [Usenix'22]

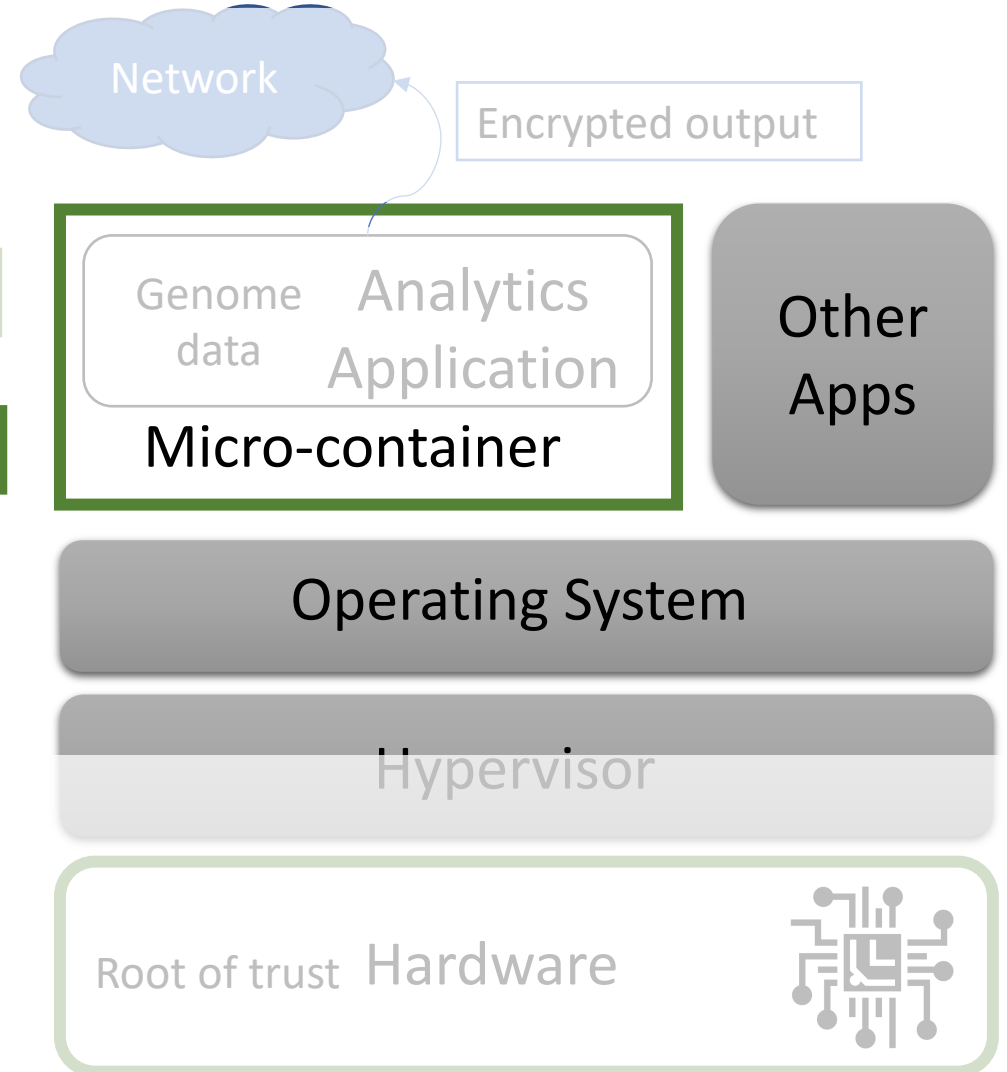
Formal verification [Usenix'20]

Attacks & Defenses [AsiaCCS'16, CCS'21]

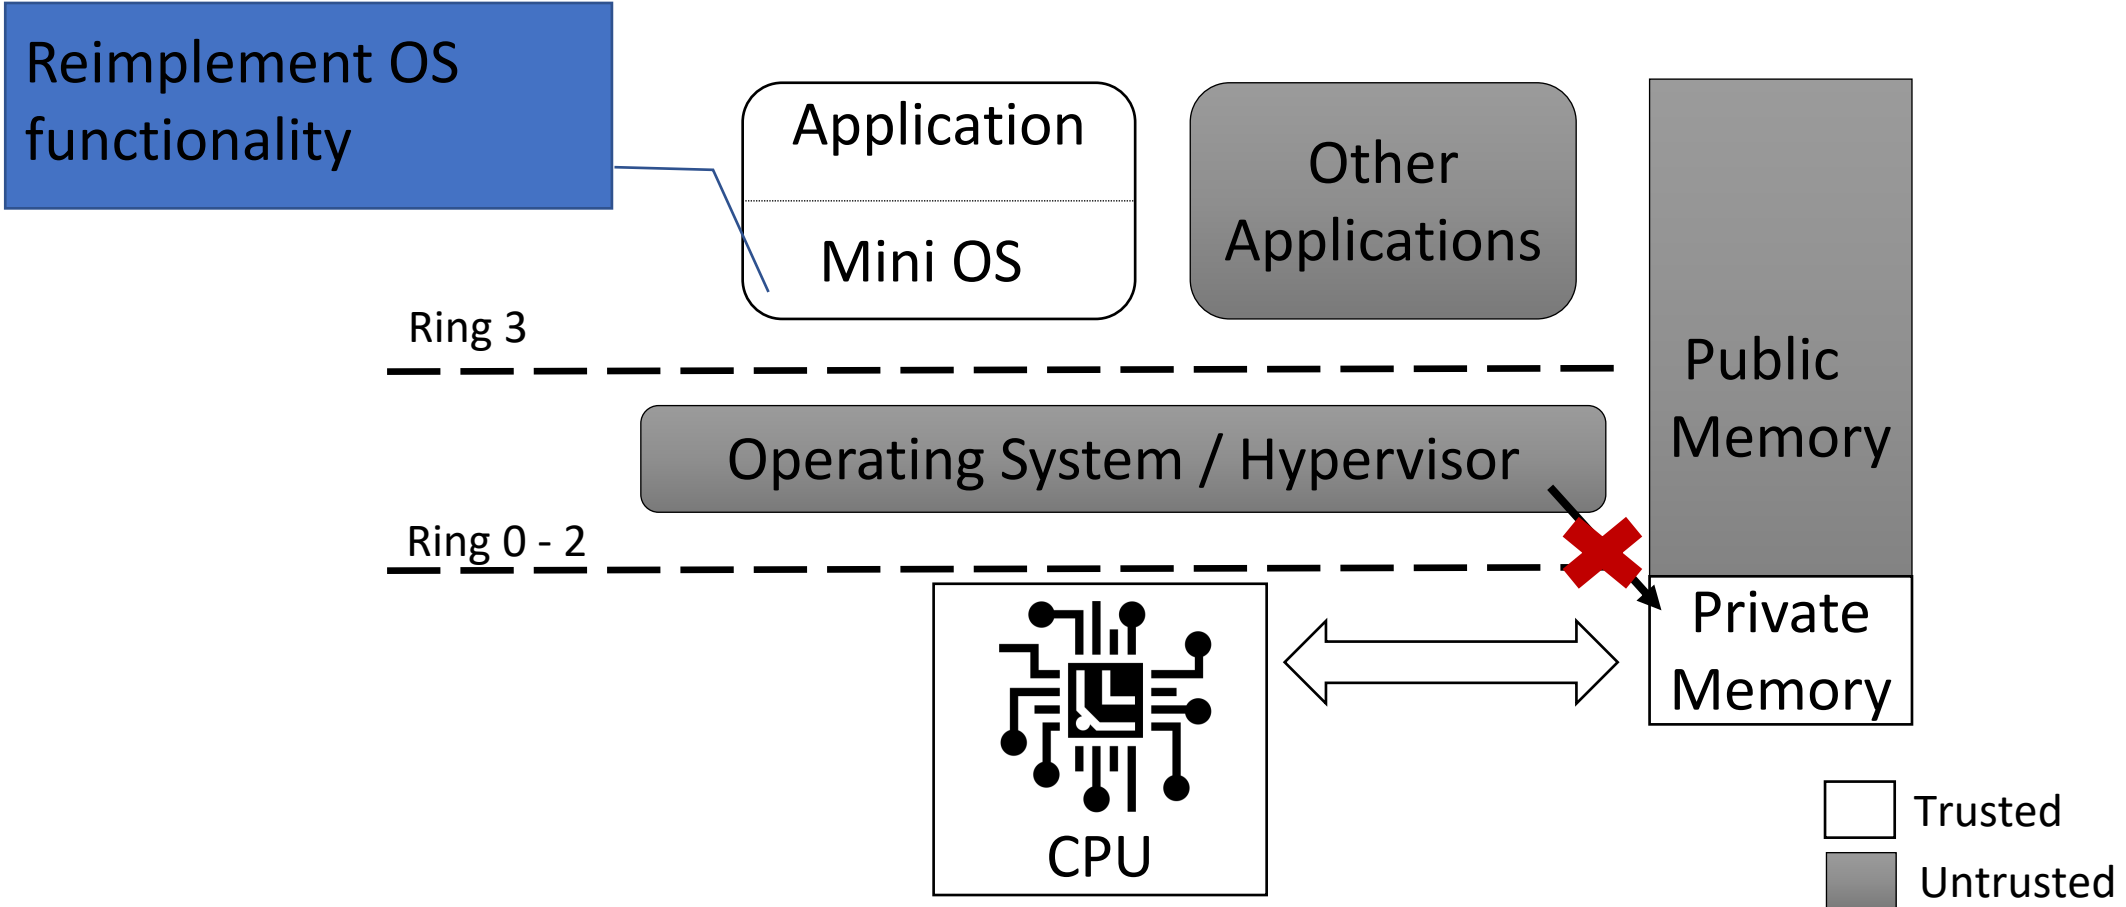
Trusted Computing Primitives
[TR'15], [Eurosys'20]

