Secure Tera-scale Data Crunching with a Small TCB

Bruno Vavala UL / CMU

Nuno Neves UL

Peter Steenkiste CMU

Ciências ULisboa

Carnegie Mellon University

47th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'17)

Goal

delivering **security guarantees** for
large-scale data processing
on untrusted hosts with a **small TCB**

security guarantees

trusted HW based

data integrity

1 TB

large-scale data processing

small code

small interface

small TCB

No HW devices

Some use cases

public cloud service provider

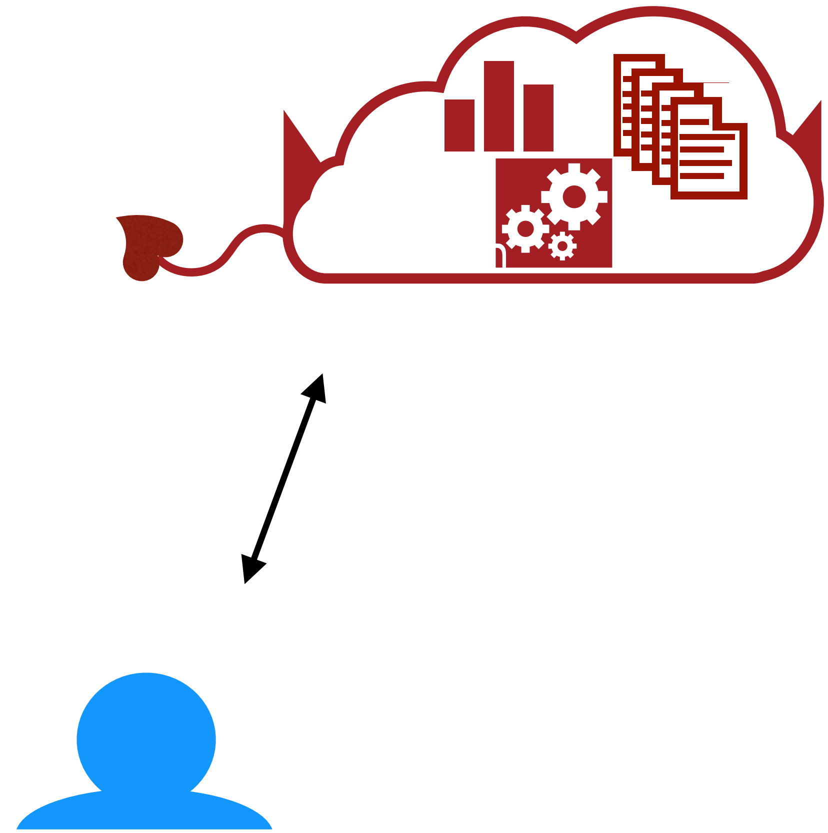

Some use cases

public cloud service provider

computational genomics

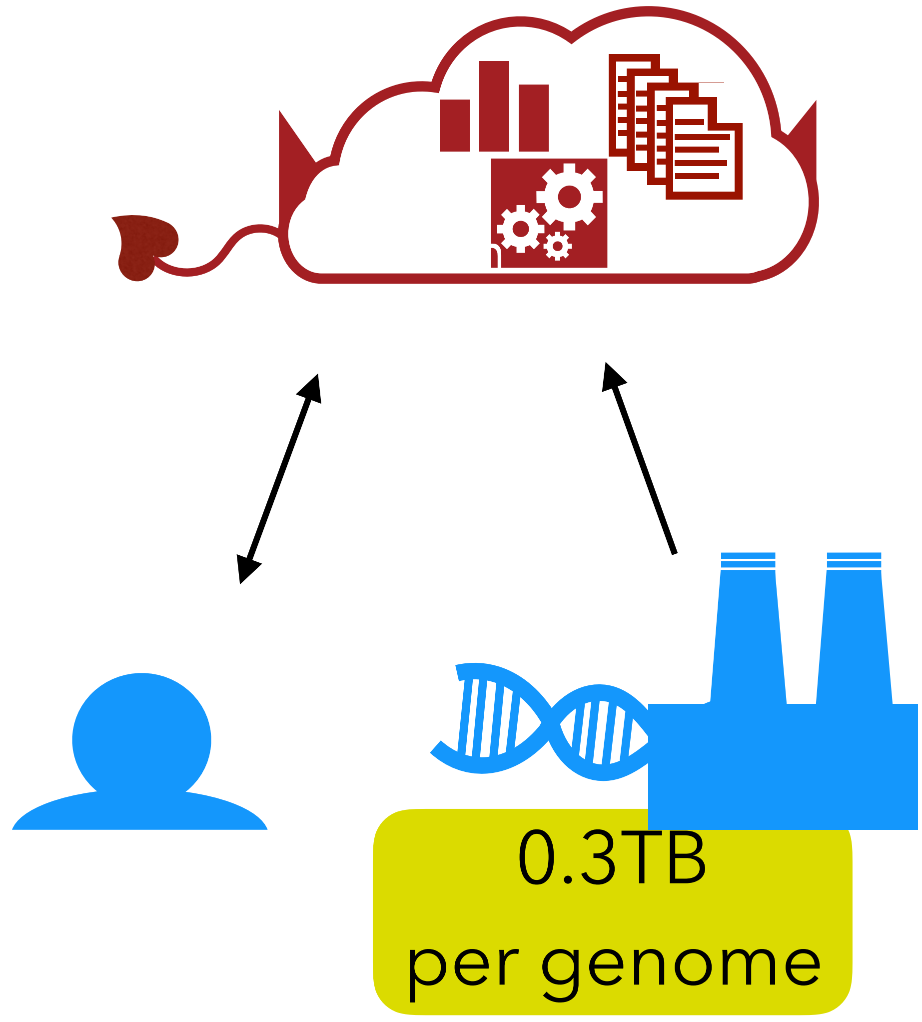

…more generally…

Model

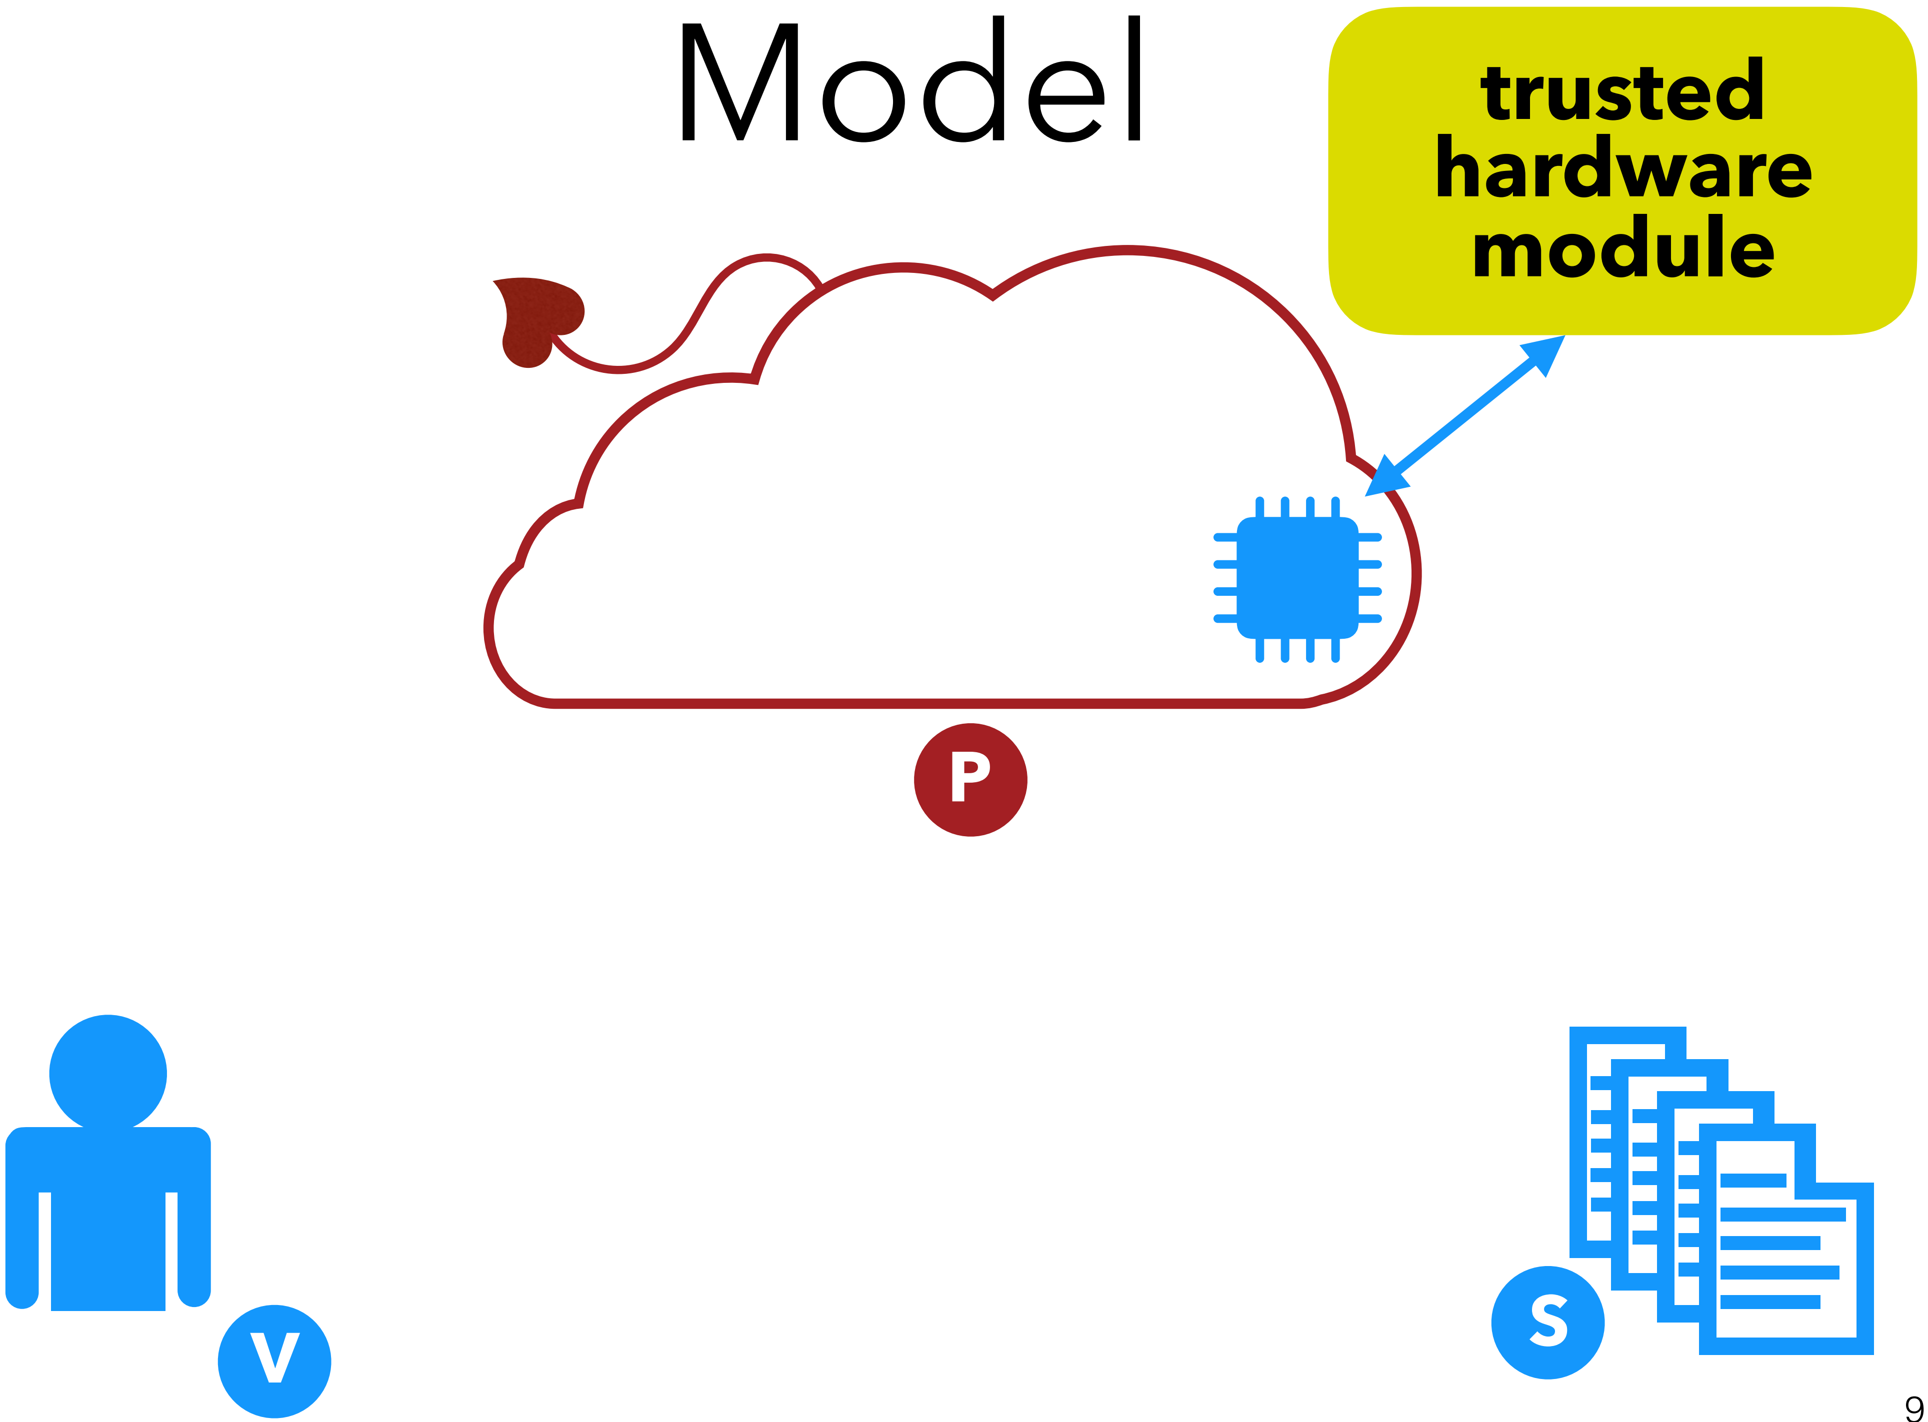

Model

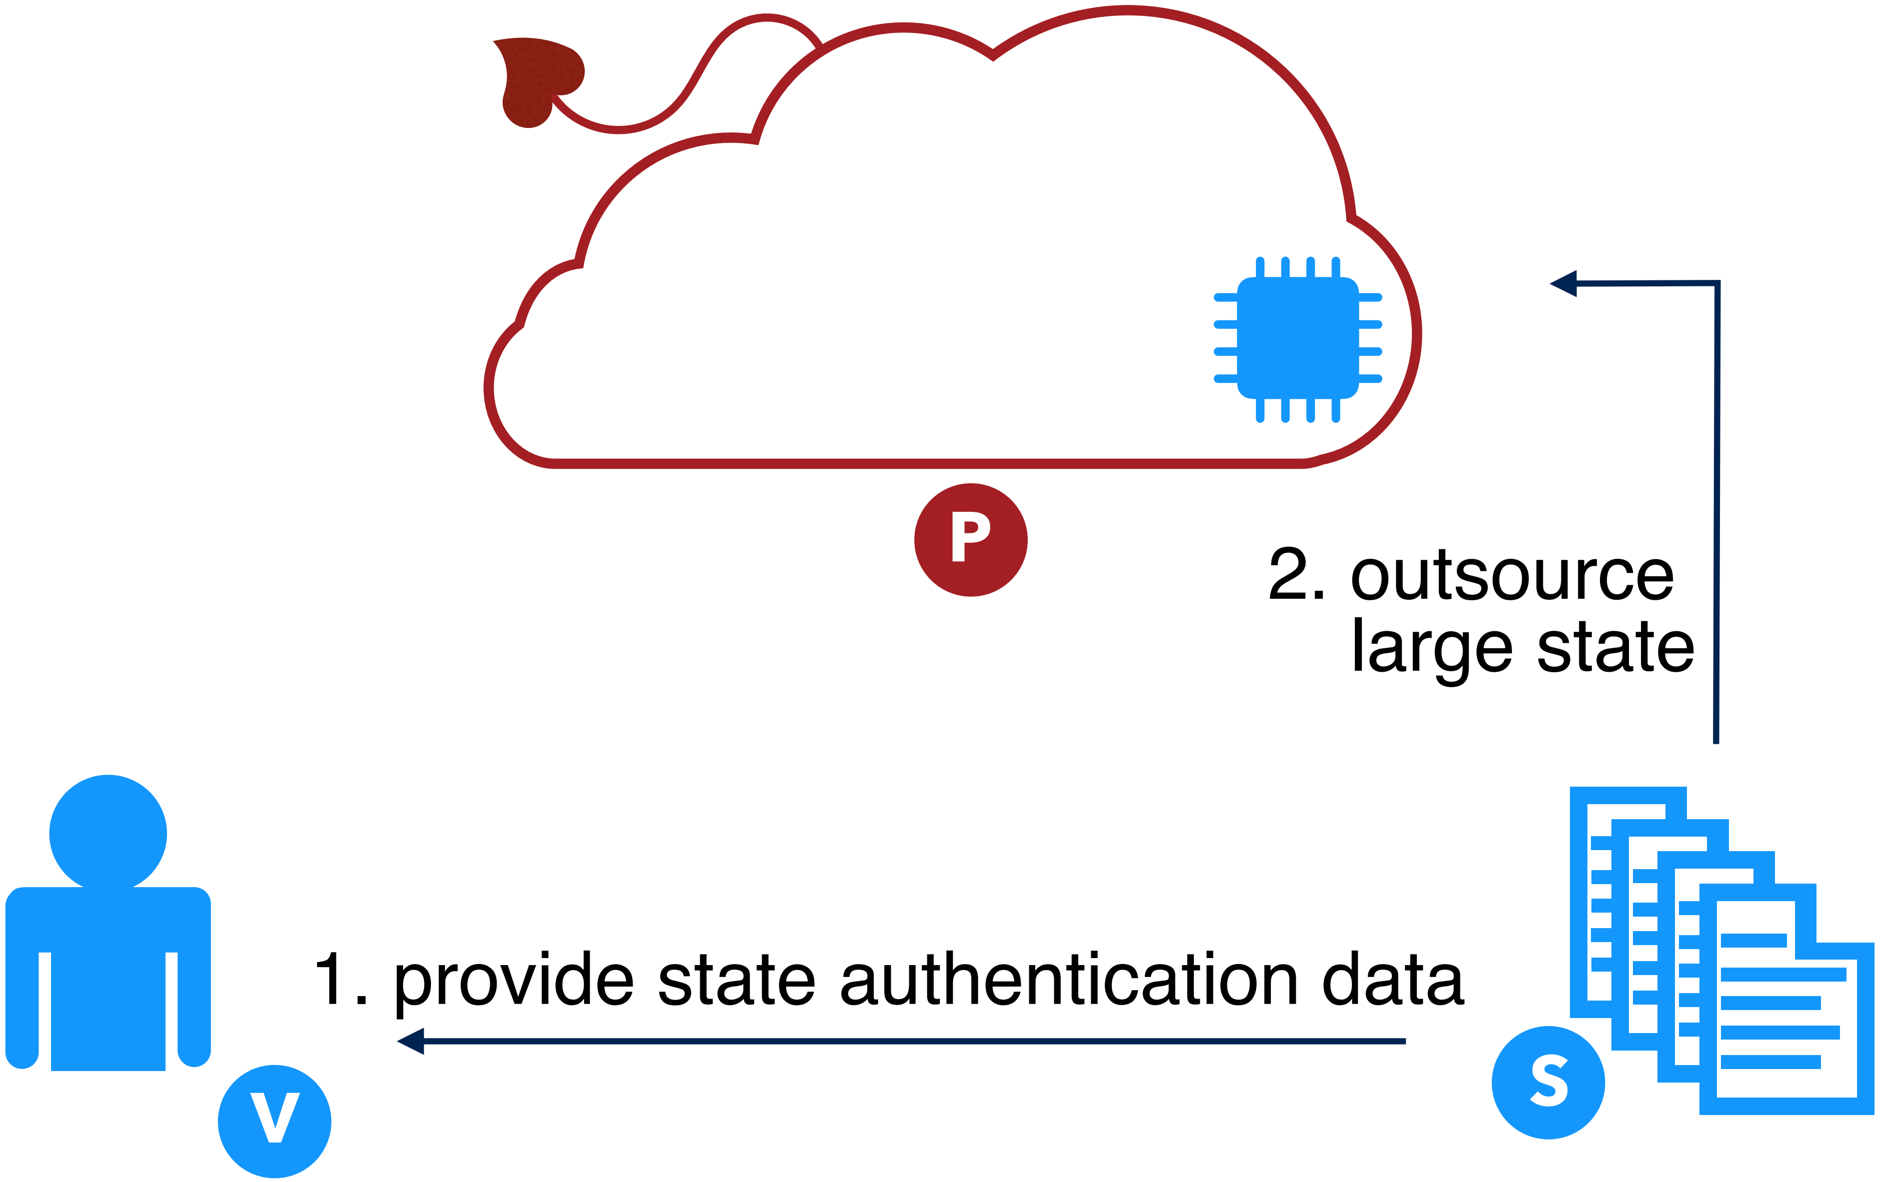
P
2. outsource
large state
1. provide state authentication data
V
S

Model

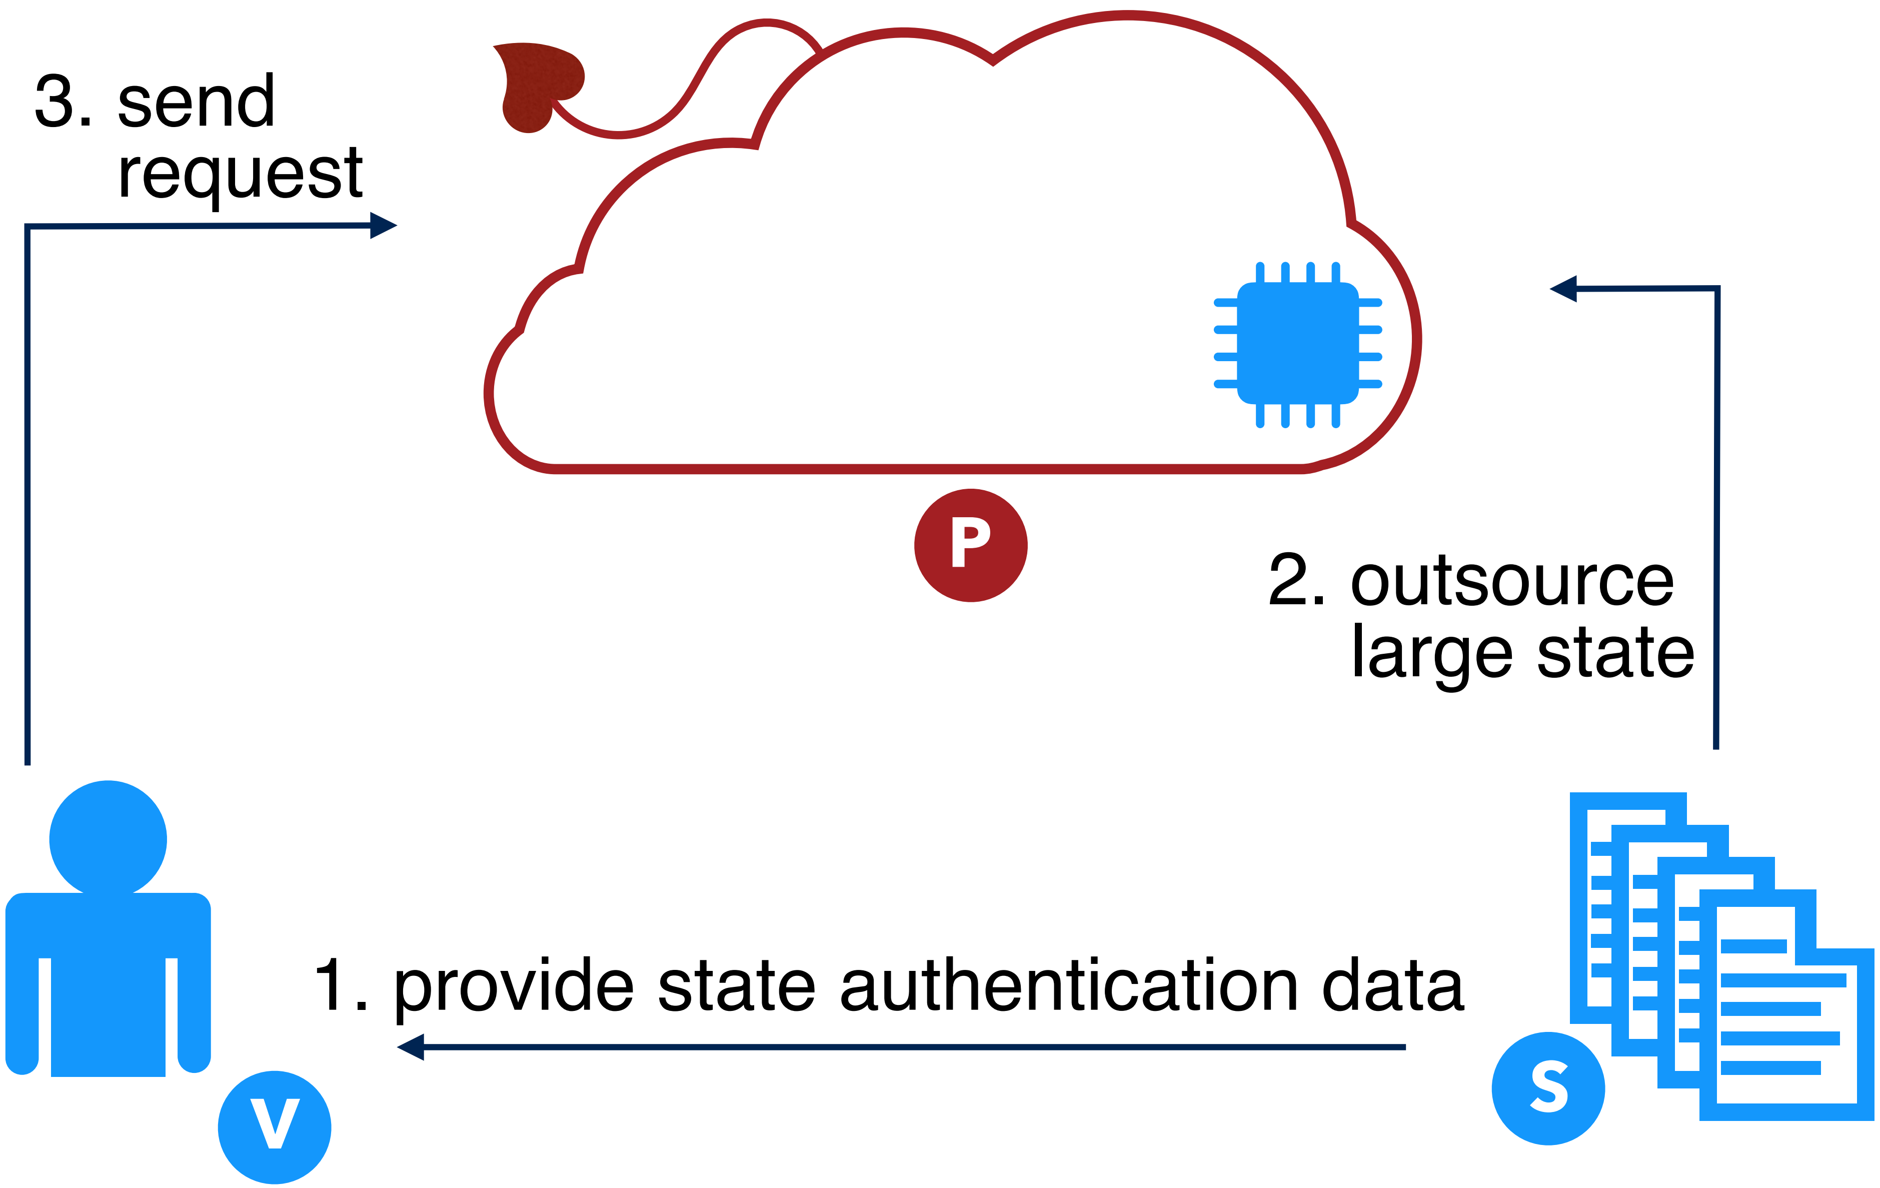
3. send
request
P
2. outsource
large state
1. provide state authentication data
V
S

