

Eleos: Exit-Less OS Services for SGX Enclaves

Meni Orenbach

Marina Minkin

Pavel Lifshits

Mark Silberstein

Accelerated Computing Systems Lab

Belgrade, Serbia
April 23-26, 2017



What do we do?

Improve performance: I/O intensive & memory demanding SGX enclaves

Why?

Cost of SGX execution for these applications is high

How?

In-enclave System Calls & User Managed Virtual Memory

Results

Eleos vs vanilla SGX

2x ↑ Throughput: memcached & face verification servers

Even for 5x ↑ available enclave memory

Available for Linux, Windows*

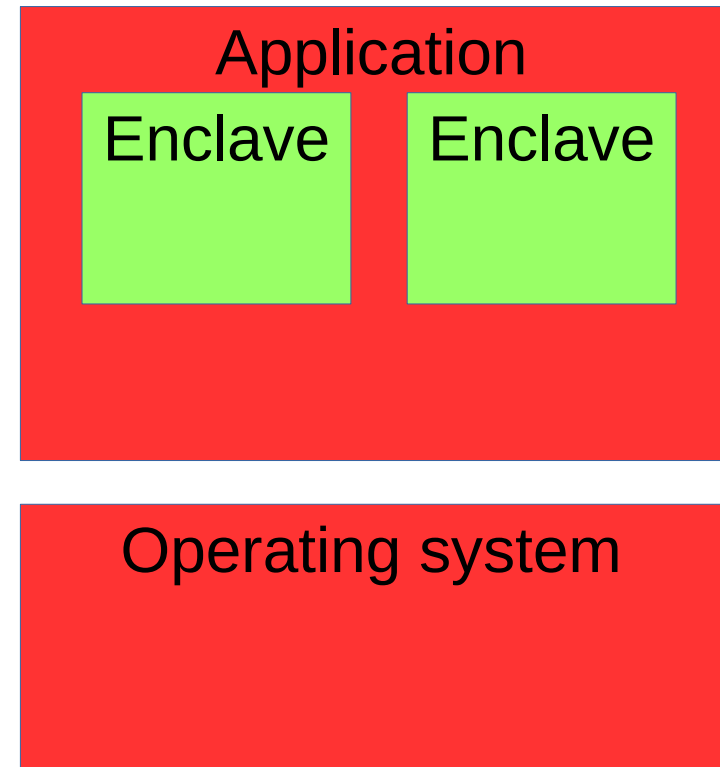
(*) Without Eleos, these applications crash in Windows enclaves

- **Background**
- Motivation
- Overhead analysis
- Eleos design
- Evaluation



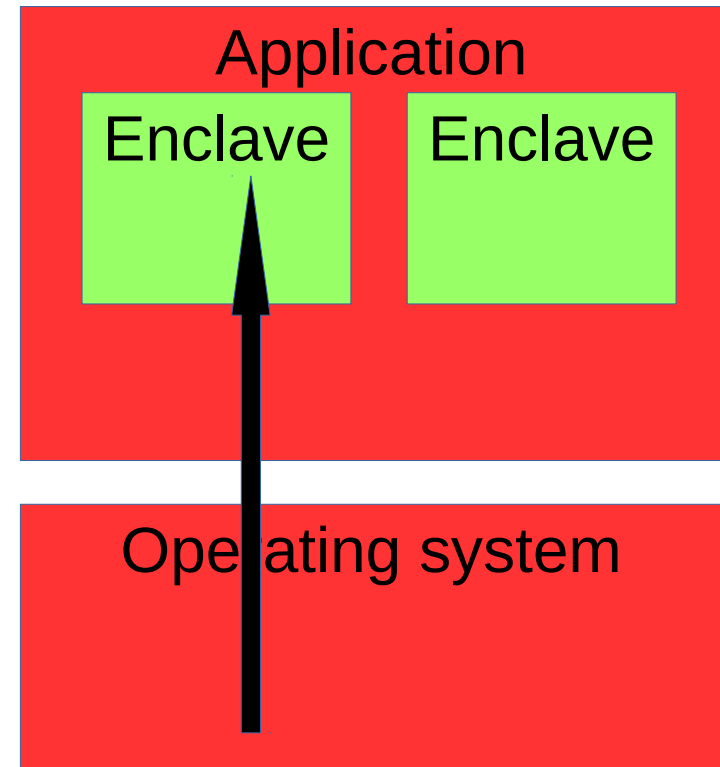
SGX enclaves are already here!

- Secured execution environment
- Reversed sandbox
- Small TCB
- Private code & data
 - Confidentiality
 - Integrity
 - Freshness
- Only CPU is trusted



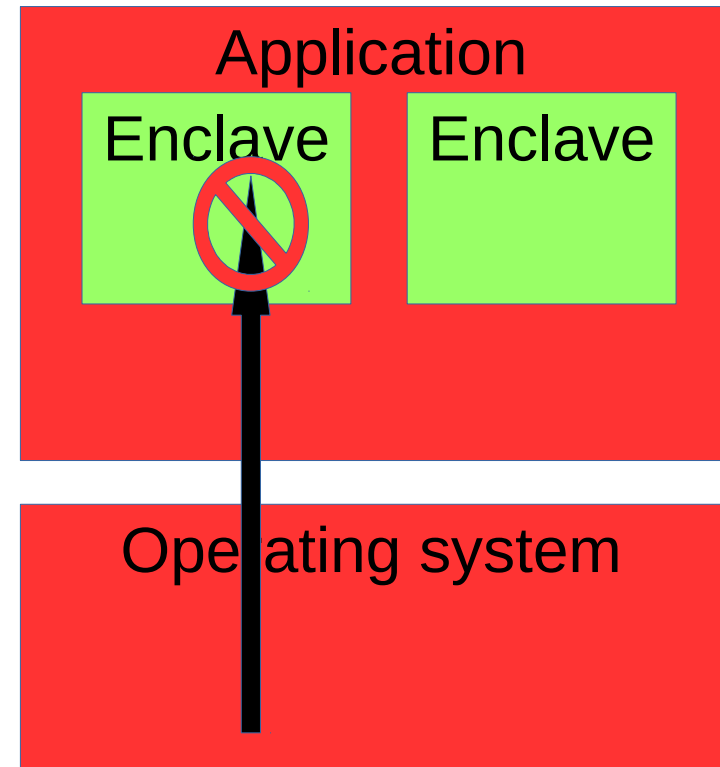
SGX enclaves are already here!

- Secured execution environment
- Reversed sandbox
- Small TCB
- Private code & data
 - Confidentiality
 - Integrity
 - Freshness
- Only CPU is trusted



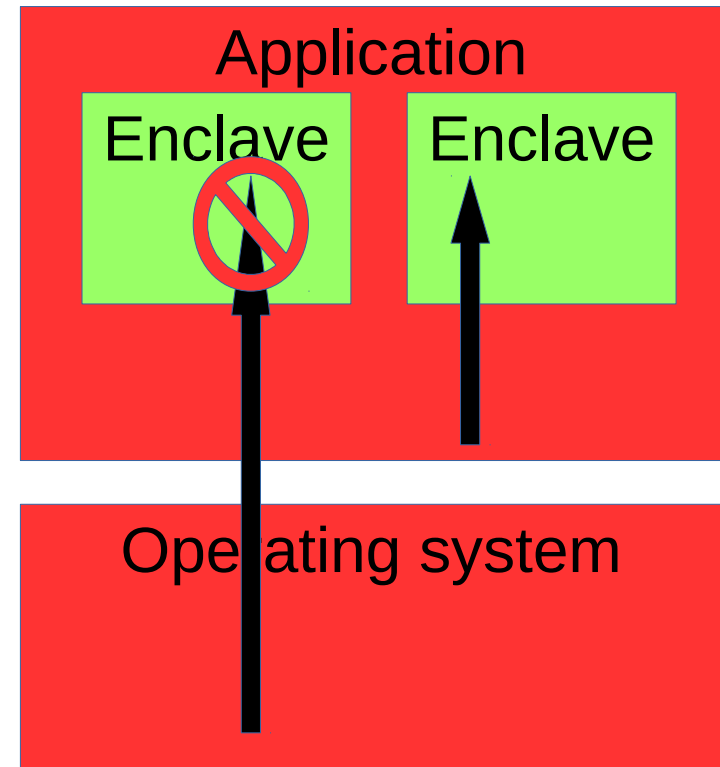
SGX enclaves are already here!

- Secured execution environment
- Reversed sandbox
- Small TCB
- Private code & data
 - Confidentiality
 - Integrity
 - Freshness
- Only CPU is trusted



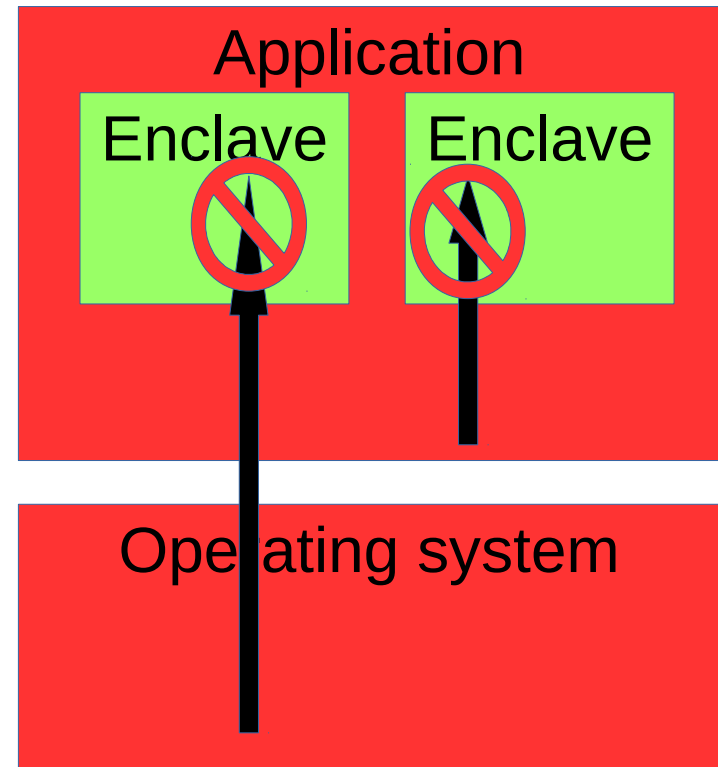
SGX enclaves are already here!

- Secured execution environment
- Reversed sandbox
- Small TCB
- Private code & data
 - Confidentiality
 - Integrity
 - Freshness
- Only CPU is trusted



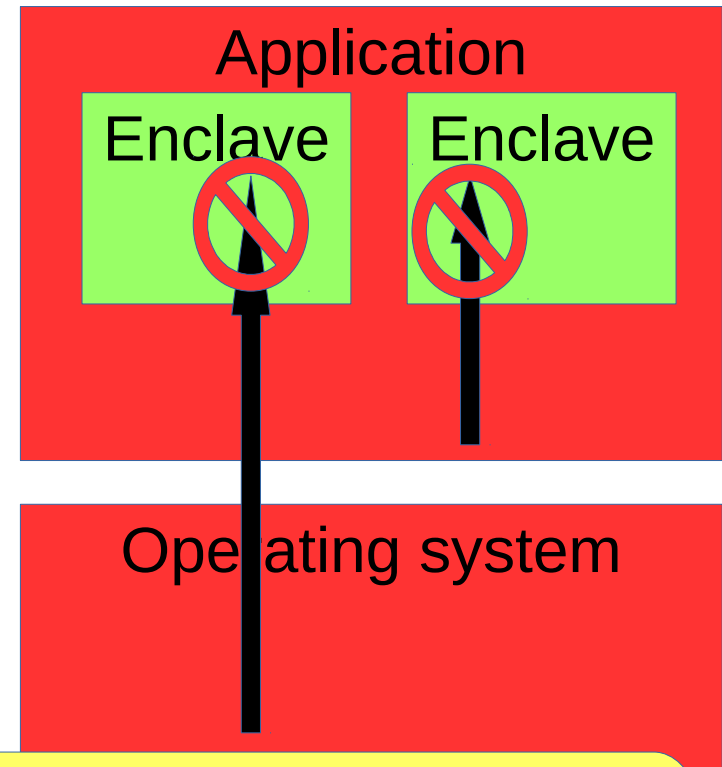
SGX enclaves are already here!

- Secured execution environment
- Reversed sandbox
- Small TCB
- Private code & data
 - Confidentiality
 - Integrity
 - Freshness
- Only CPU is trusted



SGX enclaves are already here!

- Secured execution environment
- Reversed sandbox
- Small TCB
- Private code & data
 - Confidentiality
 - Integrity
 - Freshness



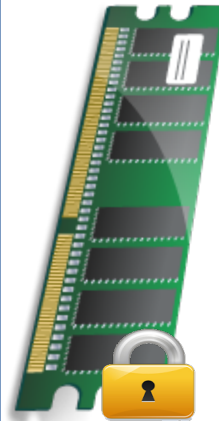
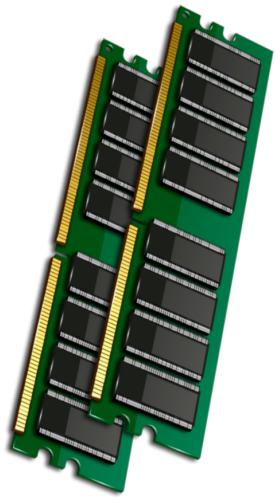
• O

Lets look at
How to secure server applications with enclaves

Background: Lifetime of a secured server

Untrusted (Host & OS)

Trusted (Enclave)

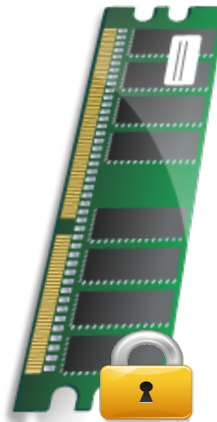


Background: Lifetime of a secured server

Untrusted (Host & OS)

Trusted (Enclave)

Untrusted memory
Unsecured access



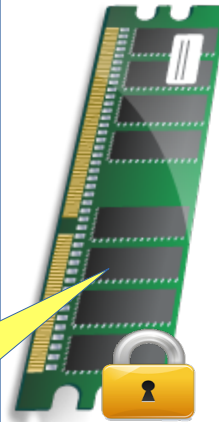
Background: Lifetime of a secured server

Untrusted (Host & OS)

Trusted (Enclave)

Untrusted memory
Unsecured access

Dedicated SGX mem
Limited to: 128 MB
Secured access



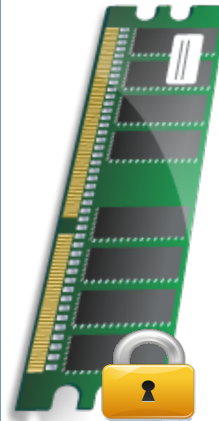
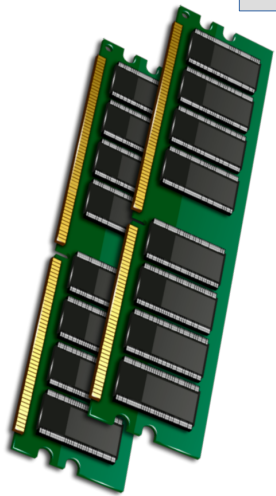
Background: Lifetime of a secured server

Untrusted (Host & OS)

Trusted (Enclave)

Host
app

Wait for network
requests



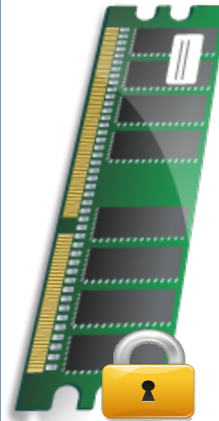
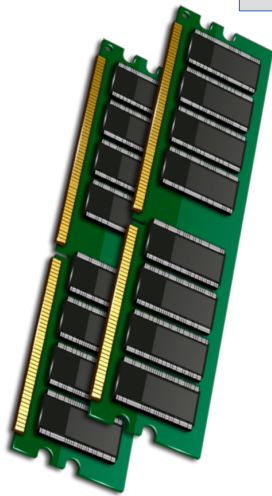
Background: Lifetime of a secured server

Untrusted (Host & OS)

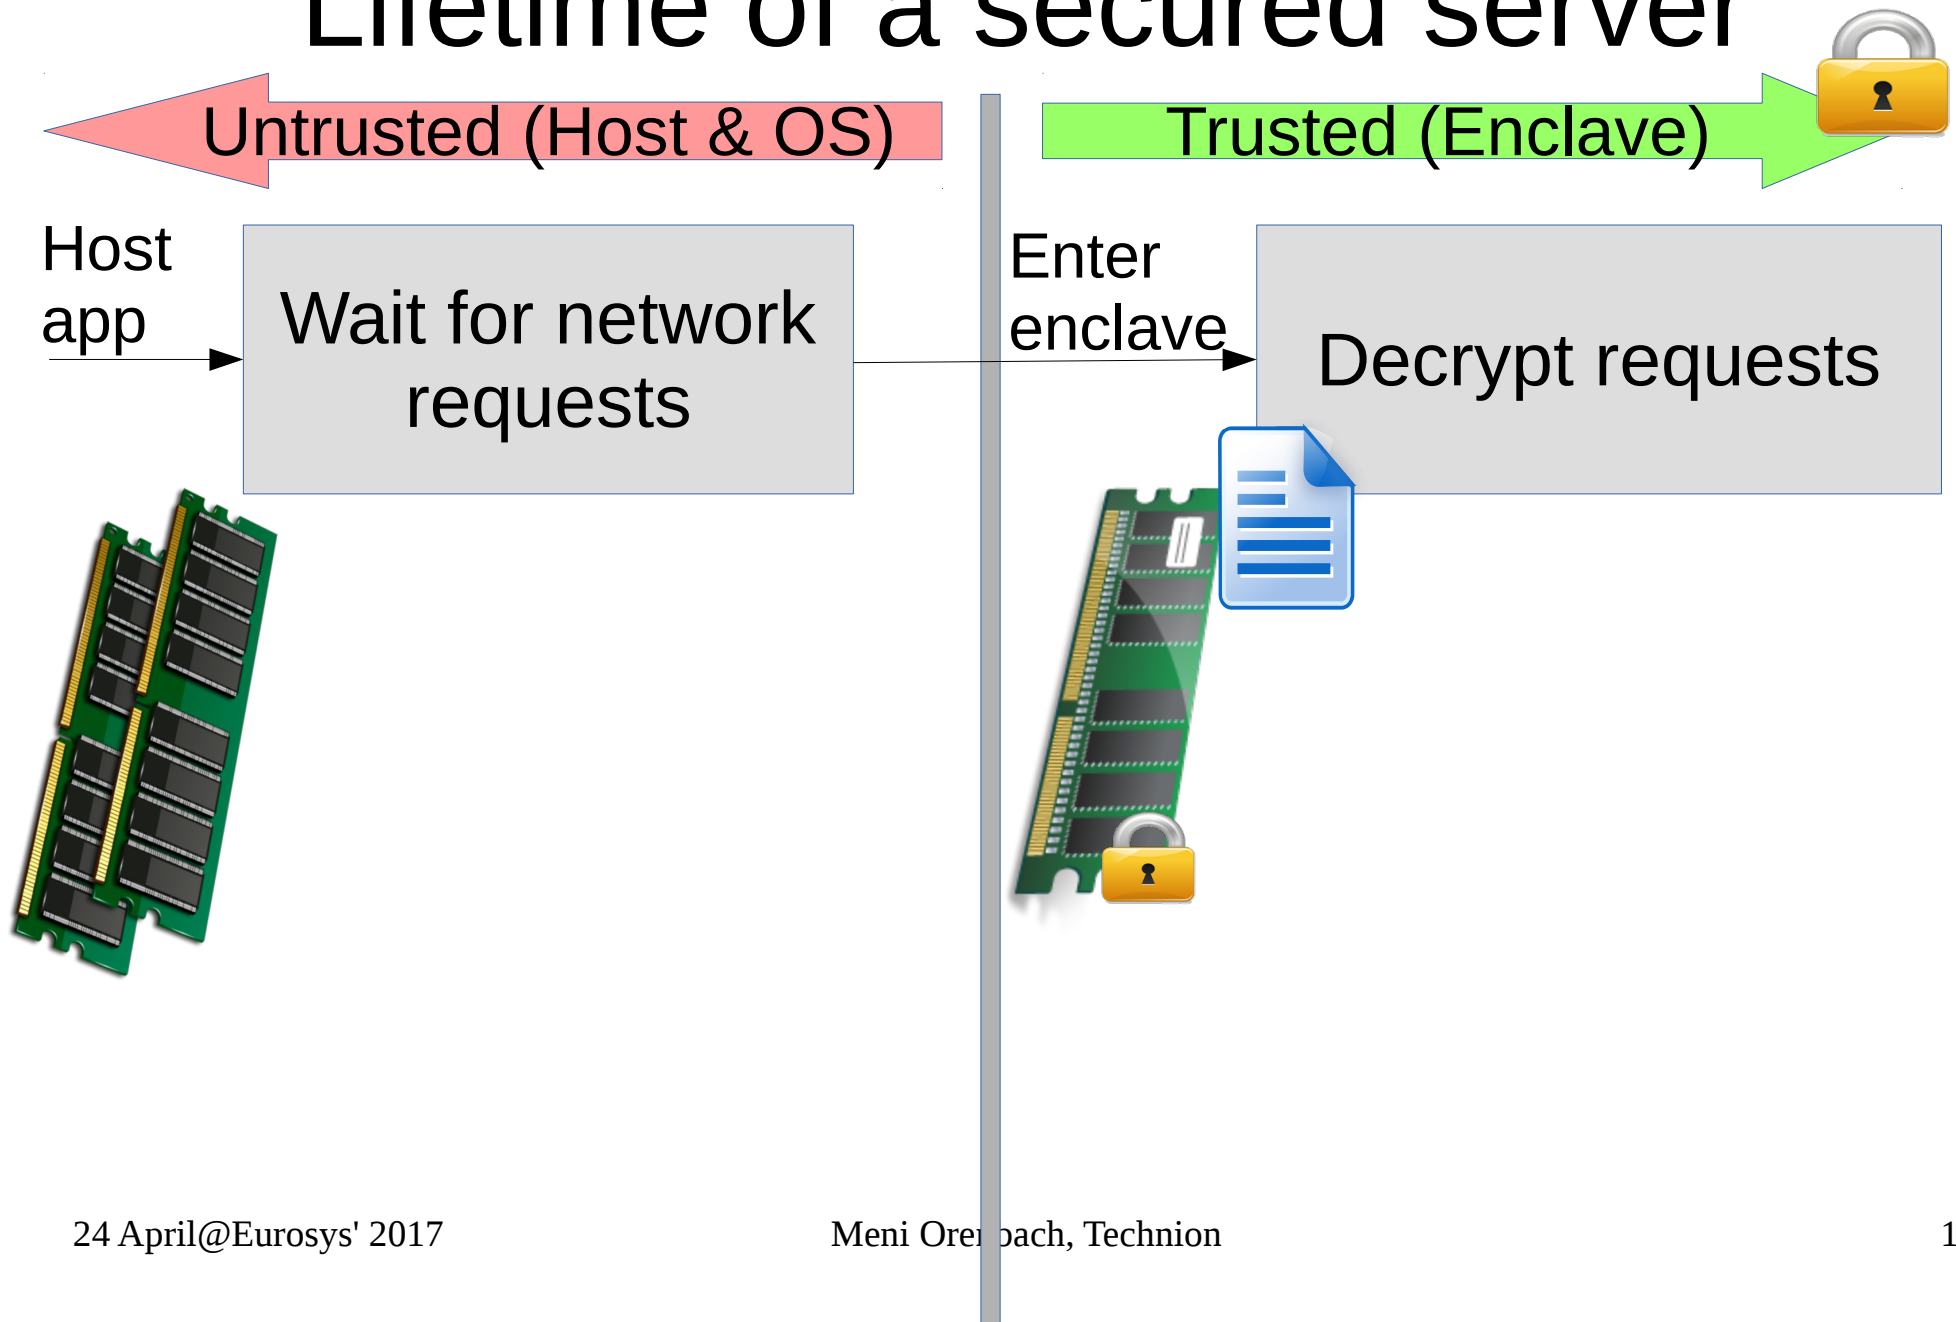
Trusted (Enclave)