~50K

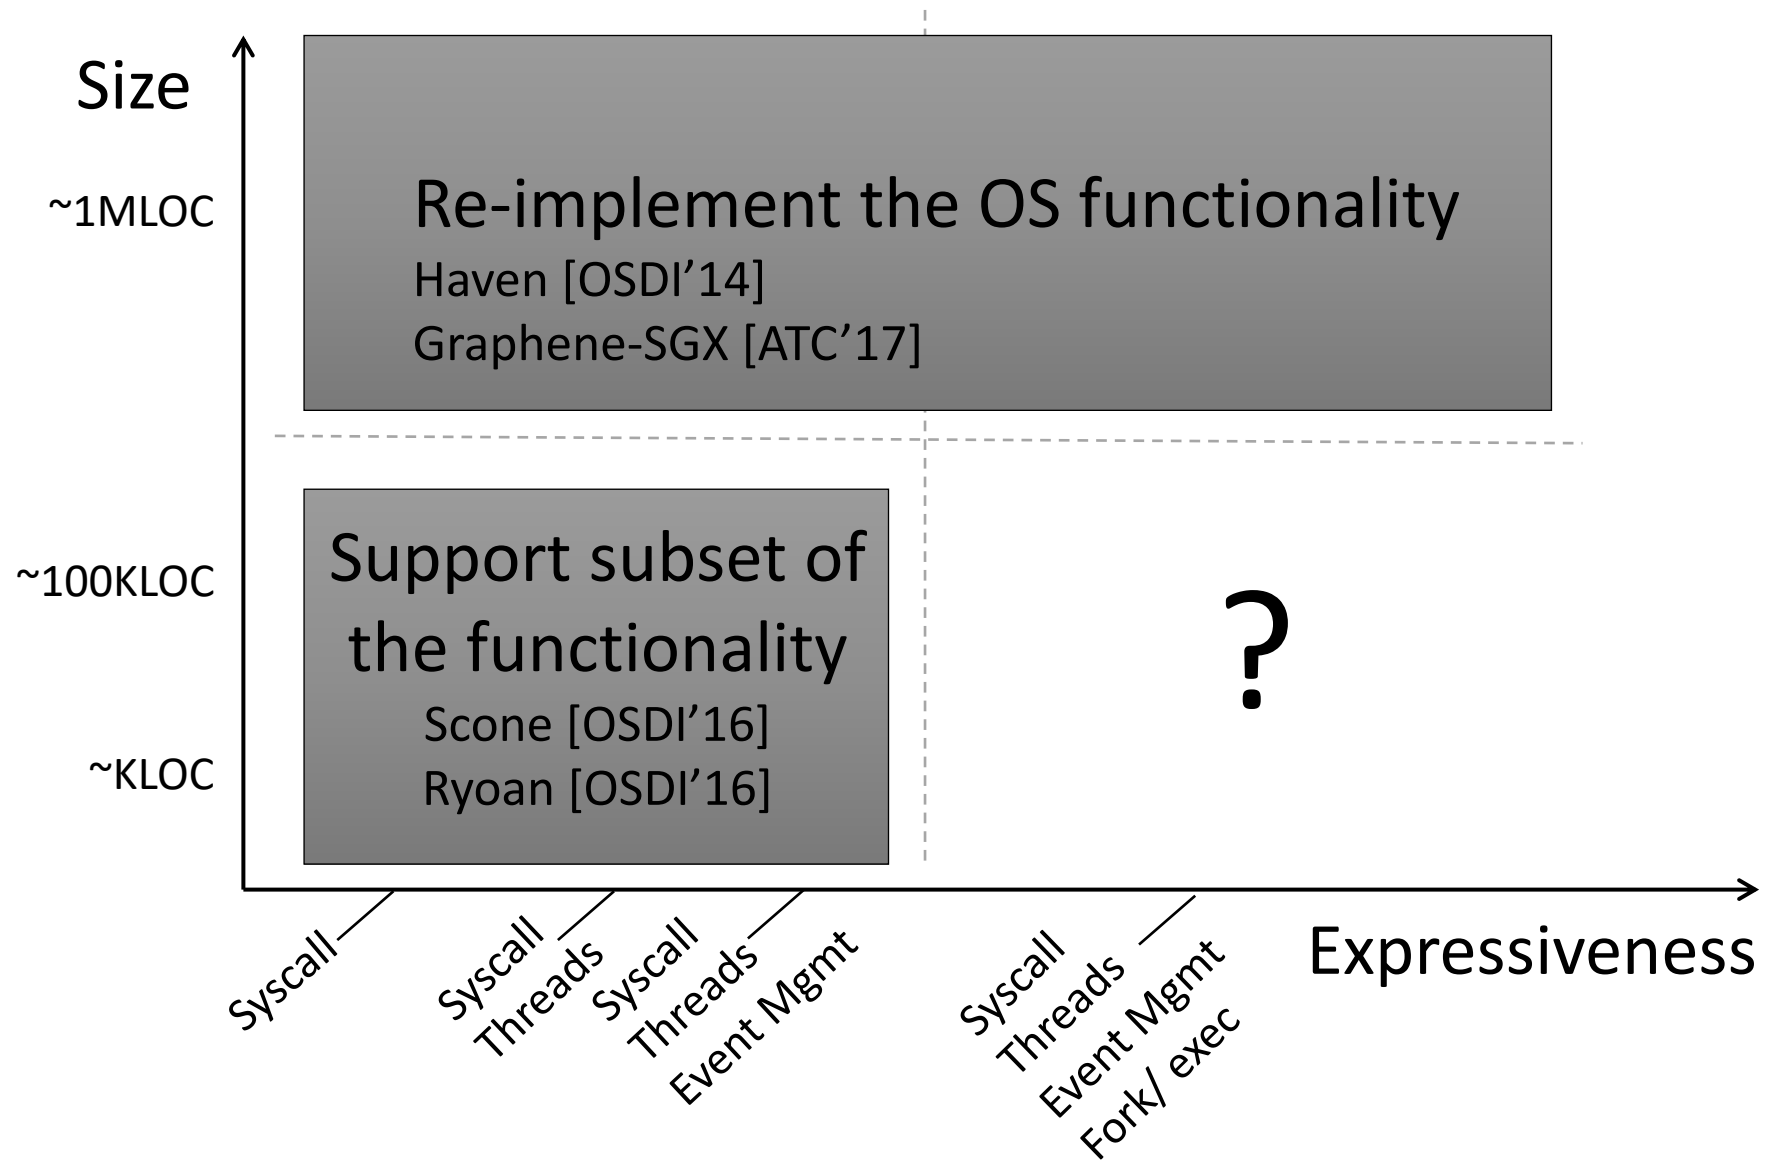
~20K



Adding Expressiveness to Commercial TEEs

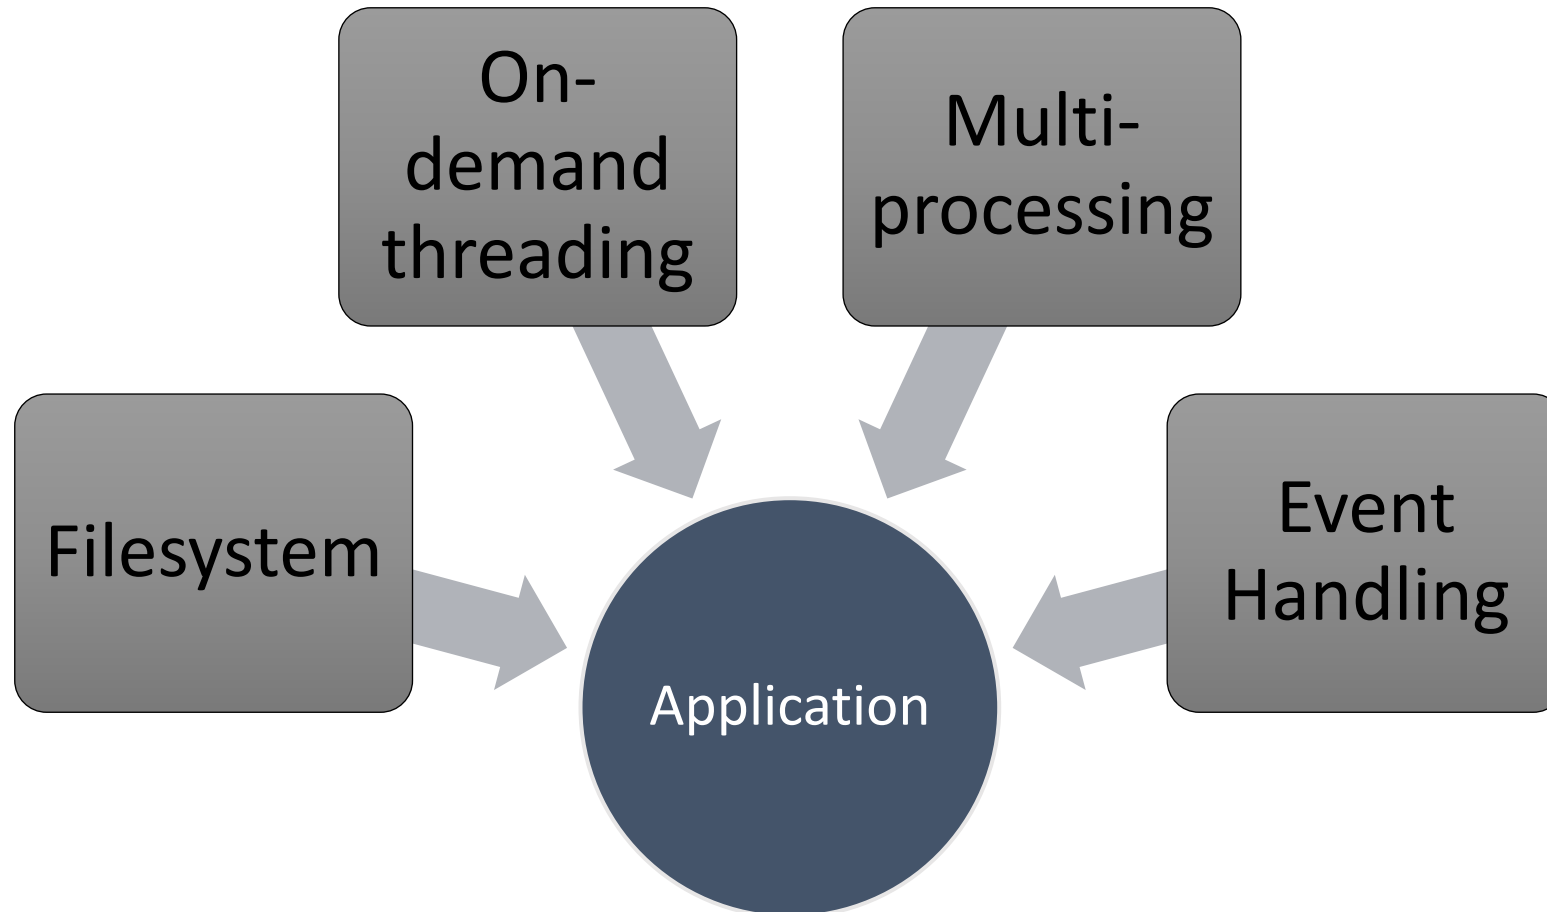


Code Size & Expressiveness Trade-off

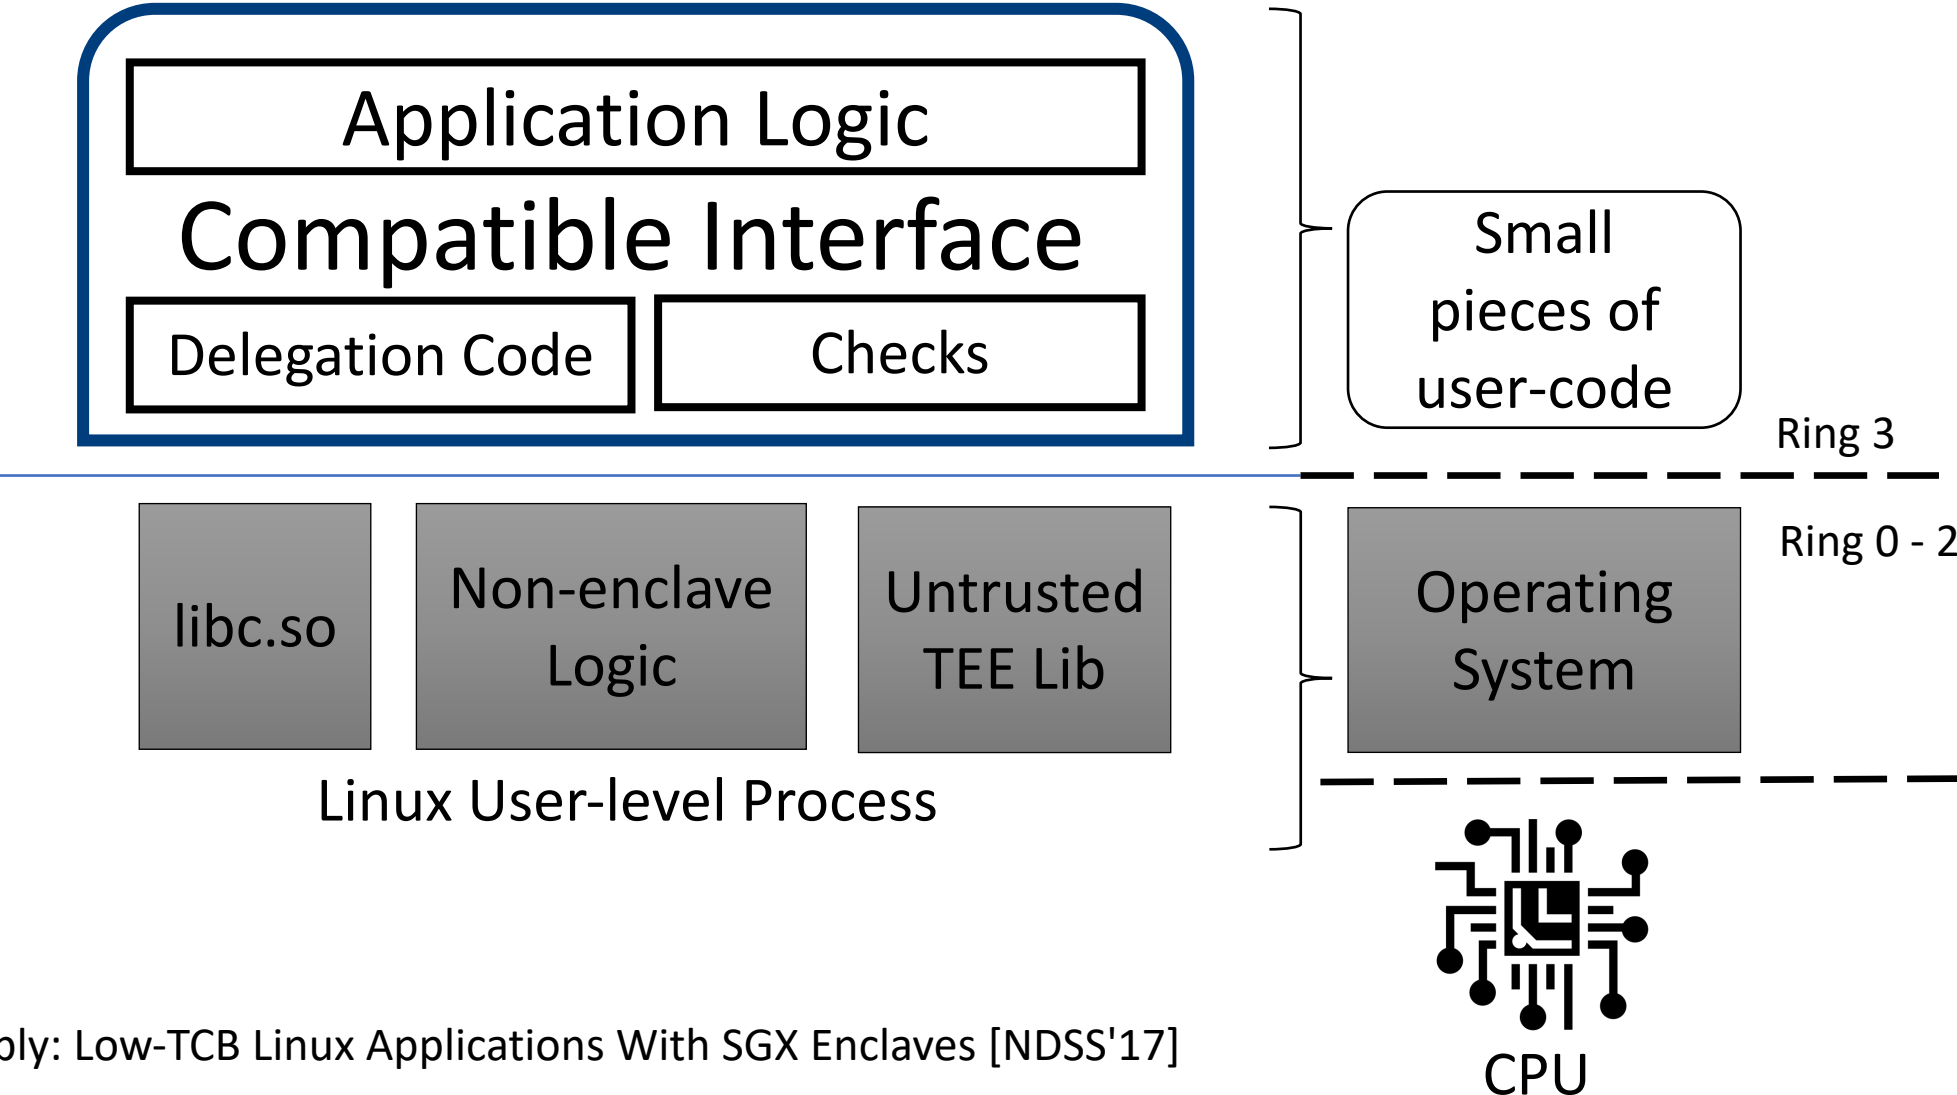


Challenge I: Expressiveness

Delegate rather than emulate