Model

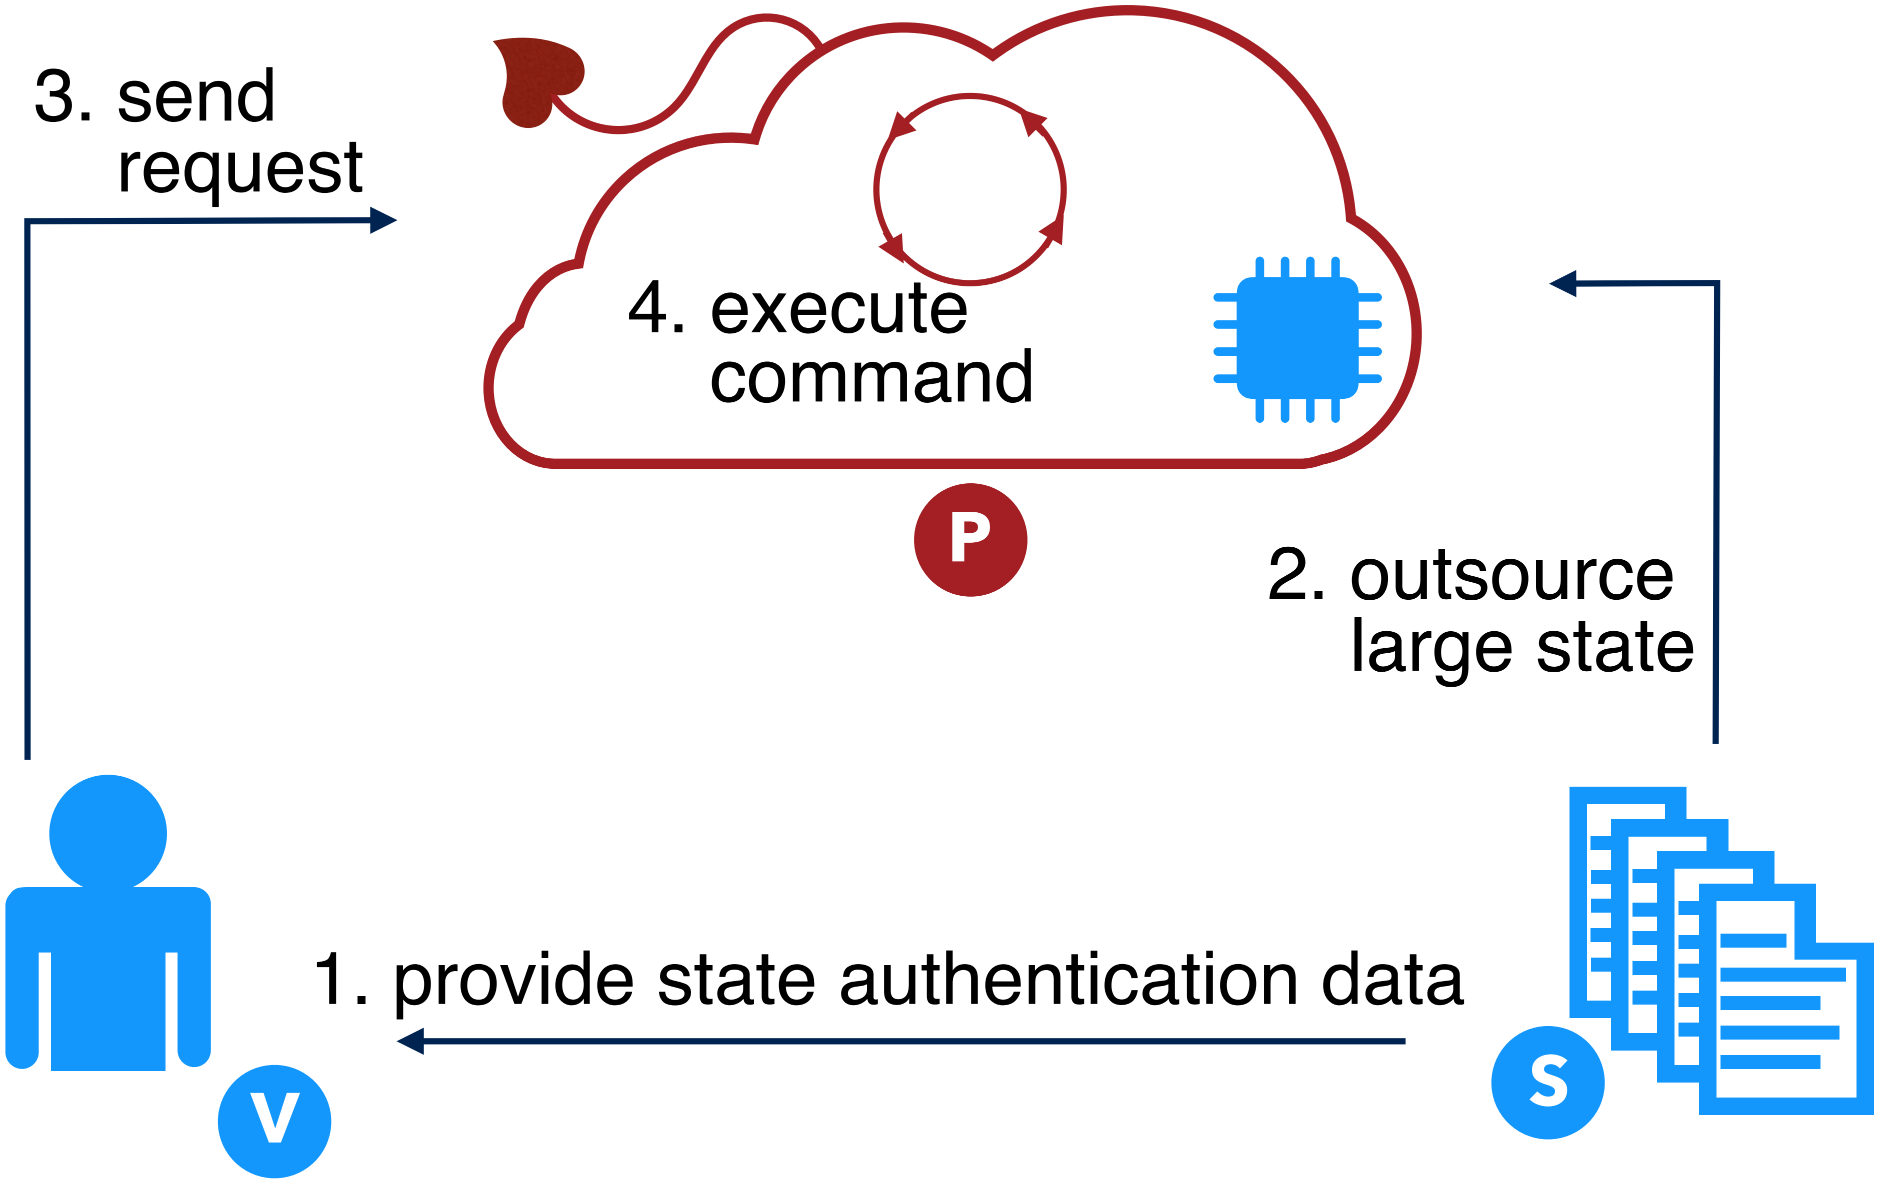
3. send request
4. execute command
P
2. outsource large state
1. provide state authentication data
V
S

Model

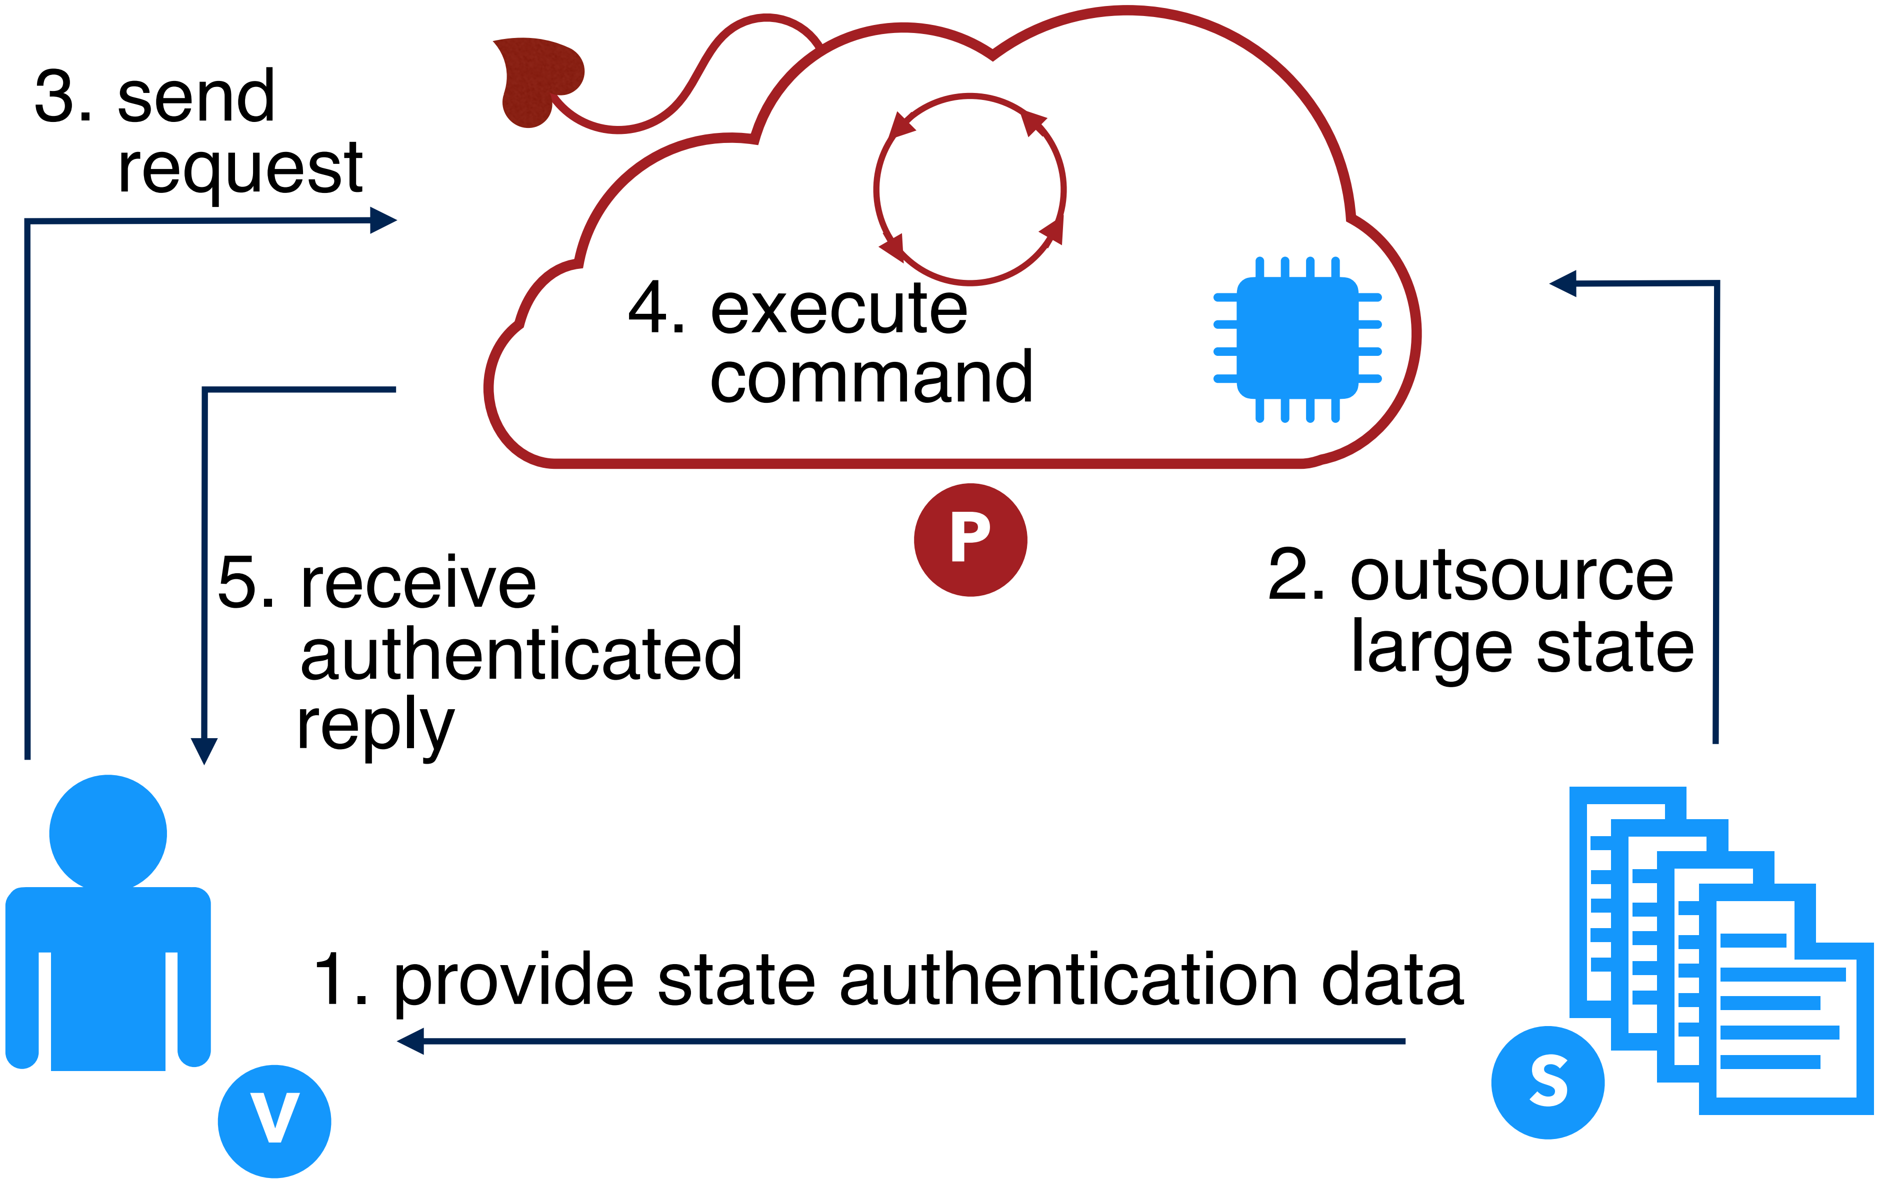
3. send request
4. execute command
P
5. receive authenticated reply
2. outsource large state
1. provide state authentication data
V
S

Outline

- **Goal**
- Previous Work
- Our solution: key ideas and overview
- Evaluation

Outline

- **Previous Work**

Haven
(OSDI'14)

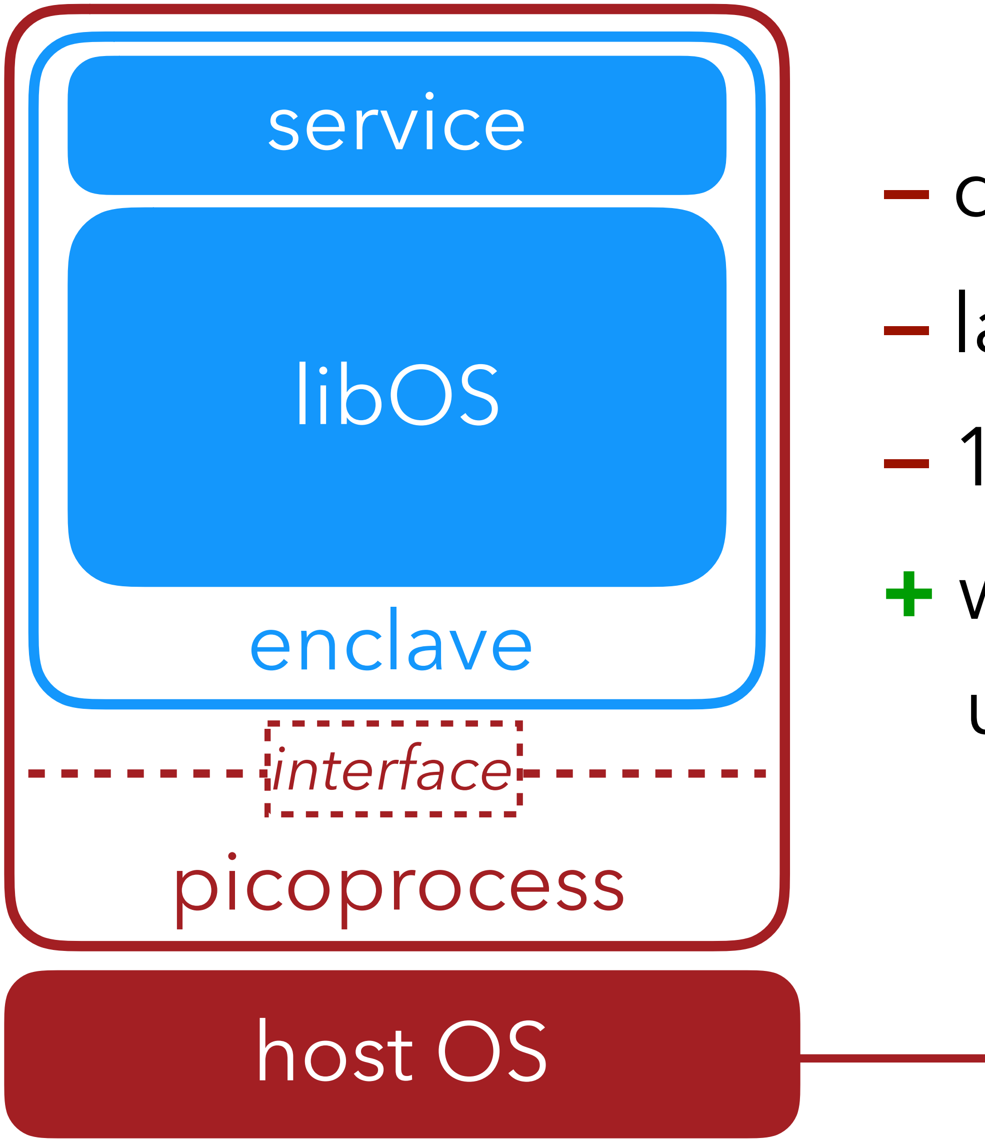

- – designed for Intel SGX
- – large TCB (due to libOS)
- – 10s of new interface calls
- + works with unmodified applications

VC^3

(IEEE S&P'15)

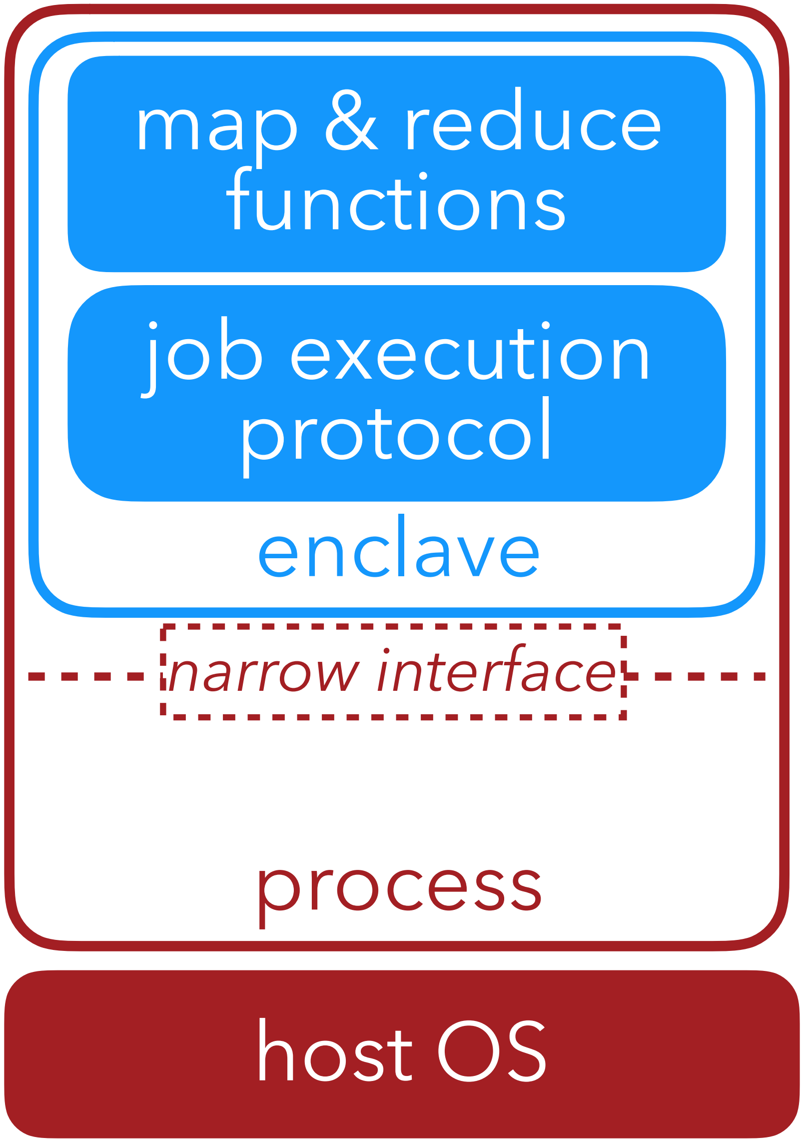

– designed for Intel SGX
– specific for Hadoop
+ small TCB
+ data confidentiality
+ can run unmodified Hadoop applications

A Niche in the State of the Art

	small TCB	Large State	Interface calls	App Specific	Trusted Computing arch.
Haven (OSDI'14)	No	Yes	tens	No	SGX
VC3 (S&P'15)	Yes	Yes MapReduce workloads	R,W	Yes	SGX
XMHF-TrustVisor (S&P'13,'10)	Yes	No	none (but Minibox has tens)	No	TPM / TXT
LaSt^{GT}	Yes	Yes	zero!	No	TV&SGX