Host
app

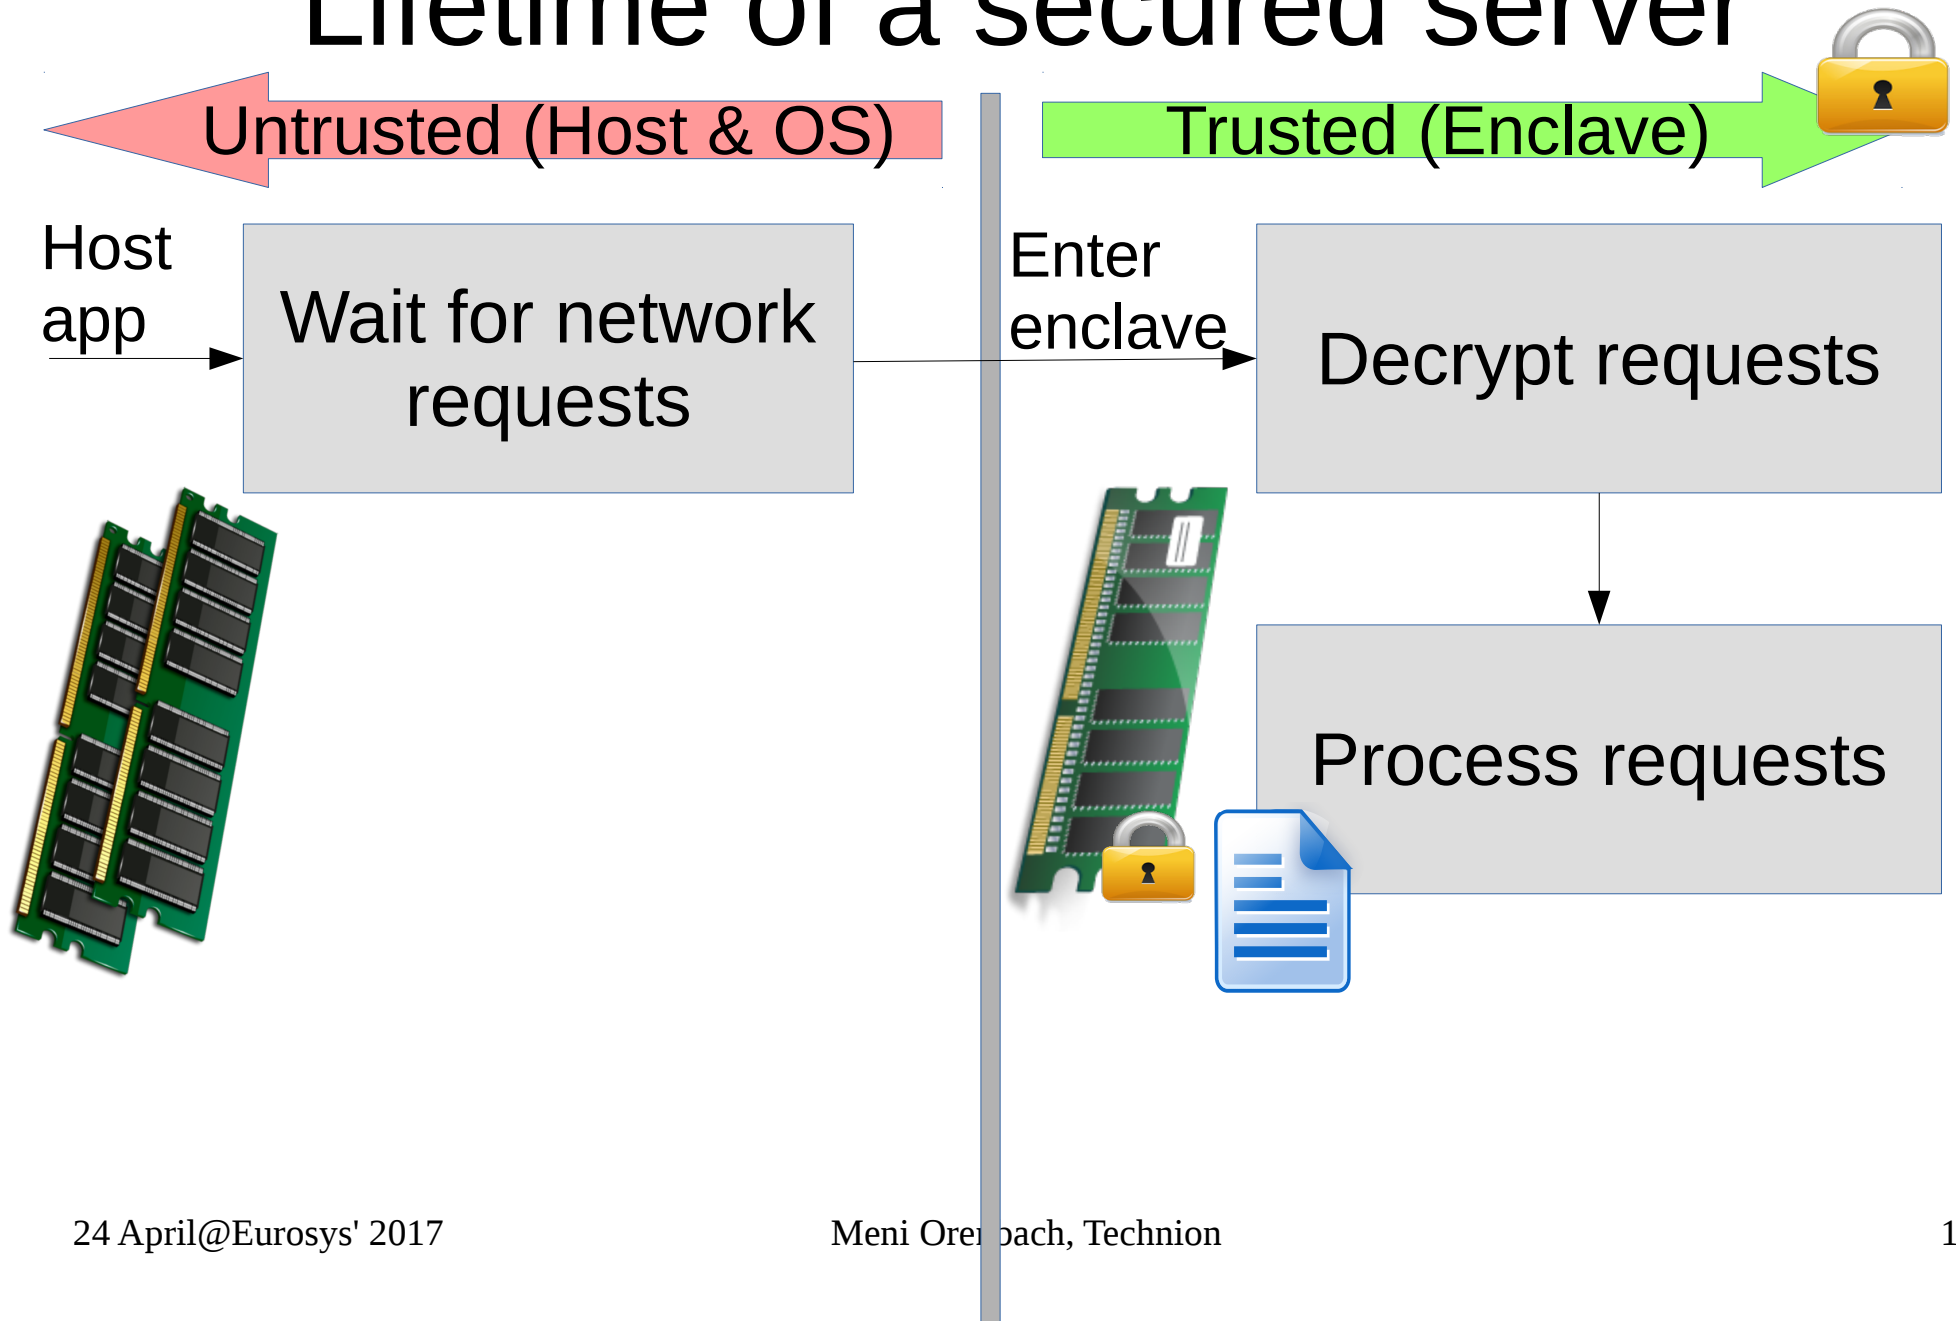
Wait for network
requests



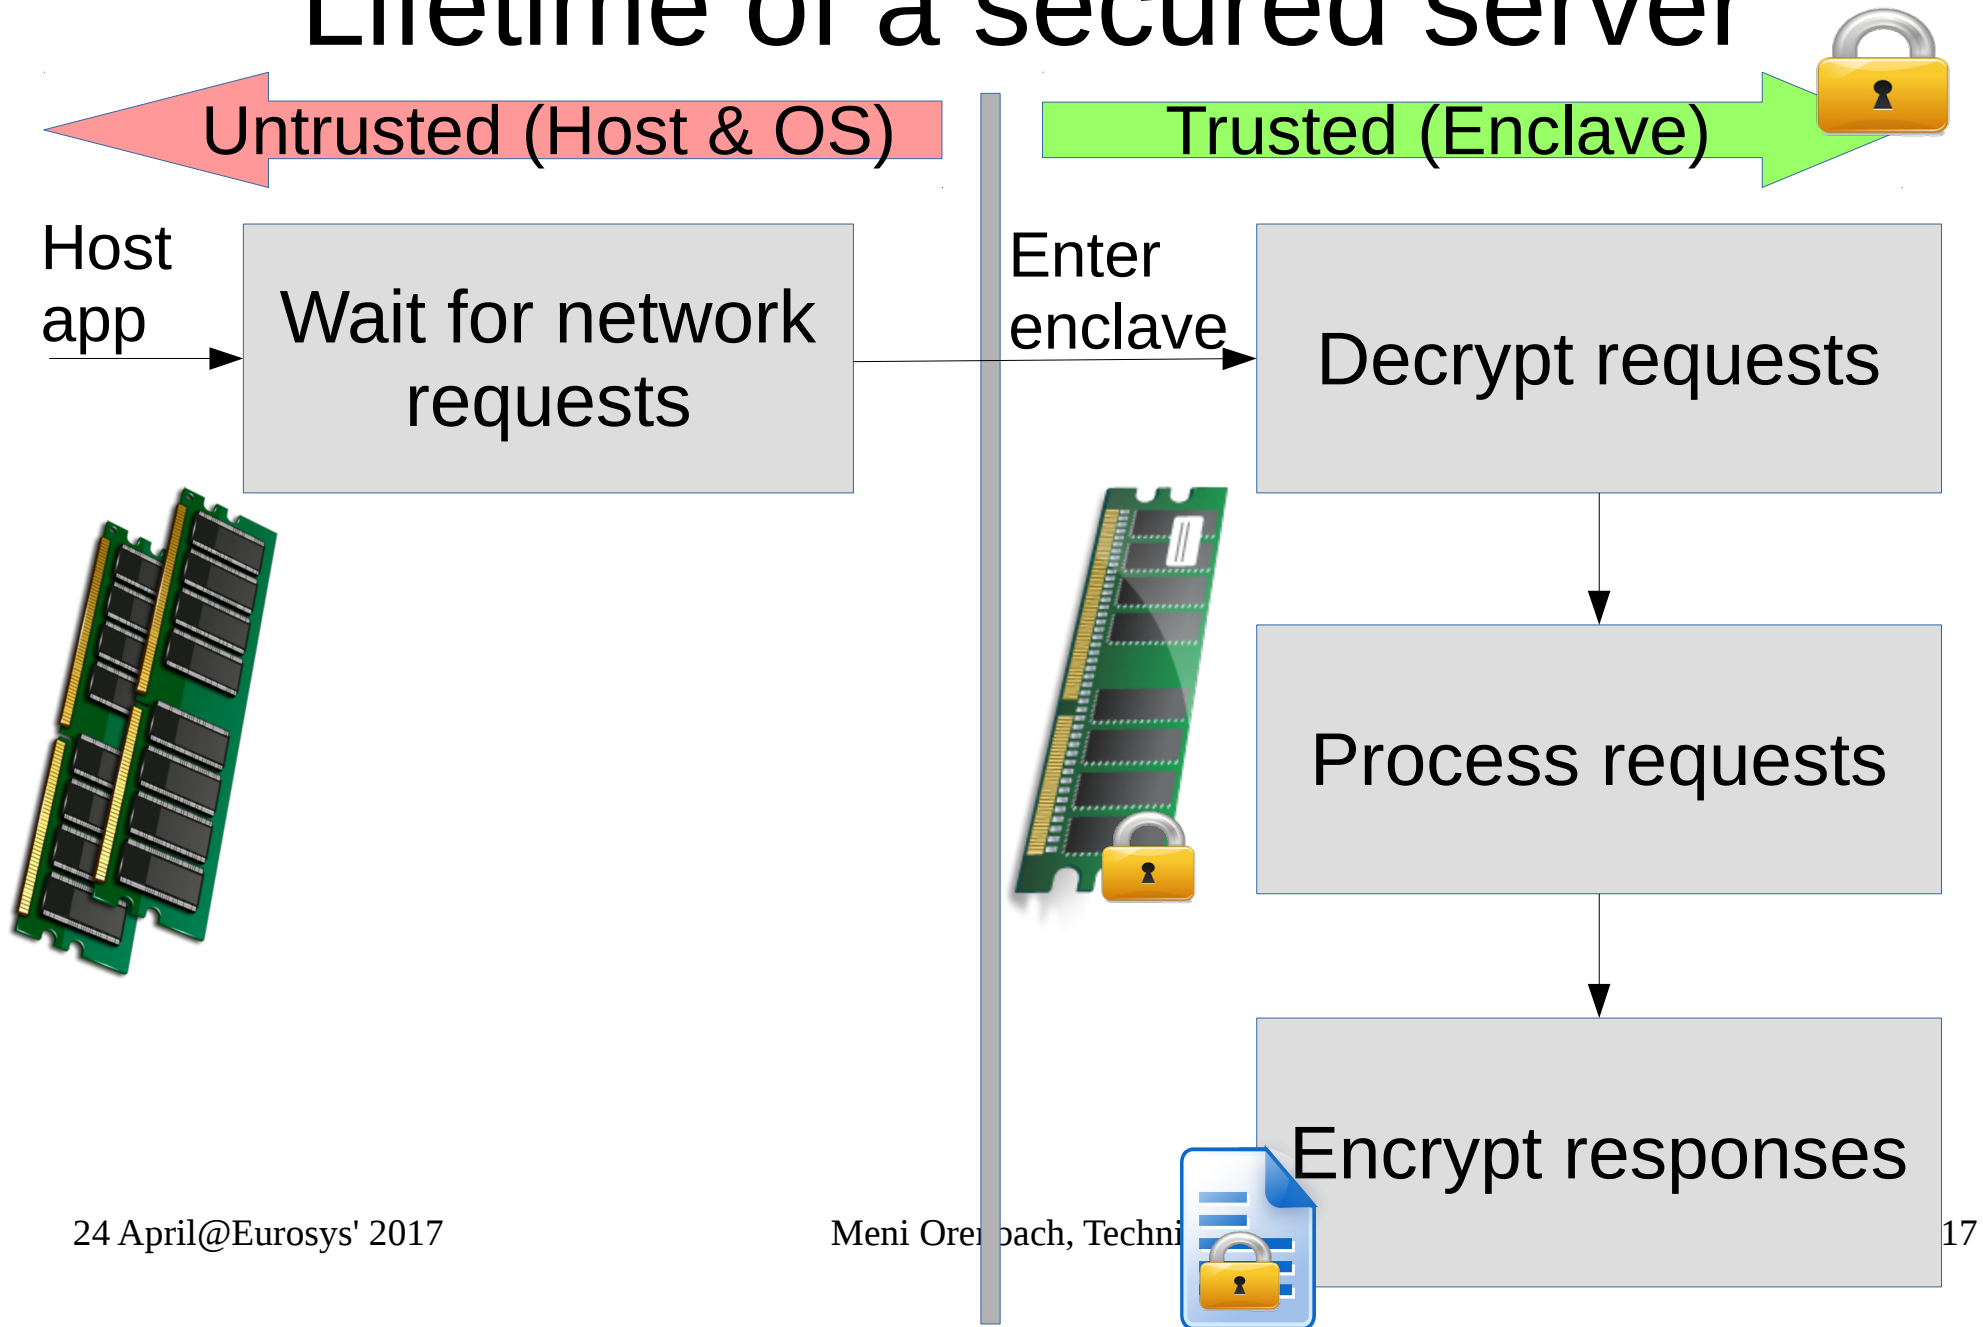
Background: Lifetime of a secured server