Building micro-container abstractions for TEEs



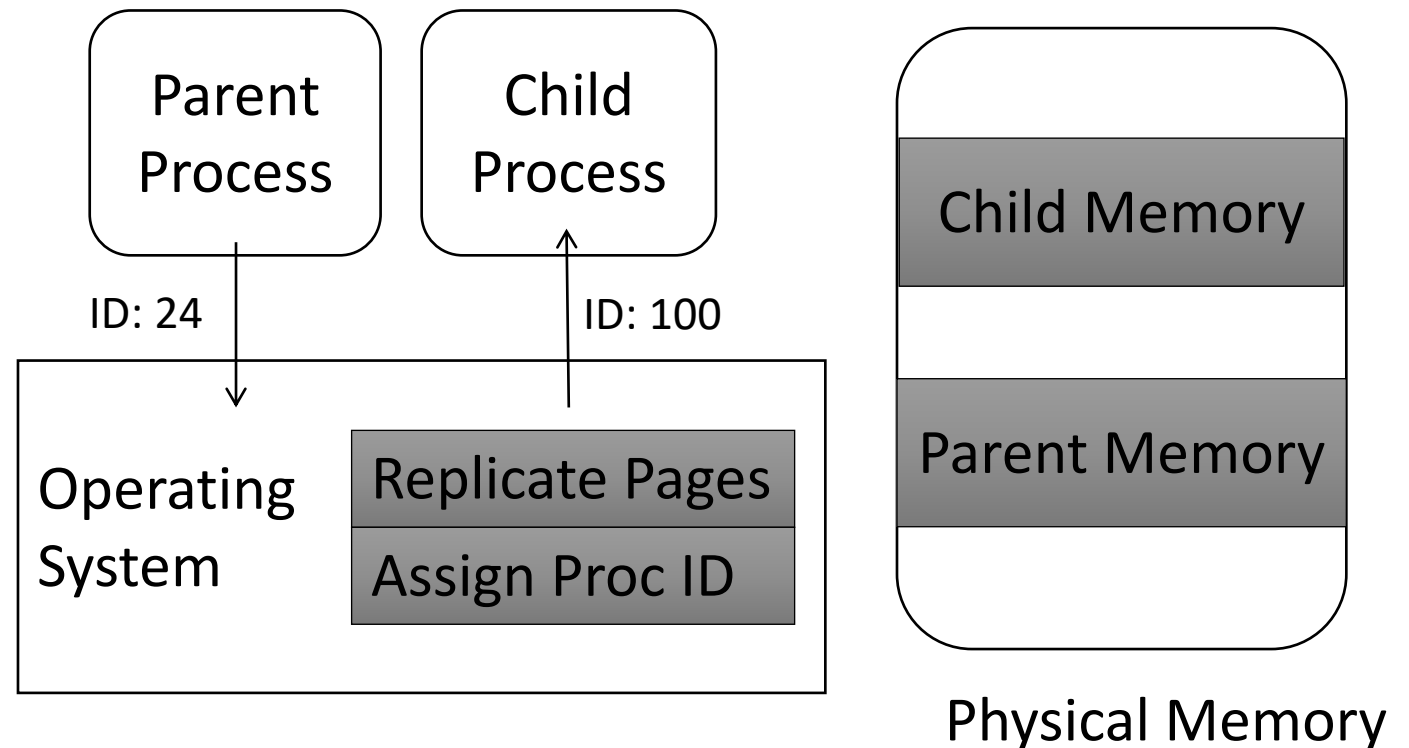
Challenge II: Delegation with isolation

- Two memory model:
 - private and public memory
- Process abstraction breaks
 - locks are in public memory
 - shared memory for processes
 - passing data to system calls