Outline

- **Our solution: key ideas and overview**

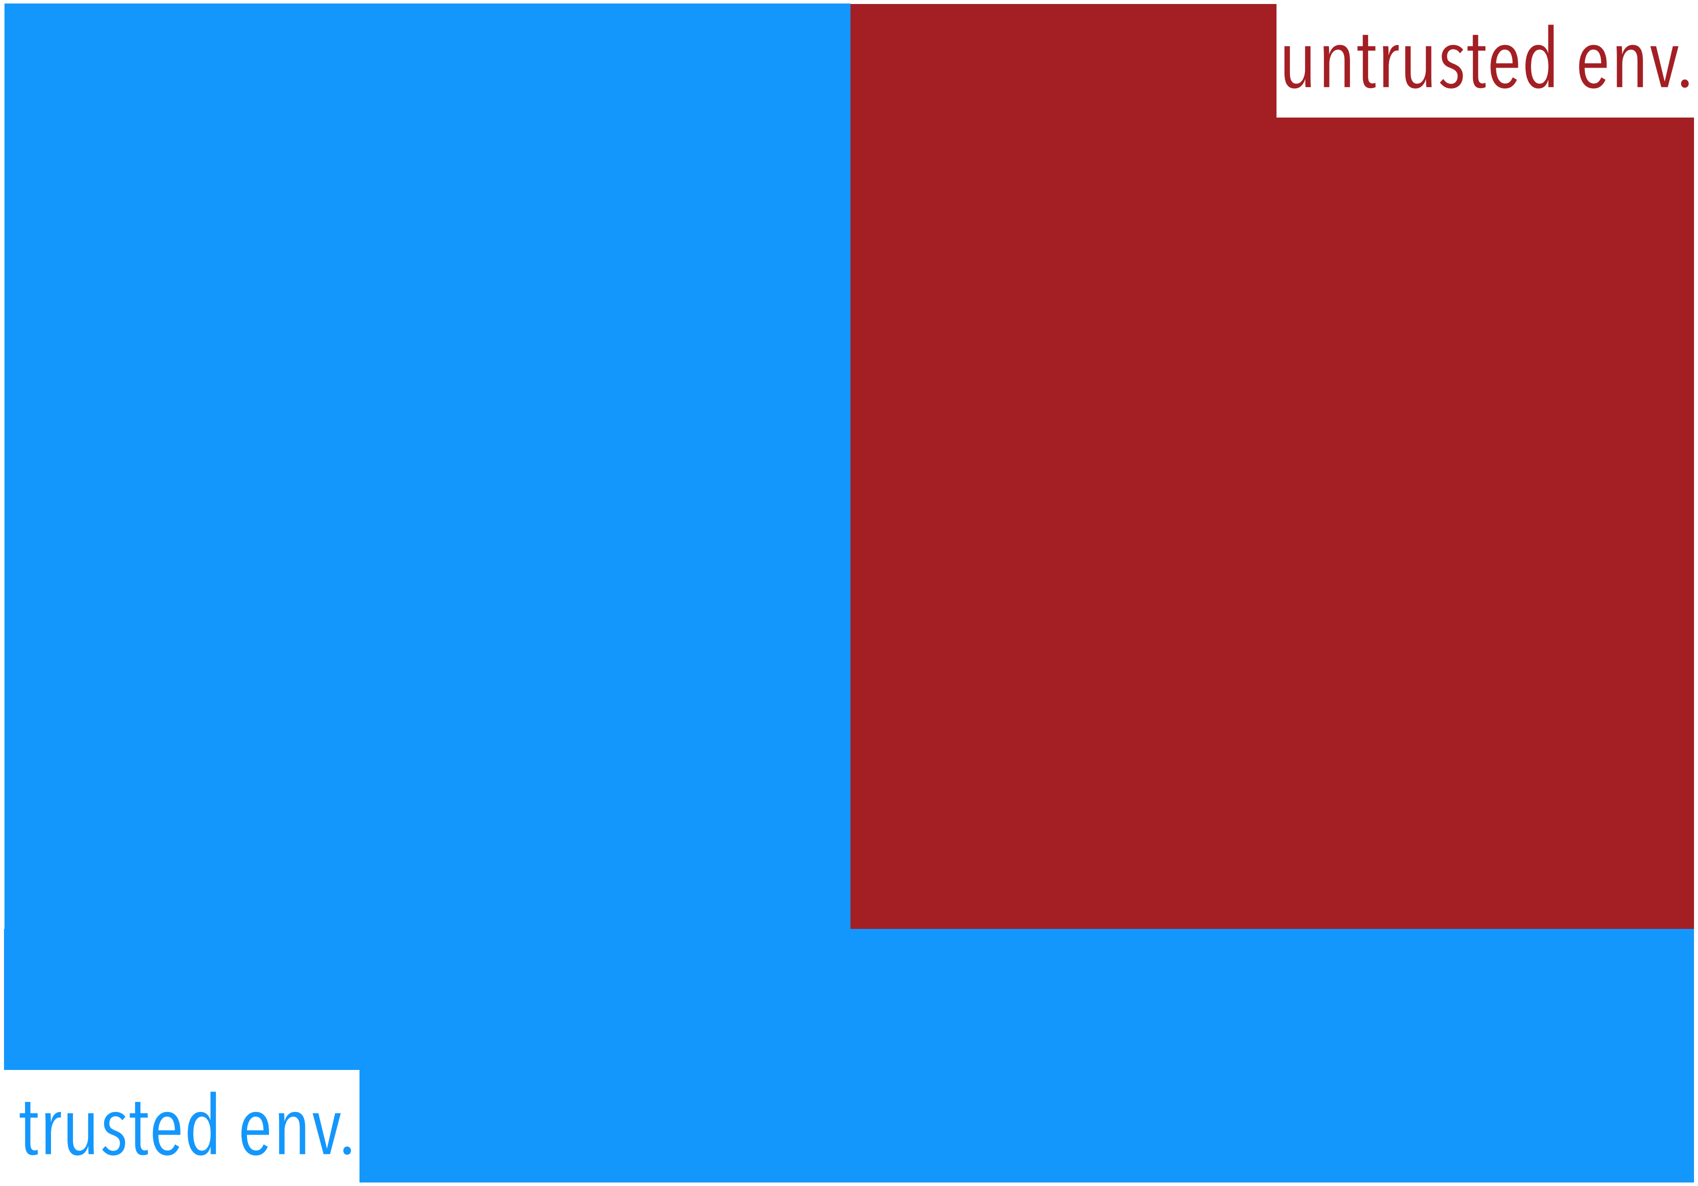

Scenario: two execution environments

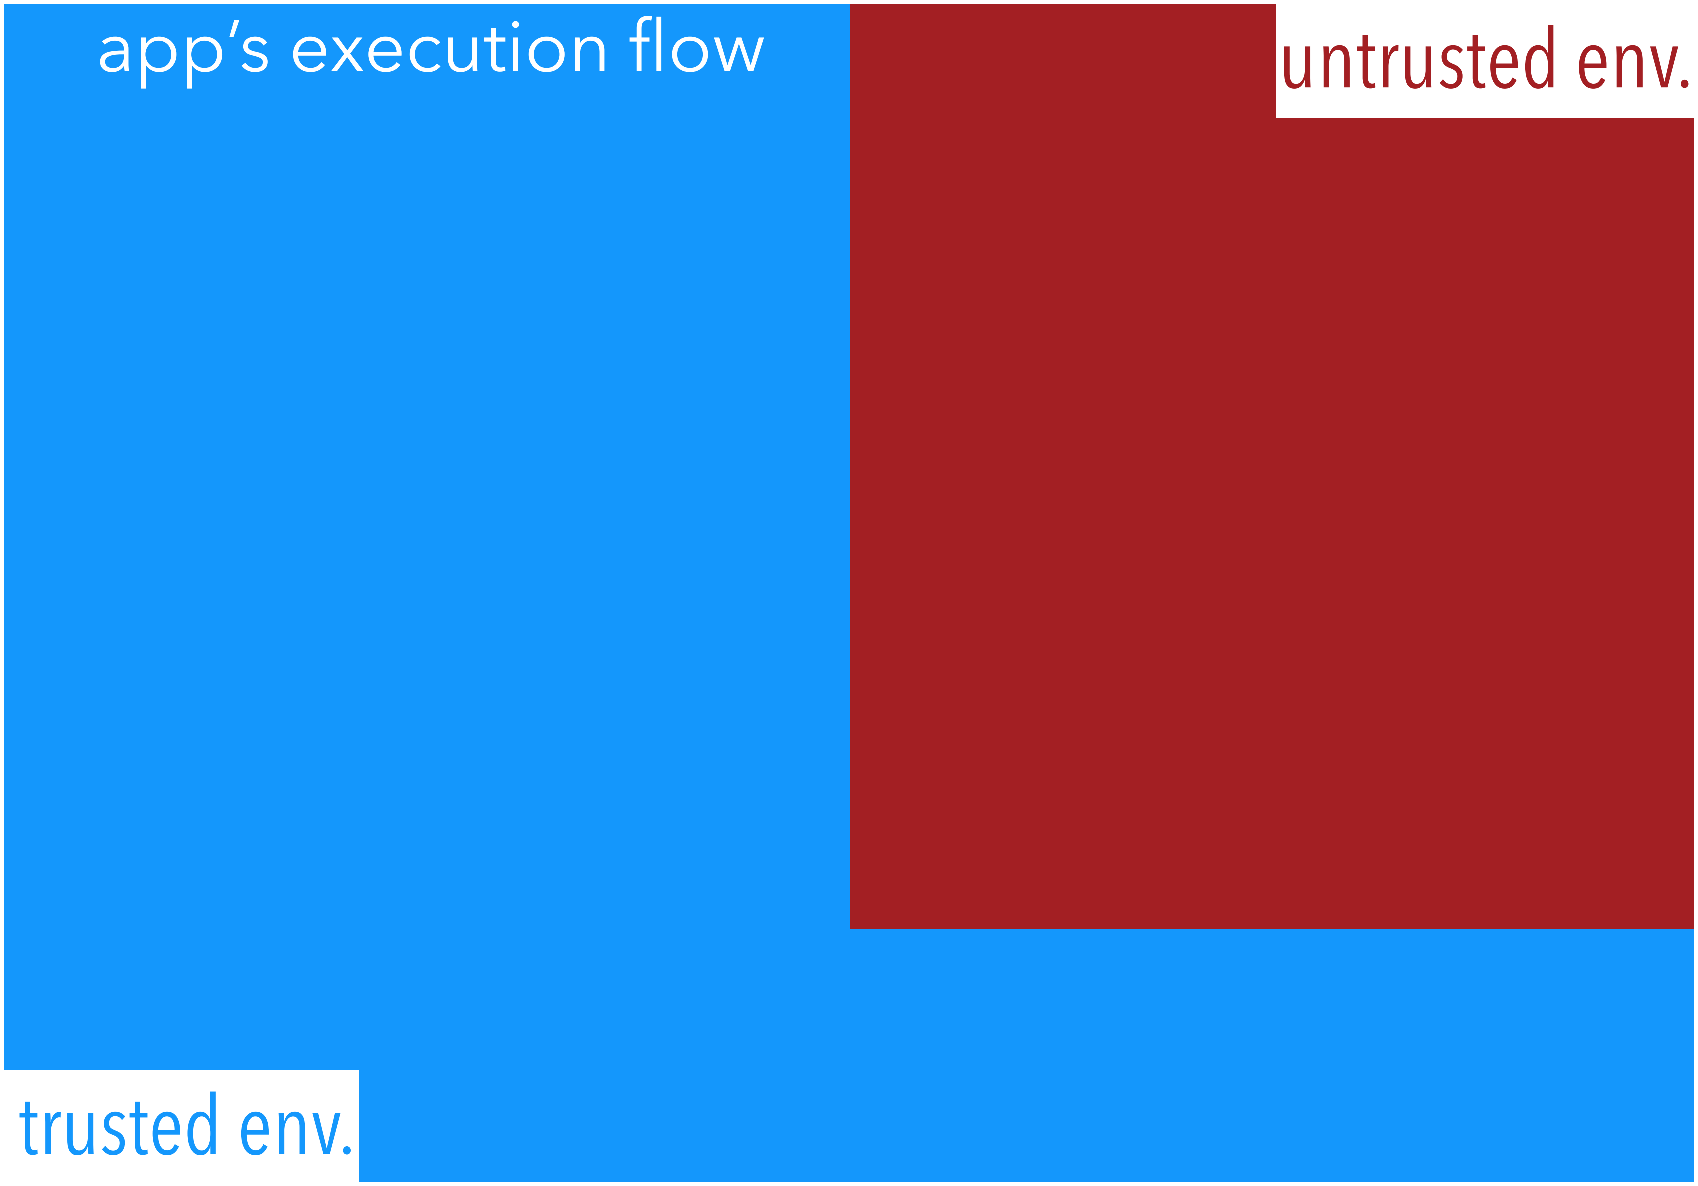

the service code is running

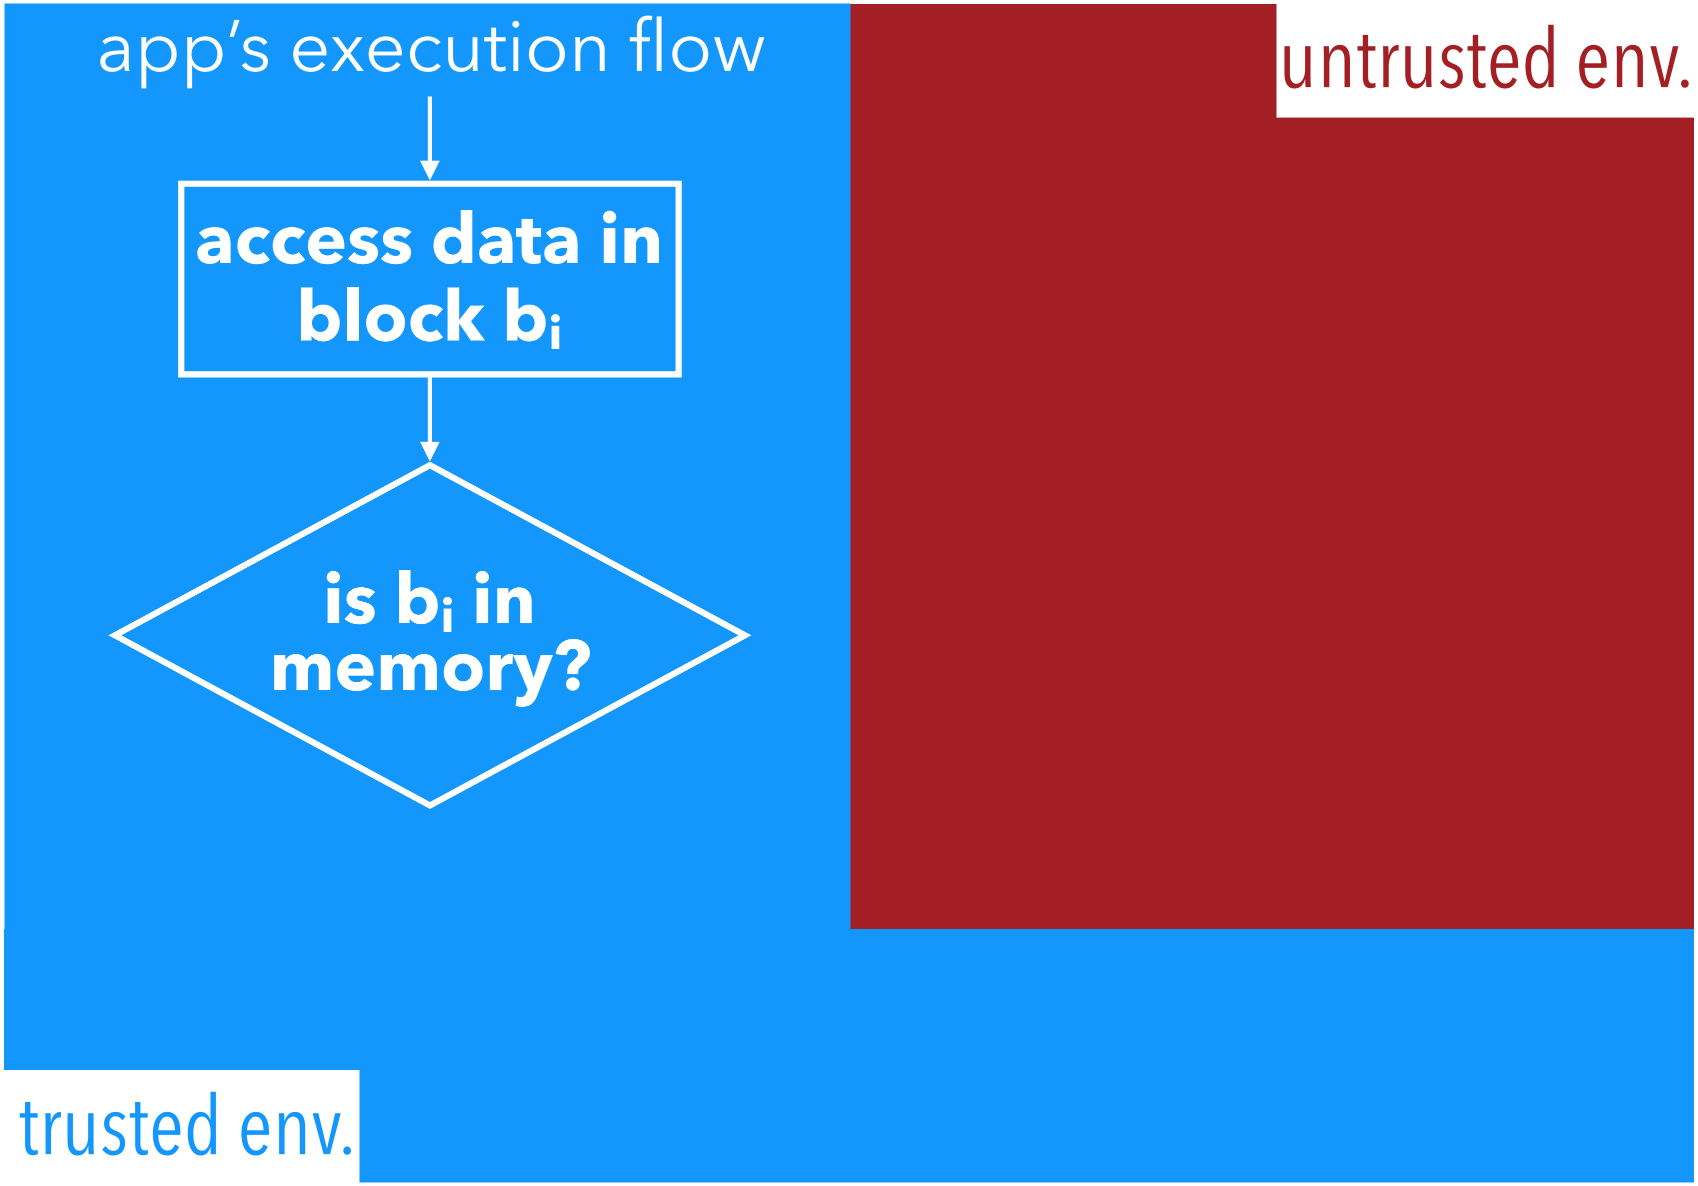

the service code accesses data in memory

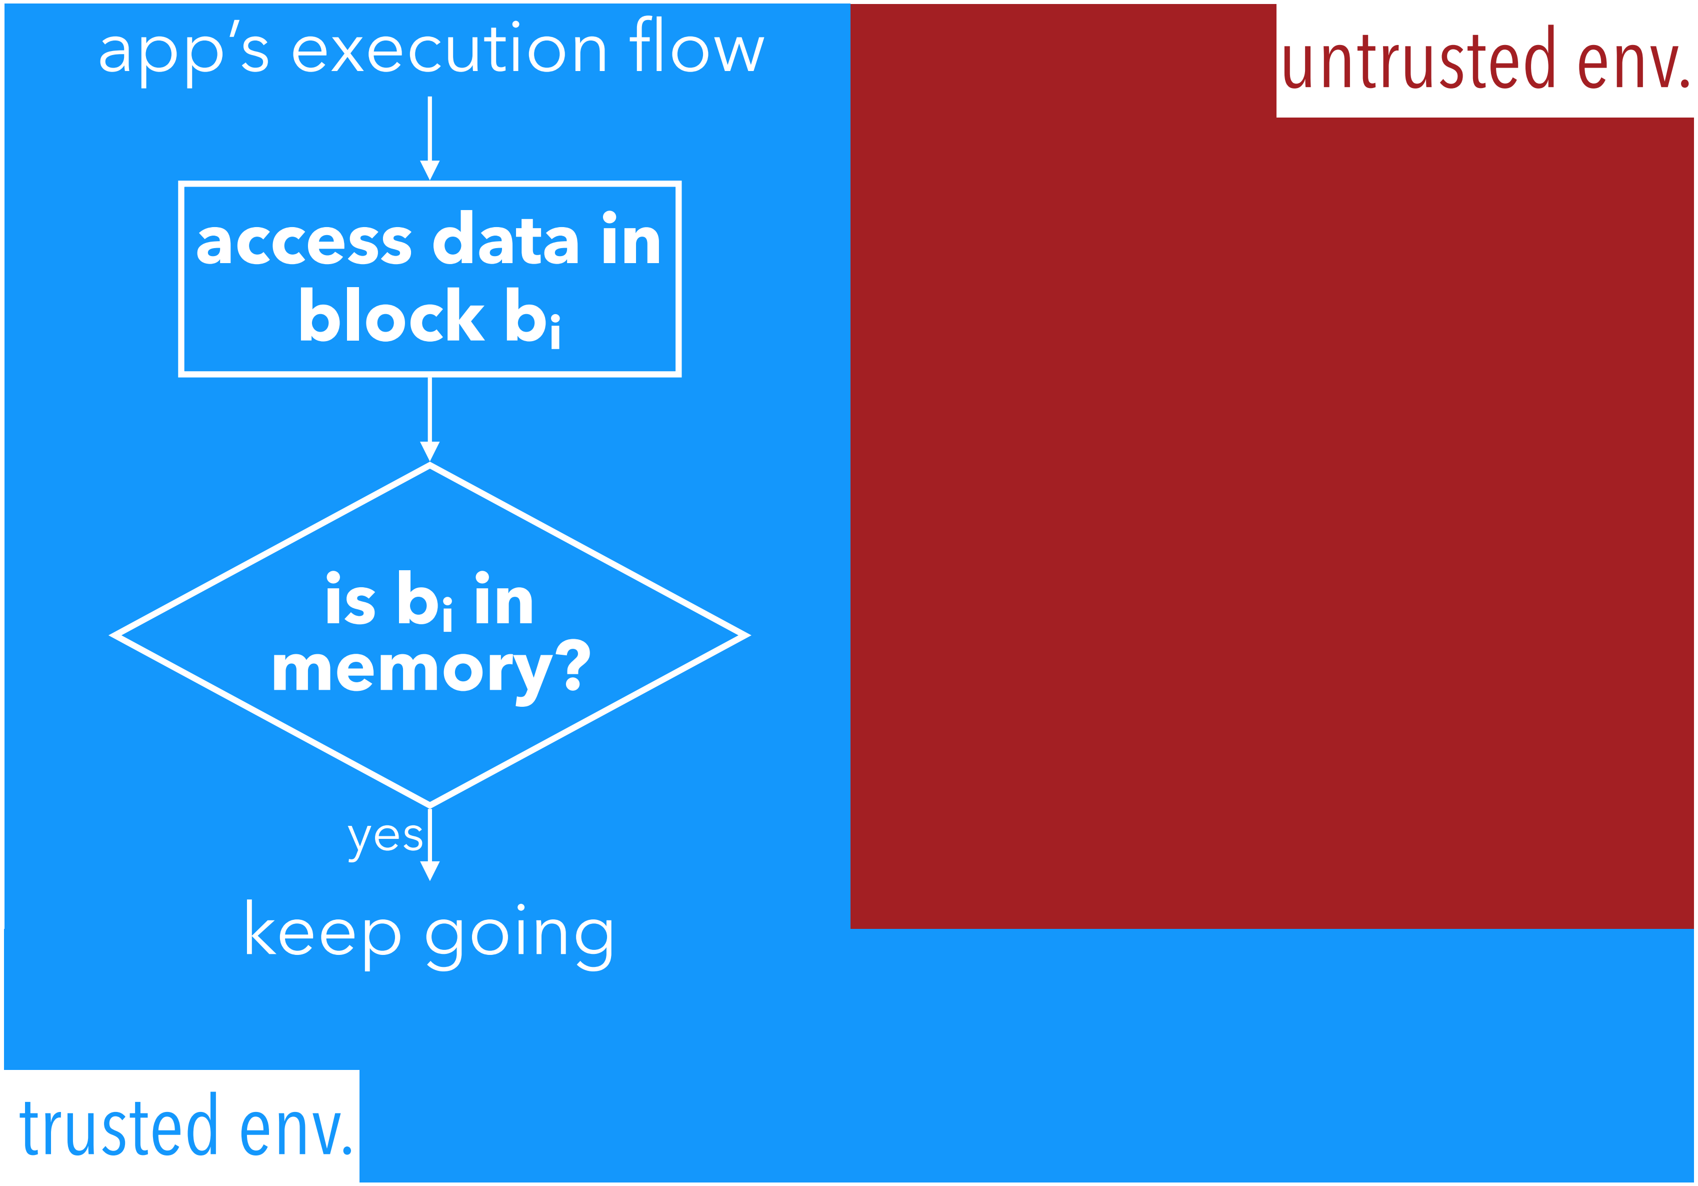

when data is available, there are no interruptions

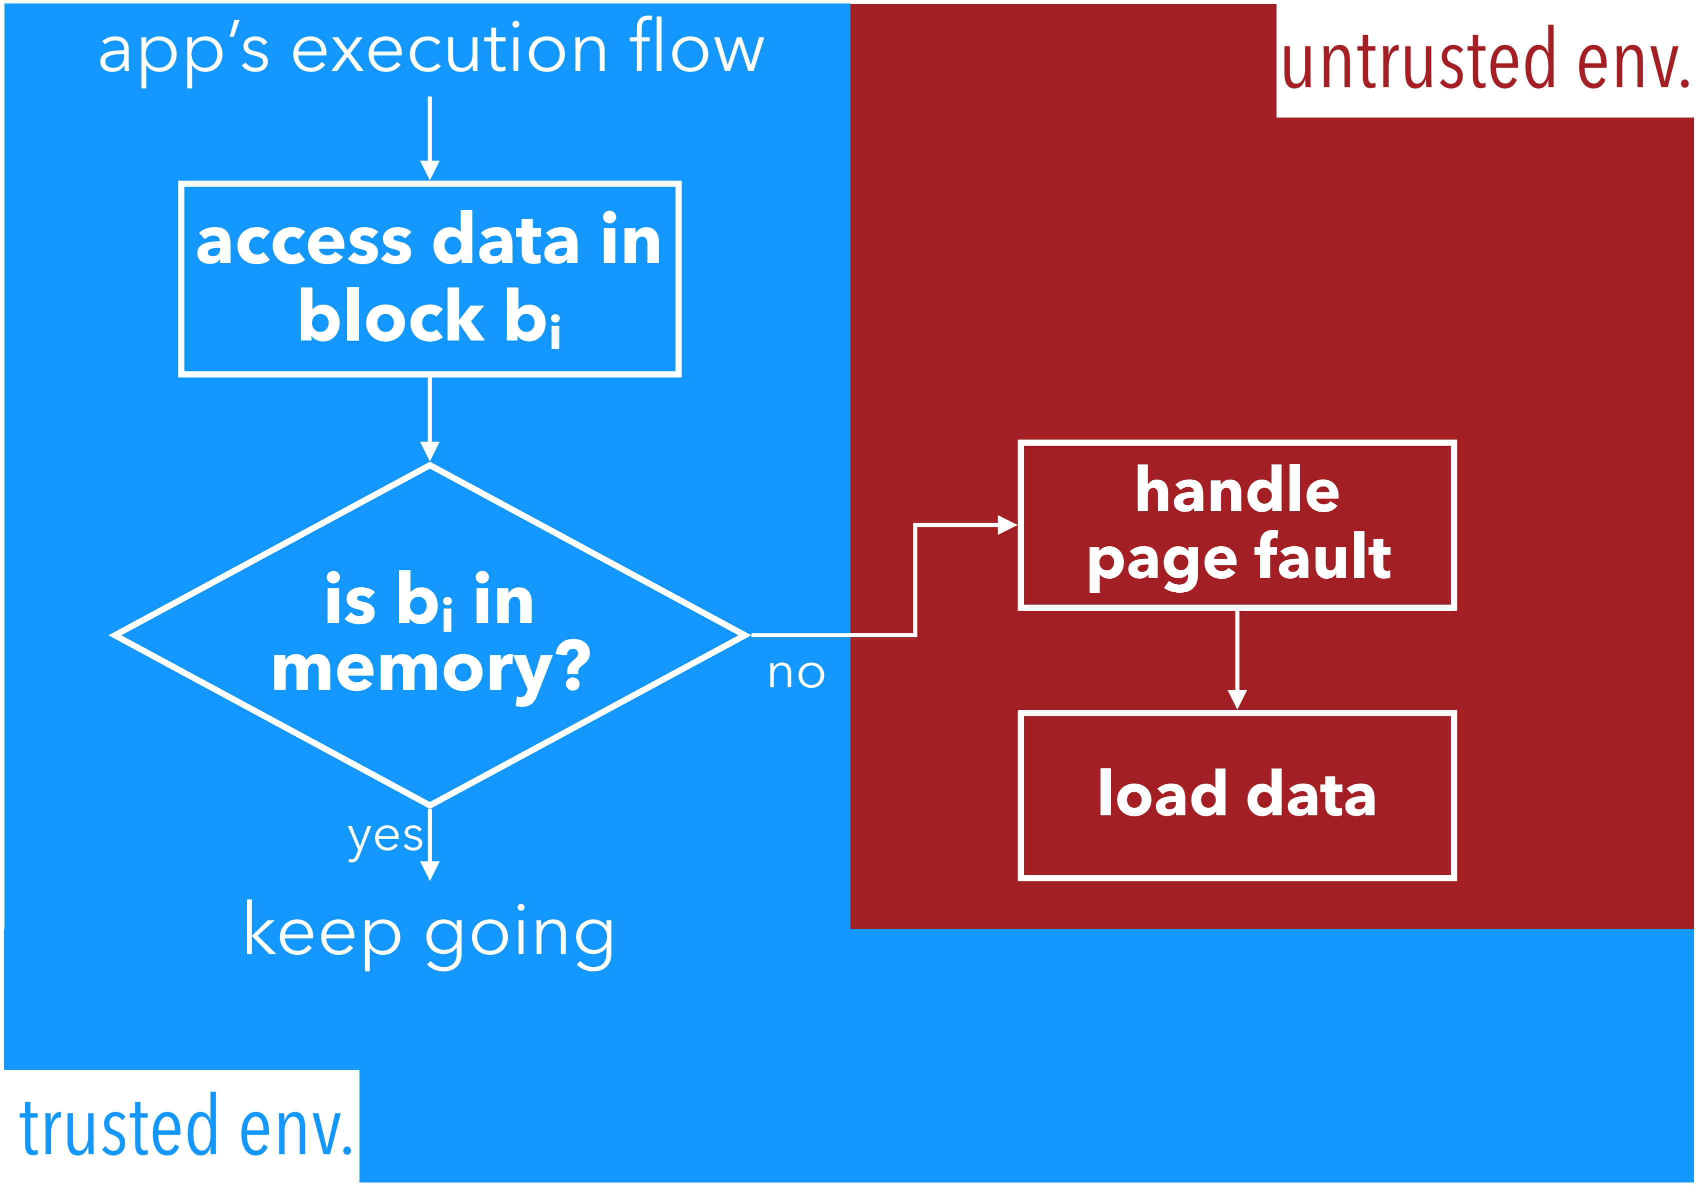

otherwise, the service is interrupted and data memory pages are loaded

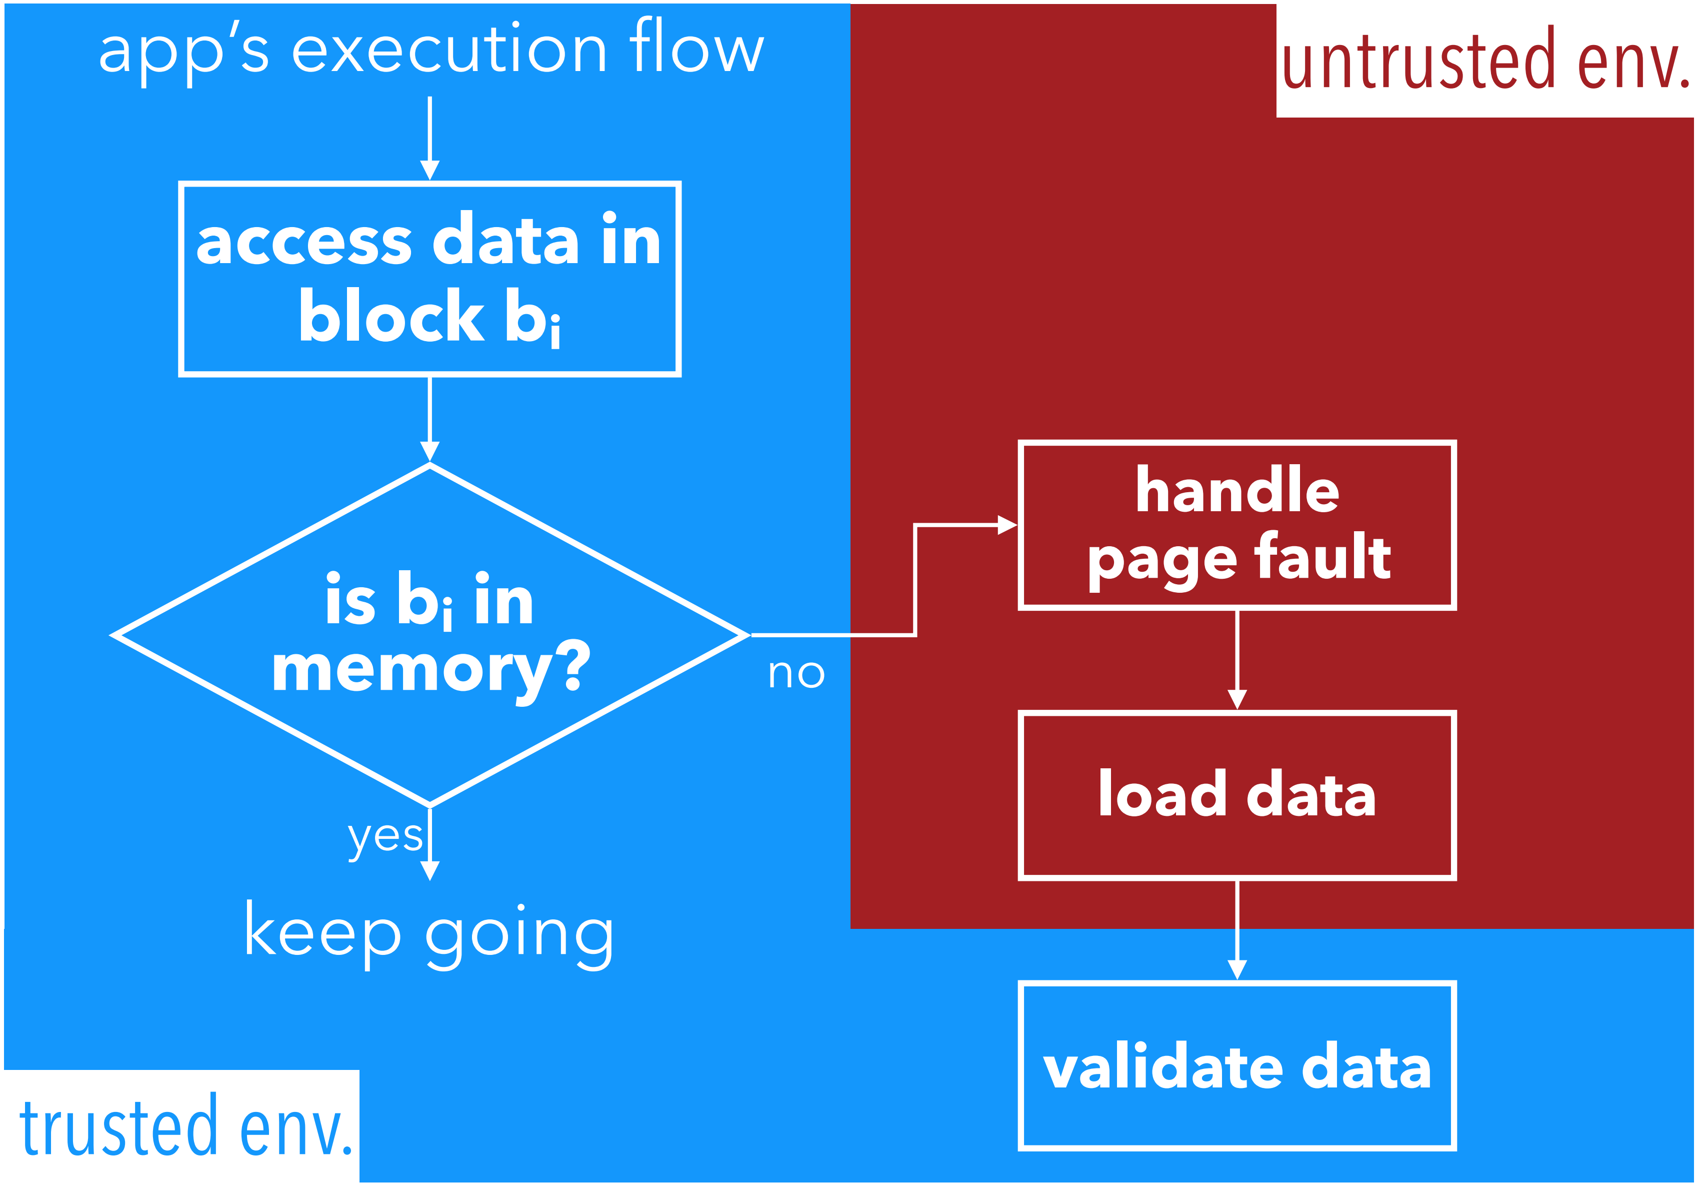