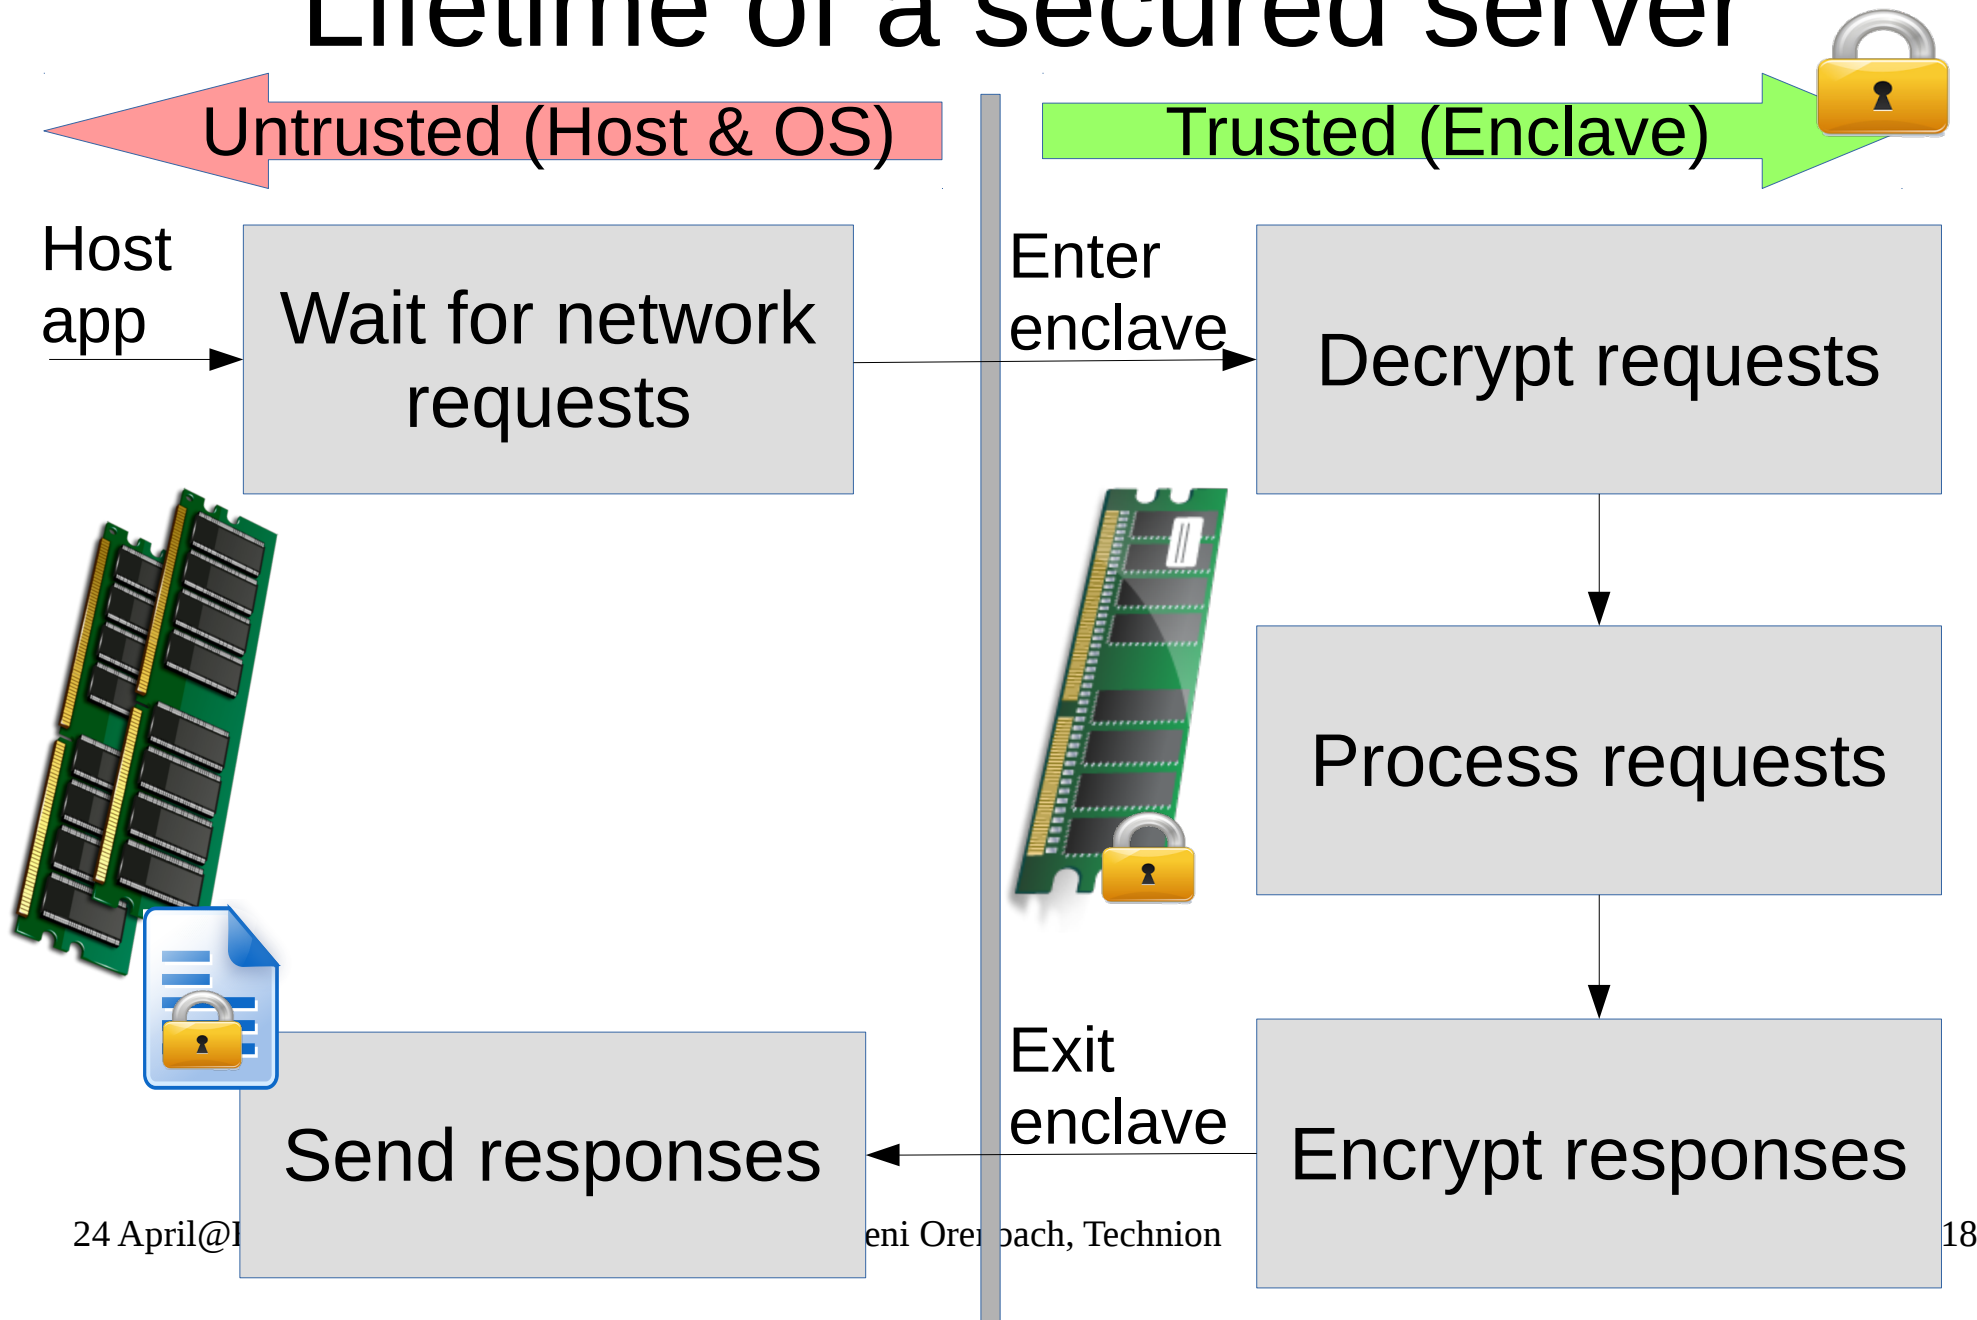
Background: Lifetime of a secured server



Background: Lifetime of a secured server

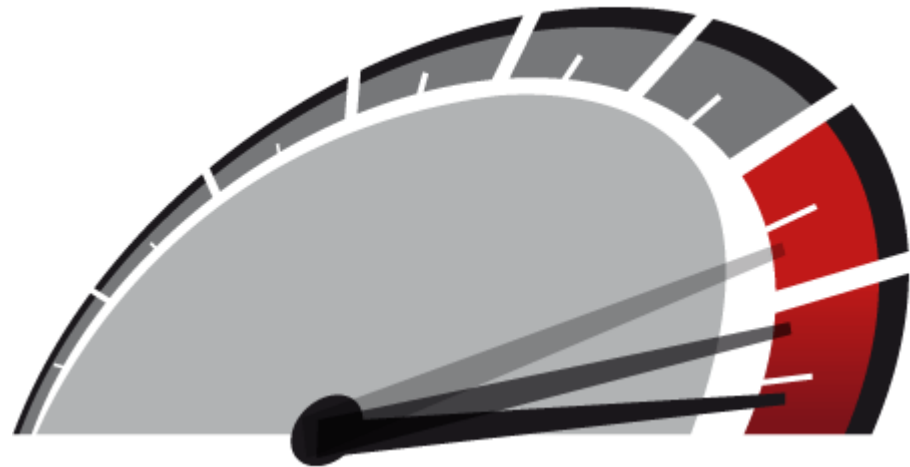


Background: Lifetime of a secured server



SGX enclaves should be fast

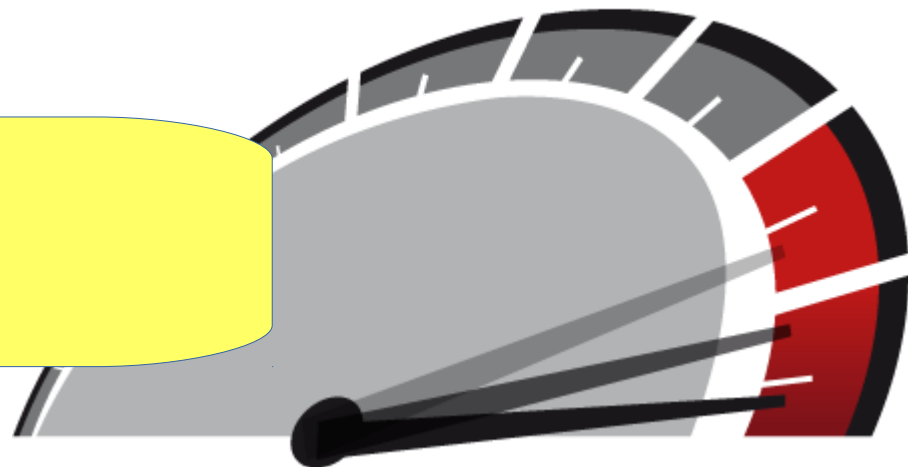
- ISA extensions
- Implemented in HW & Firmware
- Same CPU HW
- In-cache execution suffers no overheads



SGX enclaves should be fast

- ISA extensions
- Implemented in HW & Firmware
- Same CPU HW
- In-cache execution suffers no overheads

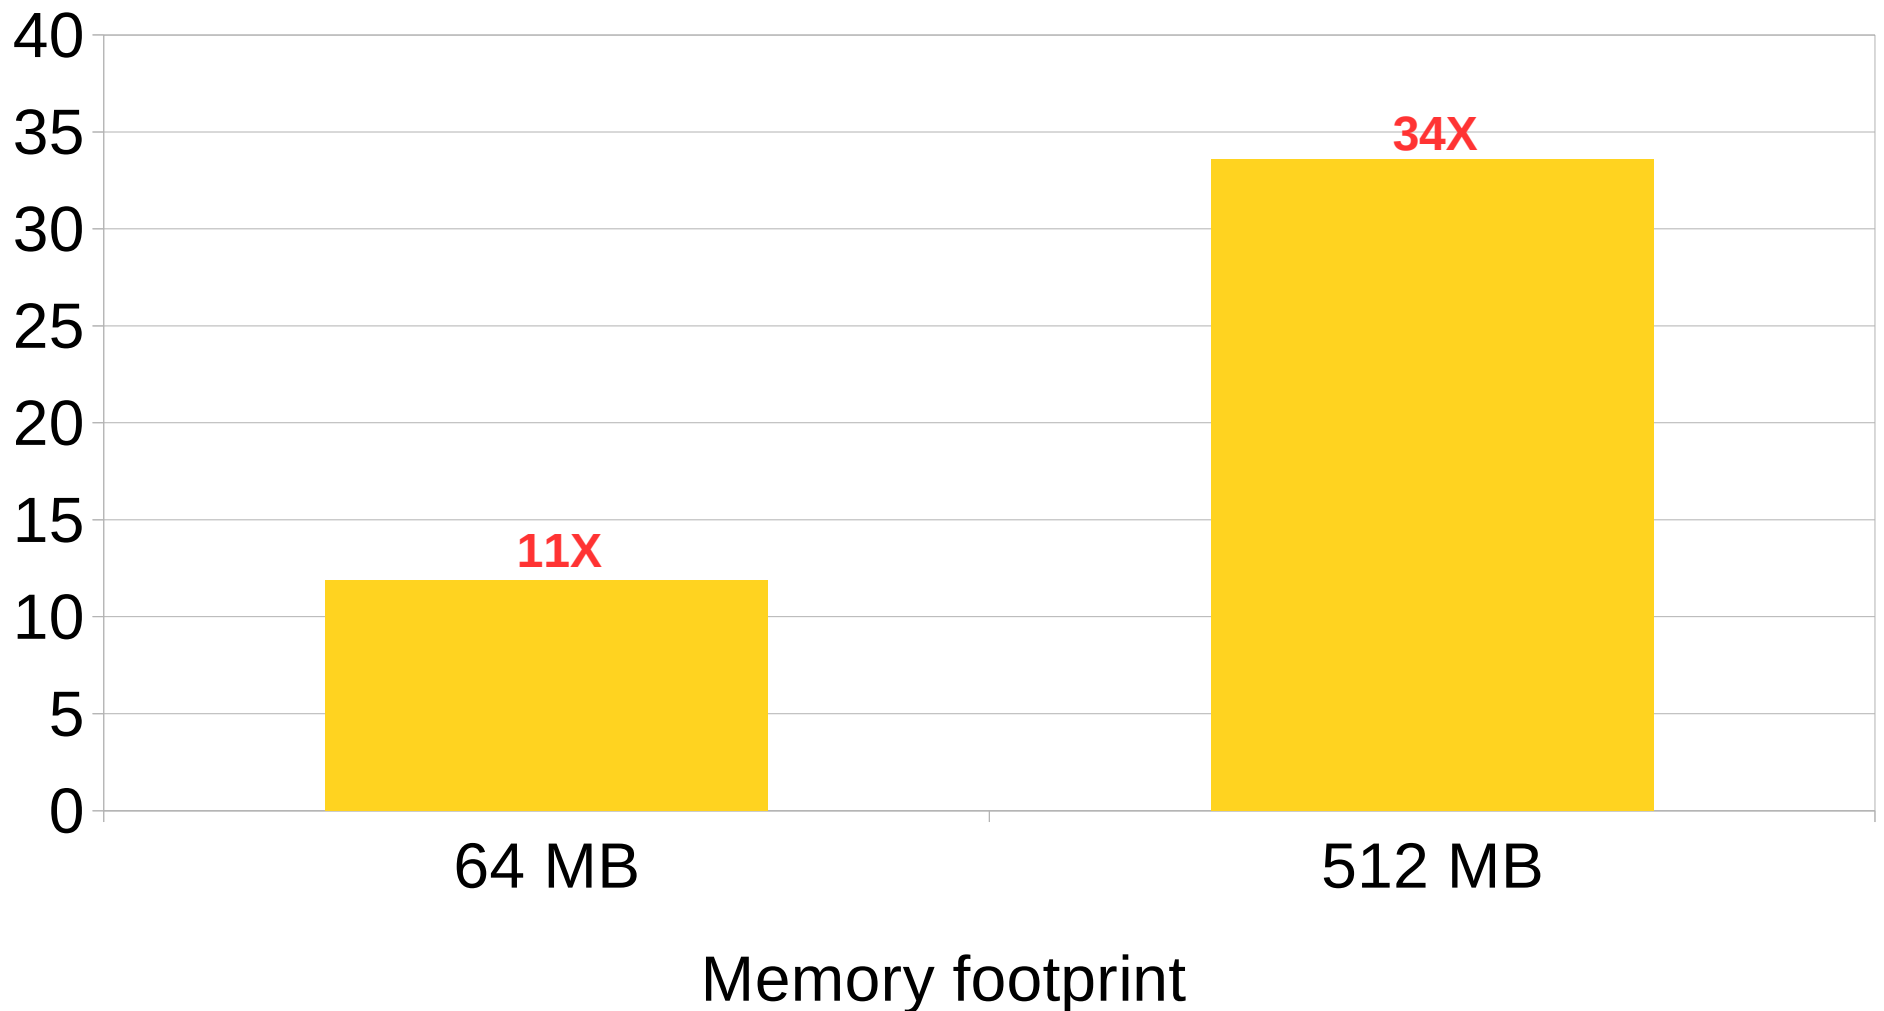
However...



Executing a Key-Value Store in enclave is slower

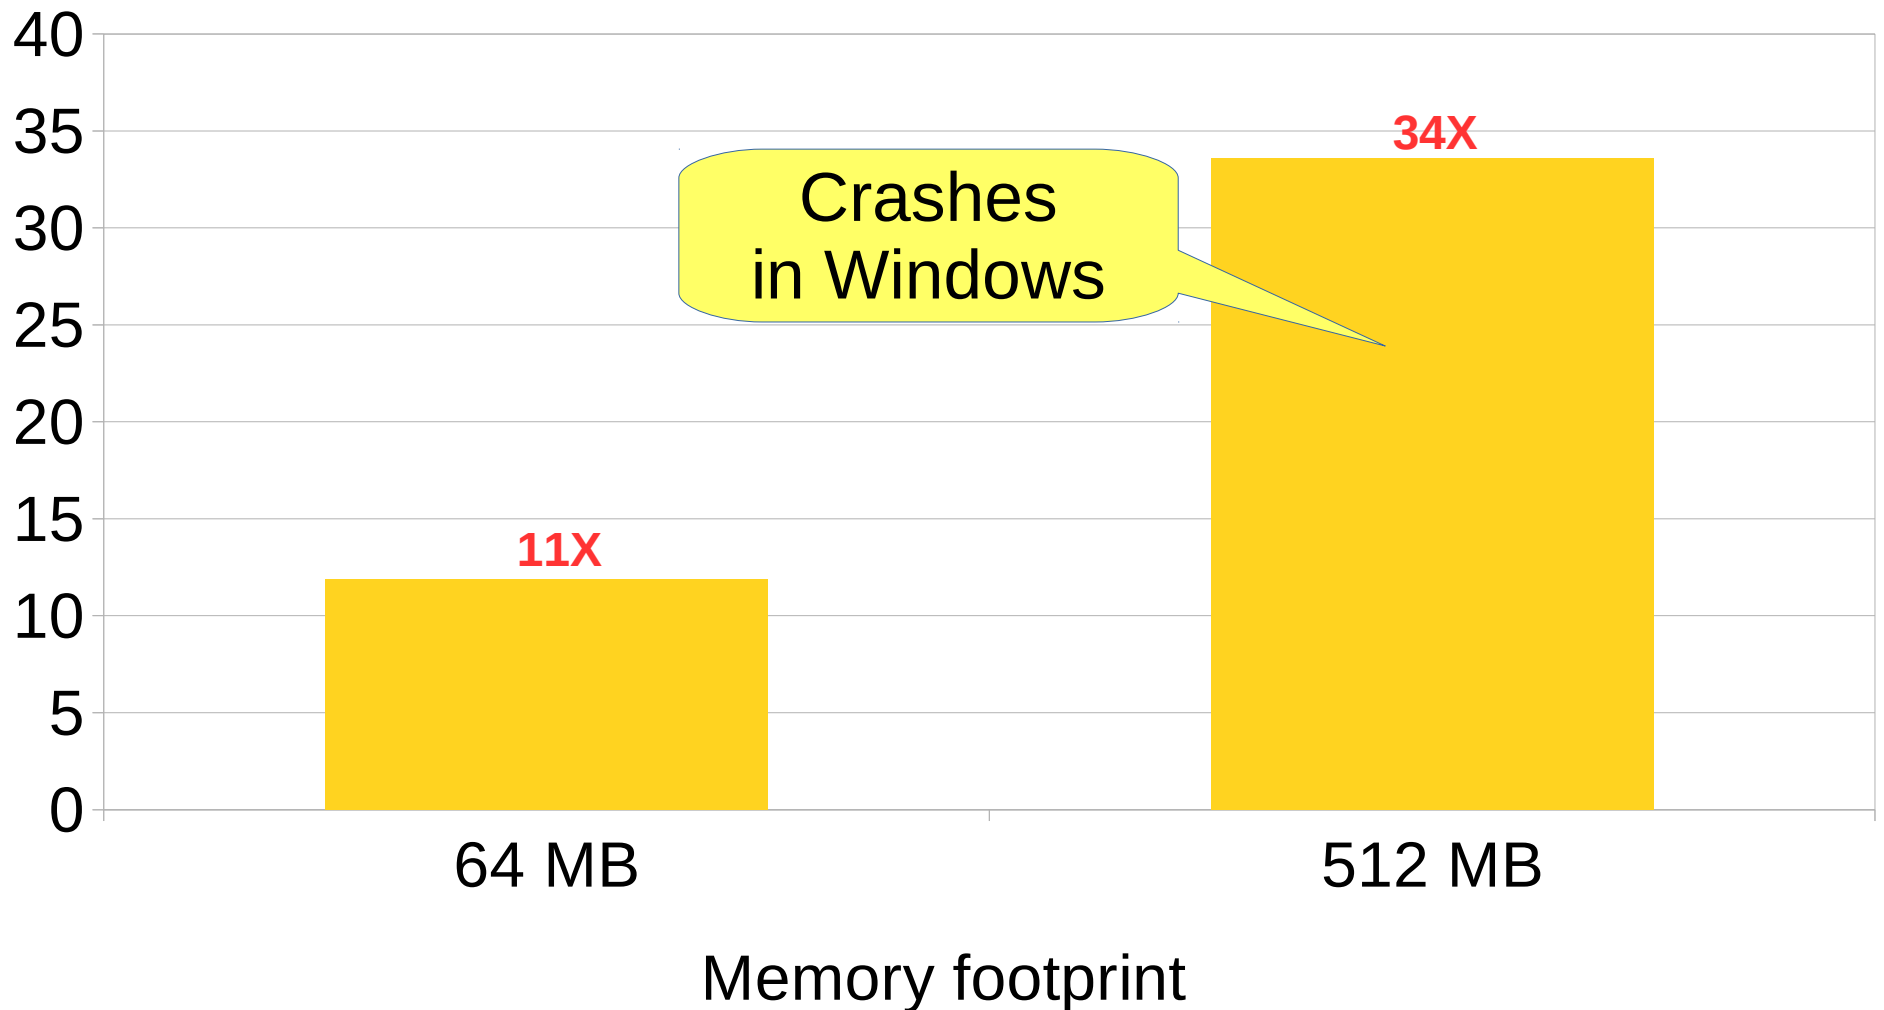
Executing a Key-Value Store in enclave is slower