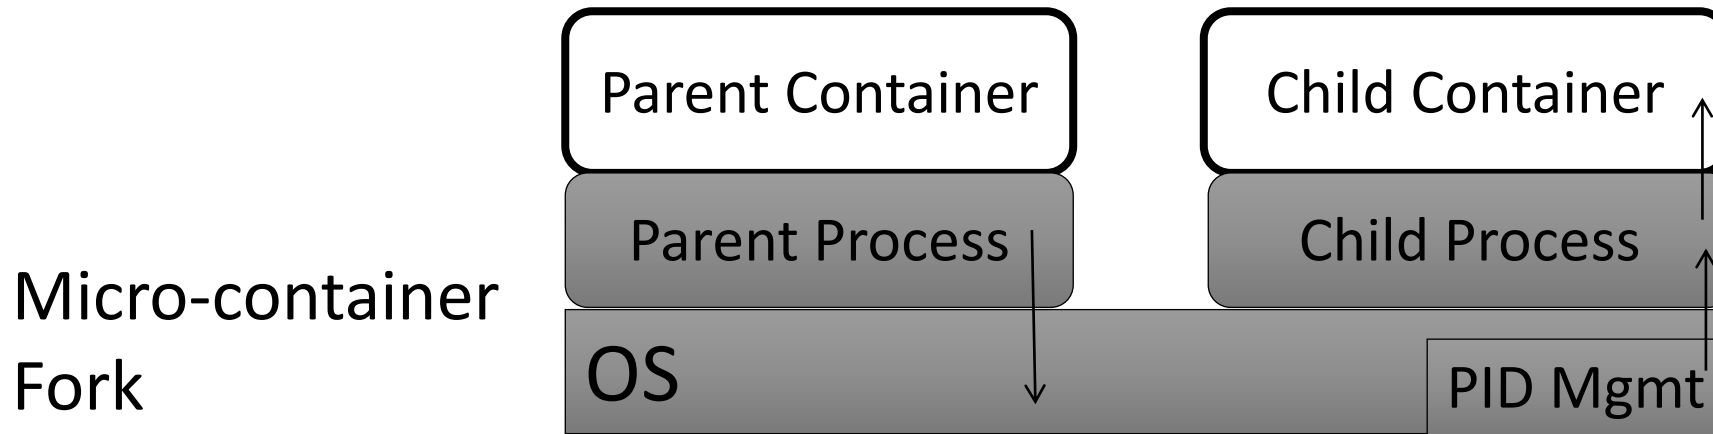
Expressiveness Example: Fork

- Fork Semantics:
 - Assigns new process id
 - Makes a memory replica
- How to maintain fork semantics if the OS cannot access private memory?



Expressiveness Example: Delegating Fork

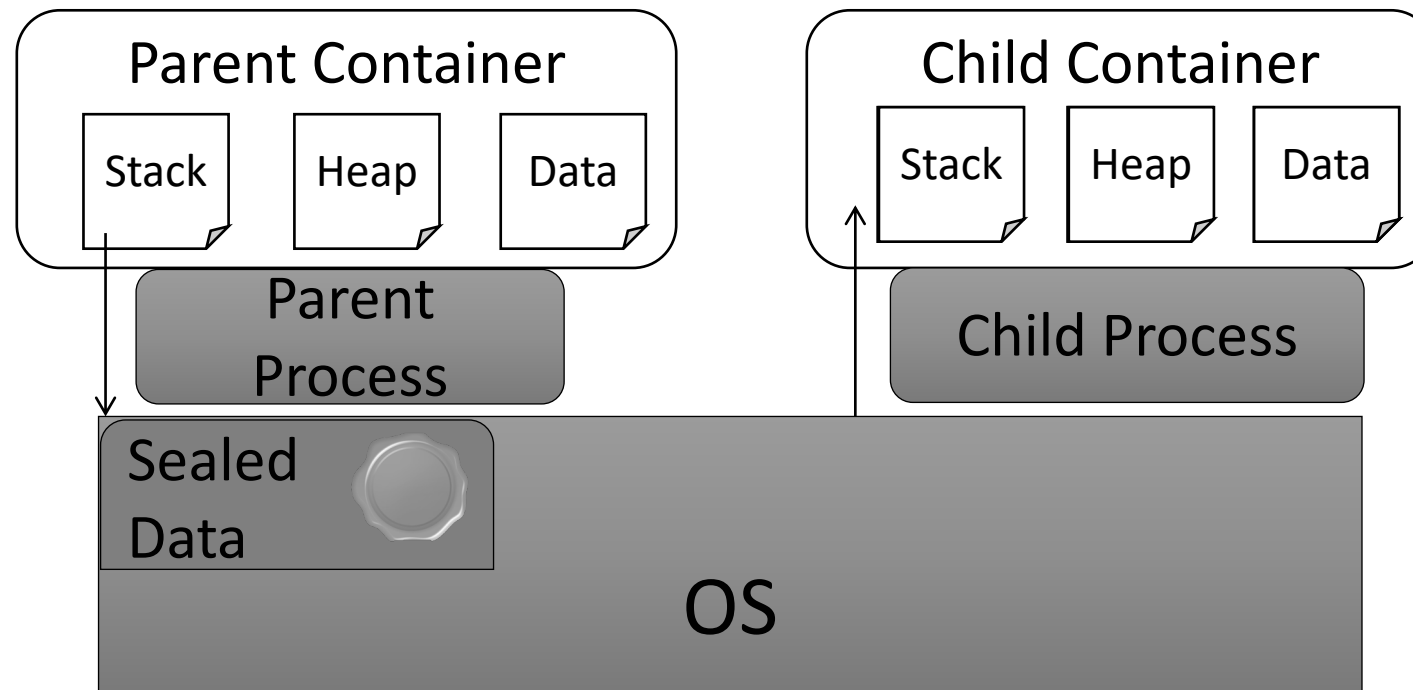
- Creating child process and child micro-container



- Child enclave has a clean memory state

Expressiveness Example: Achieving Fork Semantics

- Mirroring parent's memory in child micro-container
 - After the fork call, before resuming execution



Expressiveness: Supporting POSIX APIs

Core Services

Process Creation and Control	5
Signals	6
Timers	5
File and Directory Operations	37
Pipes	4
C Library (Standard C)	66
I/O Port Interface and Control	40

Thread Extensions


Thread Creation, Control, and Cleanup	17
Thread Scheduling	4
Thread Synchronization	10
Signal Delivery	2
Signal Handling	3

Real-time Extensions

Real-Time Signals	4
Clocks and Timers	1
Semaphores	2
Message Passing	7
Shared Memory	6
Asynchronous and Synchronous I/O	29
Memory Locking Interface	6