data is validated inside trusted environment, independently from service execution

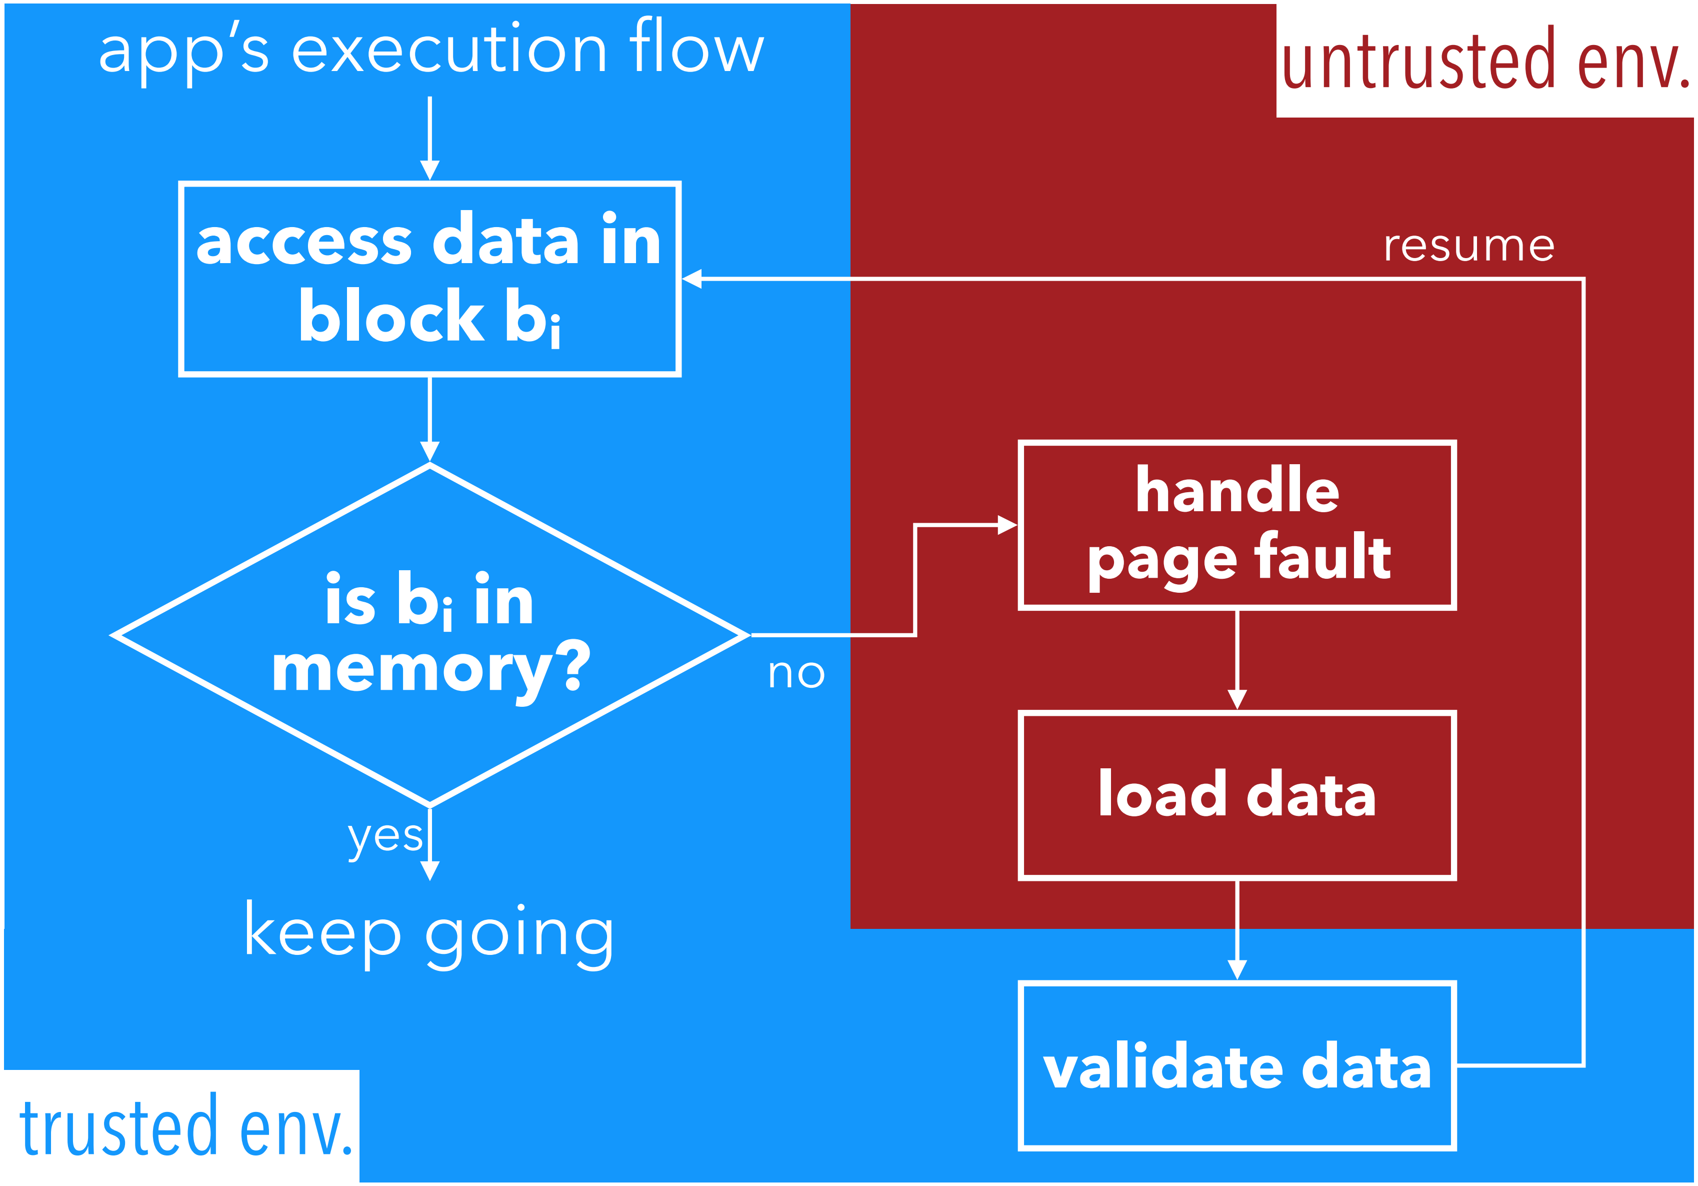

service is resumed and
only if data is valid, service can make progress

…in practice…

Architecture

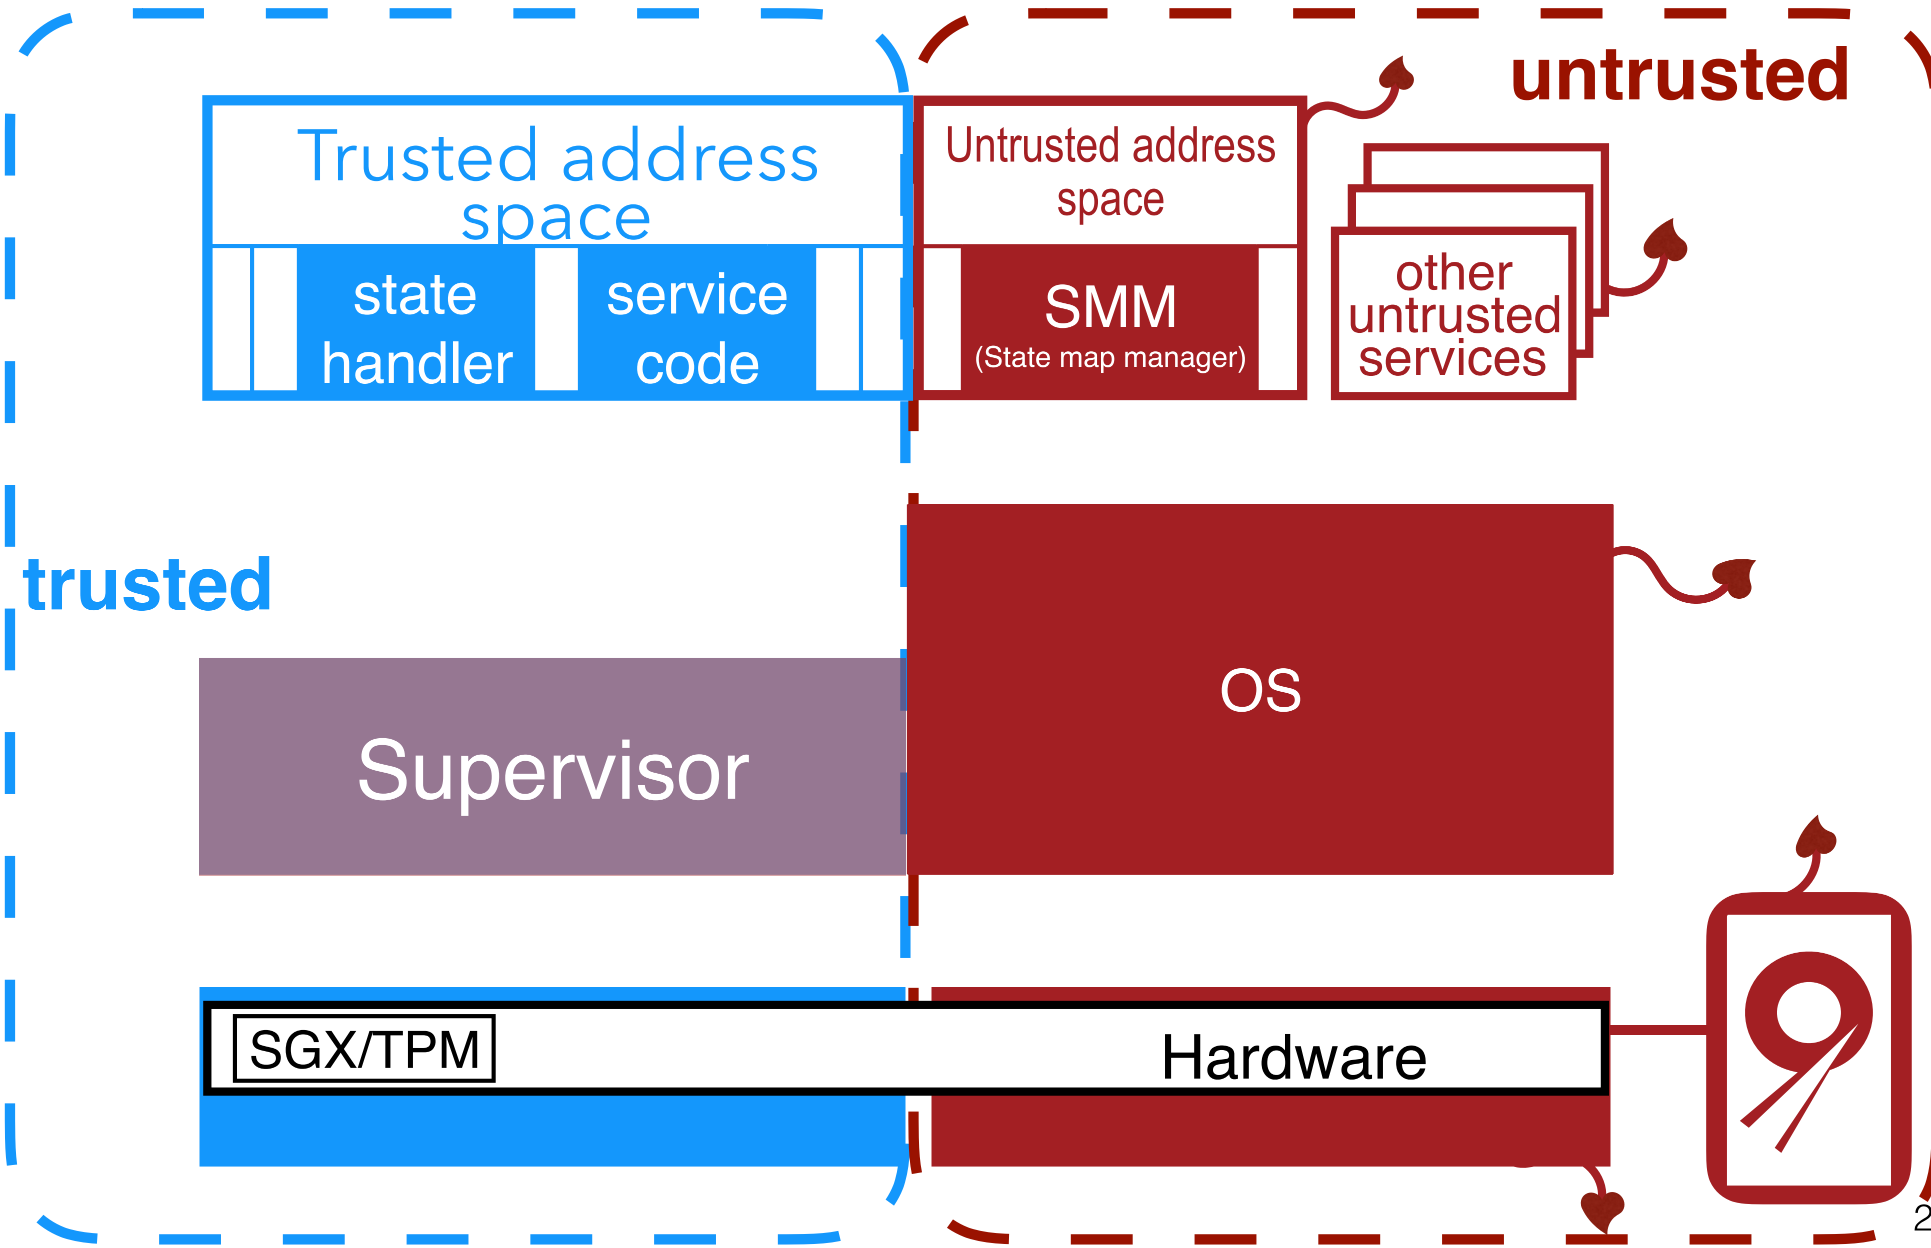

Architecture

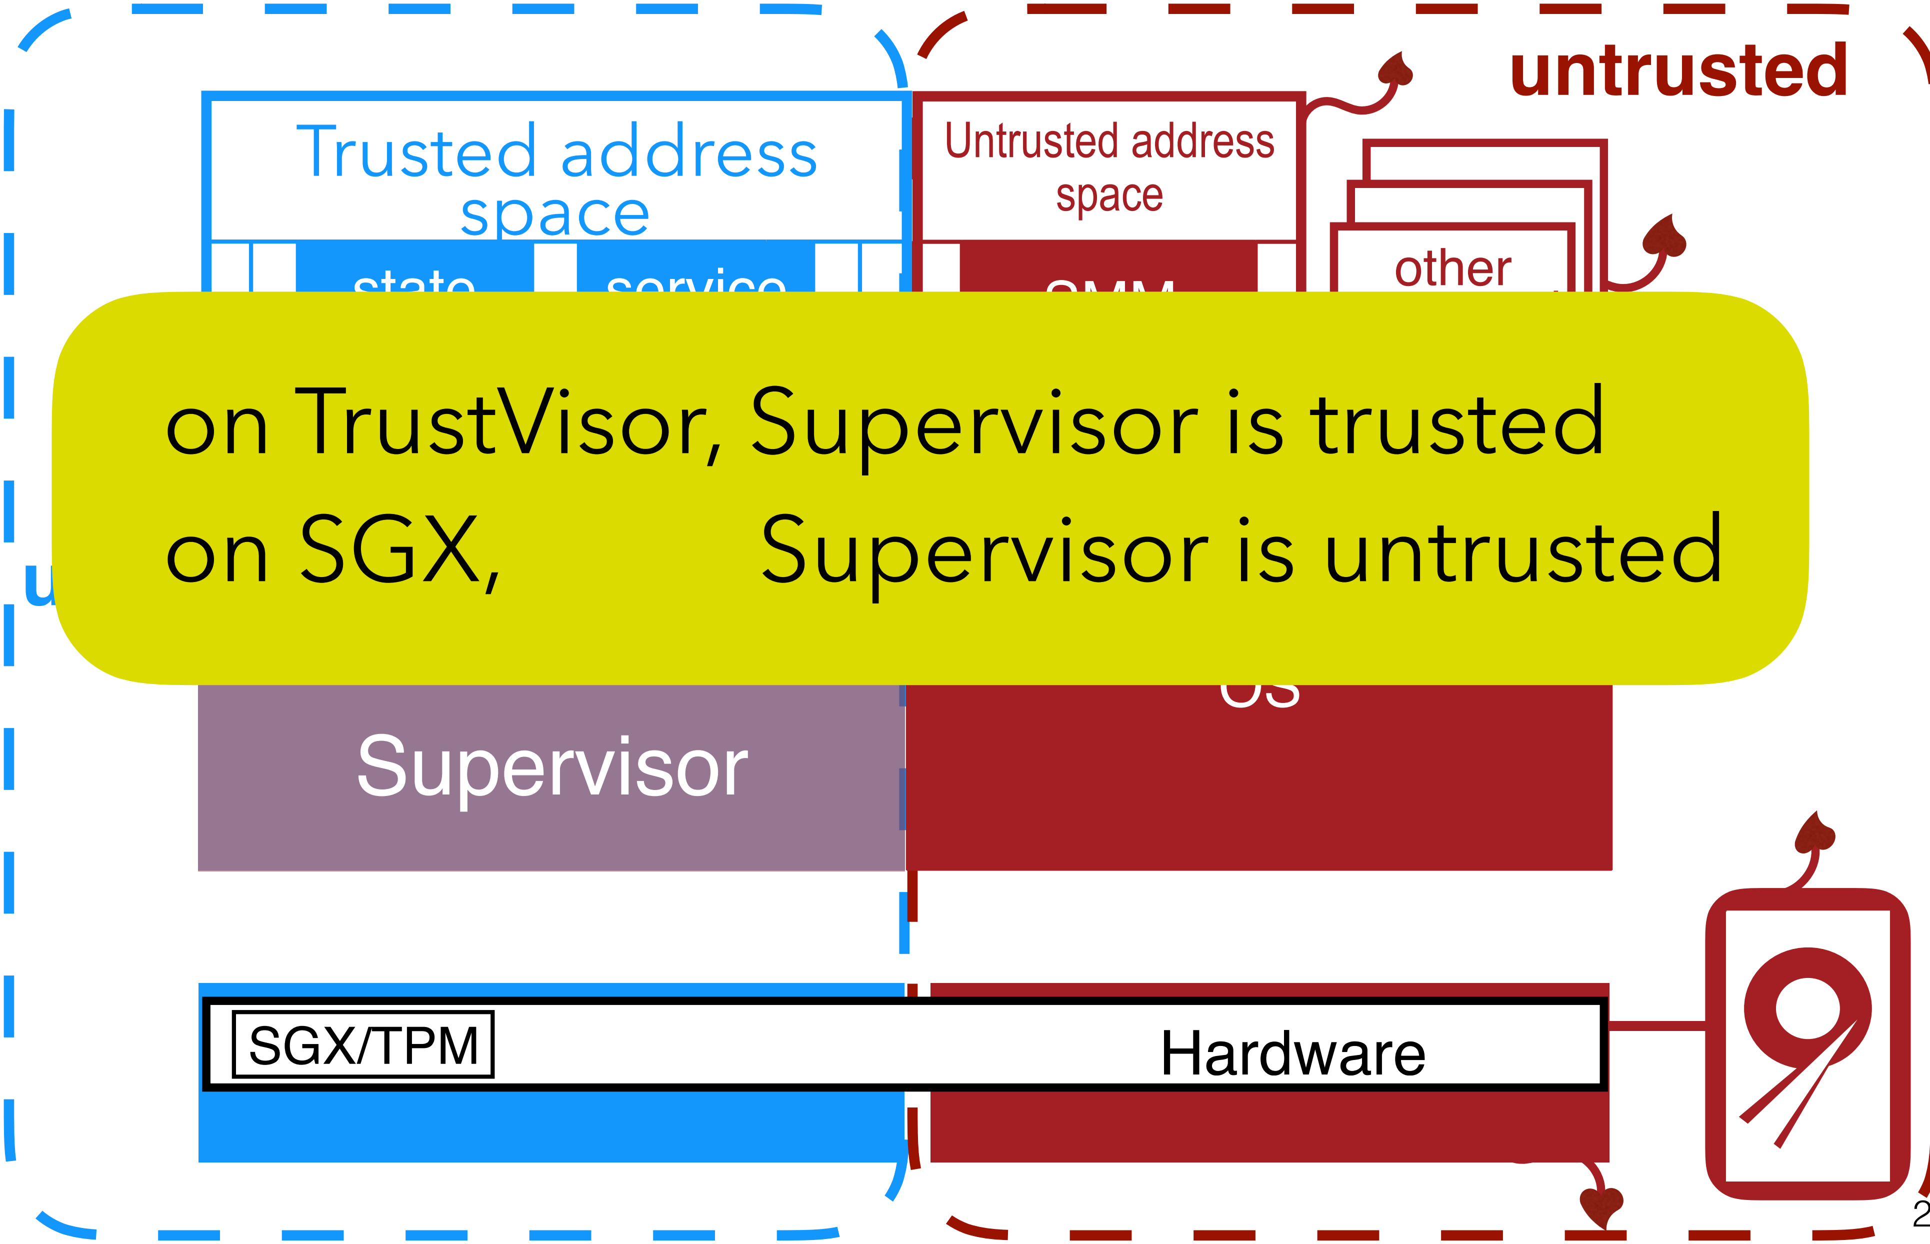

on TrustVisor, Supervisor is trusted
on SGX, Supervisor is untrusted

LaSt^{GT} in ~~5~~ 4 steps

- Offline data protection at the source
- State registration
- Data processing
- Lazy loading from memory & disk
- ~~Execution verification~~