Throughput: Slowdown factor



Executing a Key-Value Store in enclave is slower

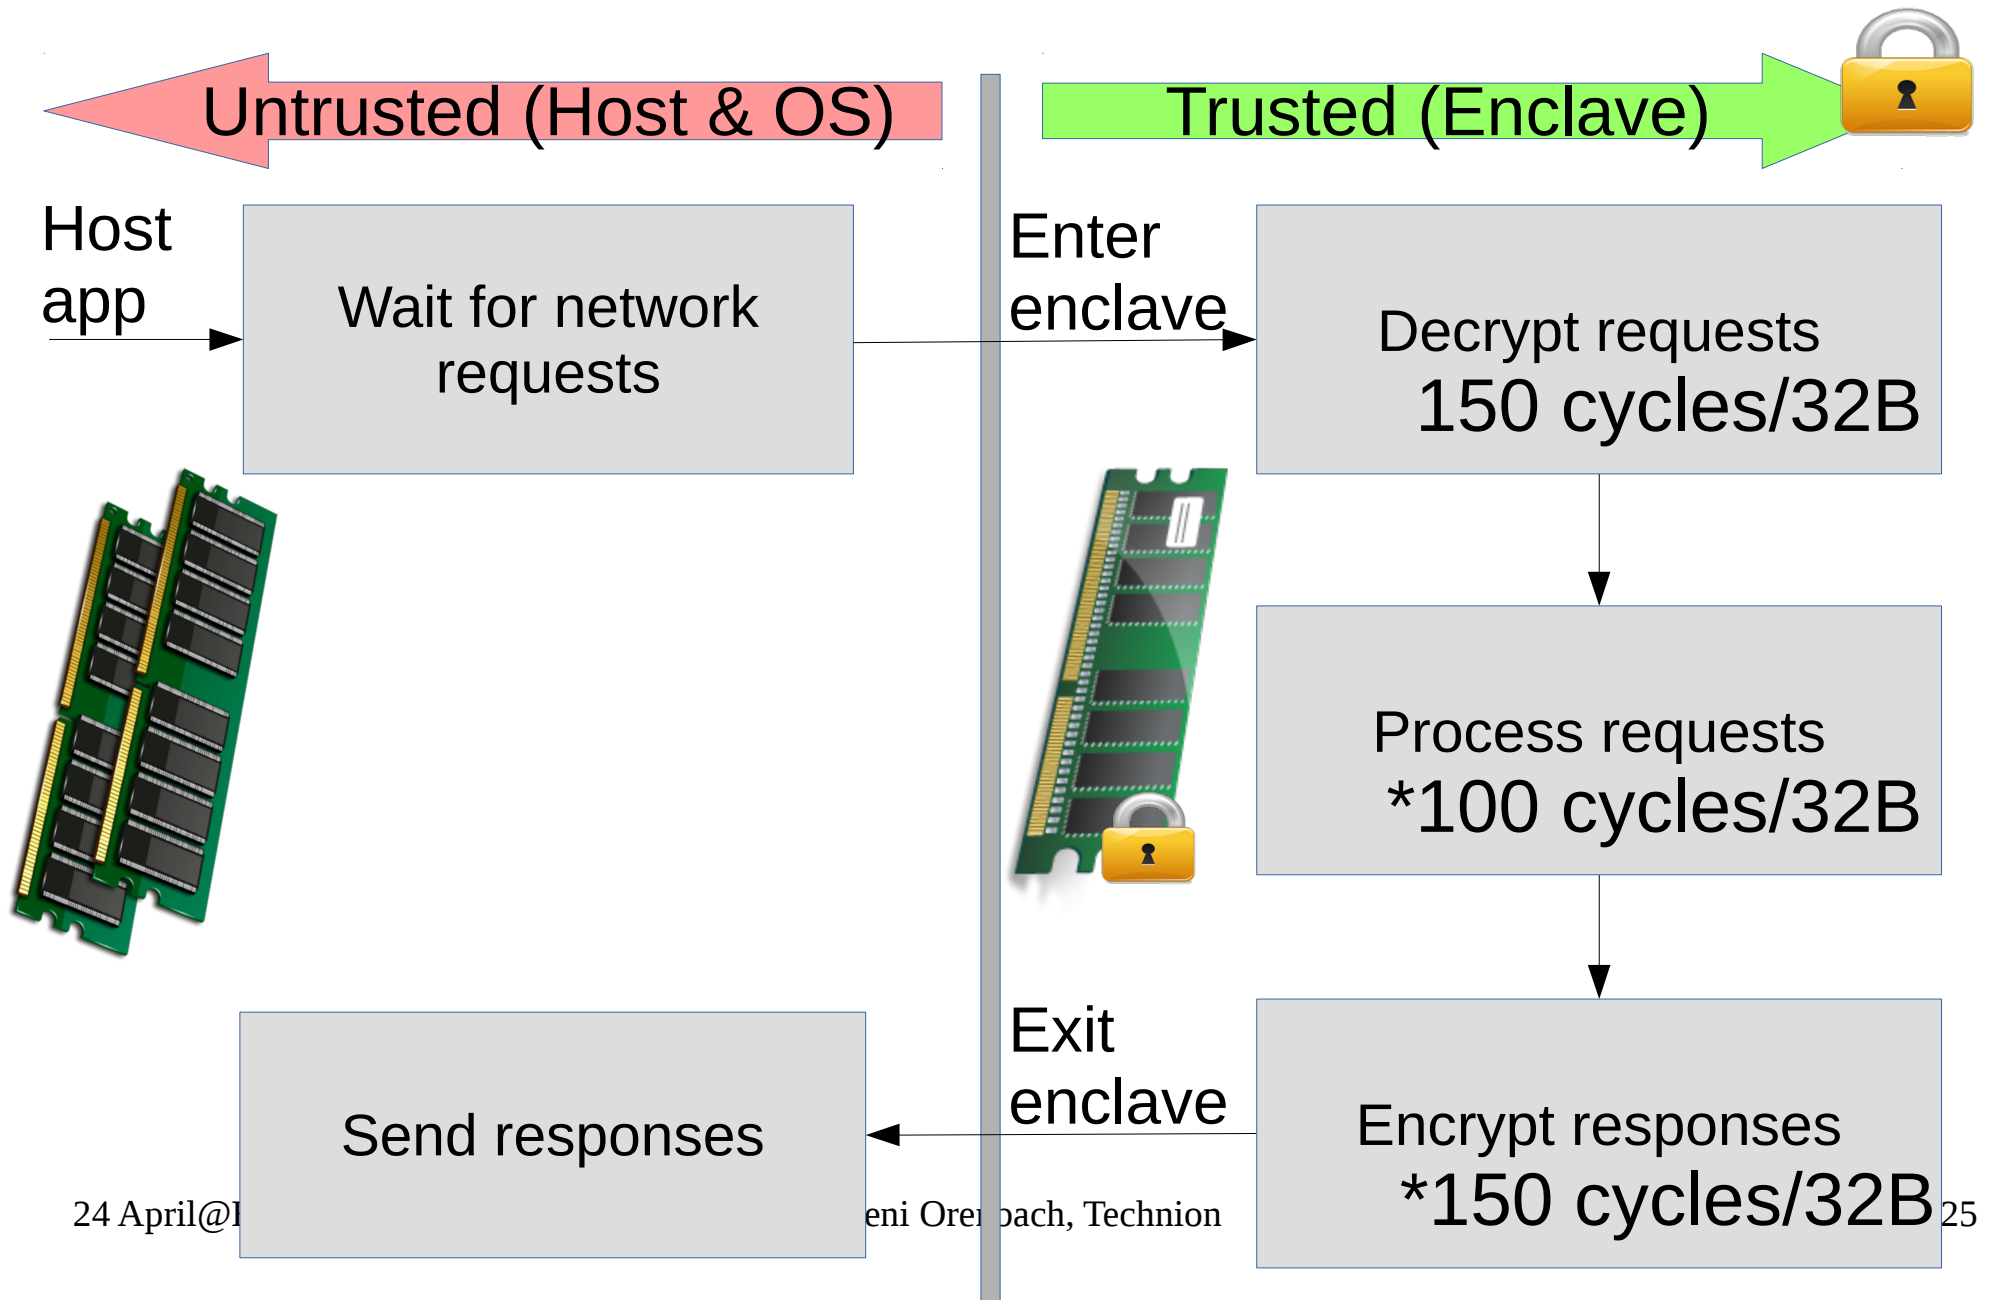
Throughput: Slowdown factor



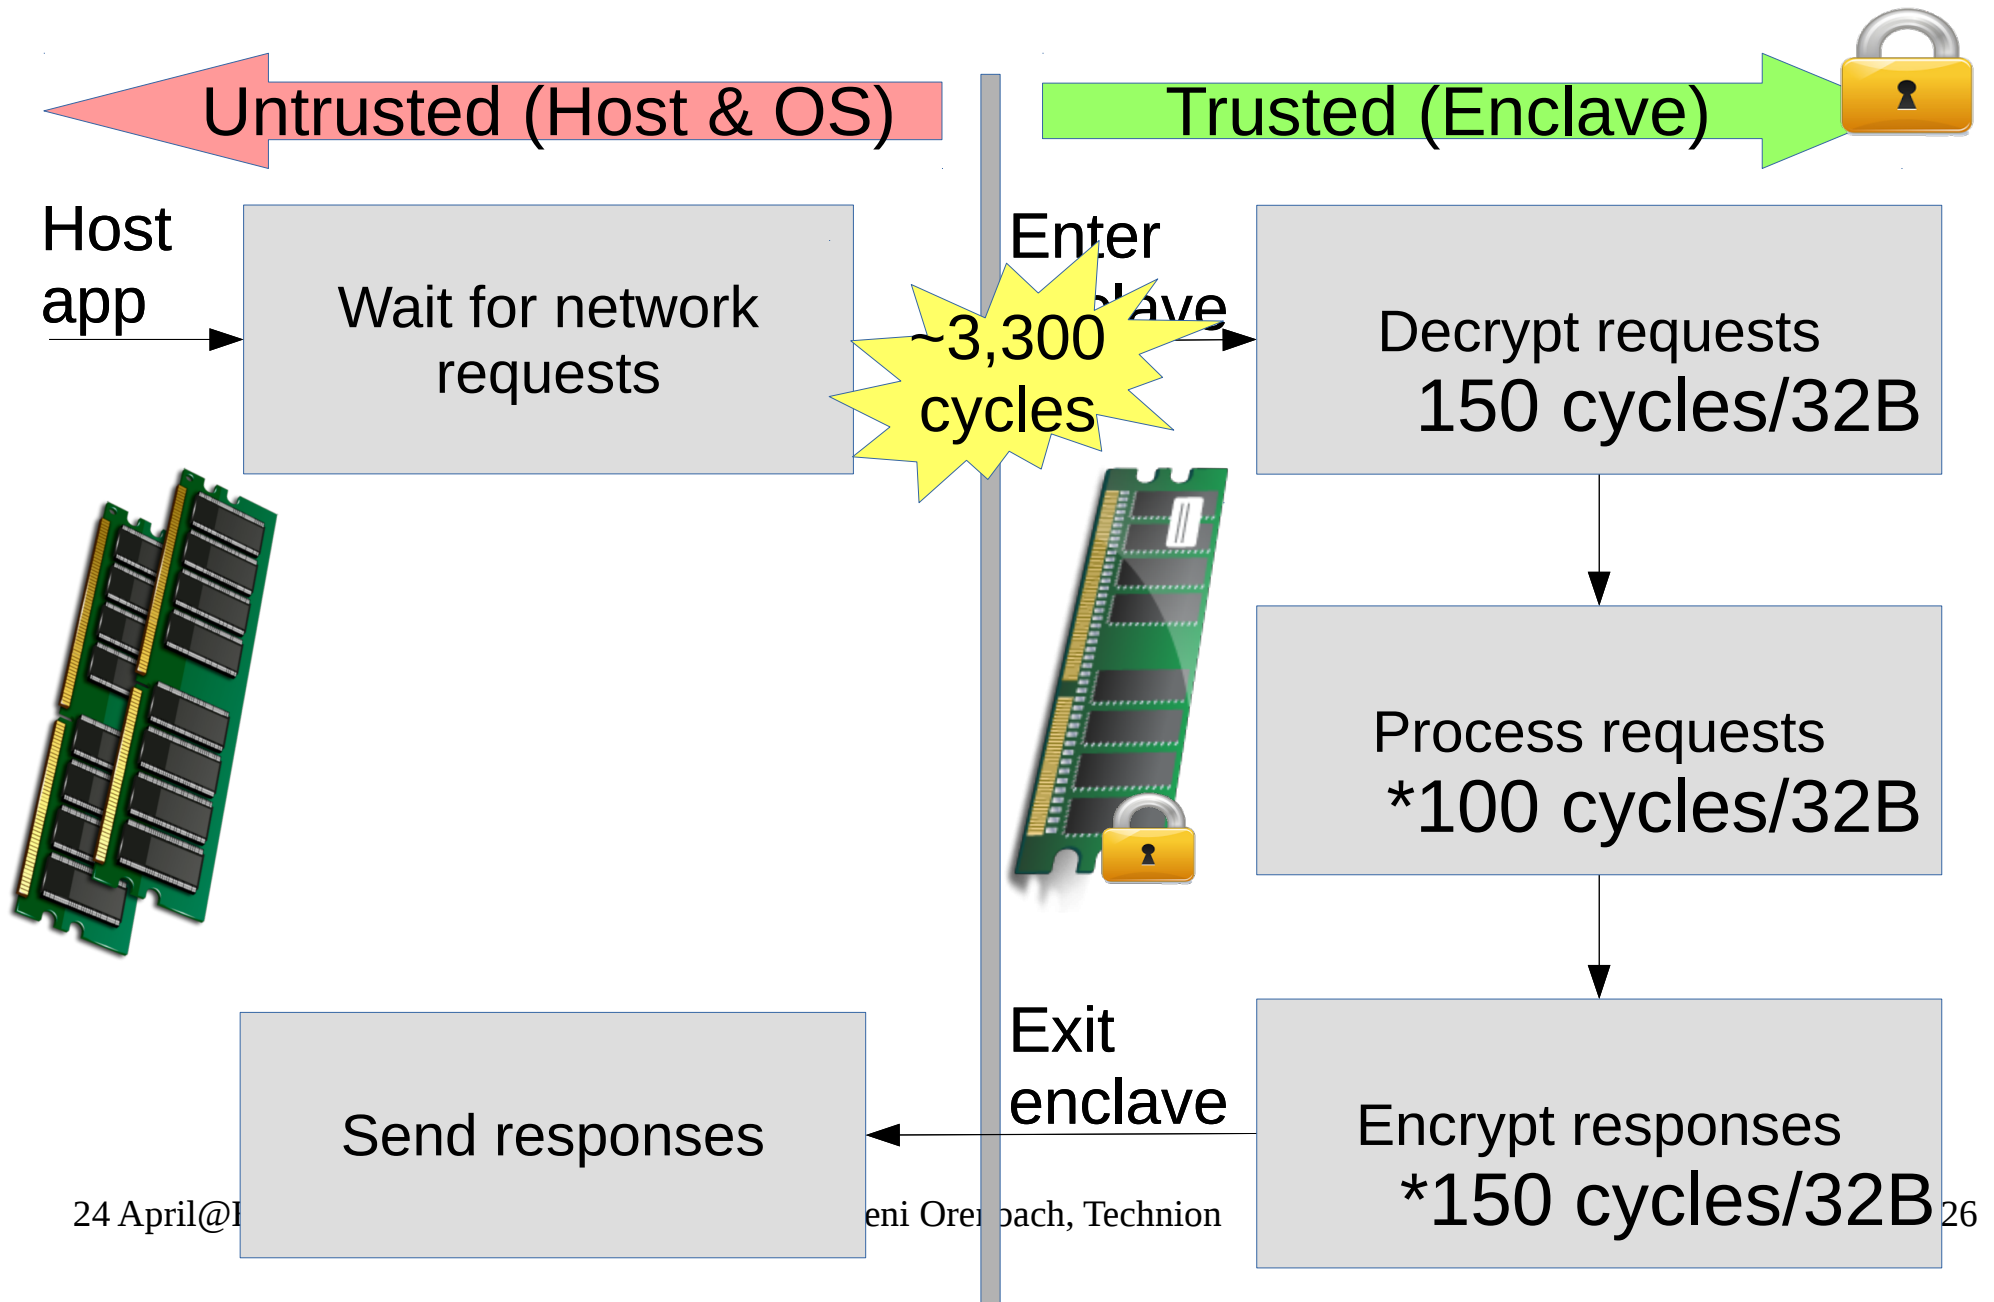
- Background
- Motivation
- **Overhead analysis**
- Eleos design
- Evaluation