**POSIX APIs
Supported for
Commodity Linux Apps**

Micro-containers execute TEE use-cases



ANONYMITY
PROTOCOLS



WEBSERVERS




DATABASE
CLIENTS



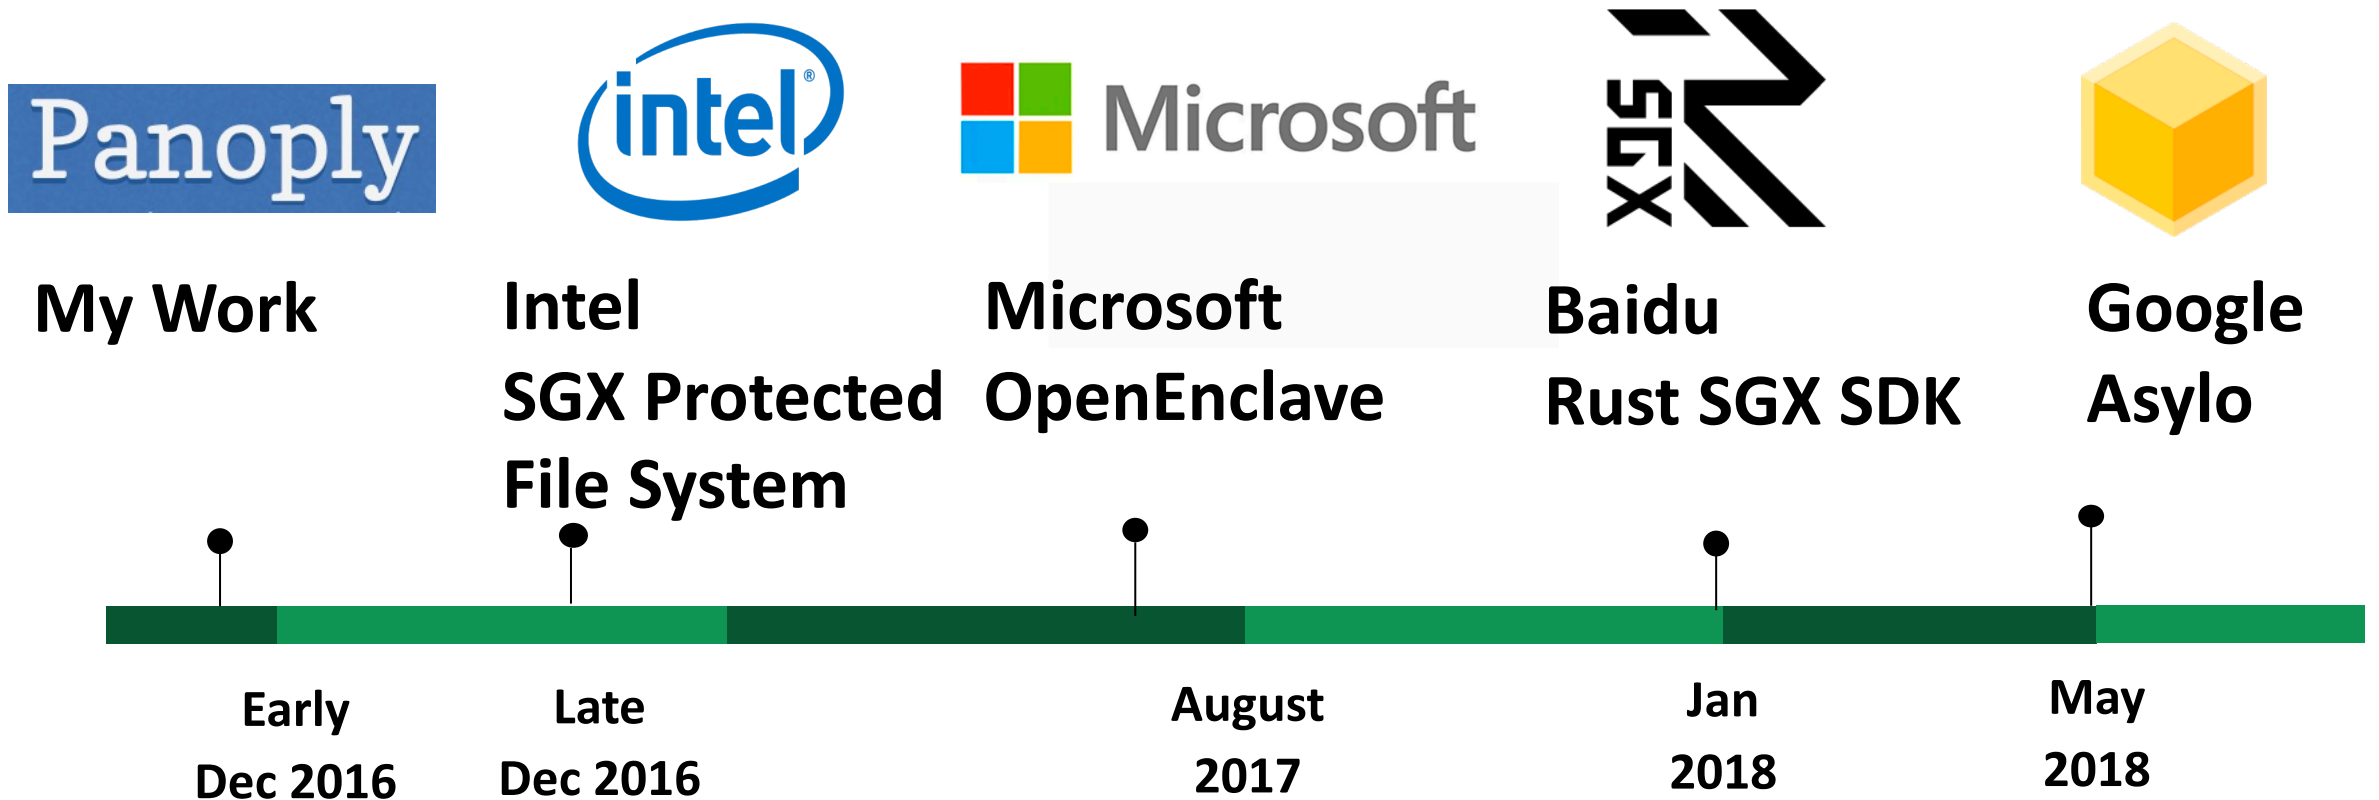
CRYPTOGRAPHIC
LIBRARIES

Performance is comparable to
importing a mini-OS

Minimize Trust to 20,000 lines of code



Adoption of the Delegation Approach



3rd component of this stack

New Applications [Arxiv'18], [ICDCS'19]
Secure Computation [CCS'13]

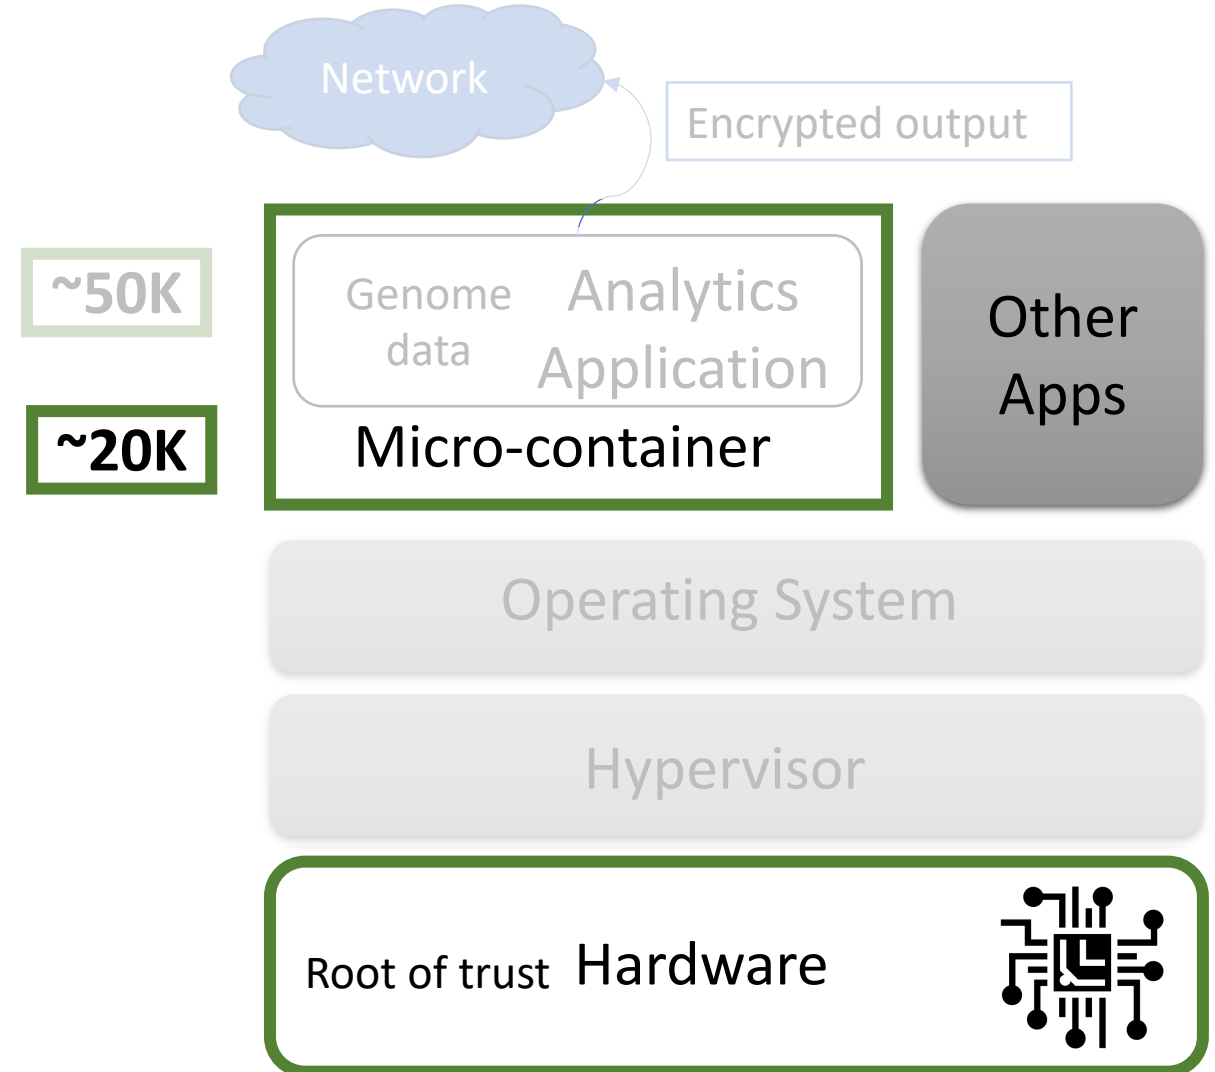
Analysis & hardening
[PLDI'14], [FSE'15], [NDSS'19], [CCS'20]

Rich functionality [NDSS'17], [Usenix'22]

Formal verification [Usenix '20]

Attacks & Defenses [AsiaCCS'16, CCS'21]

Trusted Computing Primitives
[TR'15], [Eurosys'20]



Example: What is the damage via OS interface?

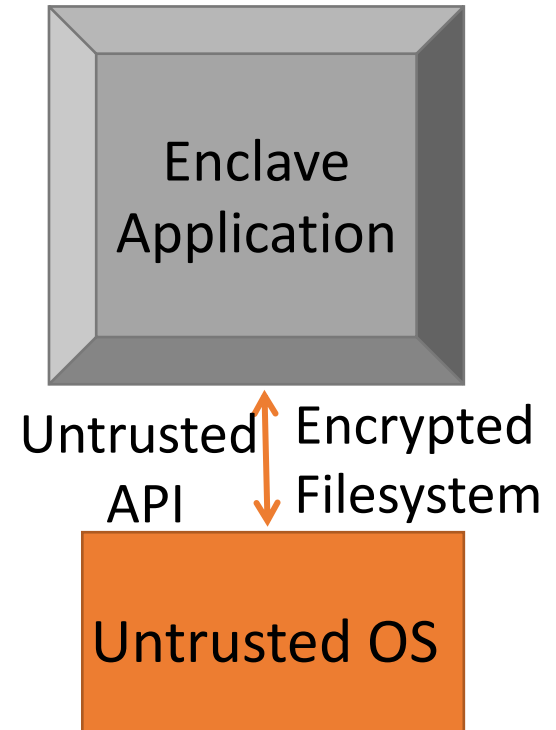
```
1 FILE* fd = fopen(fname, mode);
2 if (fd == NULL) {
3     errnum = errno;
4     if (errnum == EINVAL)
5         fd = fopen (fname, "a");
6     if (errnum == ENOENT)
7         if (fname == NULL)
8             fname = "vote.log";
9             fd = create_log(fname);
10    if (errnum == EINTR)
11        fd = fopen(fname, mode);
12 }
13 if (fd)
14     cnt = fwrite(buf, 1, len, fd);
15 return cnt;
```

Open the vote file

Failed to open the
vote file

Create new file →
overwrite previous vote

Register the vote sequence



Attacks are possible in delegation frameworks

```
9 int enc_untrusted_open(const char *path_name, int flags) {
10  uint32_t mode = 0;
11  int result;
12  sgx_status_t status = ocall_enc_untrusted_open(&result,
13  path_name, flags, mode);
14  if (status != SGX_SUCCESS) {
15    errno = EINTR;
16    return -1;
17  }
18  return result;
19 }
```

fopen: Google Asylo


```
7 static int sgx_ocall_open(void * pms) {
8   ms_ocall_open_t * ms = (ms_ocall_open_t *) pms;
9   int ret;
10  ODEBUG(OCALL_OPEN, ms);
11  ret = INLINE_SYSCALL(open, ...);
12  return IS_ERR(ret)?unix_to_pal_error(ERRNO(ret)):ret;
13 }
```

fopen: Graphene-SGX

```
4 }
5 static SGX_FILE* sgx_fopen_internal
6 (const char* filename, const char* mode) {
7   protected_fs_file* file = NULL;
8   if (filename == NULL || mode == NULL) {
9     errno = EINVAL;
10    return NULL;
11  }
12  ...
13 }
```

fopen: Intel SDK

A Formal Verification Approach: How to scale to POSIX?