- **Offline data protection at the source**

1

Data protection

Hierarchical
- Incremental as data is created

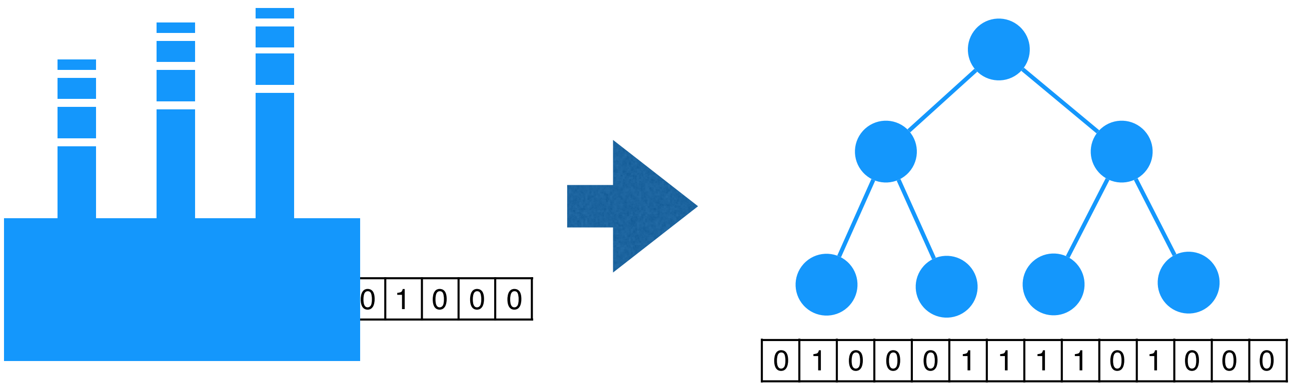

Made for:
- Incremental validation as data is loaded
- Fast verification
- Single hash tree is unsuitable

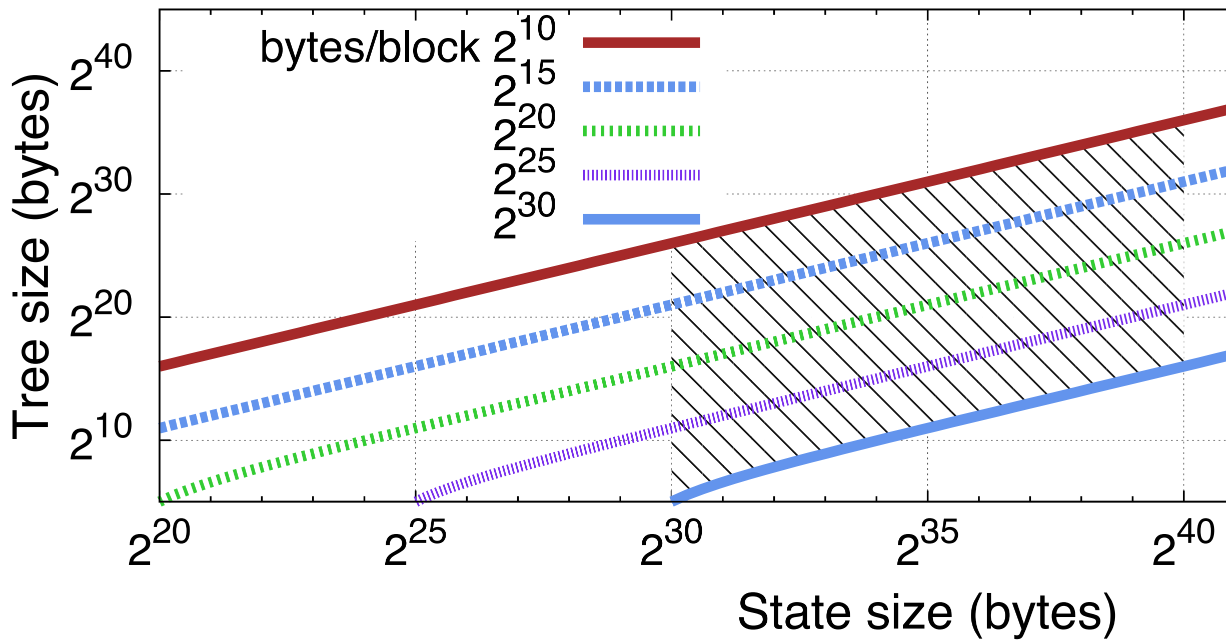

State Hierarchy

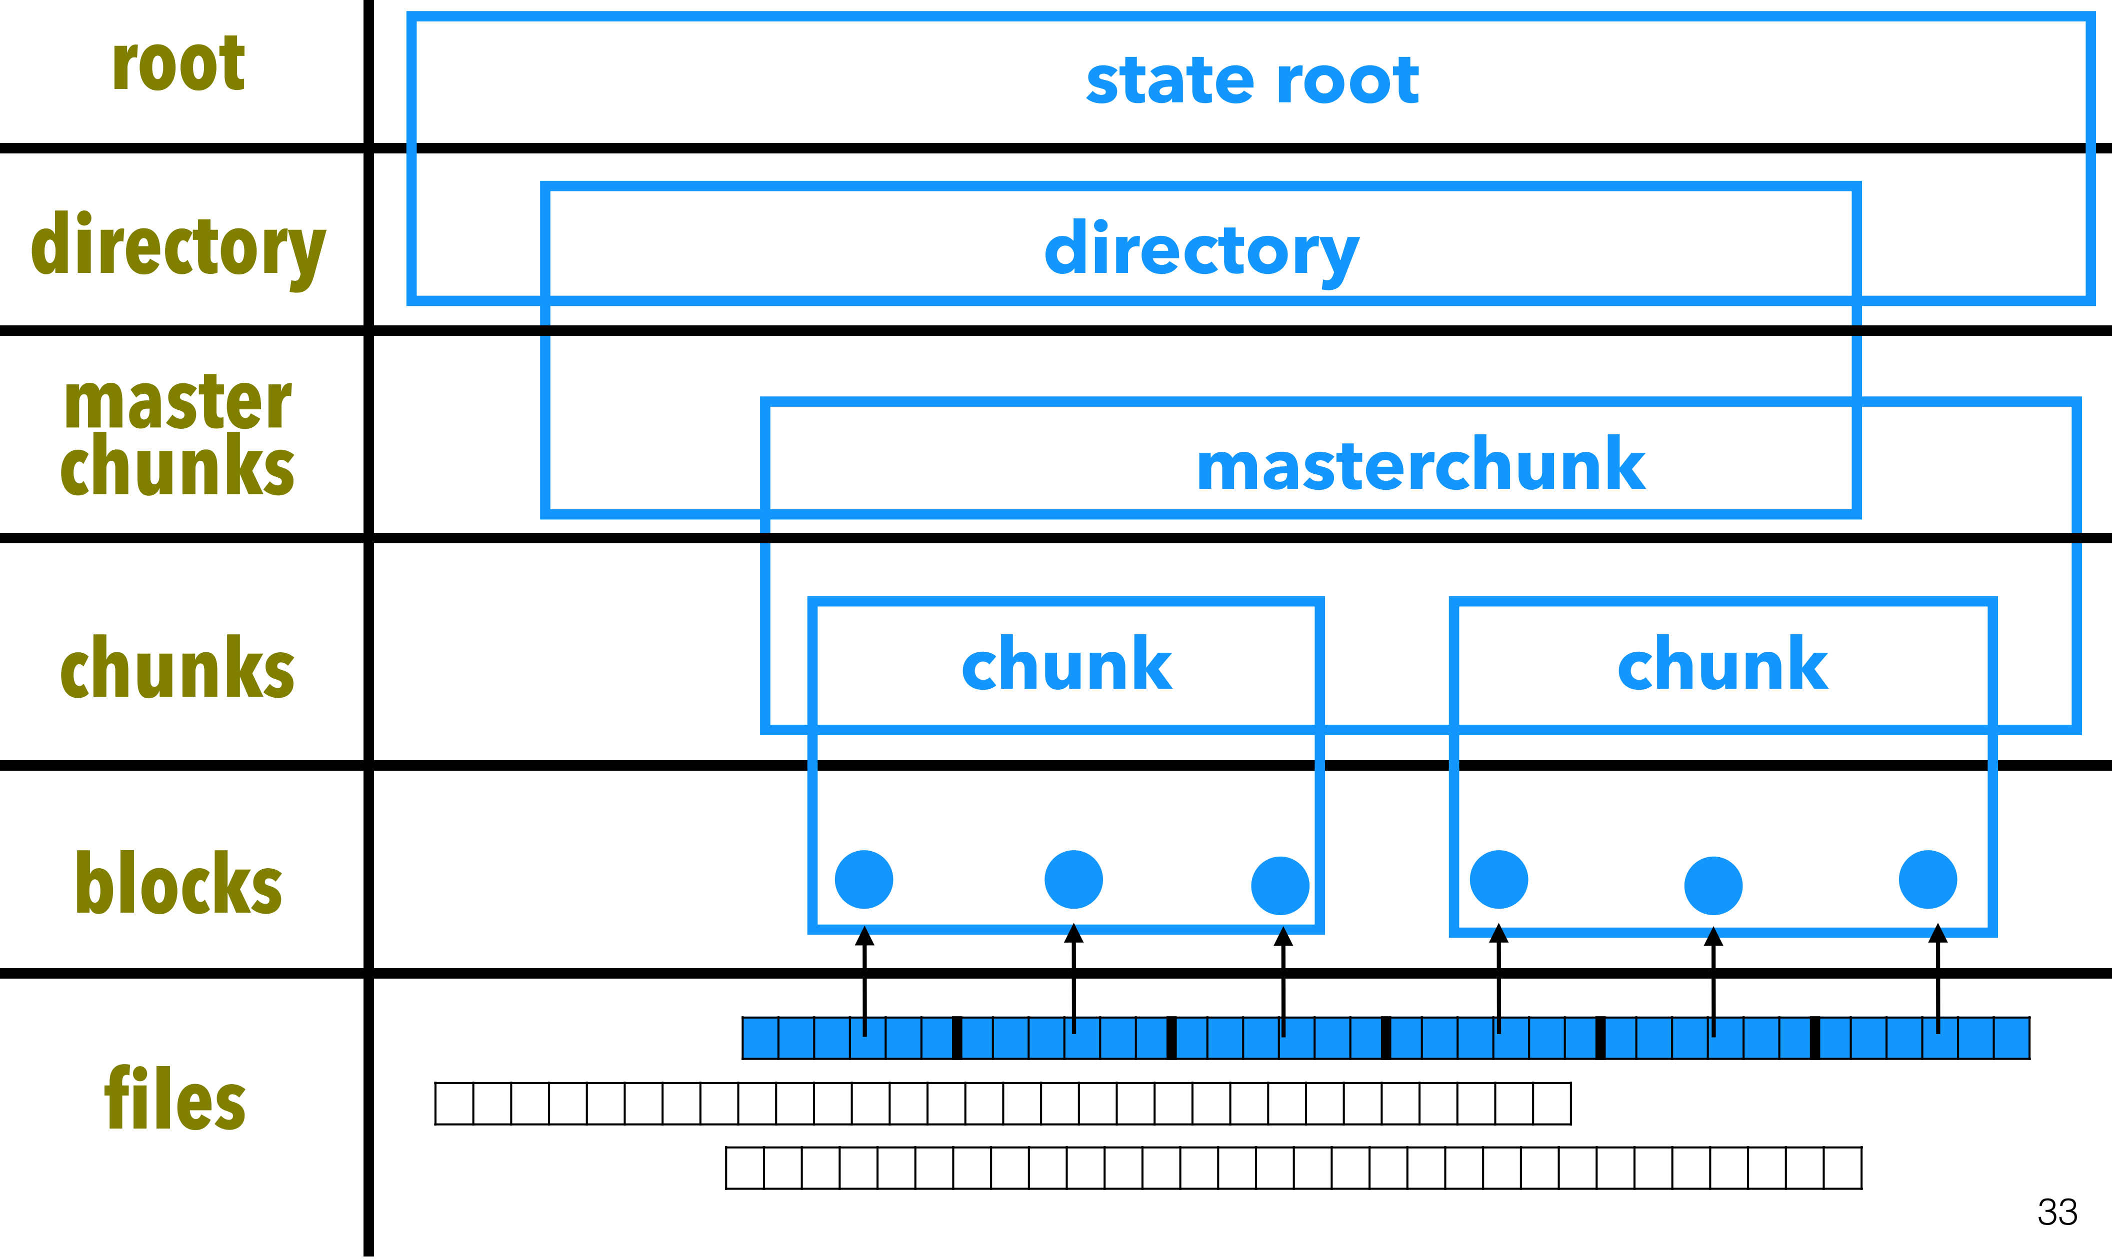

State Hierarchy

- components are loaded separately
- unneeded components not loaded in memory
- state root (1 hash) allows state validation

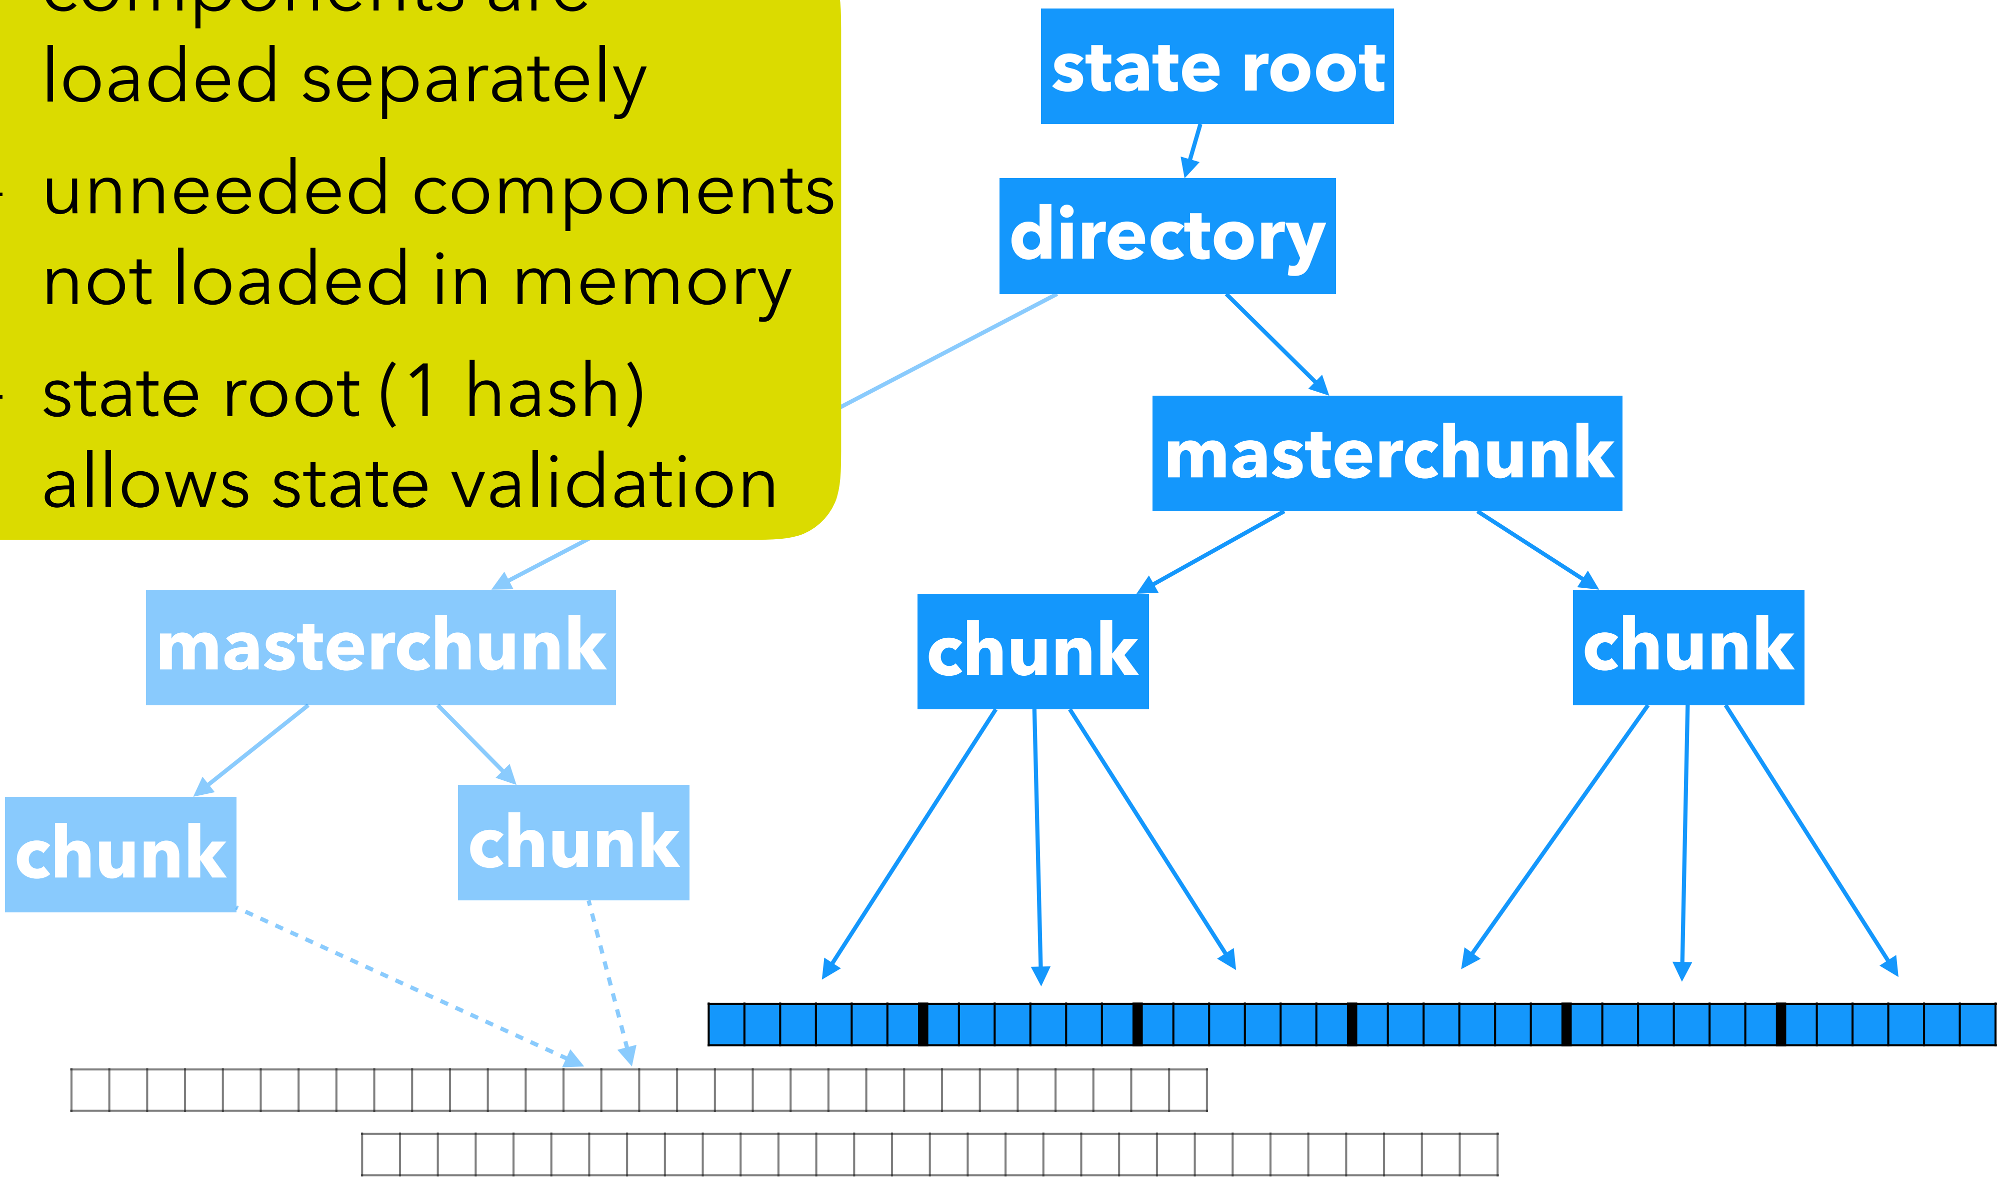

- State registration

2

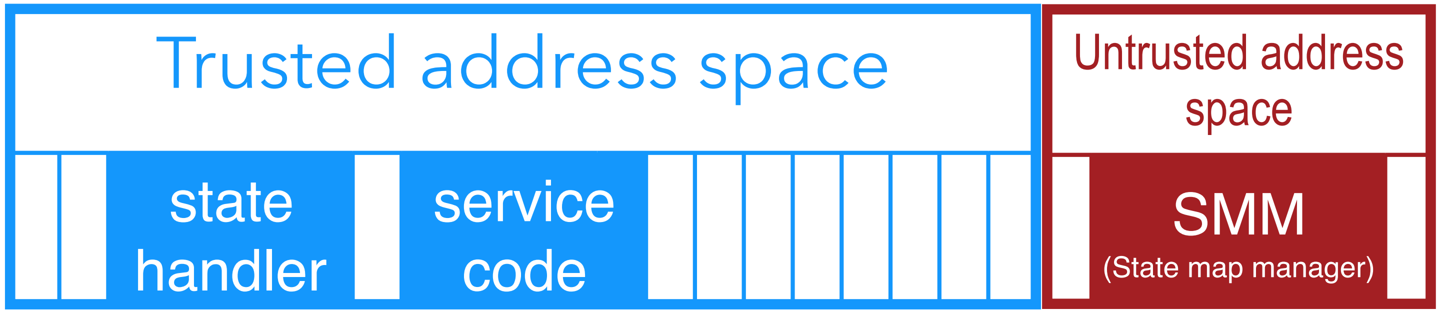

When the trusted execution environment is created, only the code is available inside

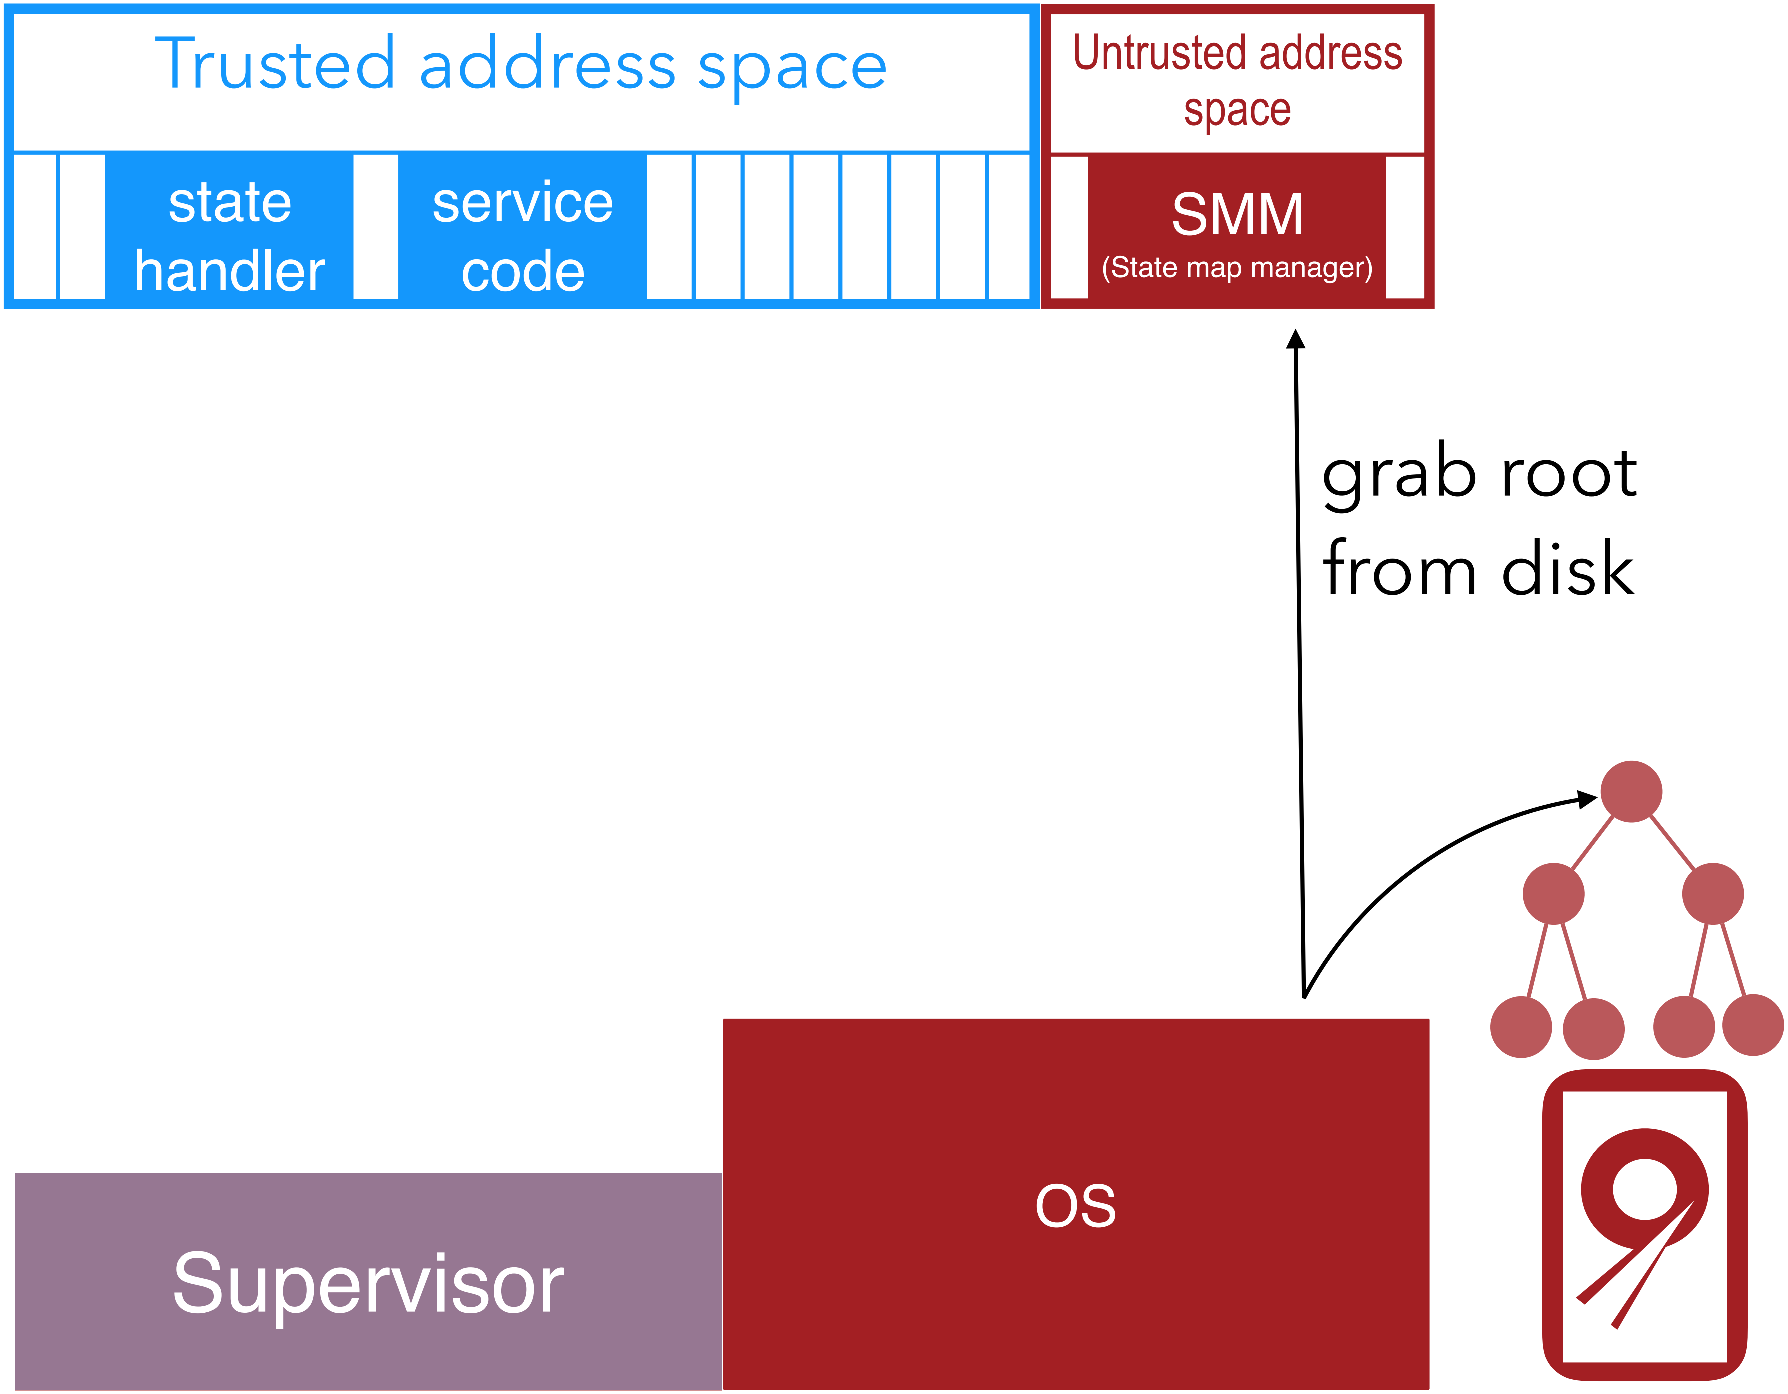

Trusted address space
state handler
service code
Untrusted address space
SMM
(State map manager)
grab root from disk
OS
Supervisor