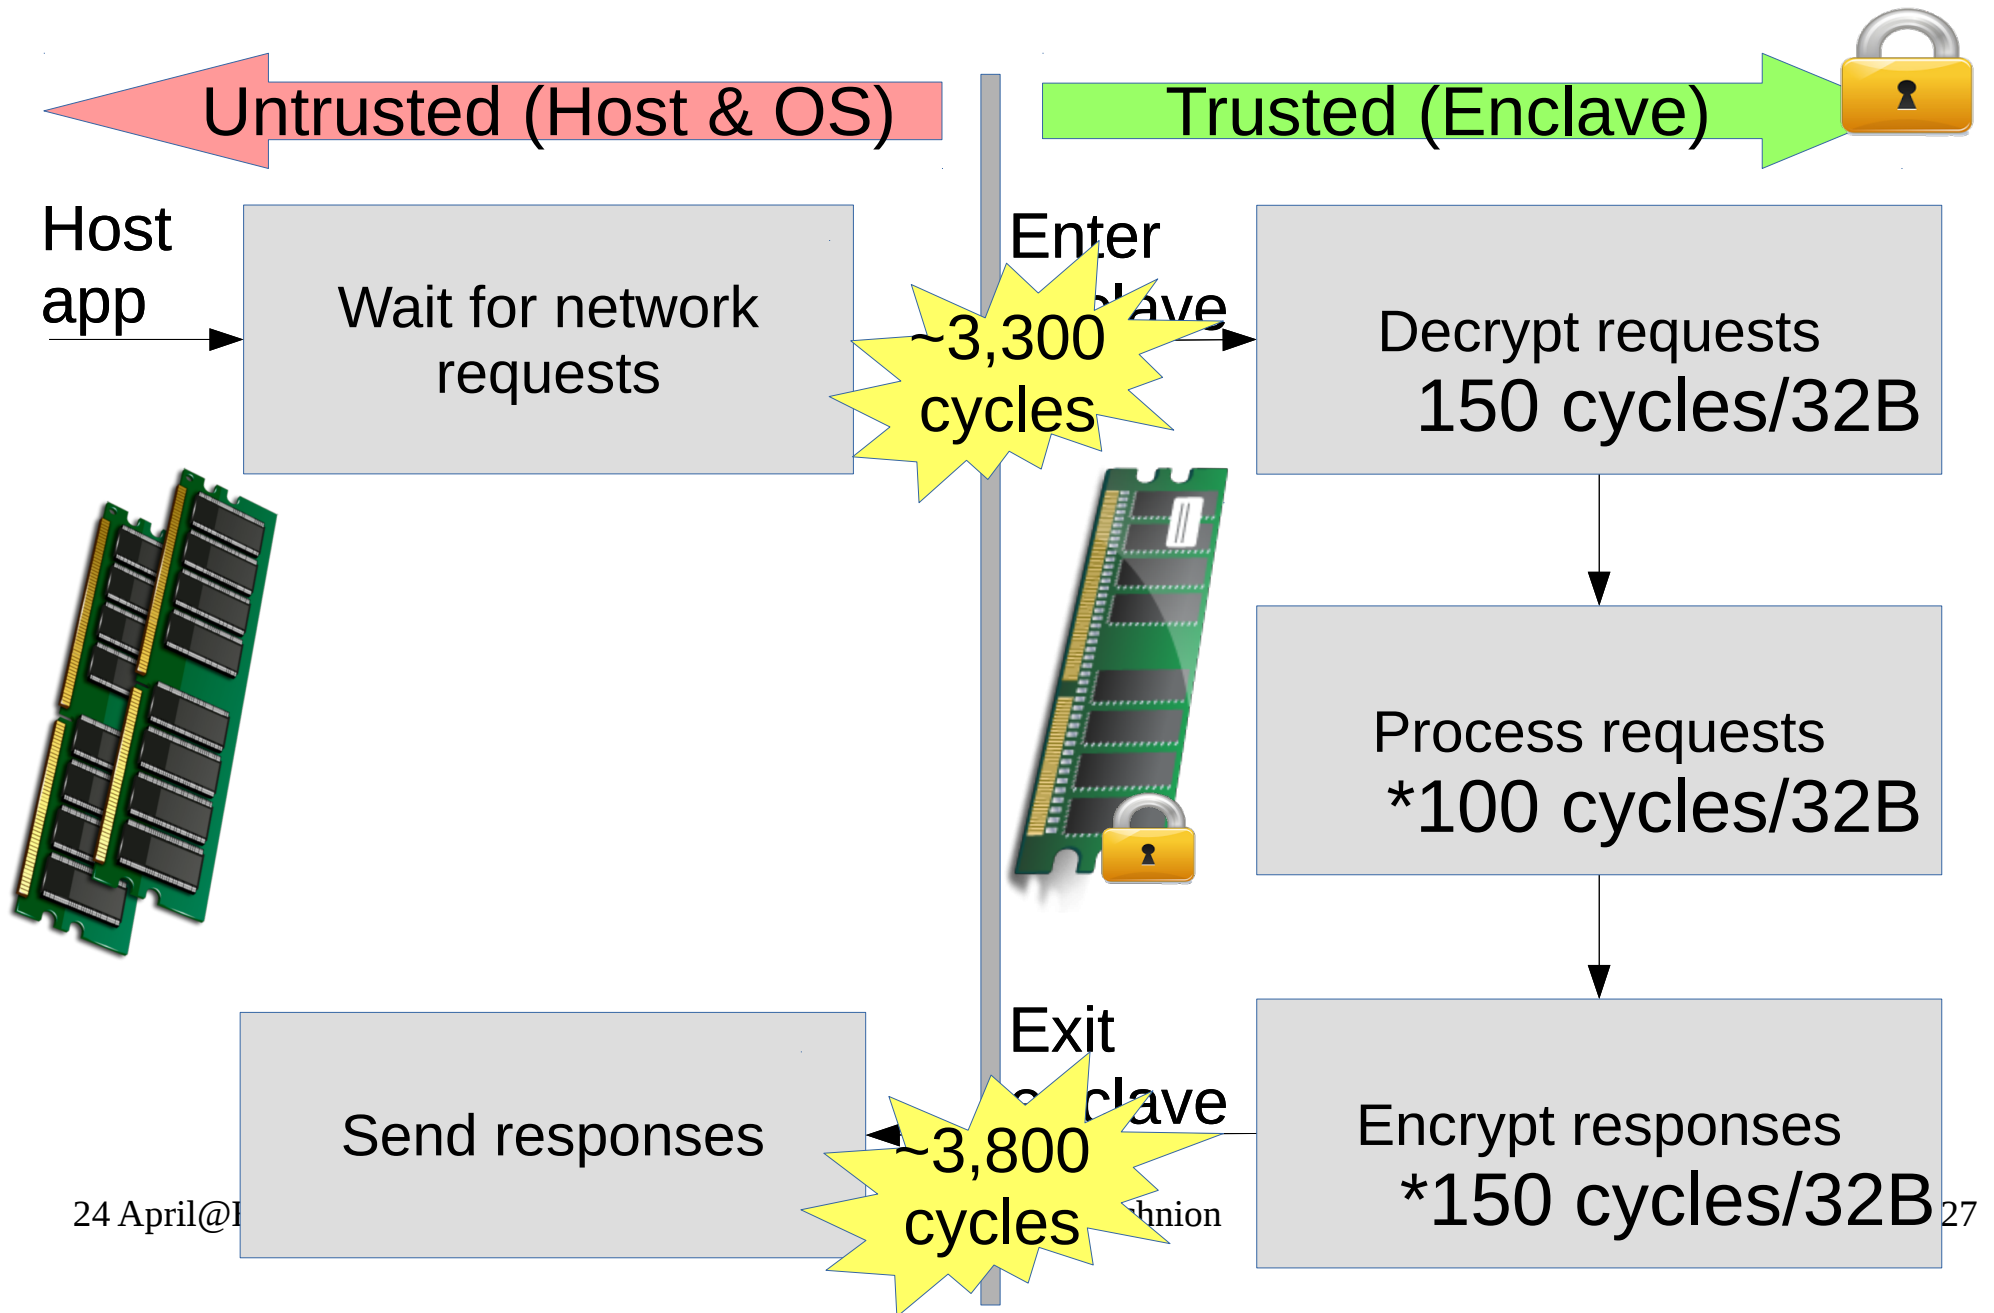
Overhead analysis



Overhead analysis




Overhead analysis

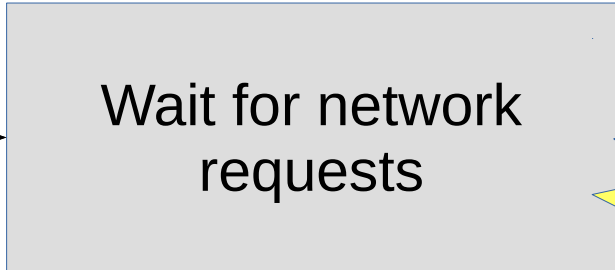


Overhead analysis

Untrusted (Host & OS)

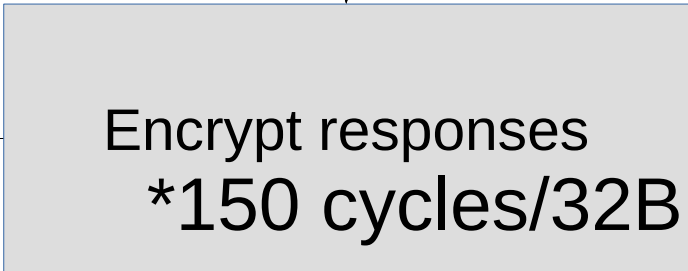
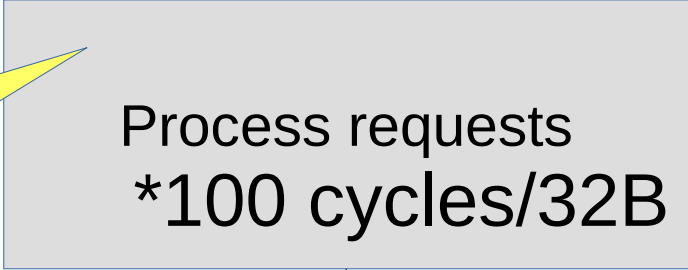
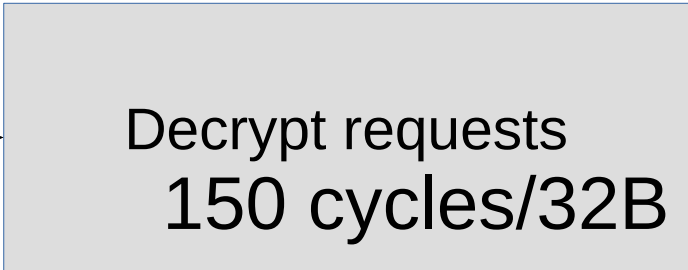
Trusted (Enclave) 

Host app

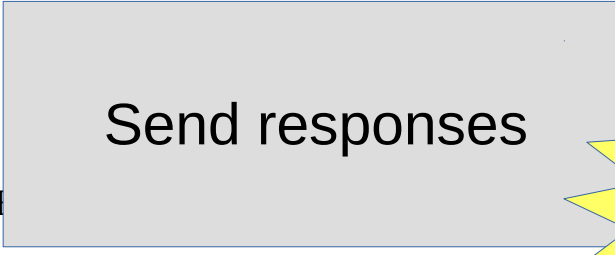


Enter enclave

~3,300 cycles



Exit enclave




~3,800 cycles

**Exits causes indirect costs:
1.5X – 5X slower execution
FlexSC [OSDI'10]
syscall analysis**

Overhead analysis

Untrusted (Host & OS)

Trusted (Enclave) 

Host app

Wait for network requests

Enter enclave

Decrypt requests
150 cycles/32B

~3,300 cycles

Process requests
*100 cycles/32B

**Exits causes indirect costs:
1.5X – 5X slower execution
FlexSC [OSDI'10]
syscall analysis**

Send responses

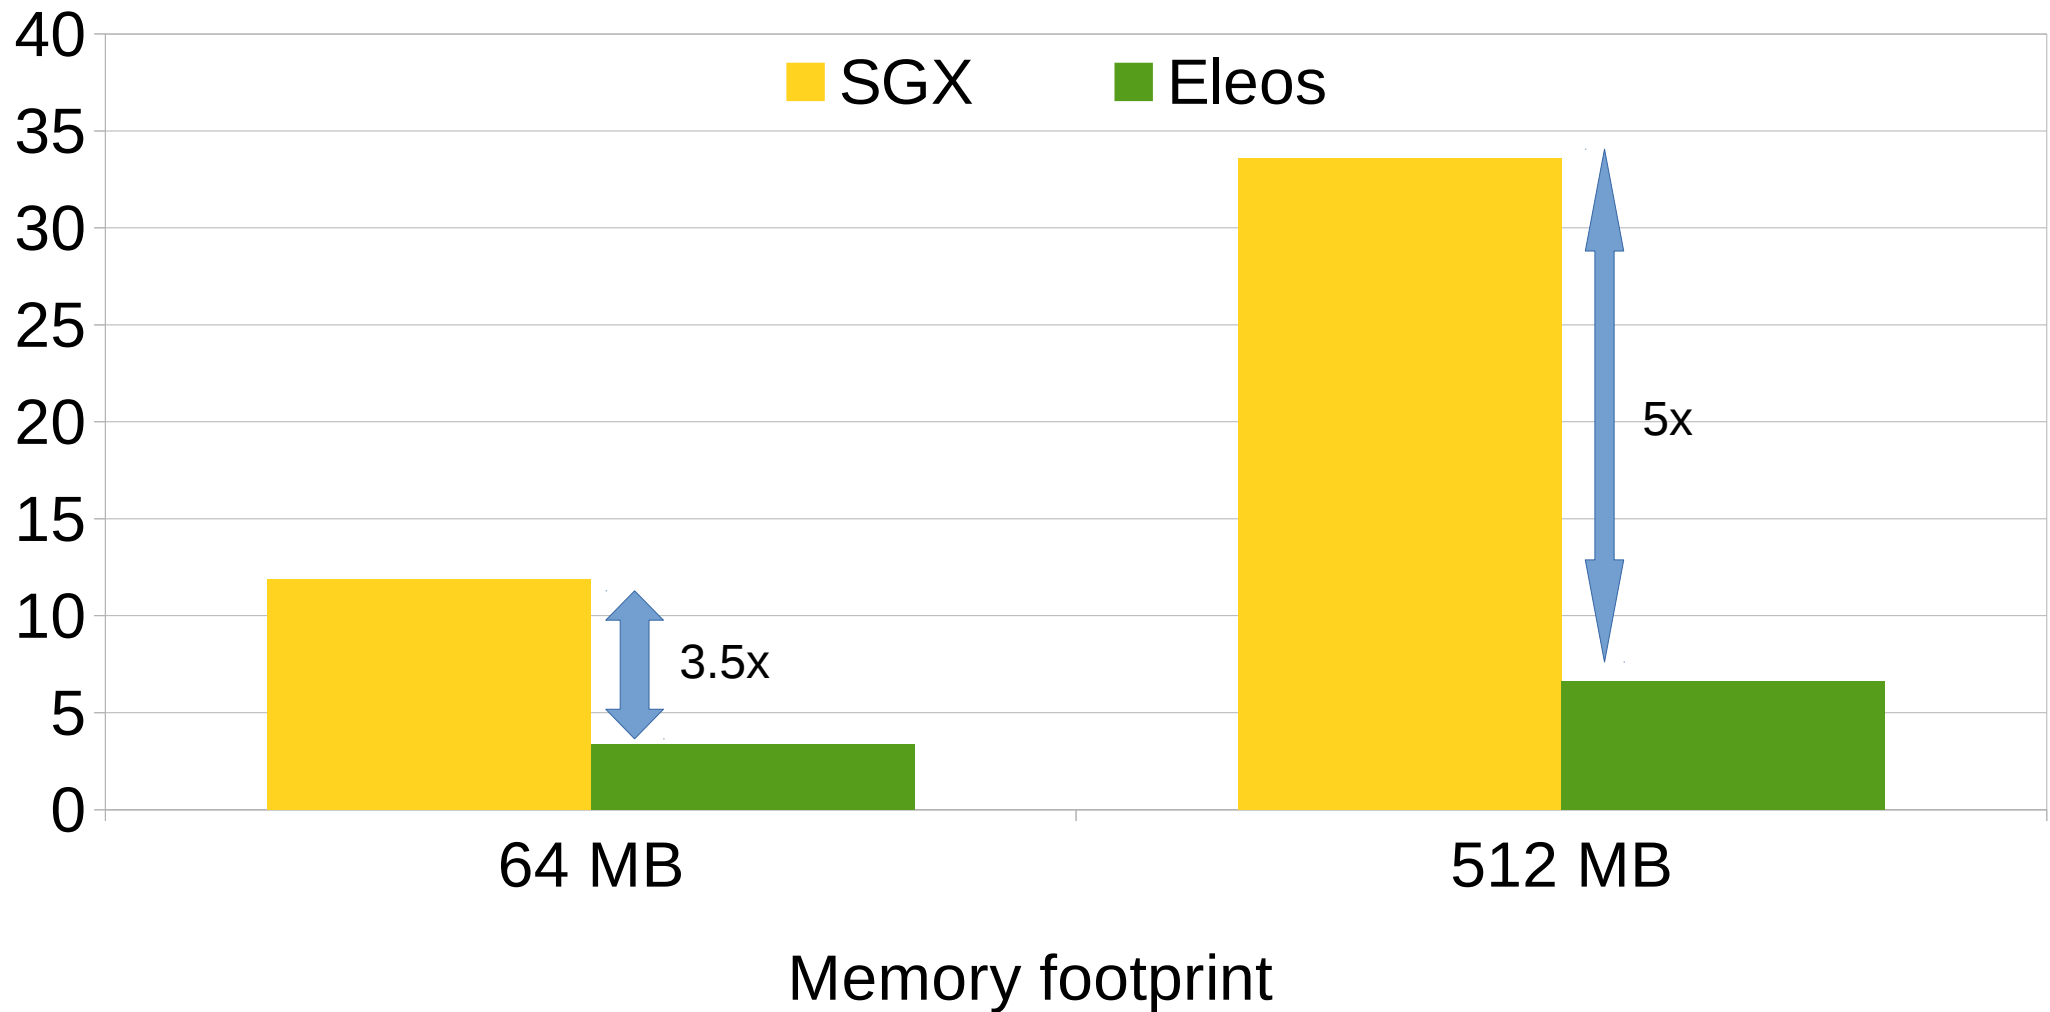
Exit enclave

Encrypt responses
*150 cycles/32B

~3,800 cycles

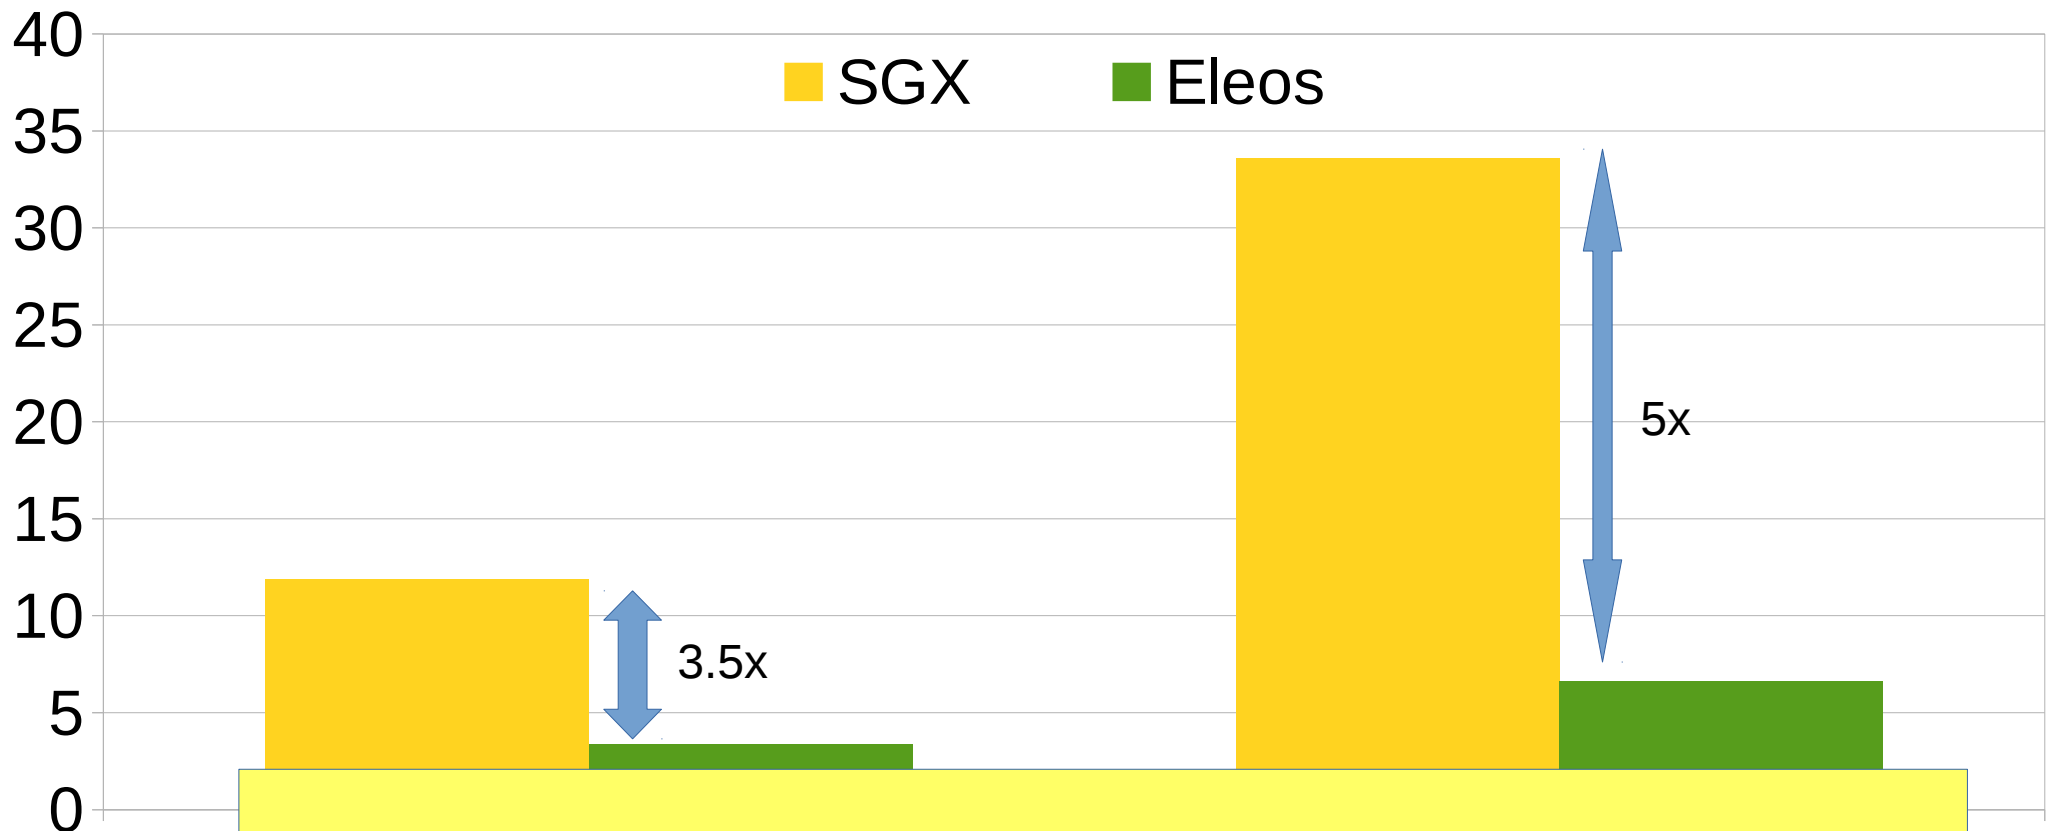
Eleos does better!