The scalability challenge:

- Specification for safe behavior for the entire POSIX API
- Proving safe implementation
 - entire libc (glibc, musl)
 - filesystem (ext4)

Designing Scalable Specification: BesFS Interface

- Our Approach
 - 15 core APIs: e.g., `open`, `close`, `read`, `write`
 - Allow to execute any sequence of these while maintaining safety property
- Can be composed to express higher-level interfaces
 - e.g., `fwrite` can be composed with `write` and `fstat`
 - Created 22 auxiliary APIs witnessed in applications

BesFS Highlights



4625 lines in Coq
167 lemmas
($< 1.5K$ in C code)



Not over restrictive
Supports all applications
from Panoply (& more)

Total 31 tested



Helped in eliminating
bugs (from Panoply, Intel
SDK, Google SDK)

Towards Next Generation Computation Stack

New Applications [Arxiv'18], [ICDCS'19]

Secure Computation [CCS'13]

Analysis & hardening

[PLDI'14], [FSE'15], [NDSS'19], [CCS'20]

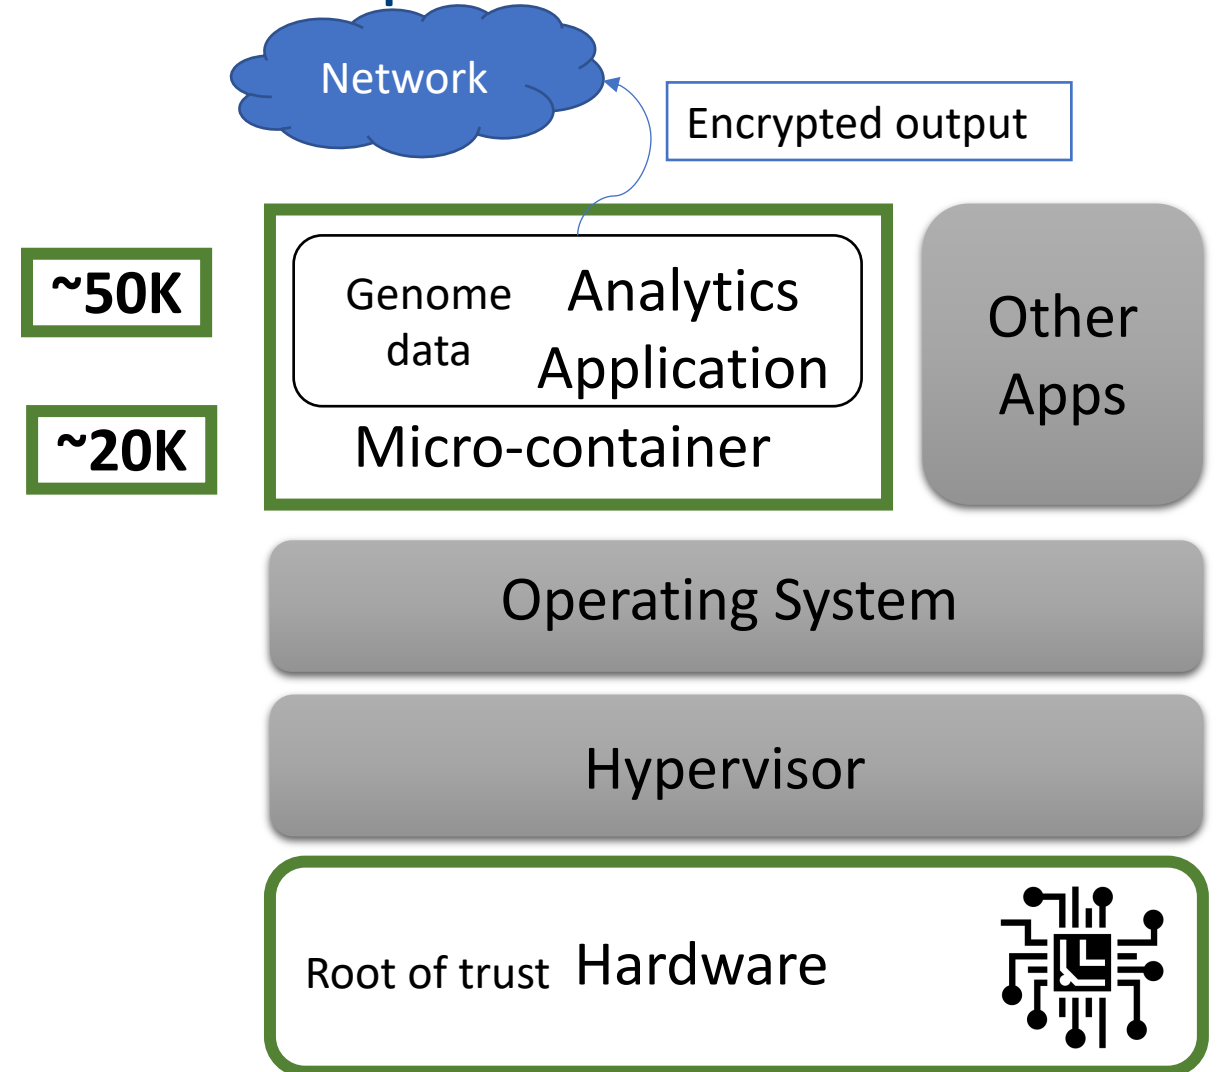
Rich functionality [NDSS'17], [Usenix '22]

Formal verification [Usenix'20]

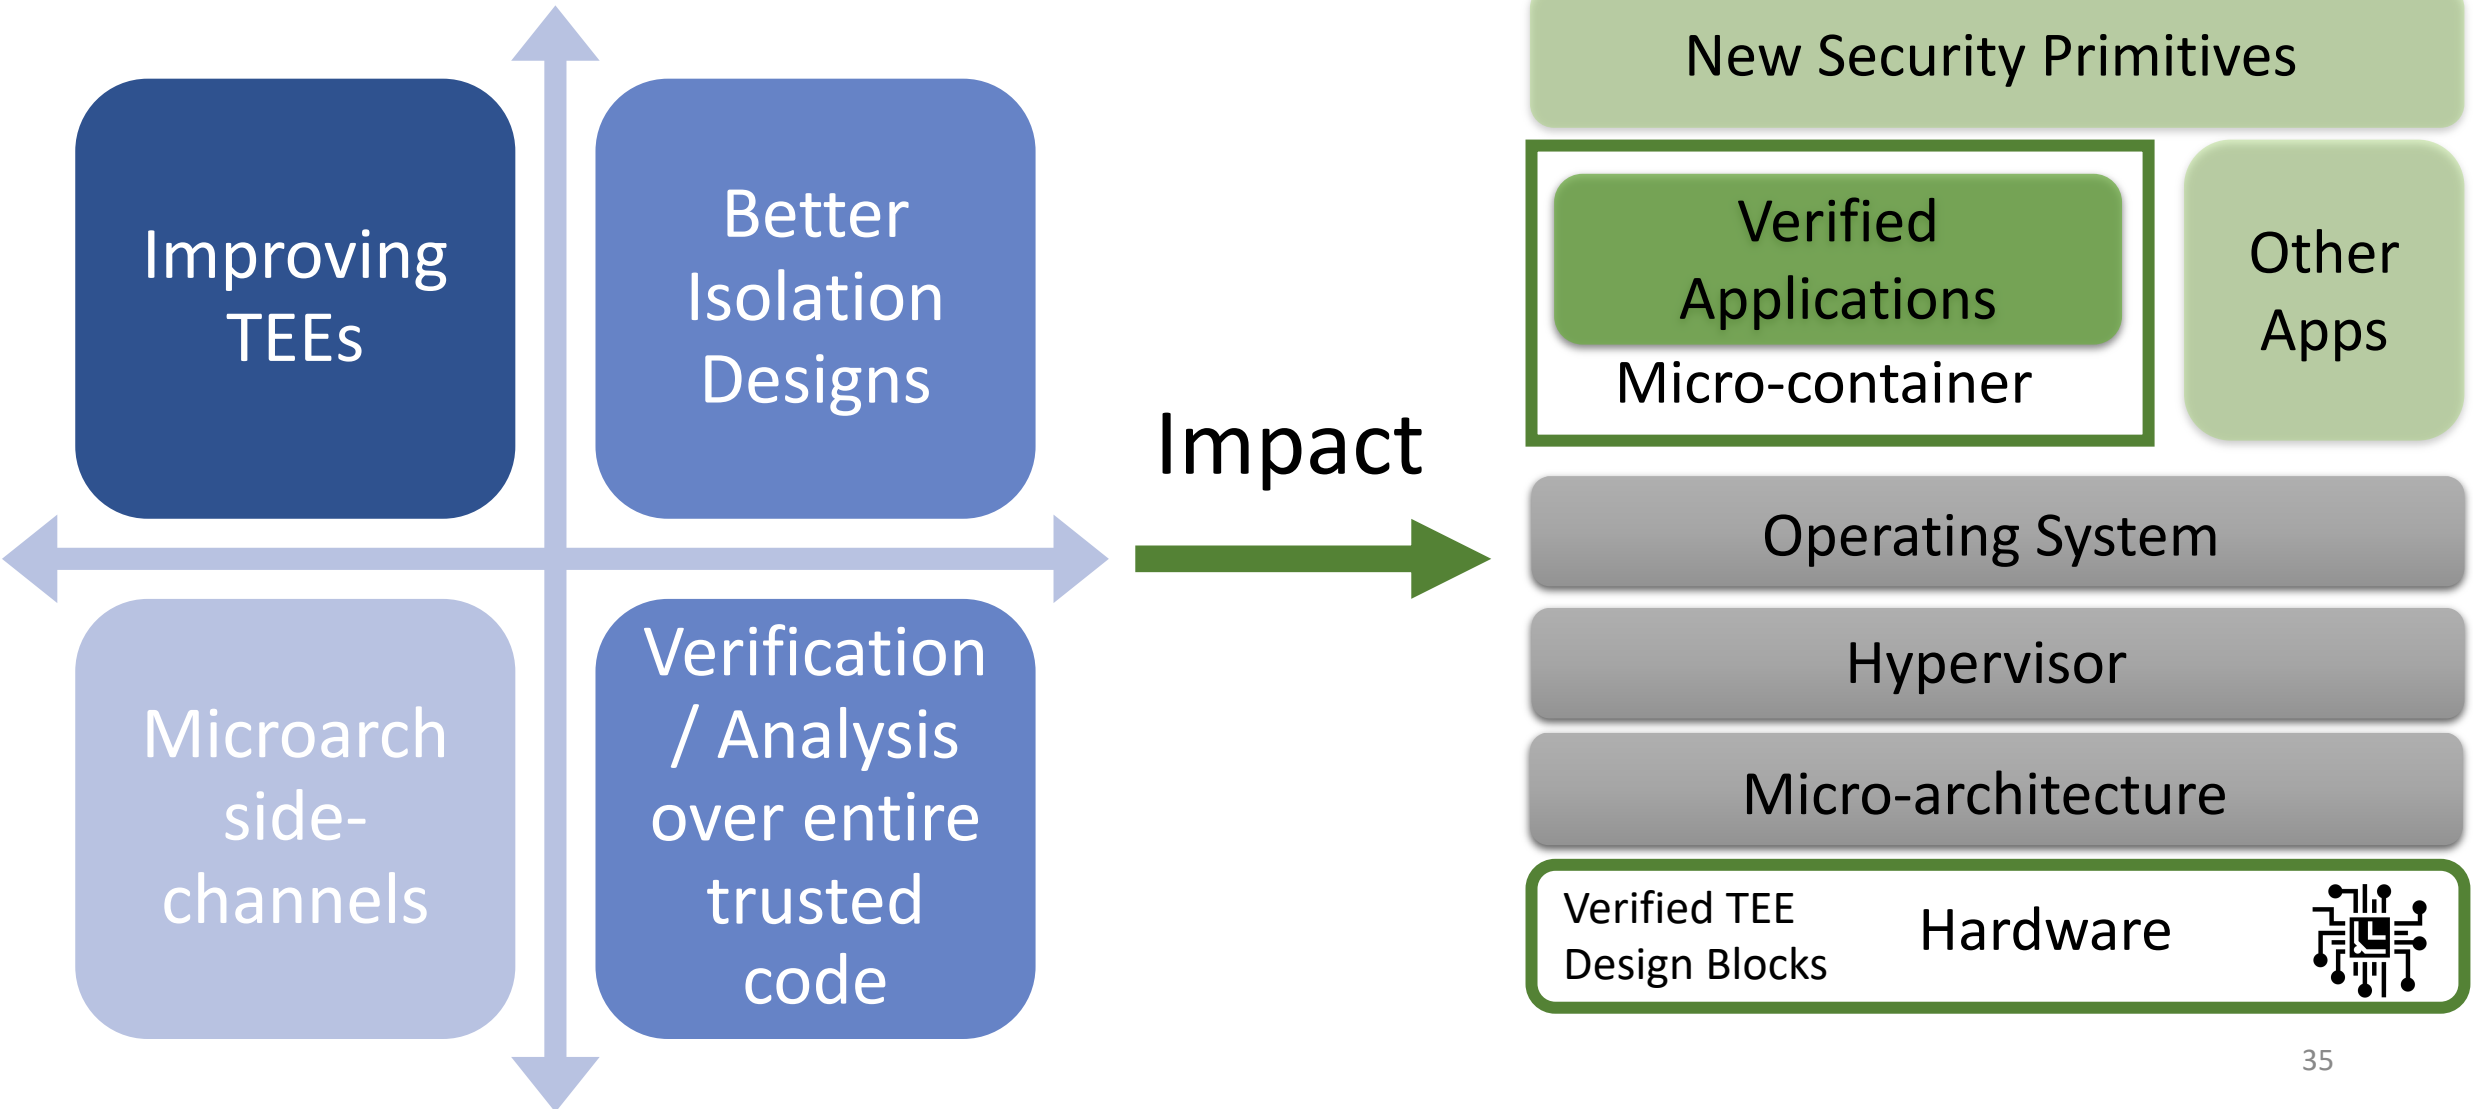
Attacks & Defenses [AsiaCCS'16, CCS'21]