- registration is the first execution
- state handler installs root
- root is trusted

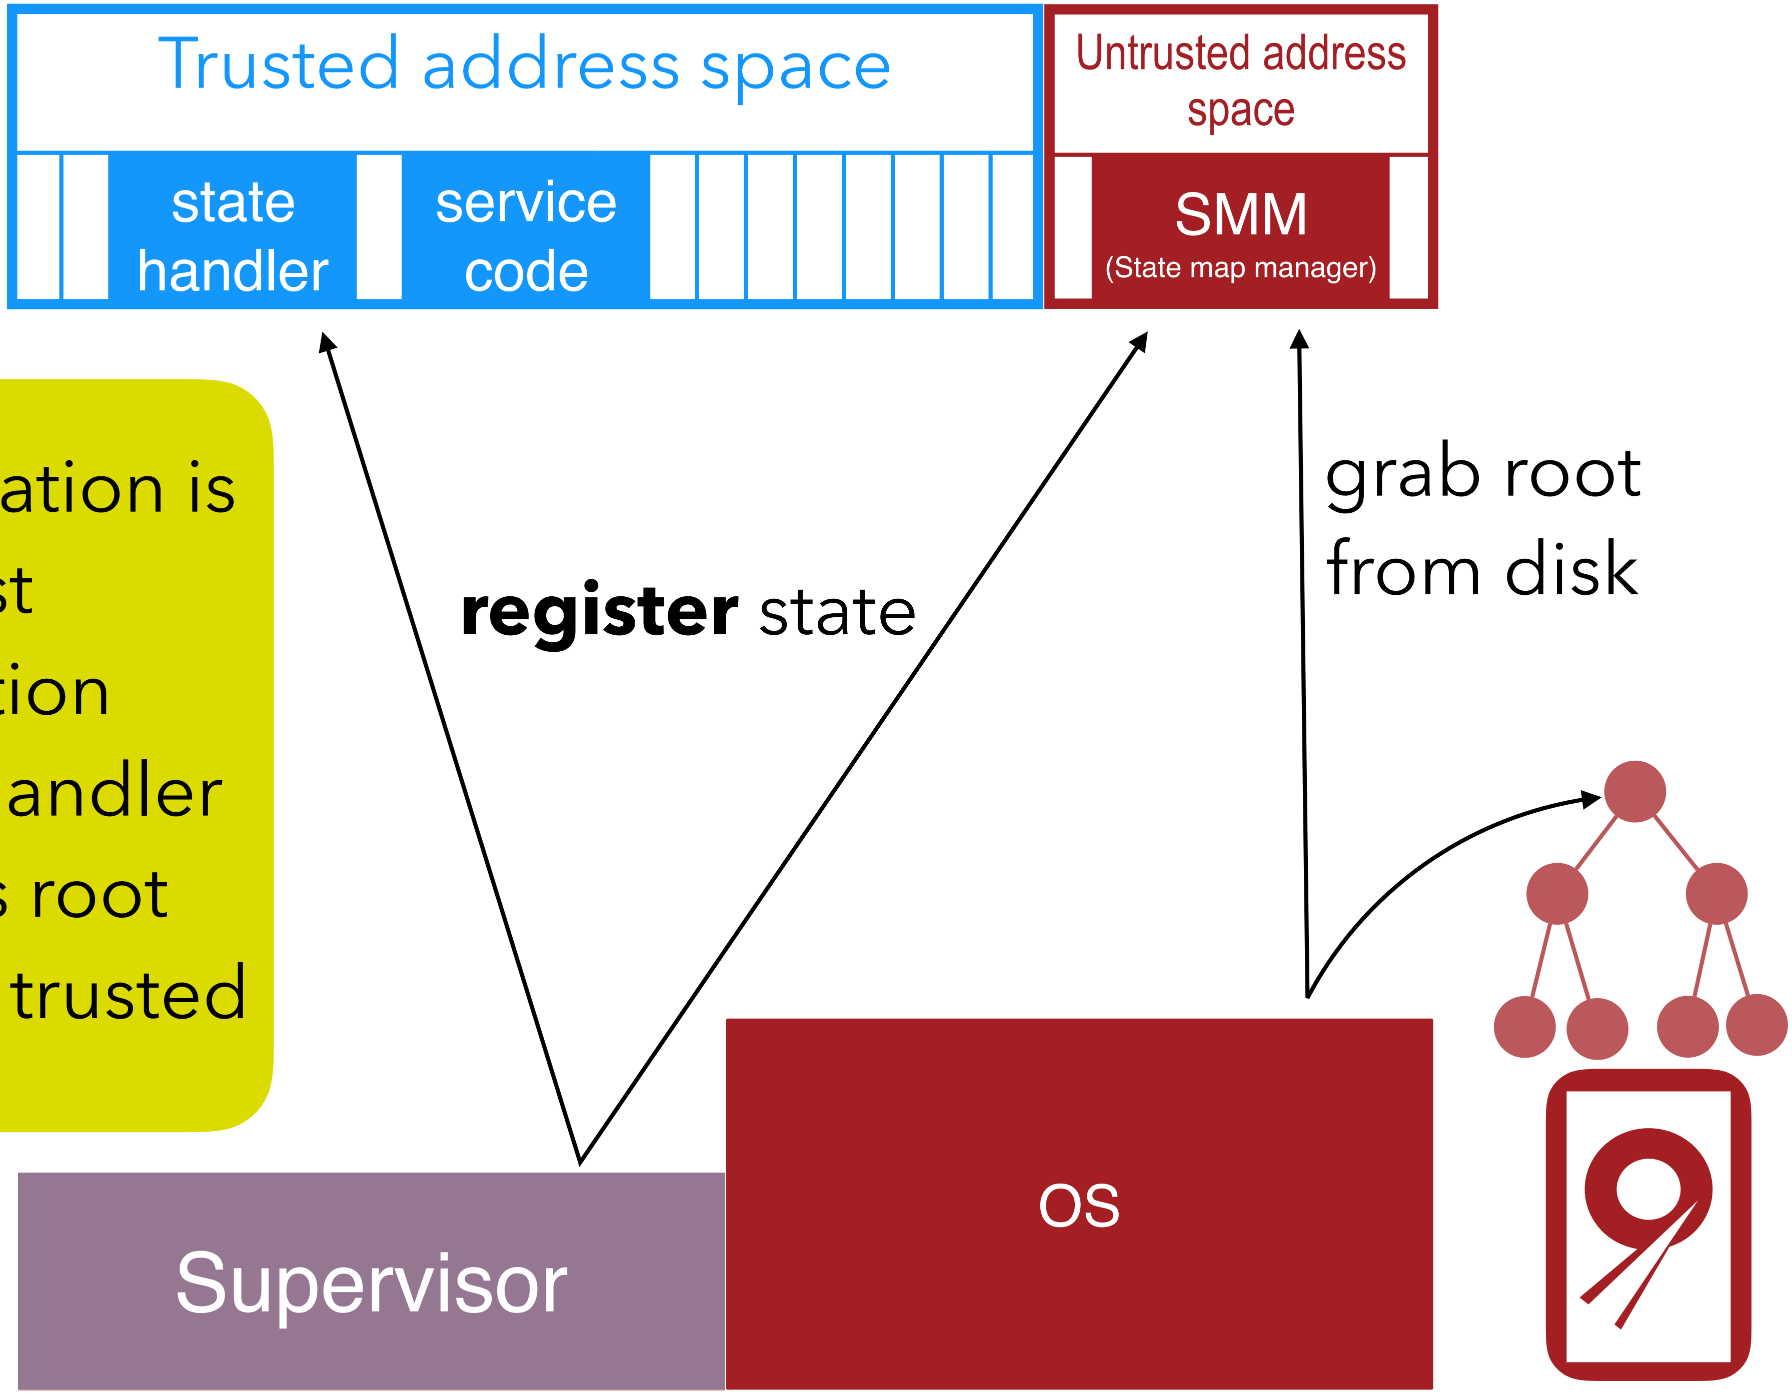

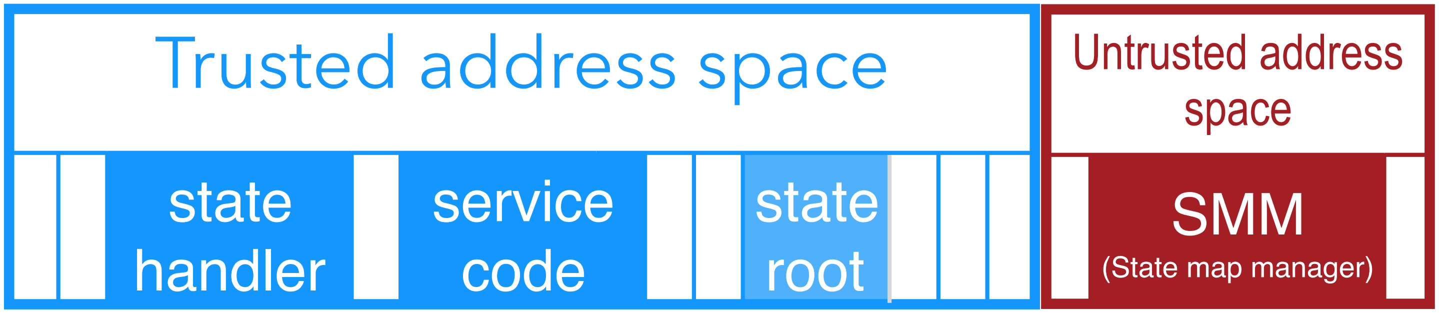

- state root is available before service code runs

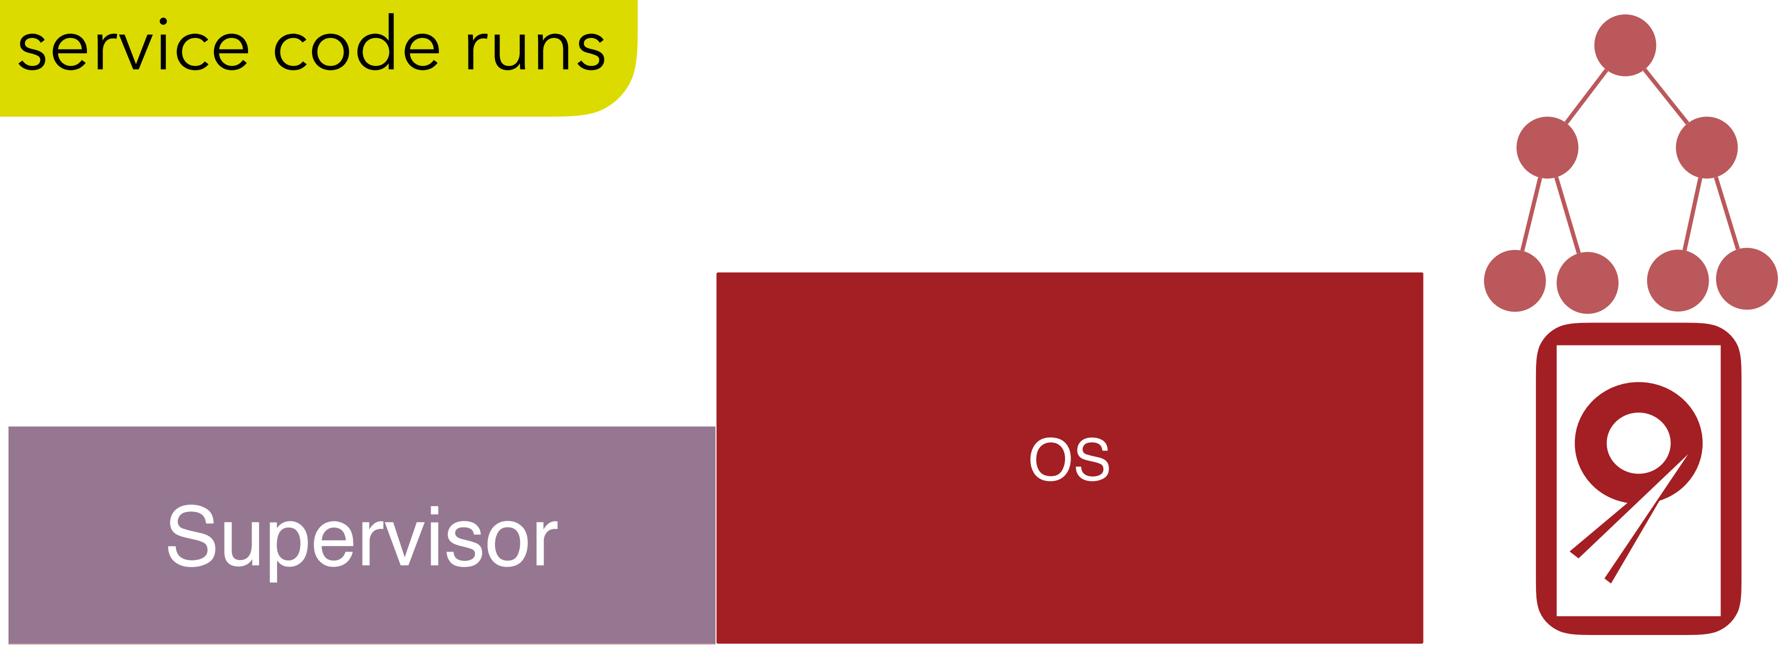

- **Data processing**

3

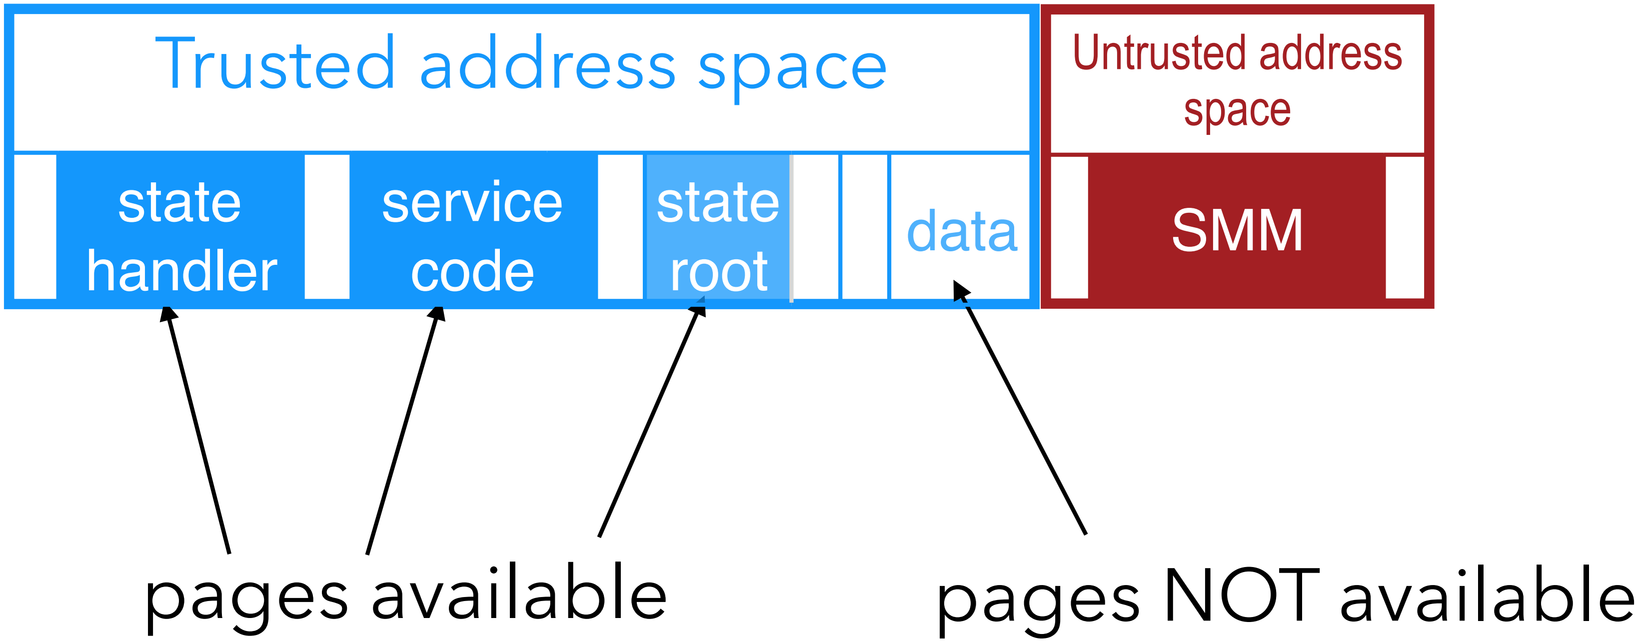

- service code has view of entire state
- state not readily available: inefficient loading it upfront

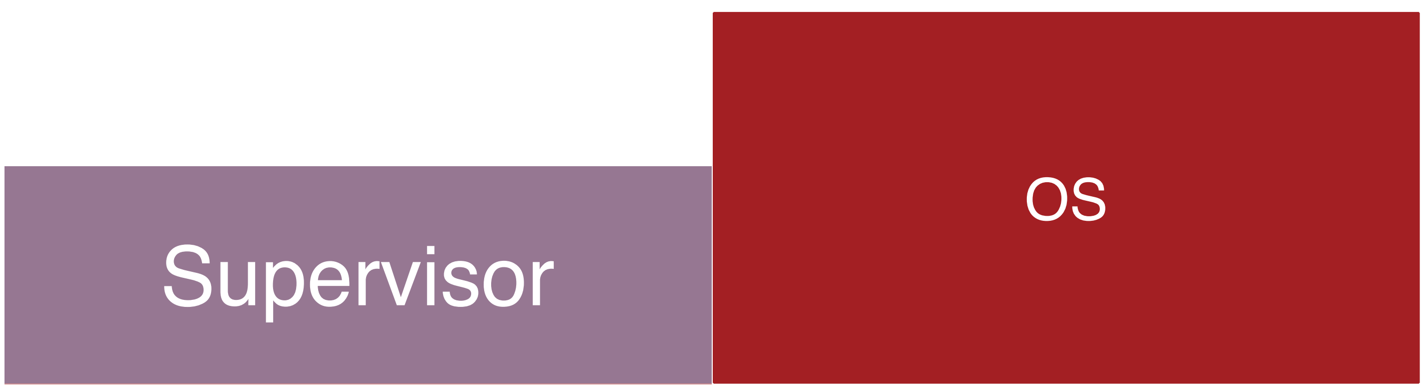

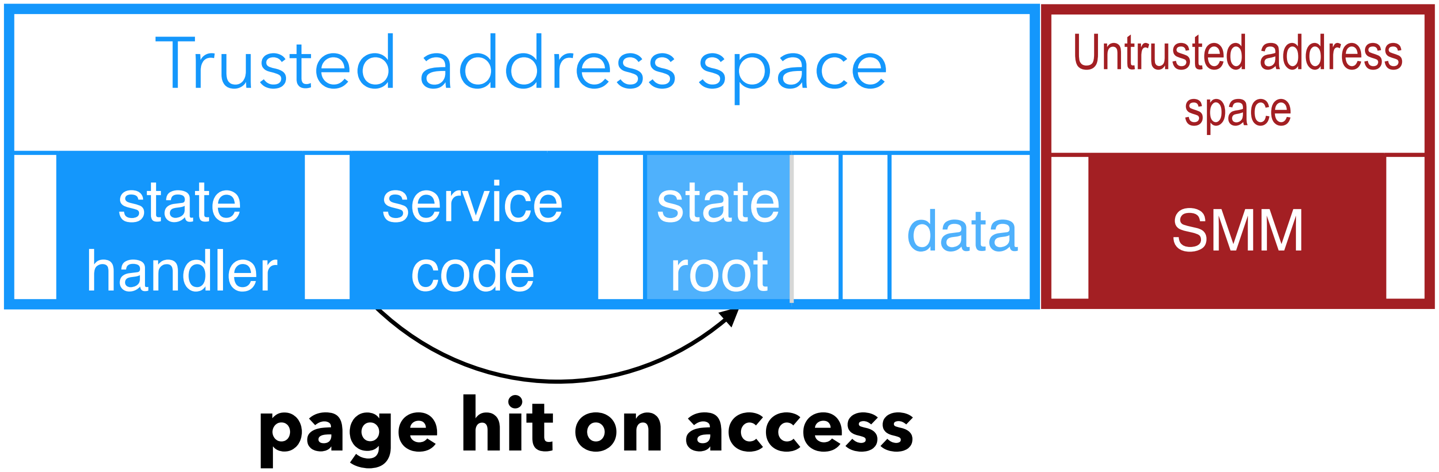

- Service code execution begins
- Service accesses data in memory
- Data retrieval is fast if data is already available