Throughput: Slowdown factor



Eleos does better!

Throughput: Slowdown factor



How does Eleos achieve this?

Eleos: **Exit-less** services

Exit-less system calls with RPC infrastructure

Exit-less SGX paging



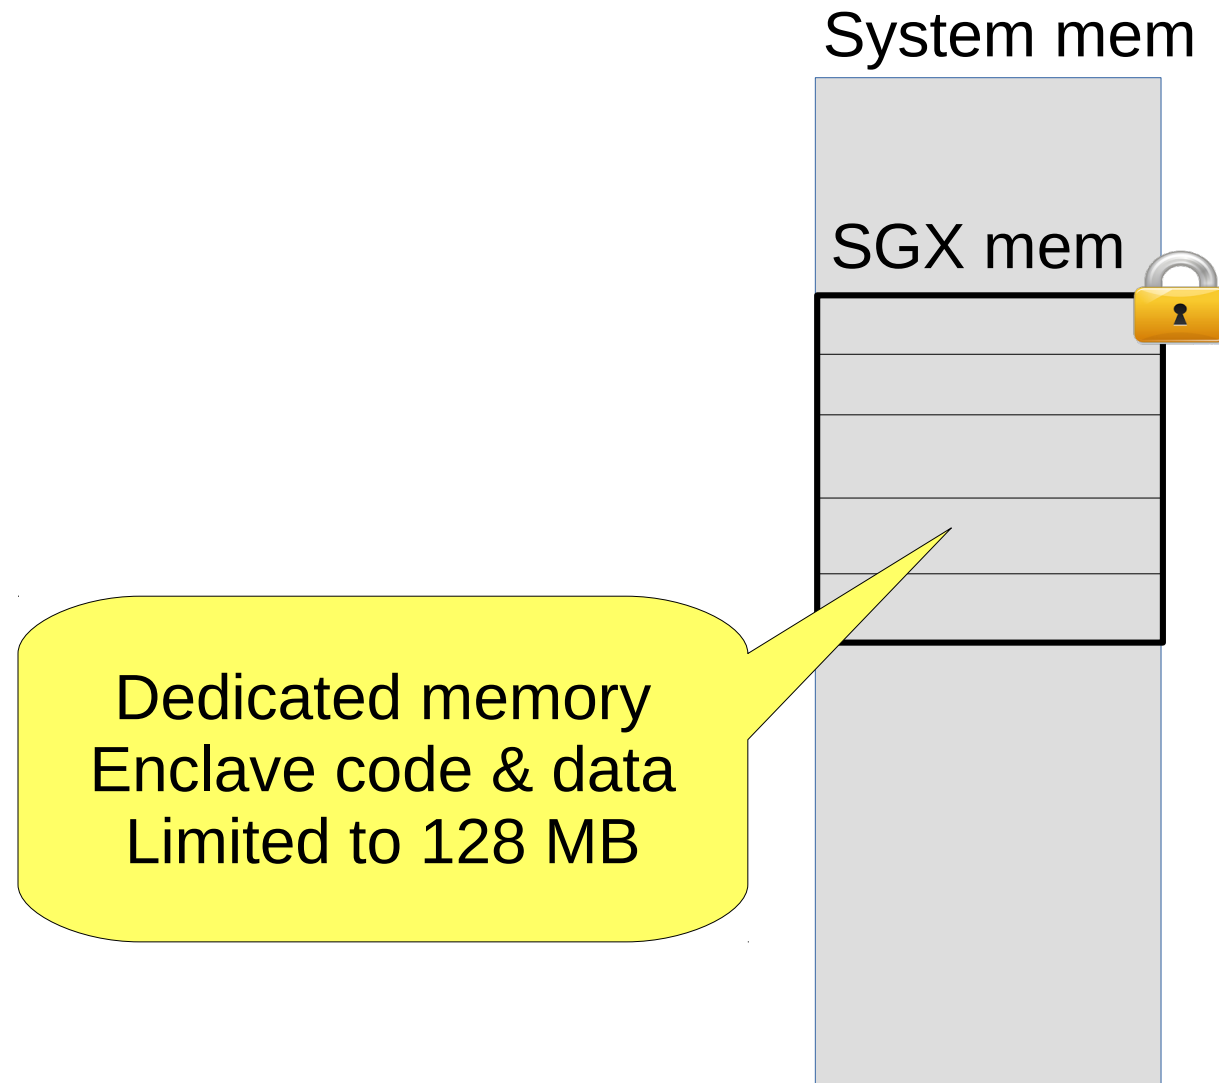
Eleos: **Exit-less** services