Trusted Computing Primitives

[TR'15], [Eurosys'20]

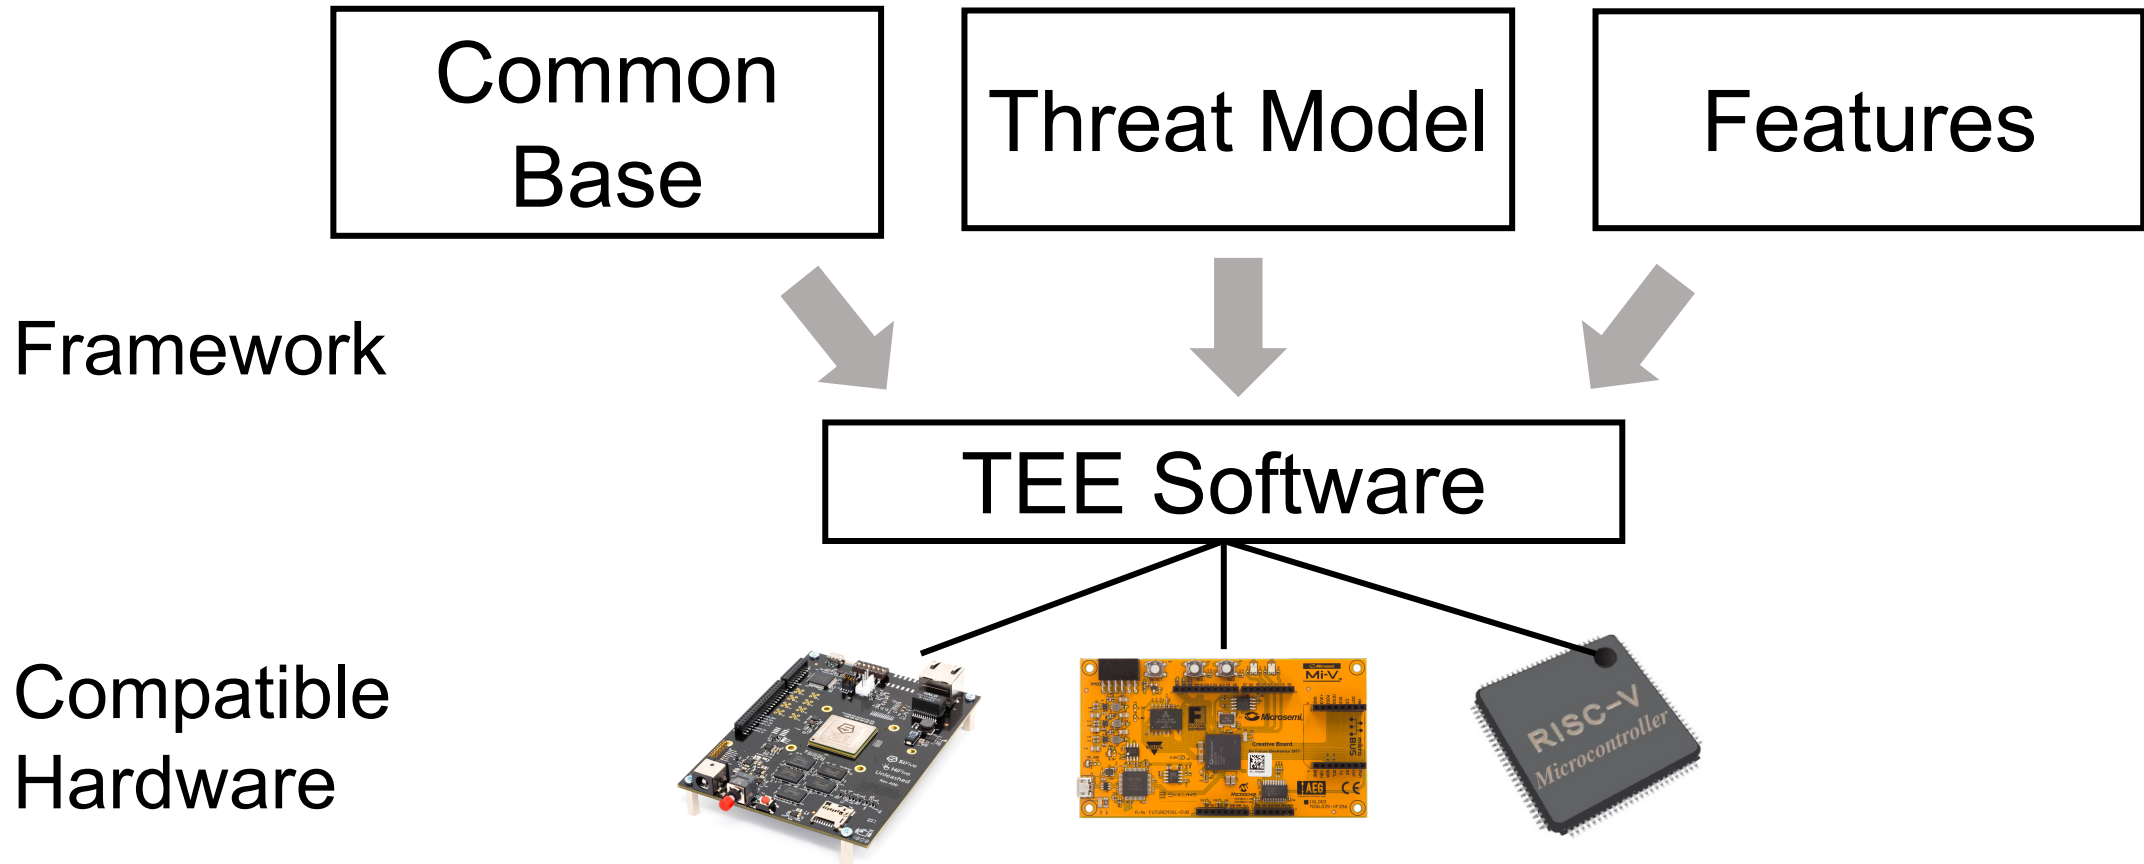


Research Directions



Customizable TEEs

- A framework that provides building blocks of TEEs
- The platform provider and the enclave developer “customizes” the TEE