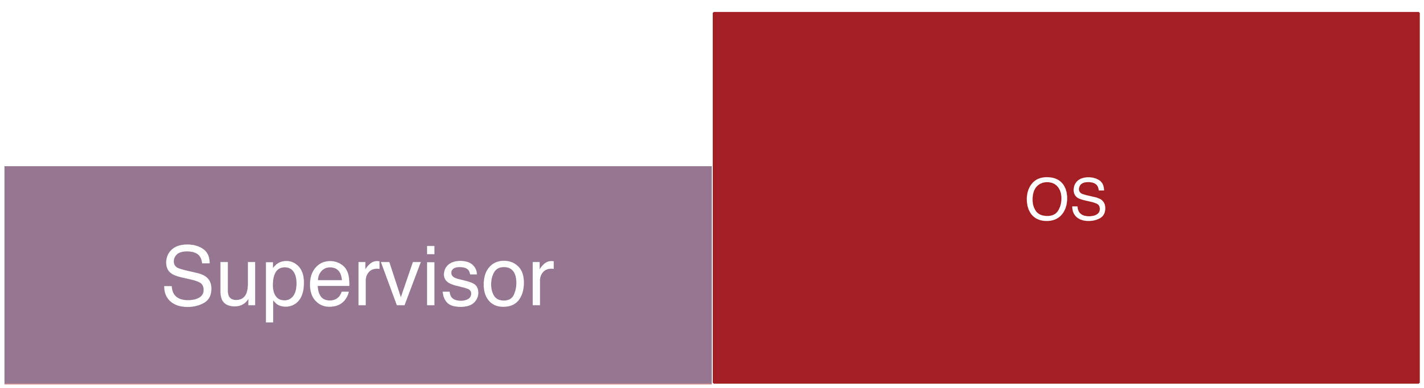

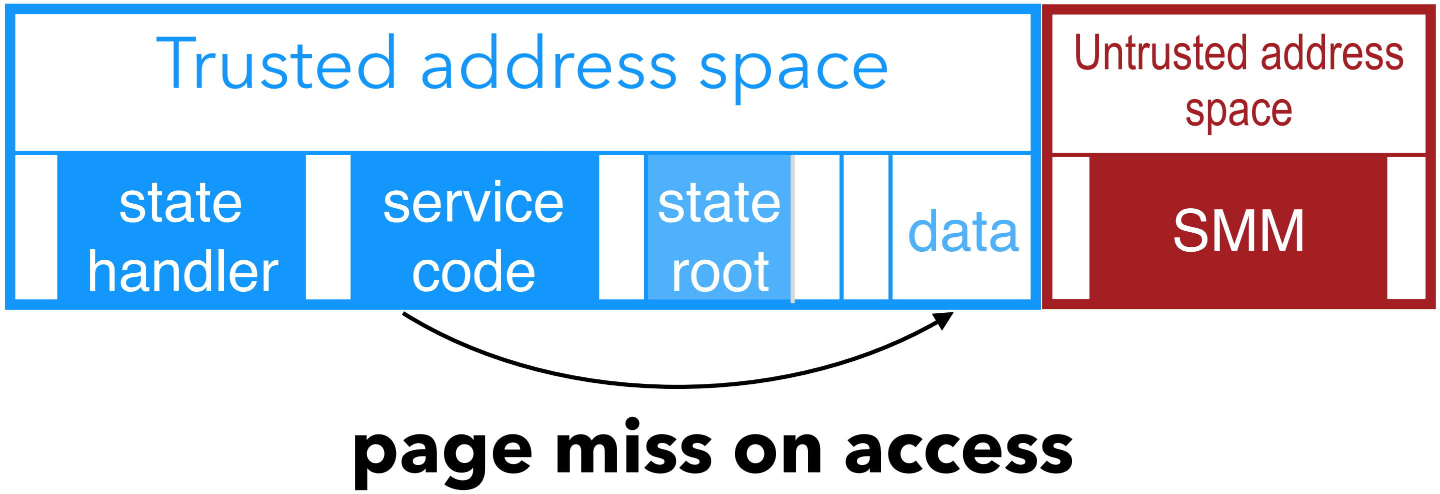

- Service code may access data on missing pages

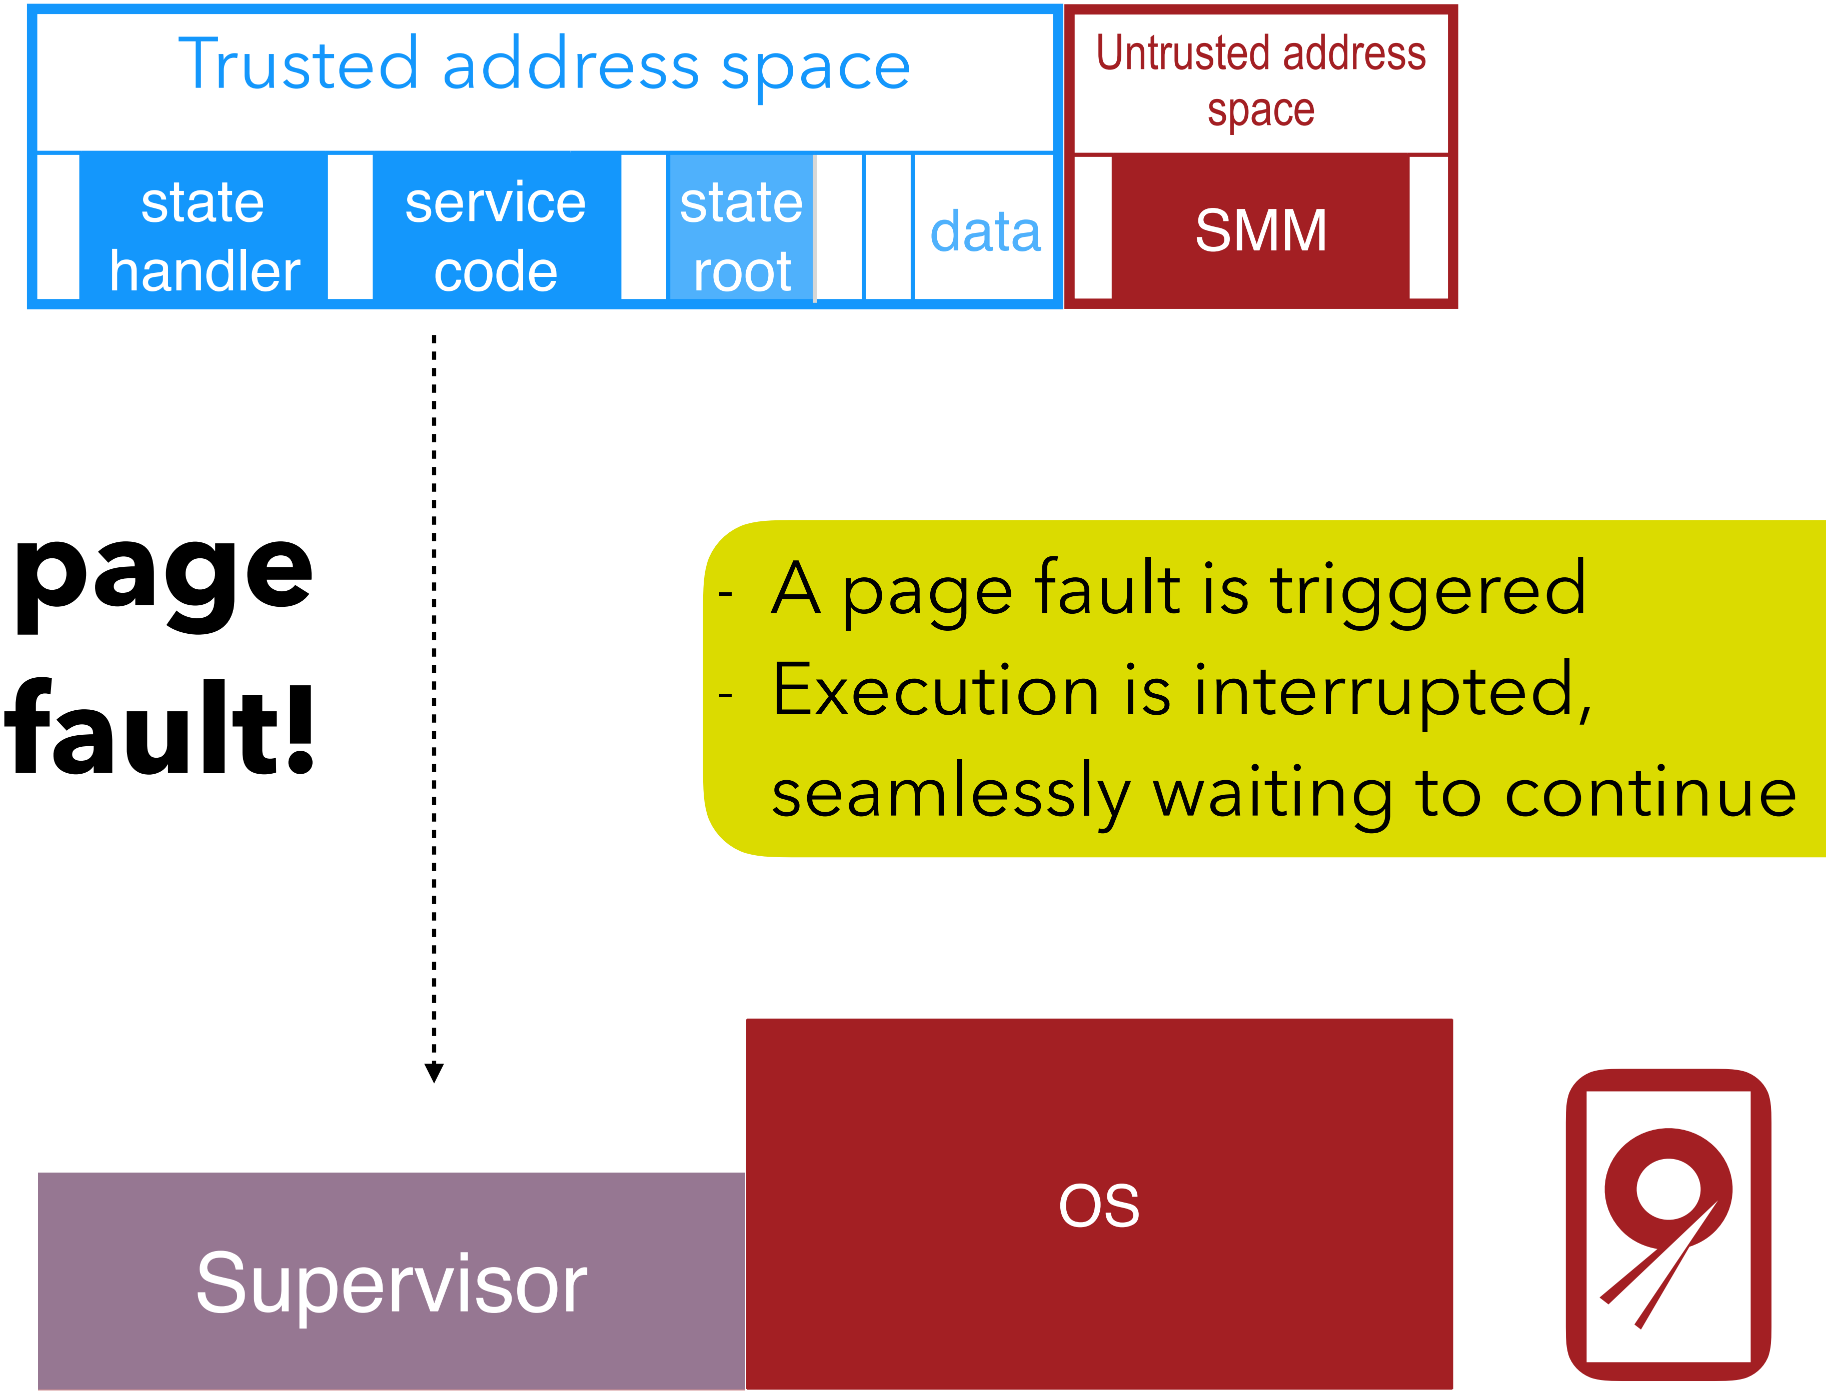
Trusted address space
state handler
service code
state root
data
Untrusted address space
SMM
page fault!
- A page fault is triggered
- Execution is interrupted, seamlessly waiting to continue
Supervisor
OS

4

- Lazy loading from memory & disk

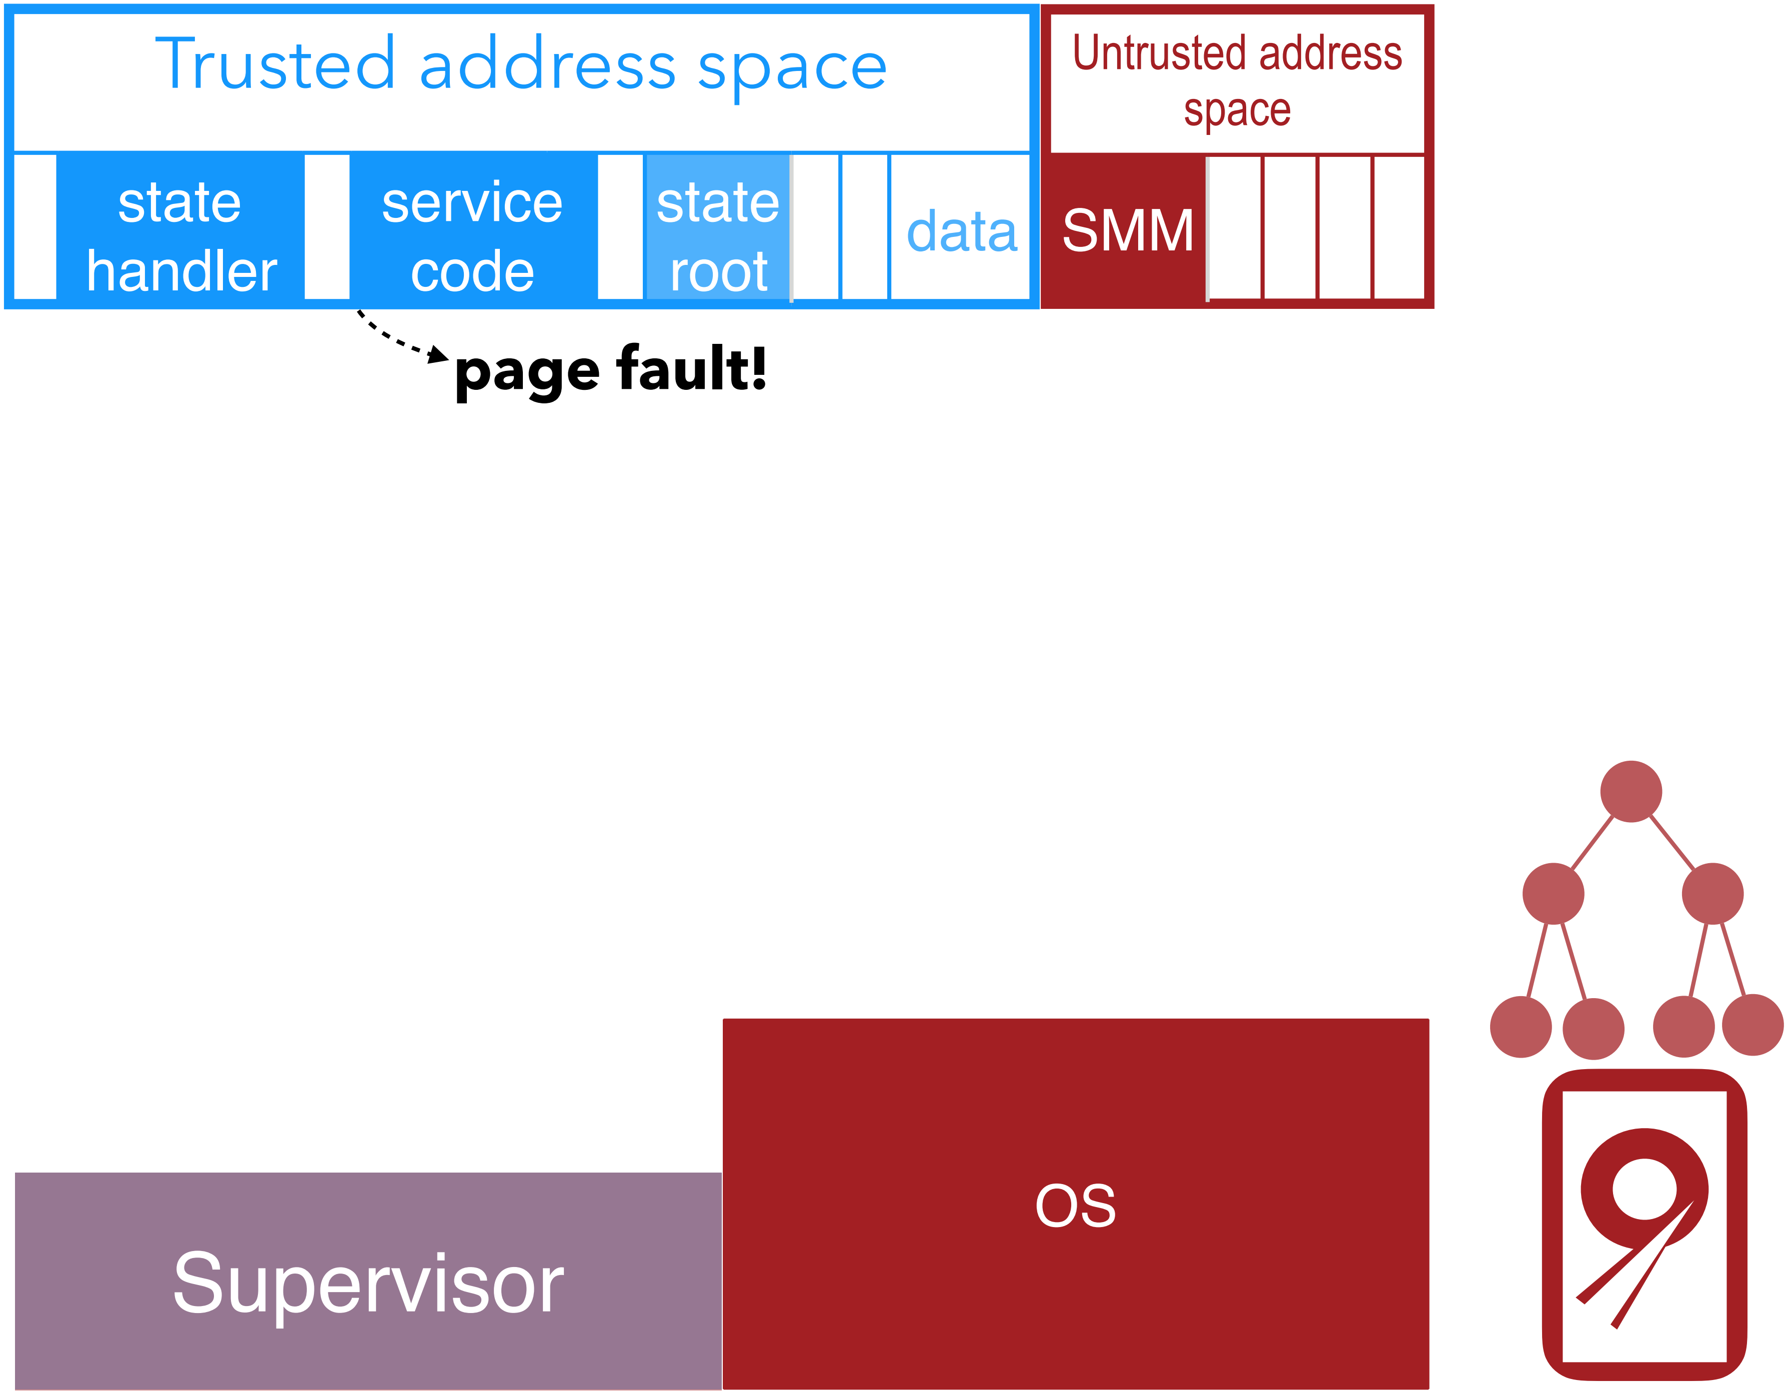
Trusted address space
state handler
service code
state root
data
Untrusted address space
SMM
page fault!
Supervisor
OS

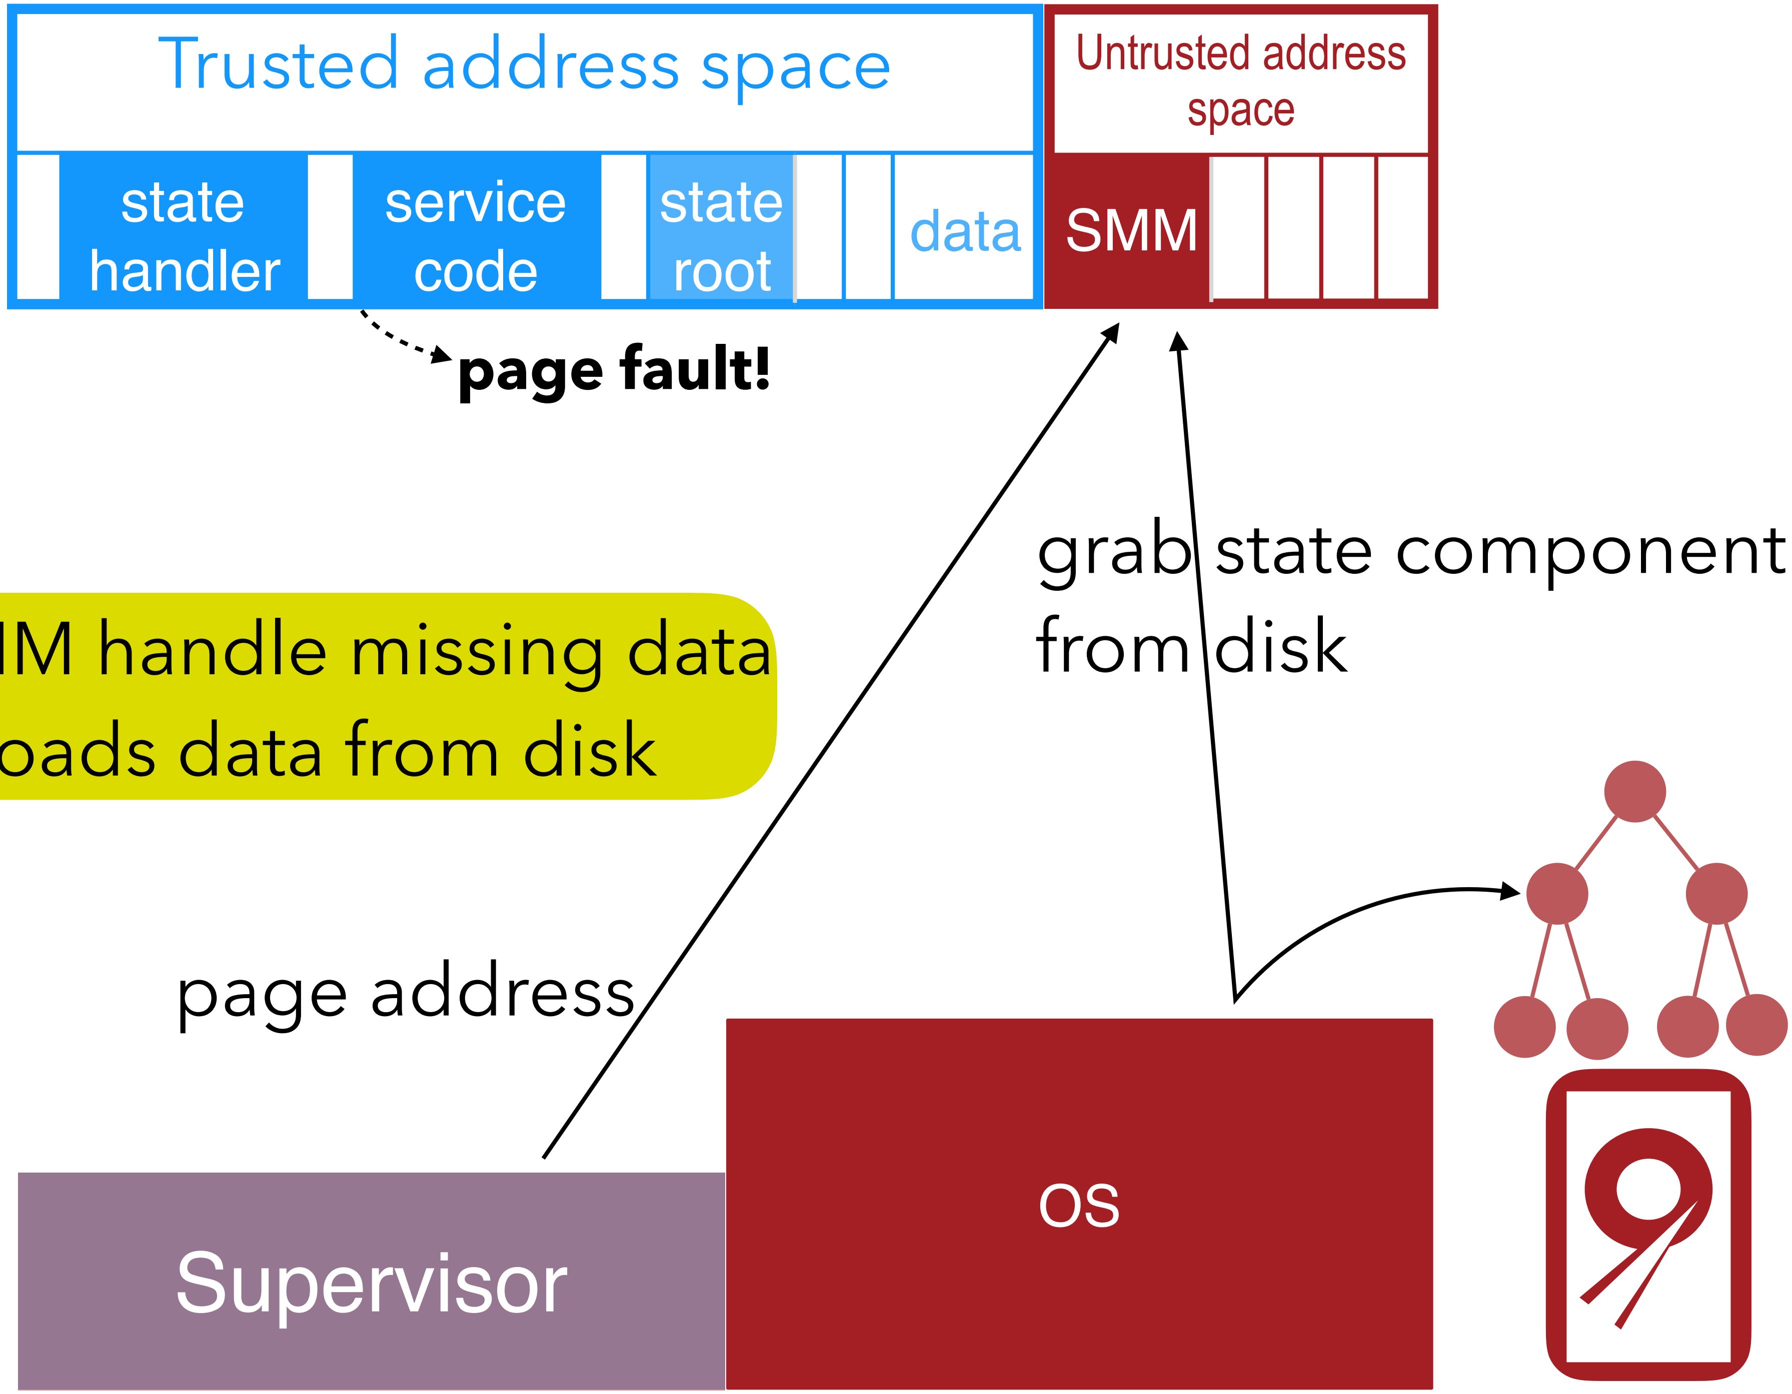

- Let SMM handle missing data
- SMM loads data from disk

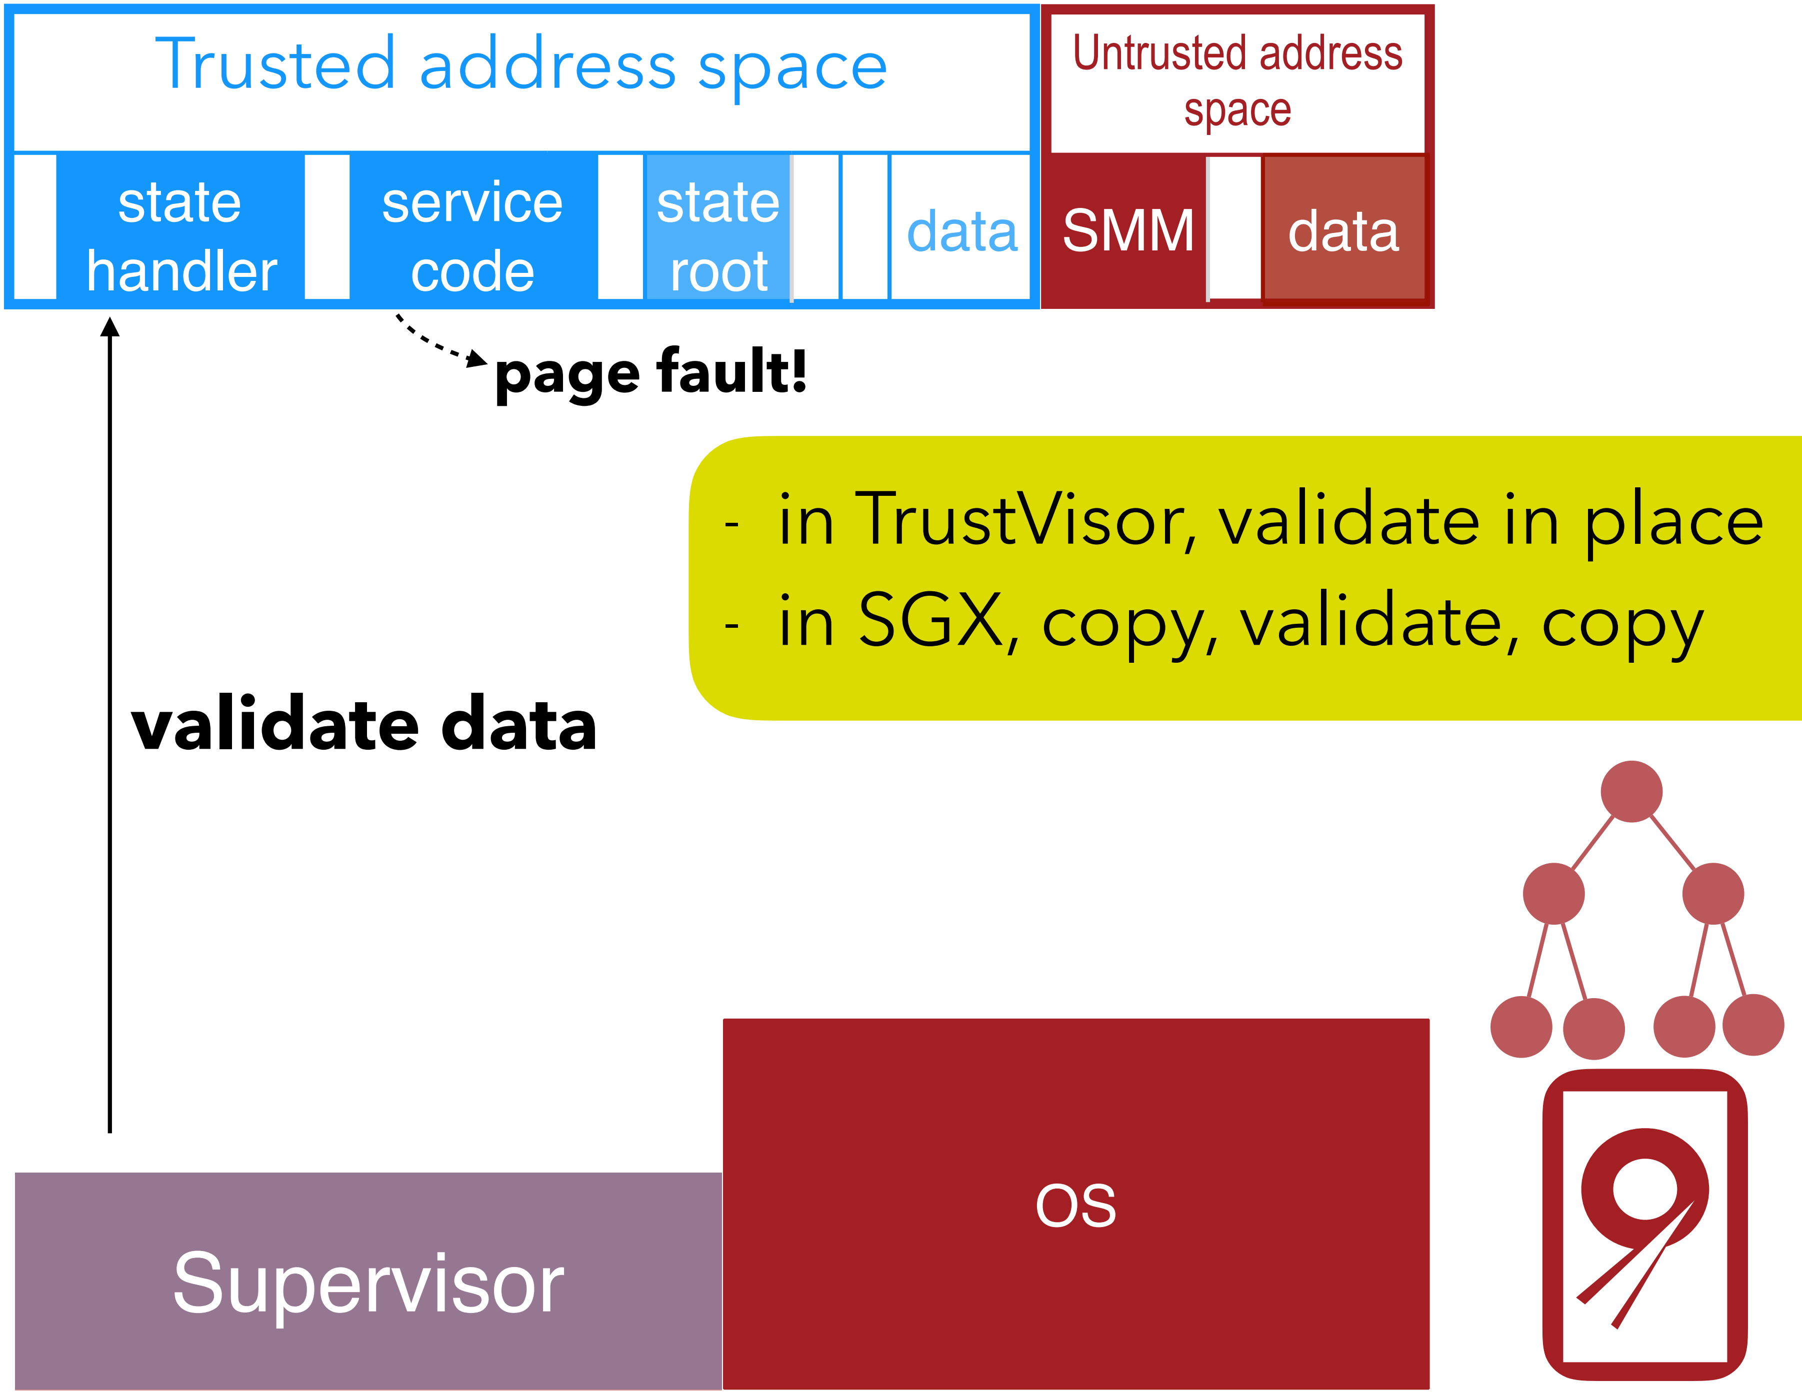
Trusted address space
state handler
service code
state root
data
Untrusted address space
SMM
data
page fault!
validate data
- in TrustVisor, validate in place
- in SGX, copy, validate, copy
Supervisor
OS