Exit-less system calls with RPC infrastructure

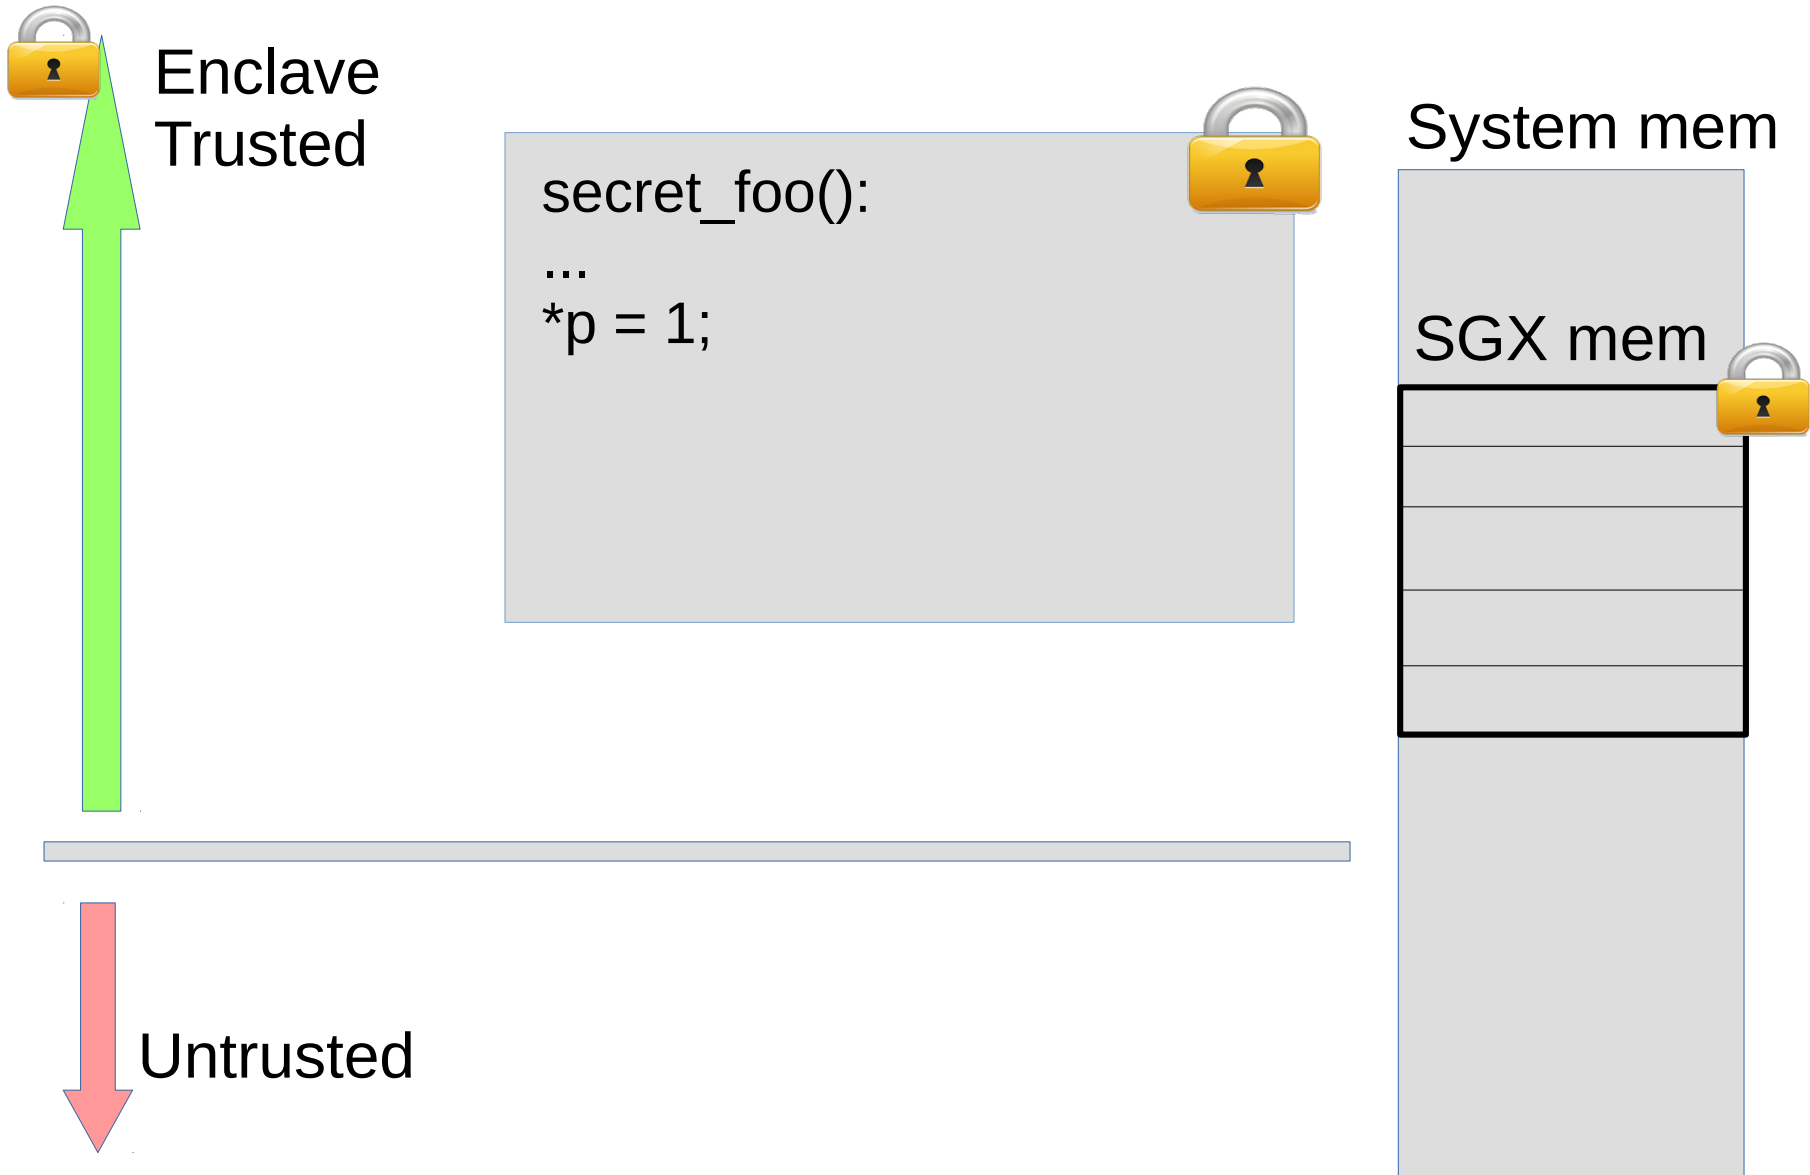
Exit-less SGX paging



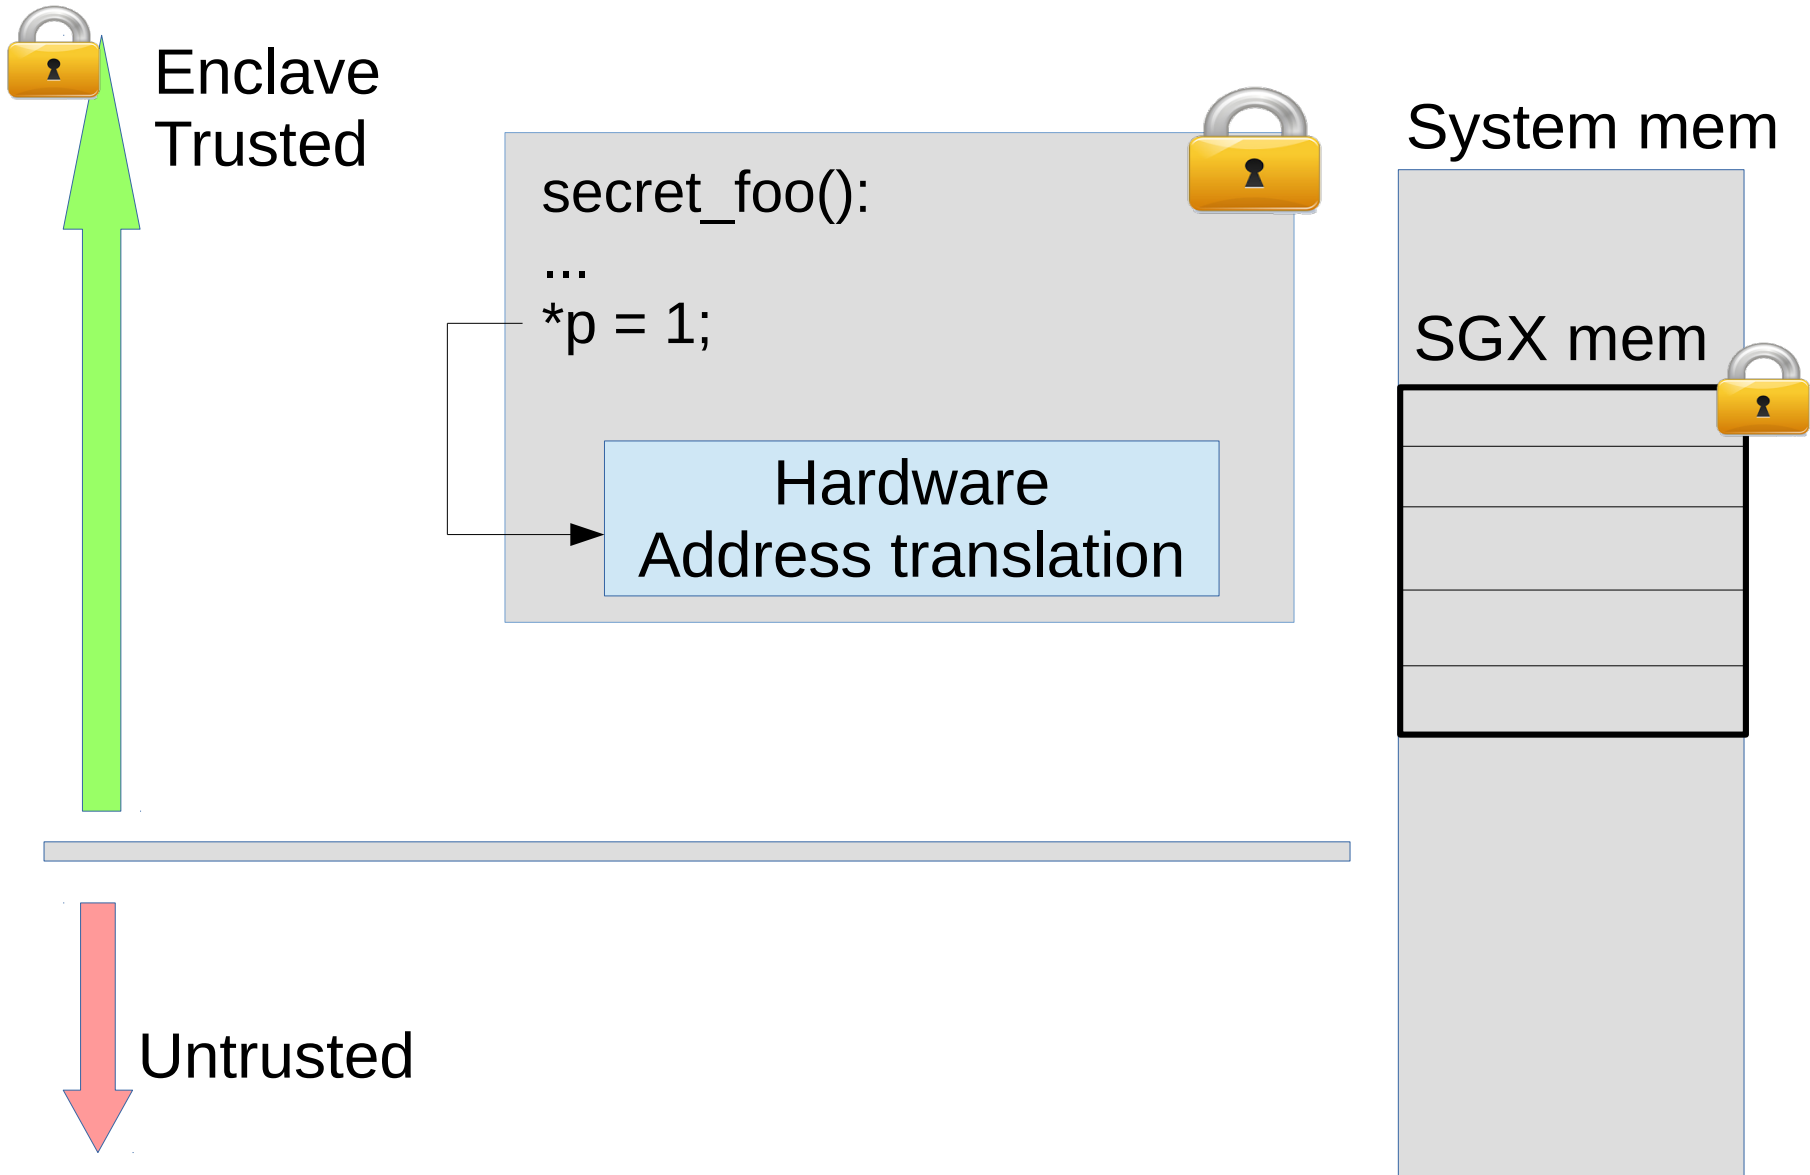
Background: SGX paging



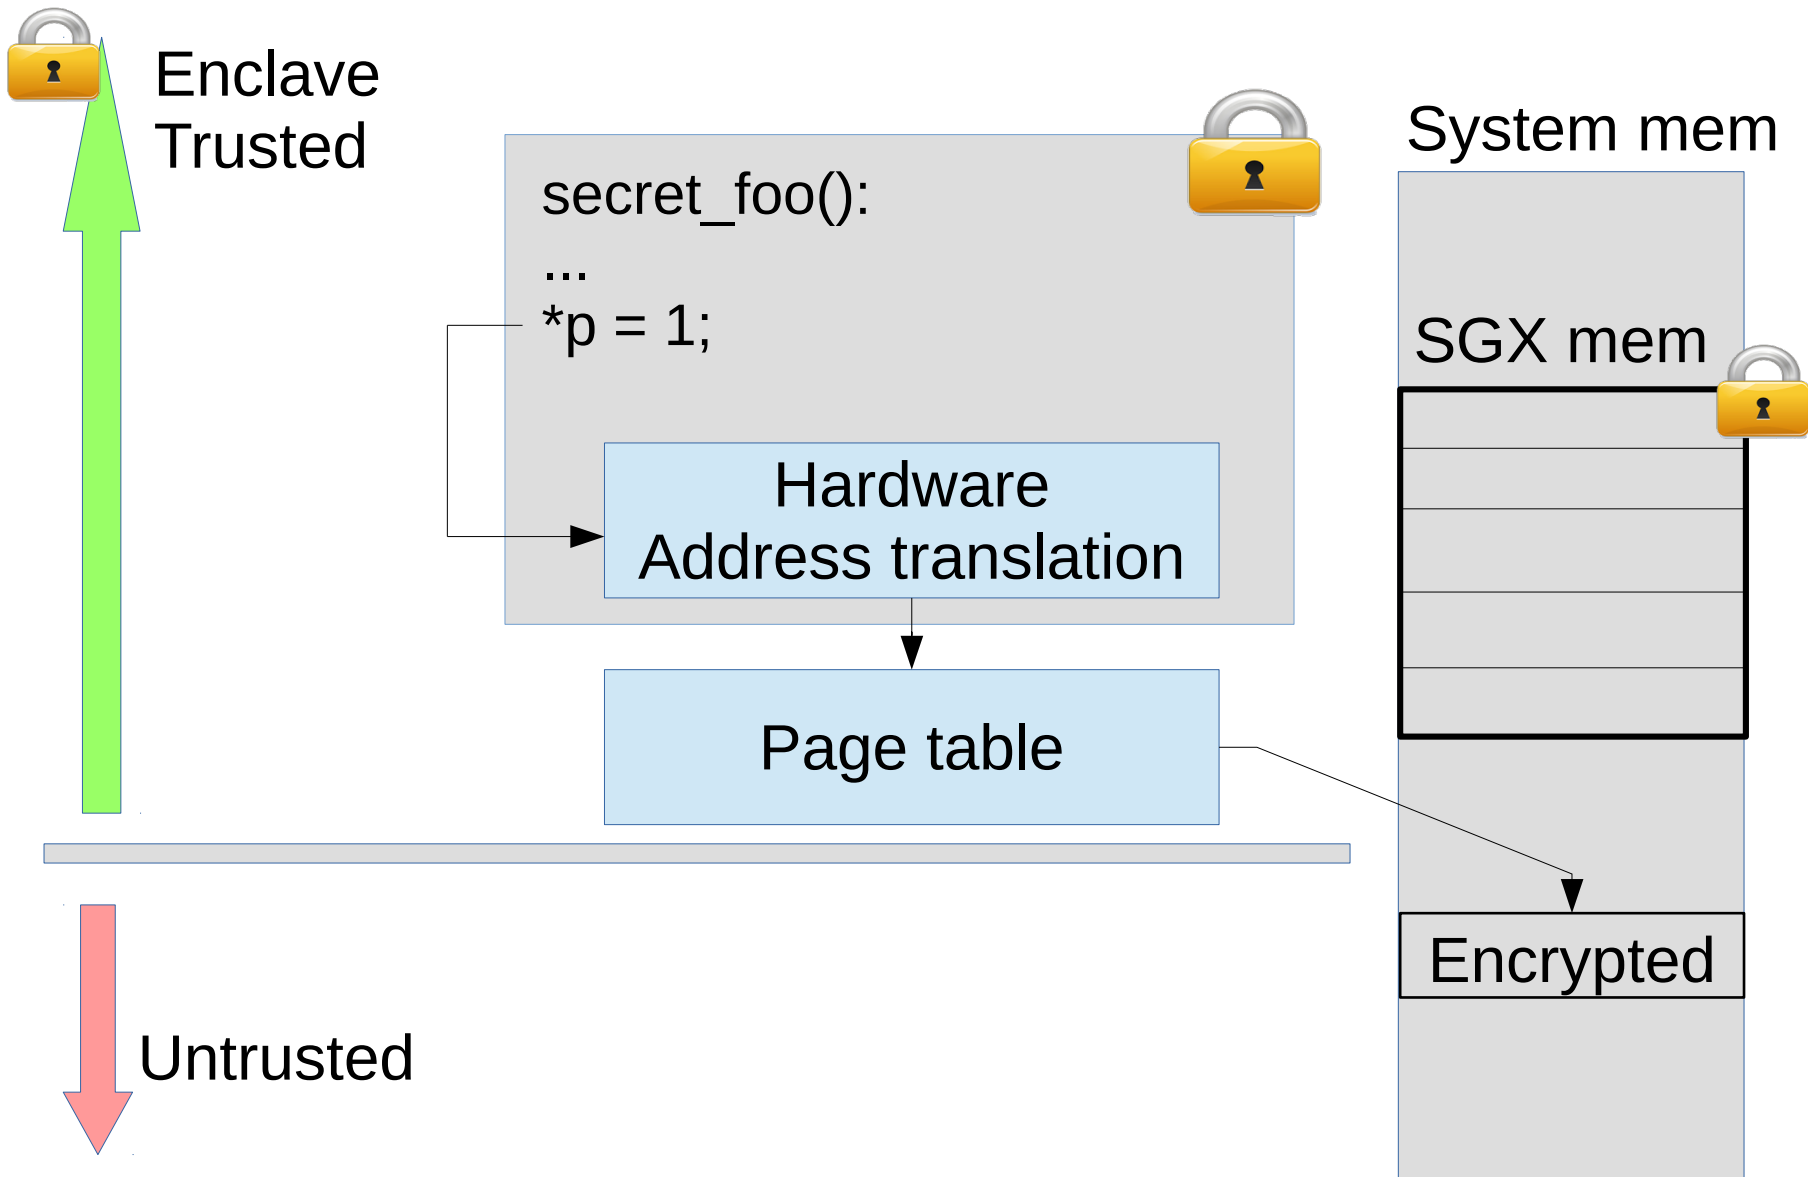
Background: SGX paging



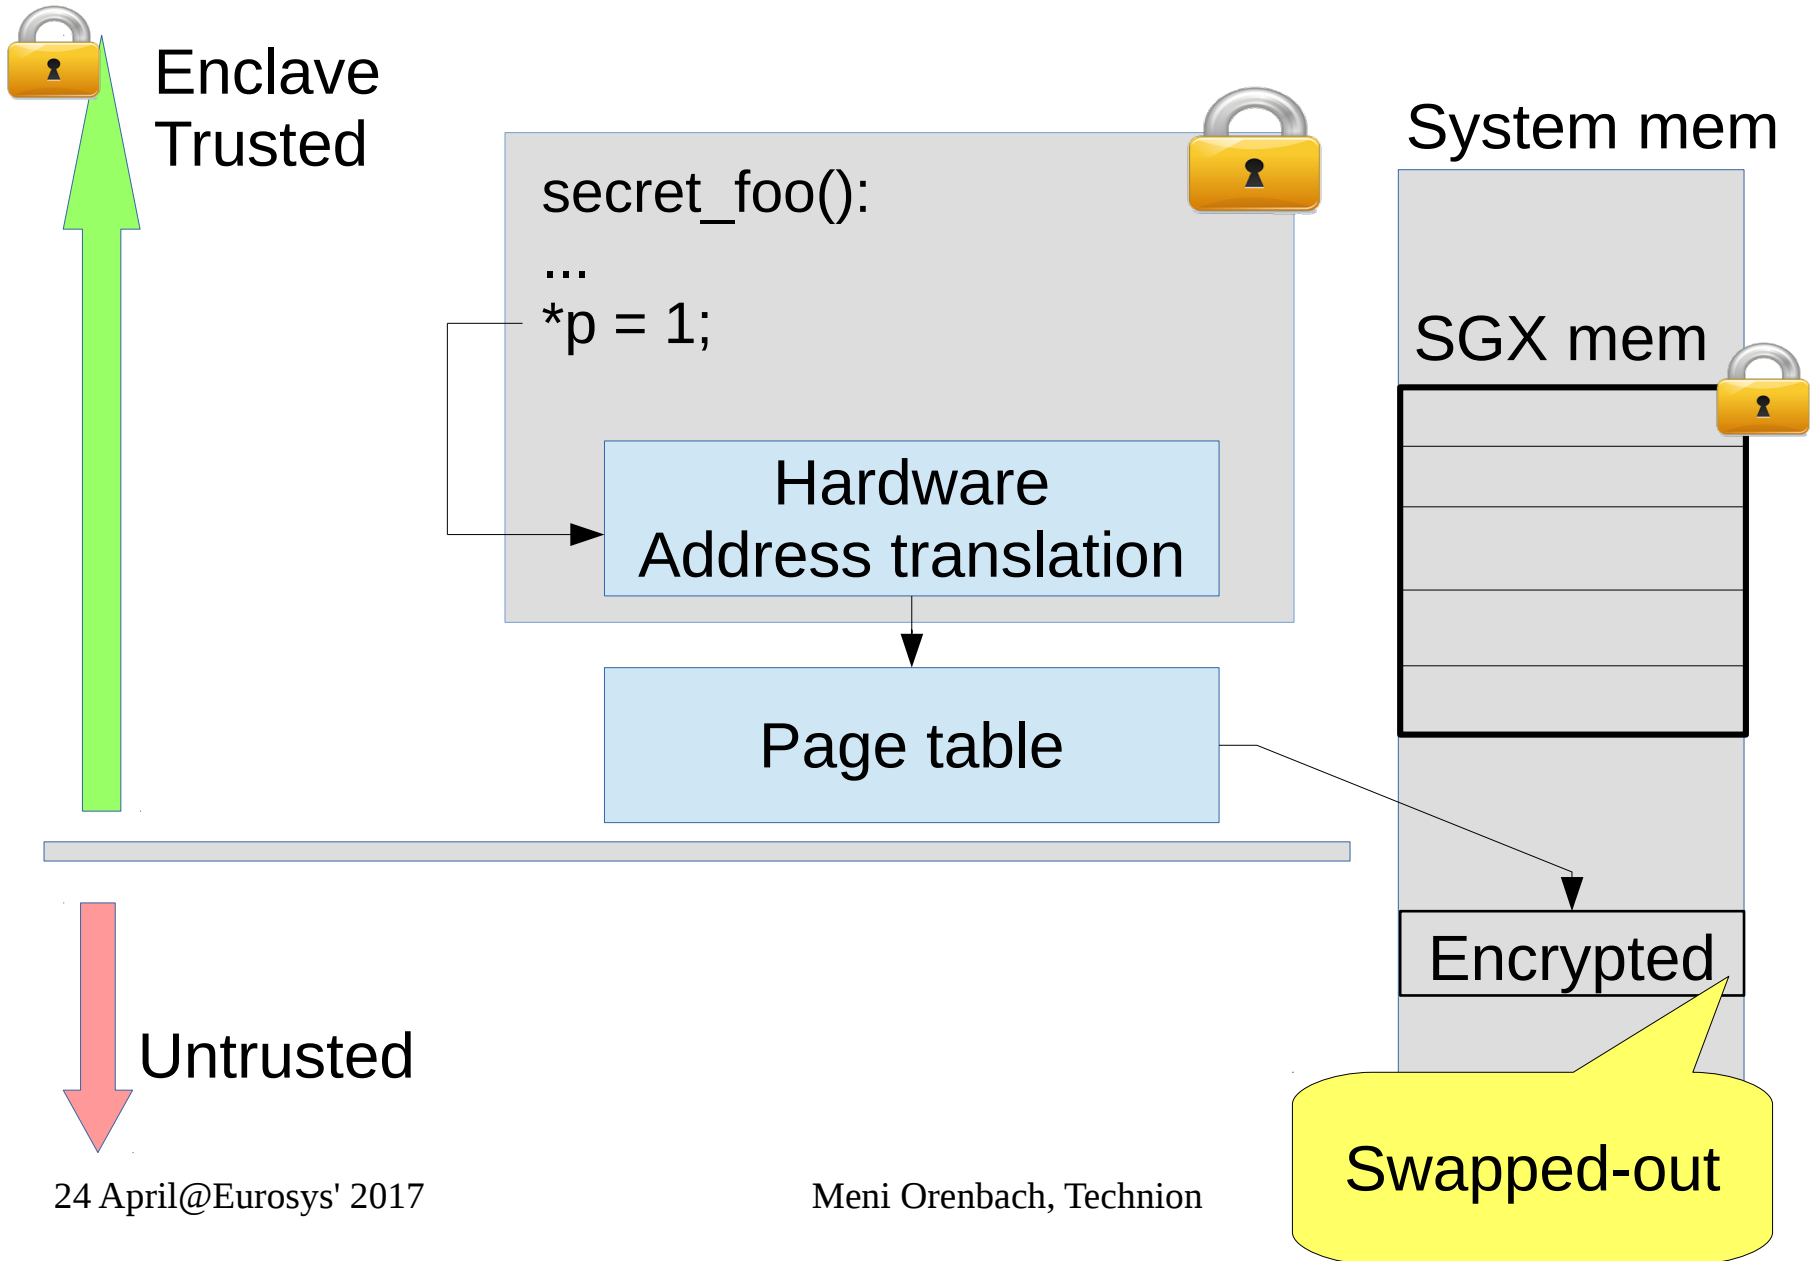