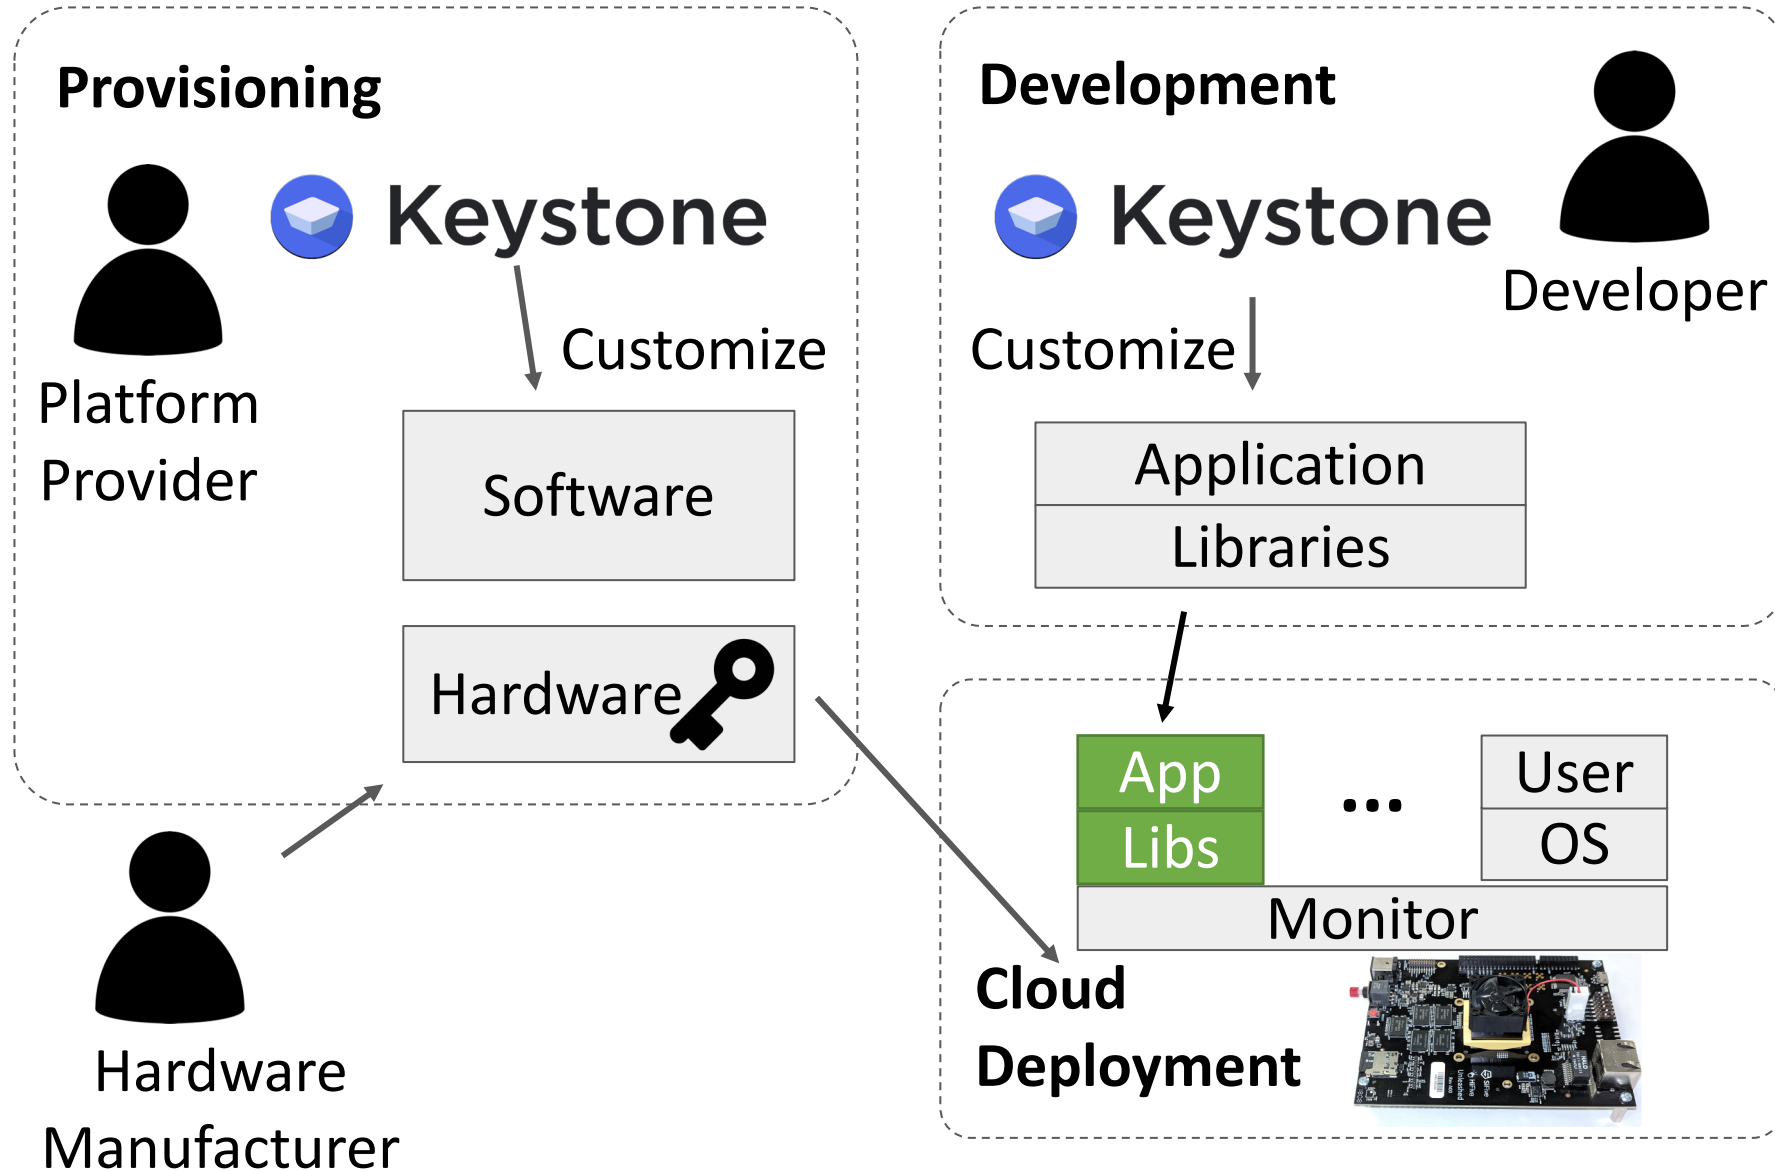
Keystone

A software framework for TEEs on RISC-V

No micro-architectural changes

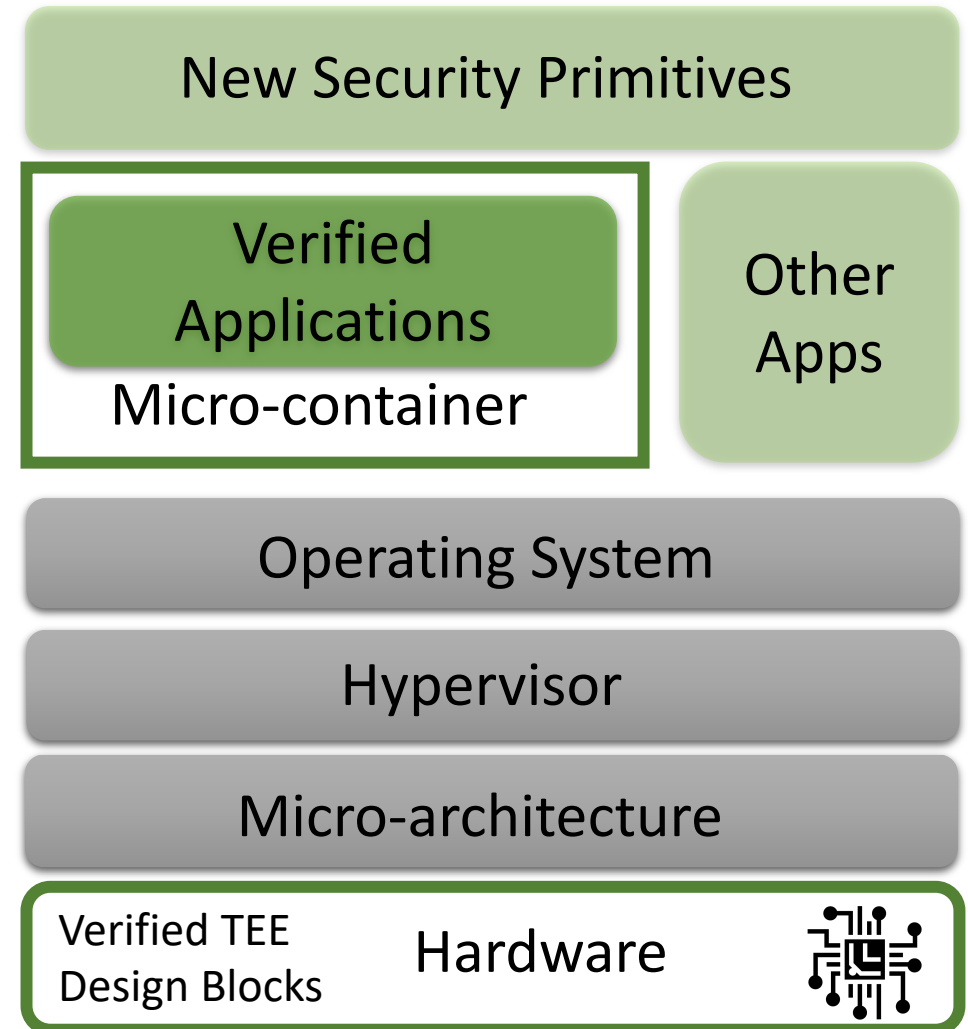
Minimal added hardware

Keystone Workflow for Customizable TEEs



Research Goals for Future TEE Platforms

- Modular TCB, easy to reduce and verify
- Binary compatibility with legacy applications
- Enable support for various backend hardware platforms
- Evolve to better hardware designs for TEE independently of the software



ETH zürich



Shweta Shinde
Assistant Professor
shweta.shivajishinde@inf.ethz.ch

ETH Zurich
Department of Computer Science



[@shw3ta_shinde](https://twitter.com/shw3ta_shinde)

<https://shwetashinde.org>
<https://www.sectrs.ethz.ch>