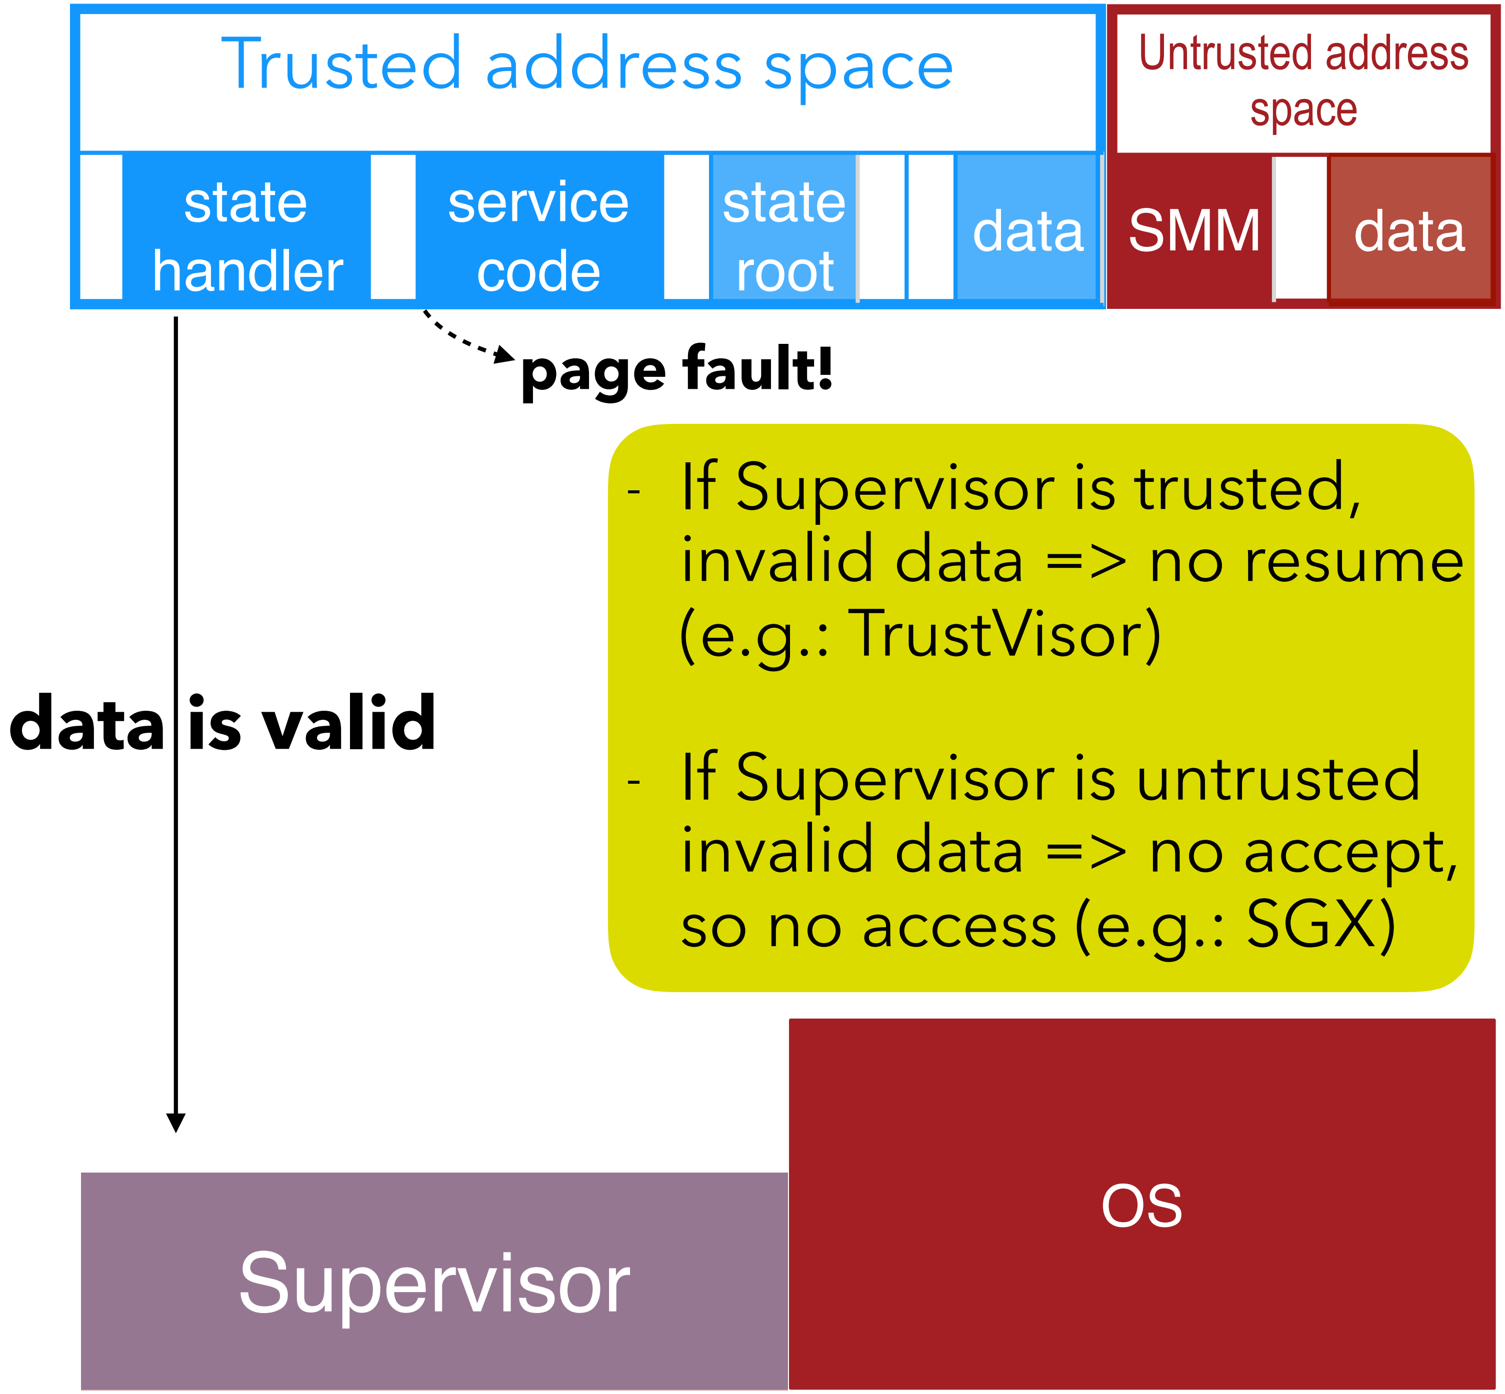
Trusted address space
Untrusted address space
state handler
service code
state root
data
SMM
data
page fault!
data is valid
- If Supervisor is trusted, invalid data => no resume (e.g.: TrustVisor)
- If Supervisor is untrusted invalid data => no accept, so no access (e.g.: SGX)
Supervisor
OS

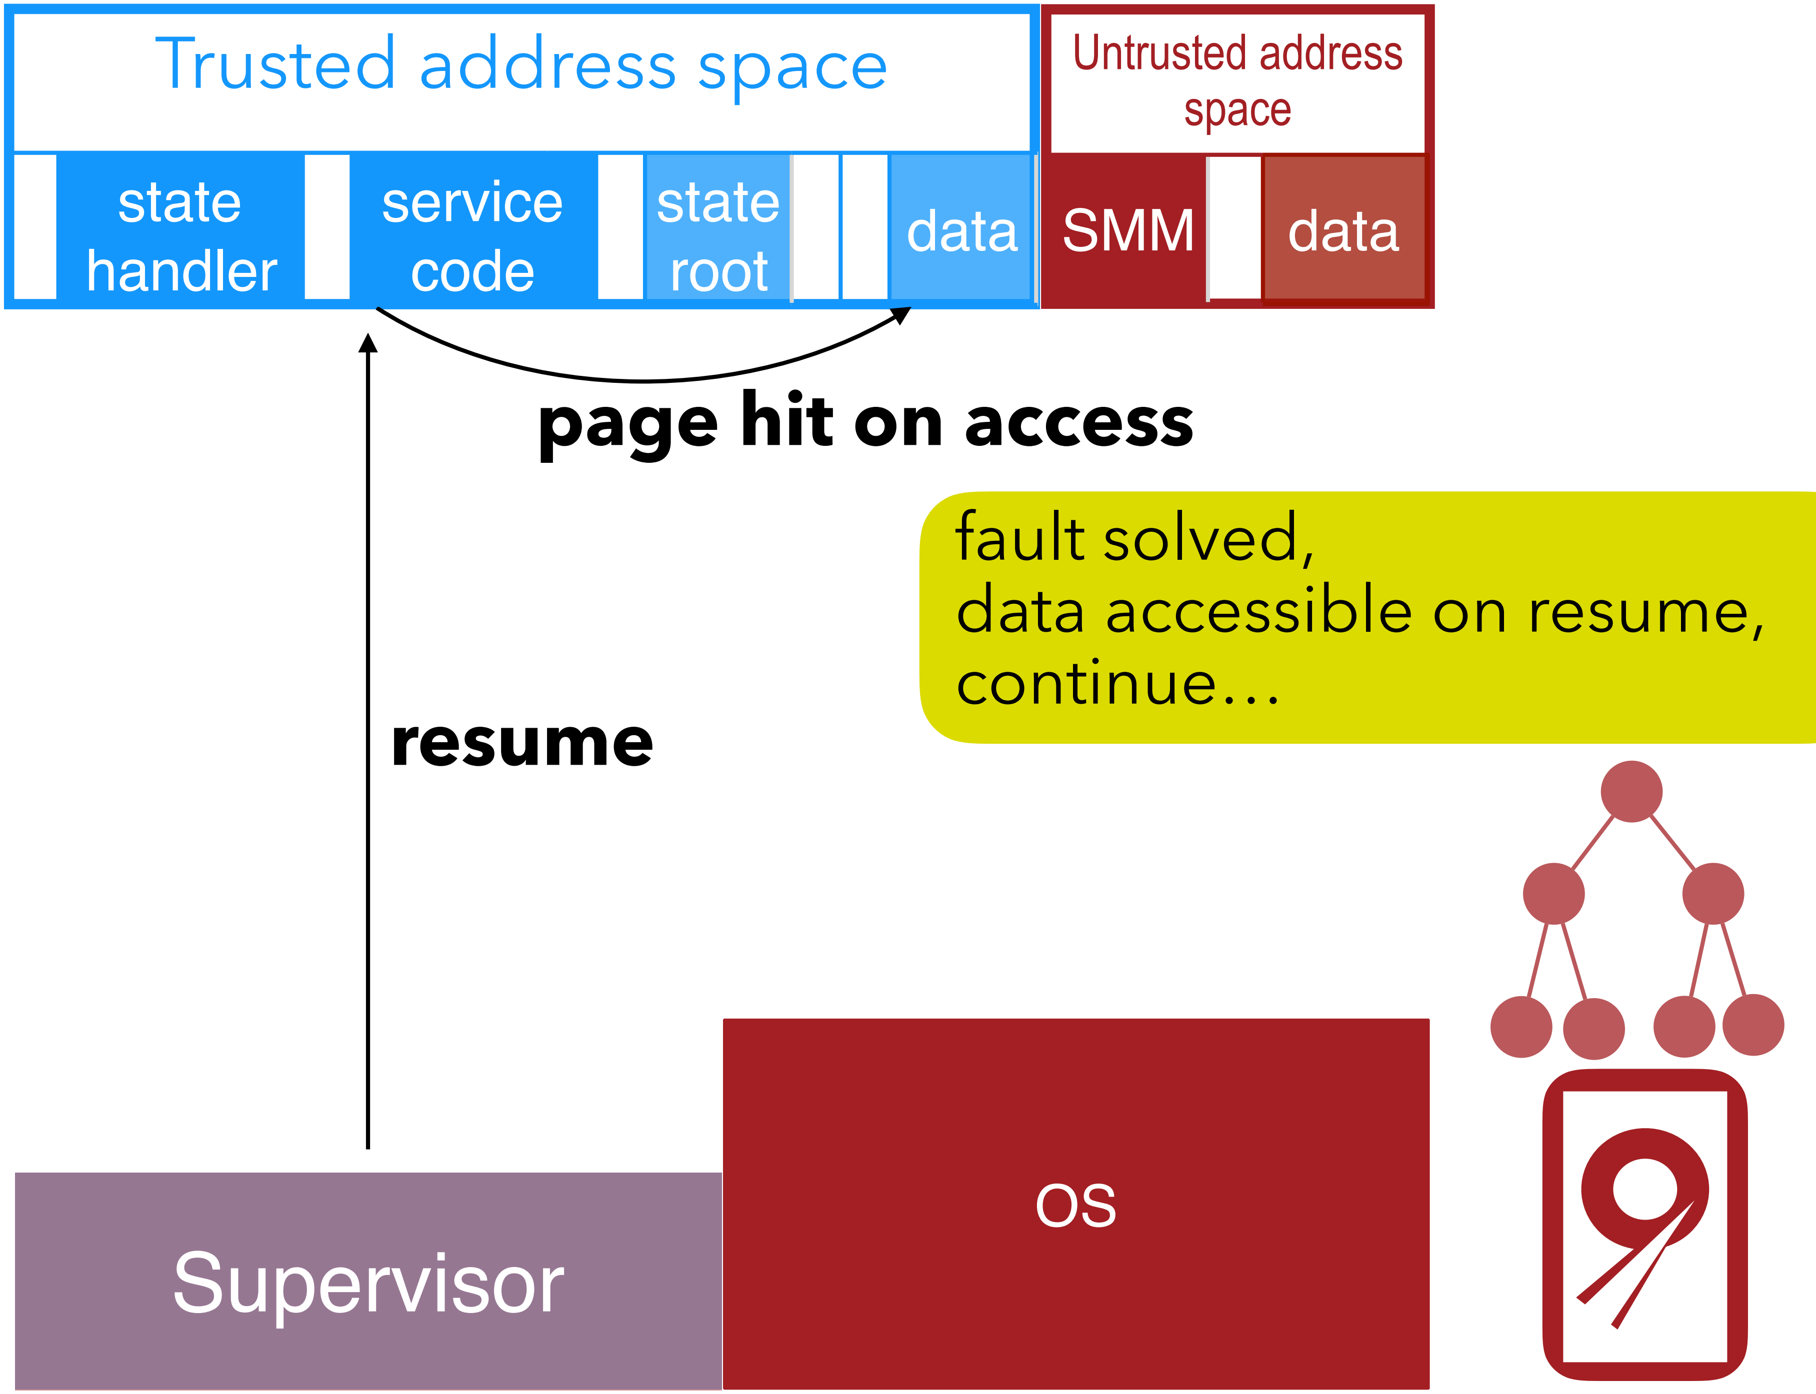

Trusted address space
state handler
service code
state root
data
Untrusted address space
SMM
data
page hit on access
fault solved,
data accessible on resume,
continue…
resume
Supervisor
OS

- HW-based attestation of code identity, including input request, state root, output reply, nonce
- Client checks validity of attestation and intended identities/hashes

5

- ~~**Execution verification**~~

Outline

- **Evaluation**

TCB size

SGX-based

TPM/TXT based

	VC^3	Haven	LaSt^{GT}		
			hypervisor	library	SQLite (example)
KSLoC (lines of code x 1000)	9.2	O(10^3)	17	7.7	92.6

library is small compared to real service

Comparison

XMHF-TrustVisor vs. LaSt^{GT}

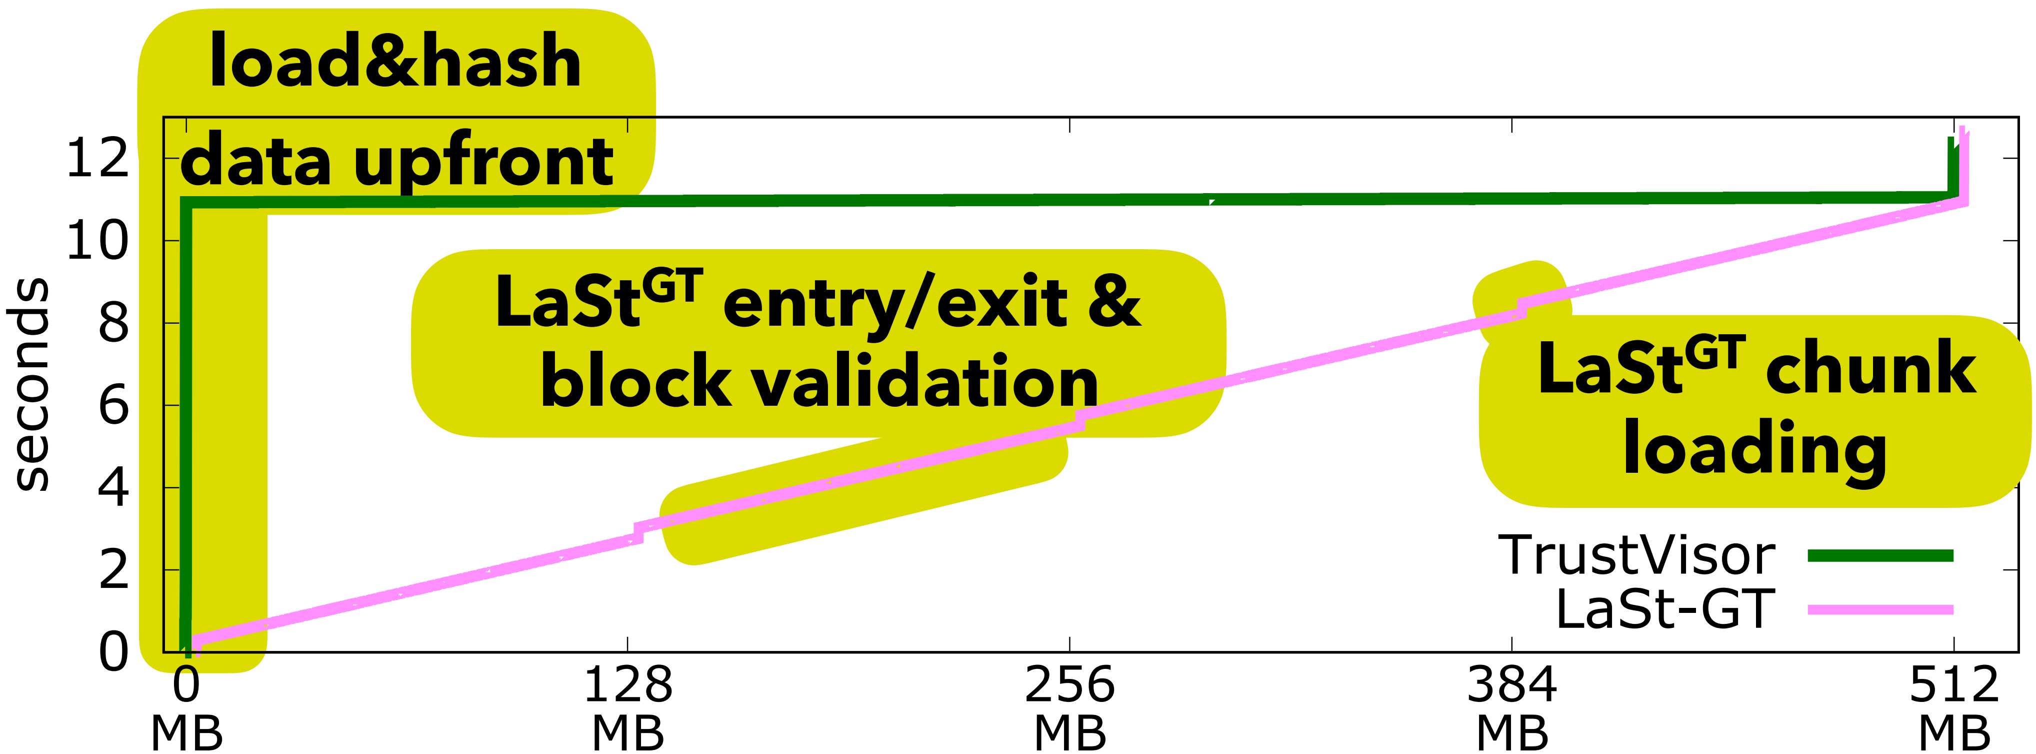

LaSt^{GT} is Incremental, Faster & Scalable

SQLite on LaSt^{GT}

- First large-scale experiment on hypervisor
- Data I/O can be optimized through state hierarchy
- SGX expected to improve substantially

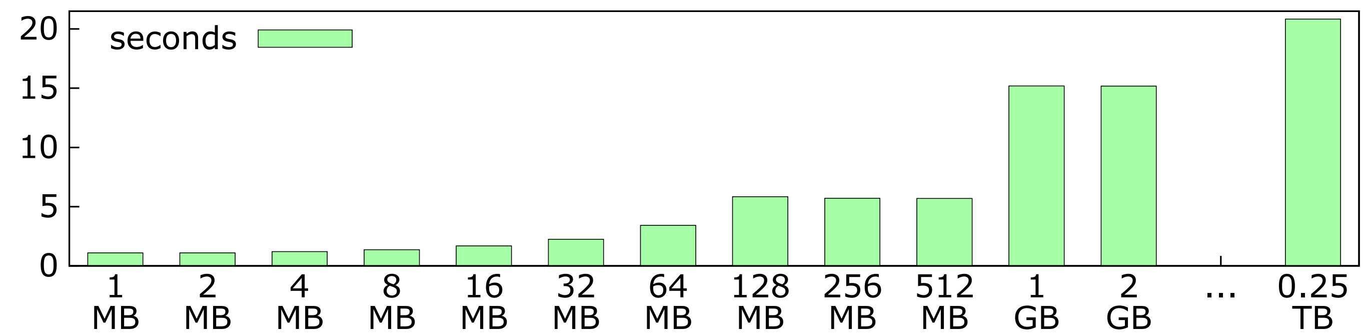

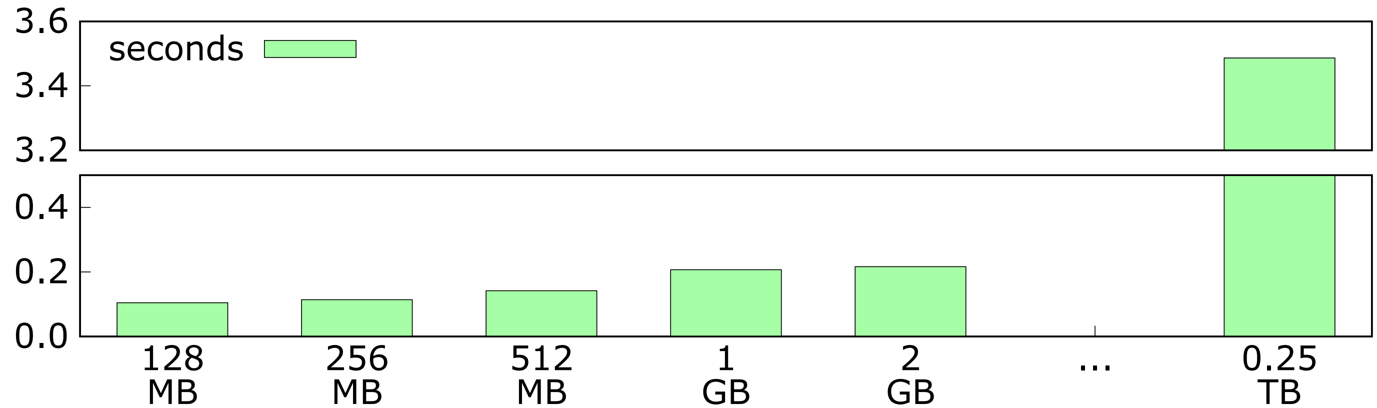

Conclusions

- **Security for large-scale data processing can be guaranteed with a small TCB**
- **Virtual memory-based data handling**
  => zero interface
- **No change to source code**
  => easy integration
- **One design can fit diverse HW & SW**

ad maiora.

Secure Tera-scale Data Crunching with a Small TCB

Bruno Vavala[1,2], Nuno Neves[1], Peter Steenkiste[2]

[1]*LaSIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal* [2]*CSD, Carnegie Mellon University, U.S.*

***Abstract*—Outsourcing services to third-party providers comes with a high security cost—to fully trust the providers. Using trusted hardware can help, but current trusted execution environments do not adequately support services that process very large scale datasets. We present LaSt^{GT}, a system that bridges this gap by supporting the execution of self-contained services over a large state, with a small and generic trusted computing base (TCB). LaSt^{GT} uses widely deployed trusted hardware to guarantee integrity and verifiability of the execution on a remote platform, and it securely supplies data to the service through simple techniques based on virtual memory. As a result, LaSt^{GT} is general and applicable to many scenarios such as computational genomics and databases, as we show in our experimental evaluation based on an implementation of LaSt^{GT} on a secure hypervisor. We also describe a possible implementation on Intel SGX.**

support the execution of either small pieces of code and data [10], or large code bases [11], or specific software like database engines [12] or MapReduce applications [13]. Recent work [14] has shown how to support unmodified services. However, since "the interface between modern applications and operating systems is so complex" [30], it relies on a considerable TCB that includes a library OS. In addition, the above systems are specific for TPMs [10], [15], secure coprocessors [12], or Intel SGX [13]. Hence, porting them to alternative architectures (e.g., the upcoming AMD Secure Memory Encryption and Secure Encrypted Virtualization [36], [37]) requires significant effort. Clearly, it is desirable to design a generic system "not relying on idiosyncratic features of the hardware" [16].

We present LaSt^{GT}, a system that can handle a LArge STate on a Generic Trusted component with a small TCB.

(blank)

“Never trust a computer you can’t throw out the window.”

Steve Wozniak

“No computer system can be absolutely secure.”

(excerpt from)

Intel’s Legal Desclaimer