Background: SGX paging



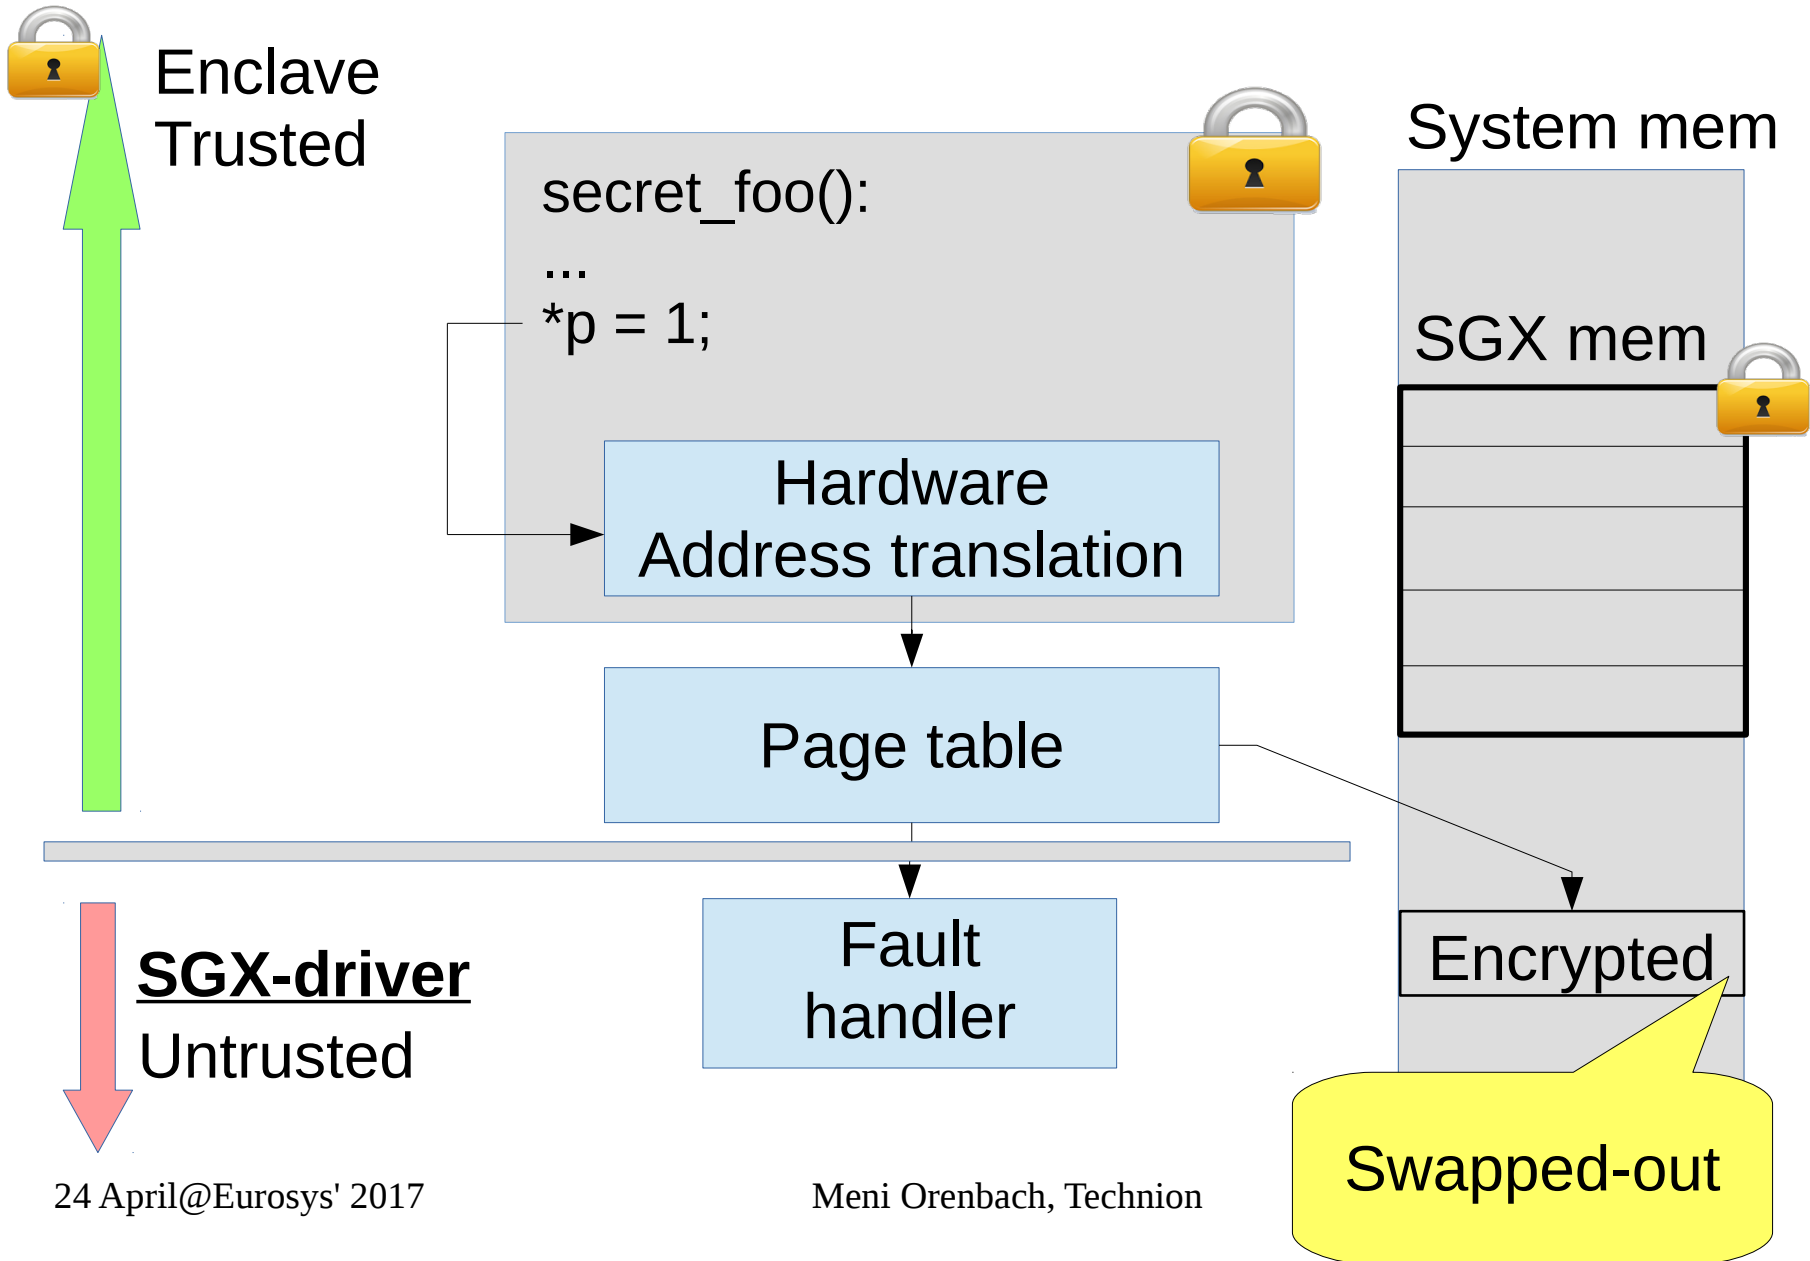
Background: SGX paging



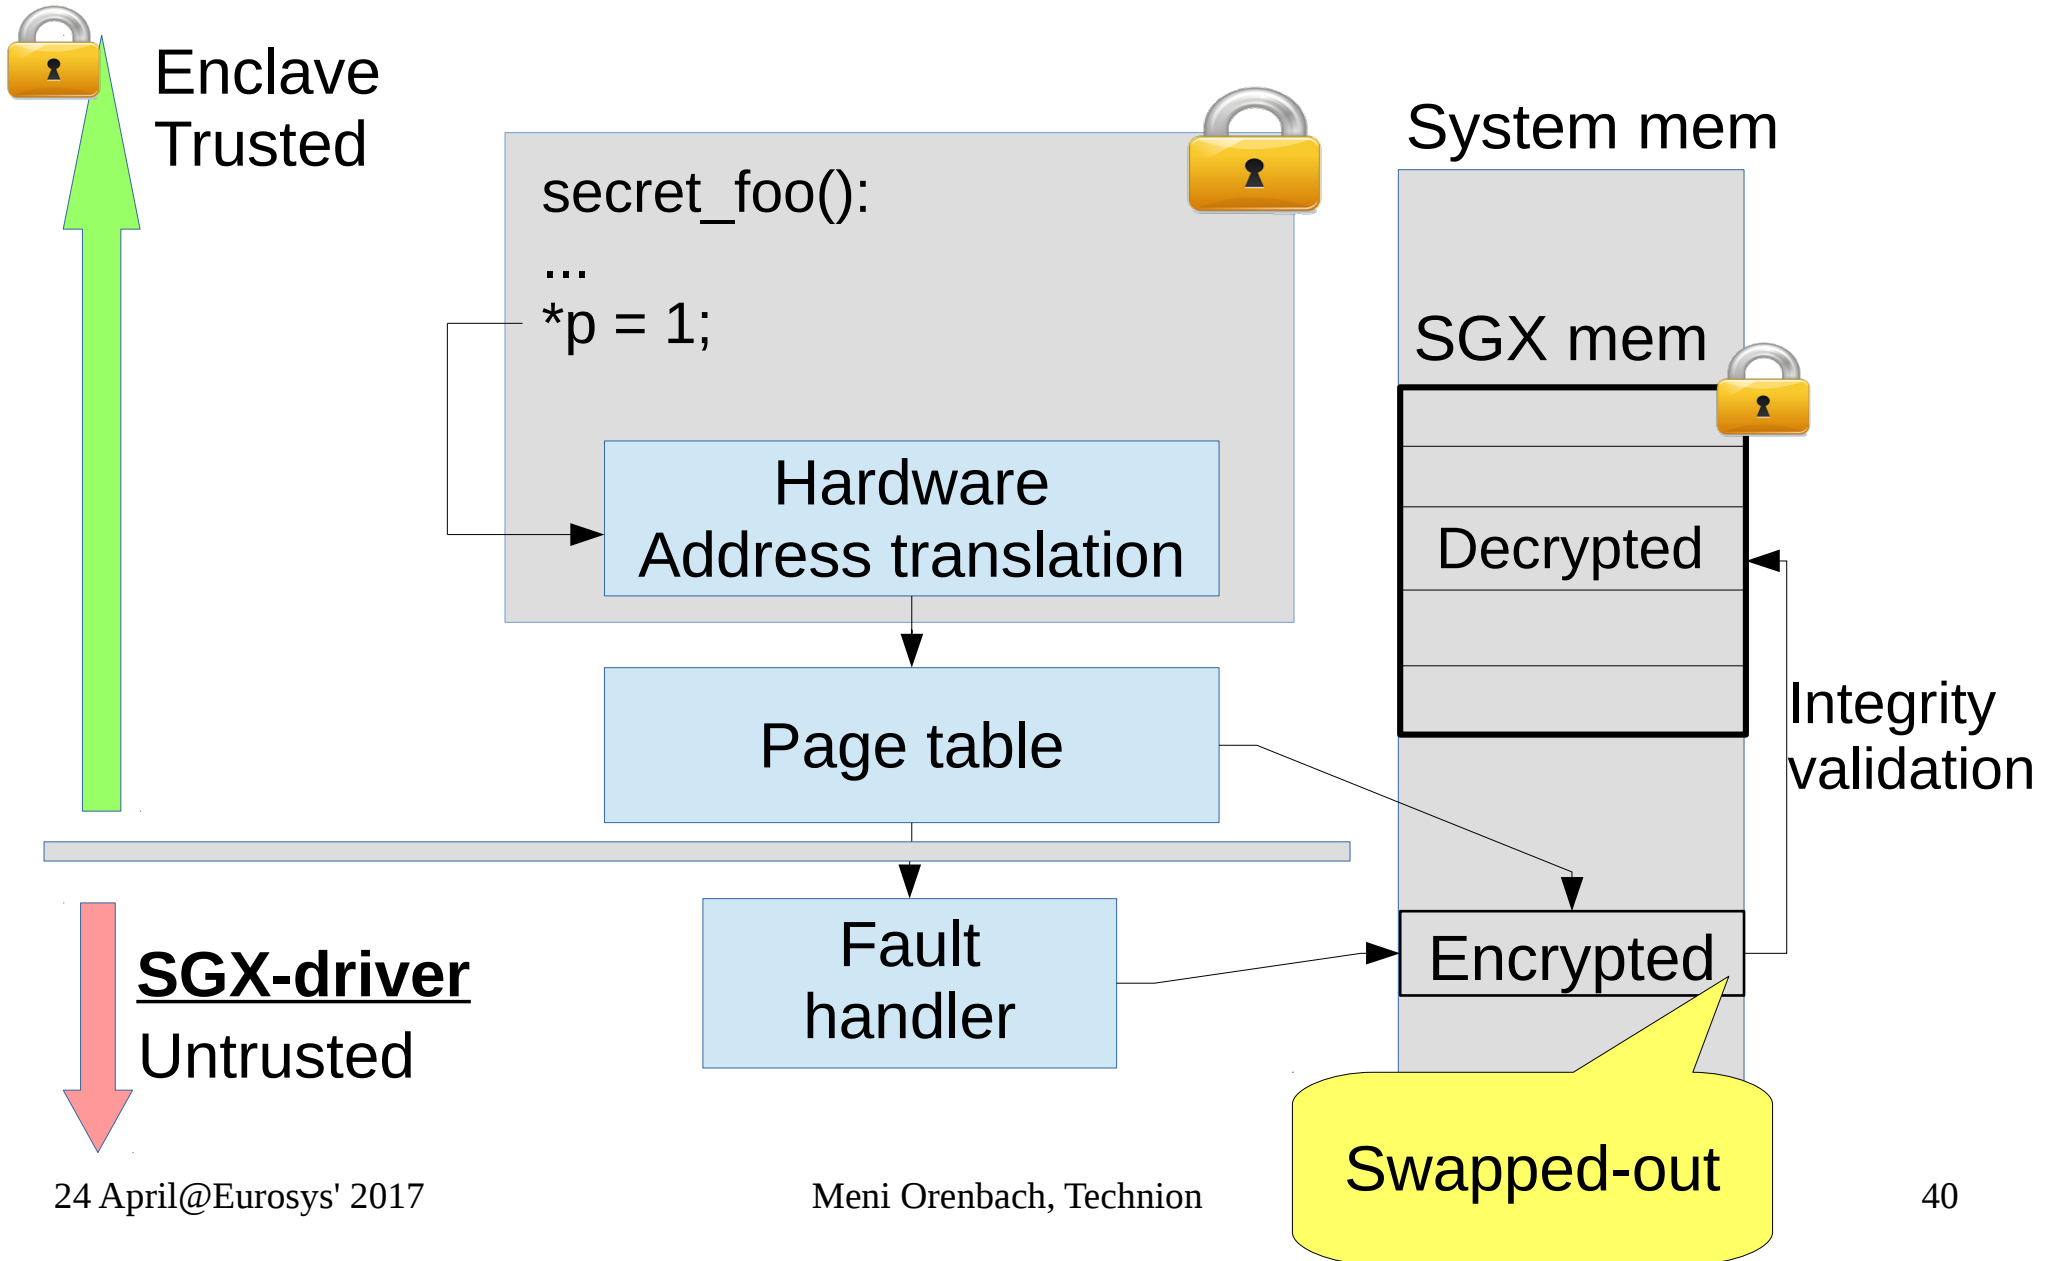
Background: SGX paging



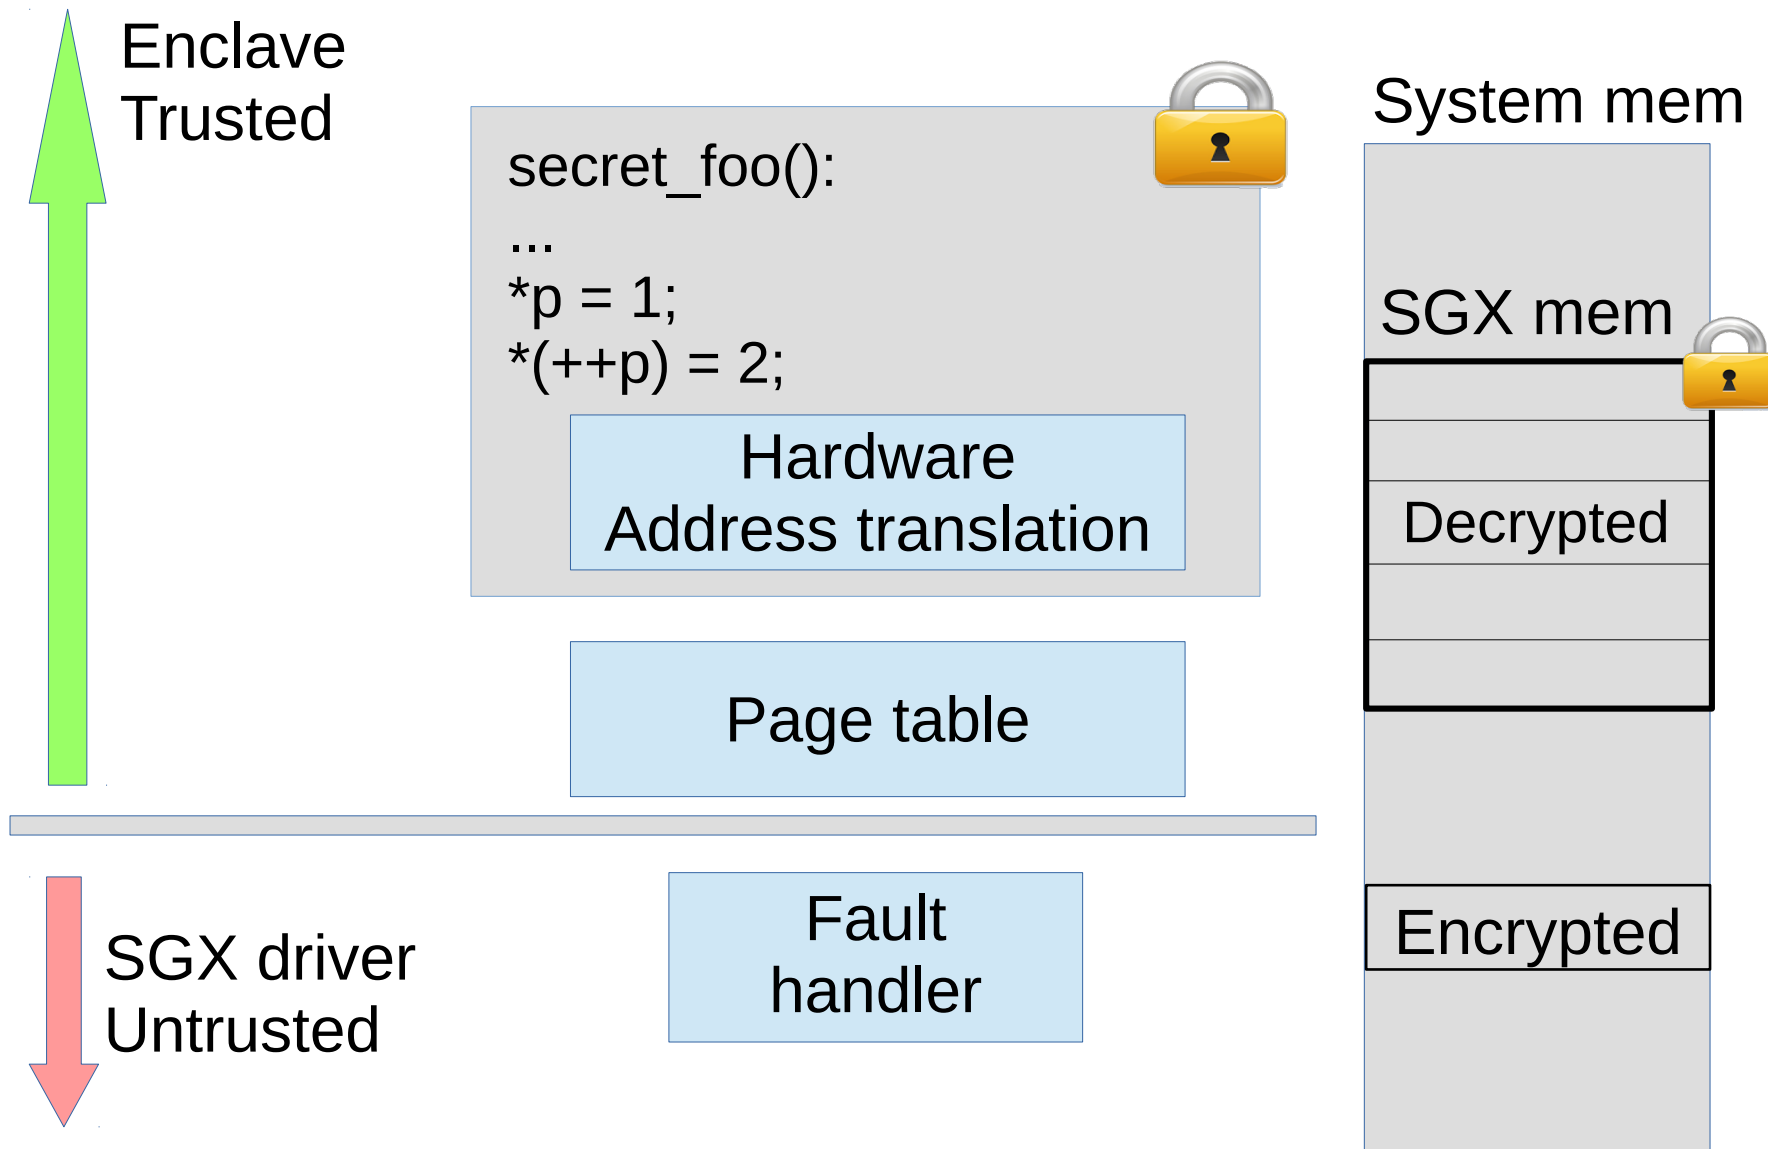
Background: SGX paging



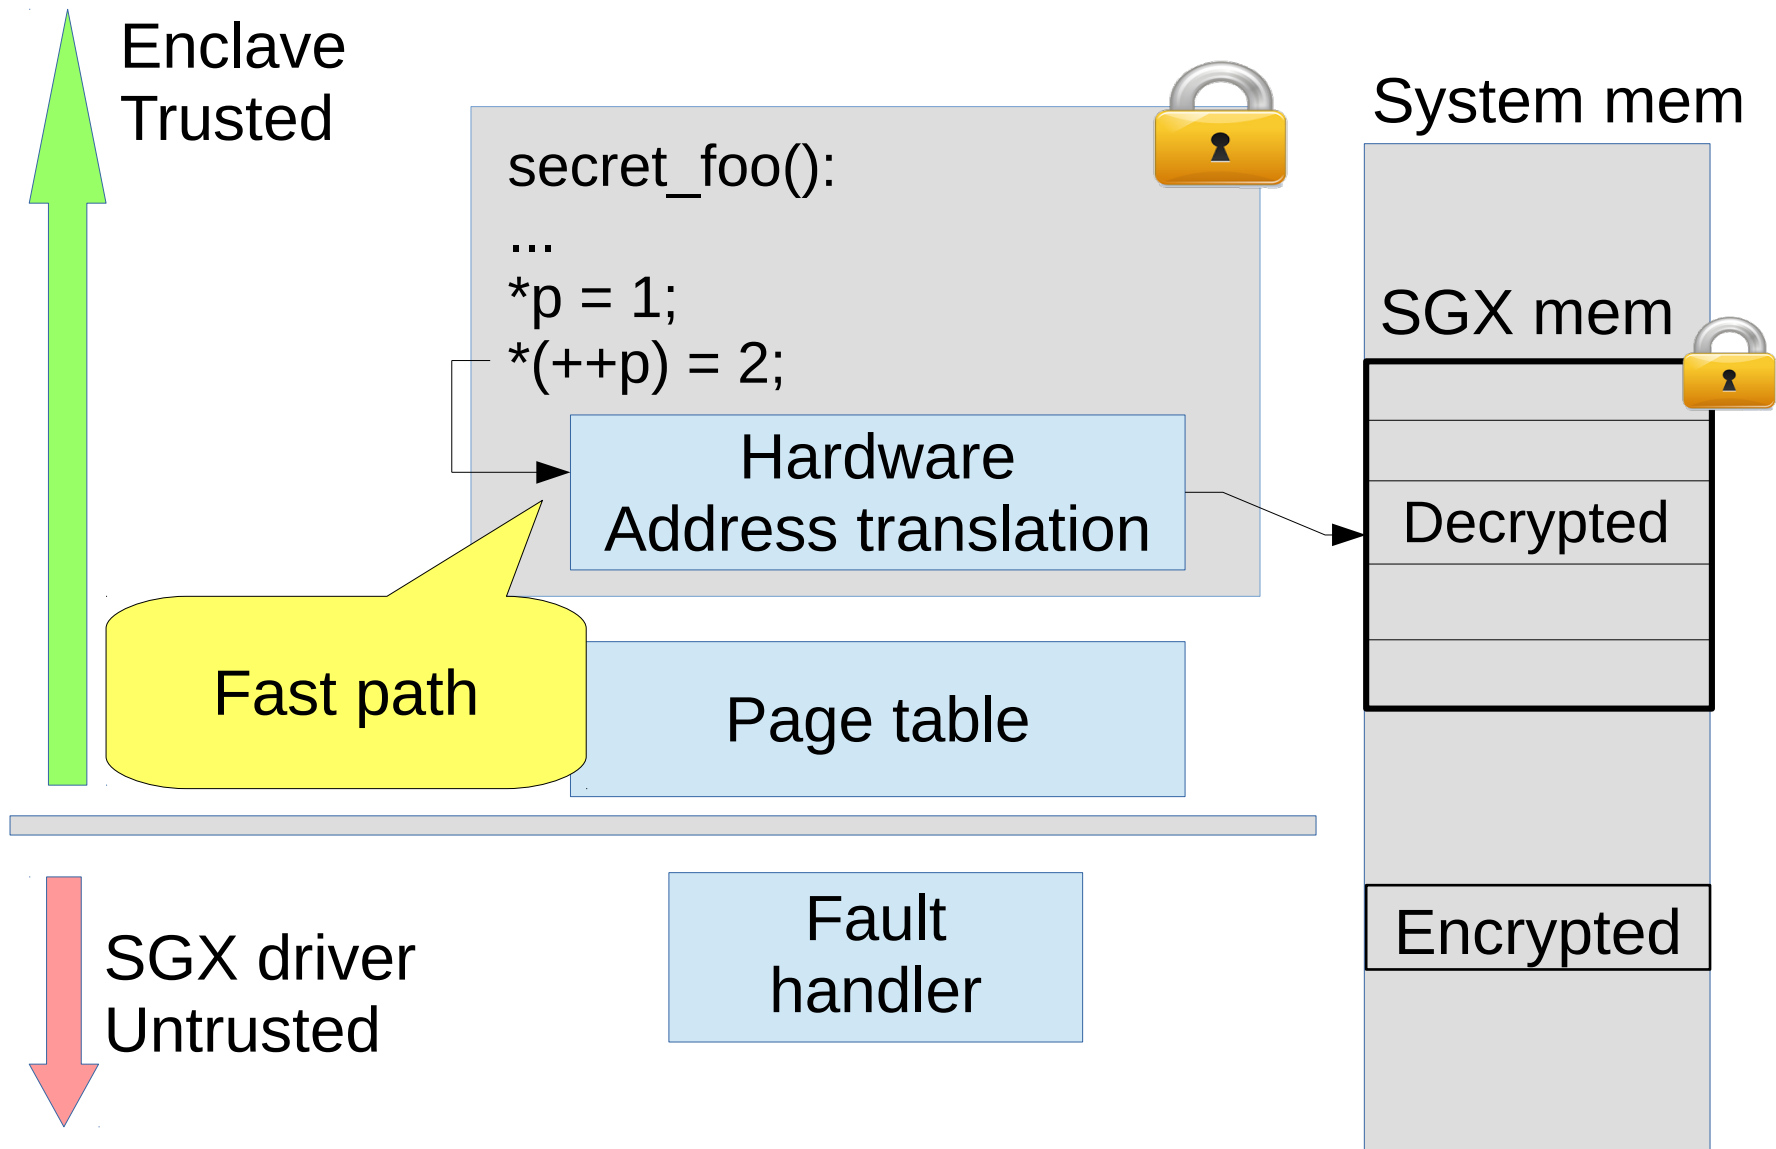
Background: SGX paging



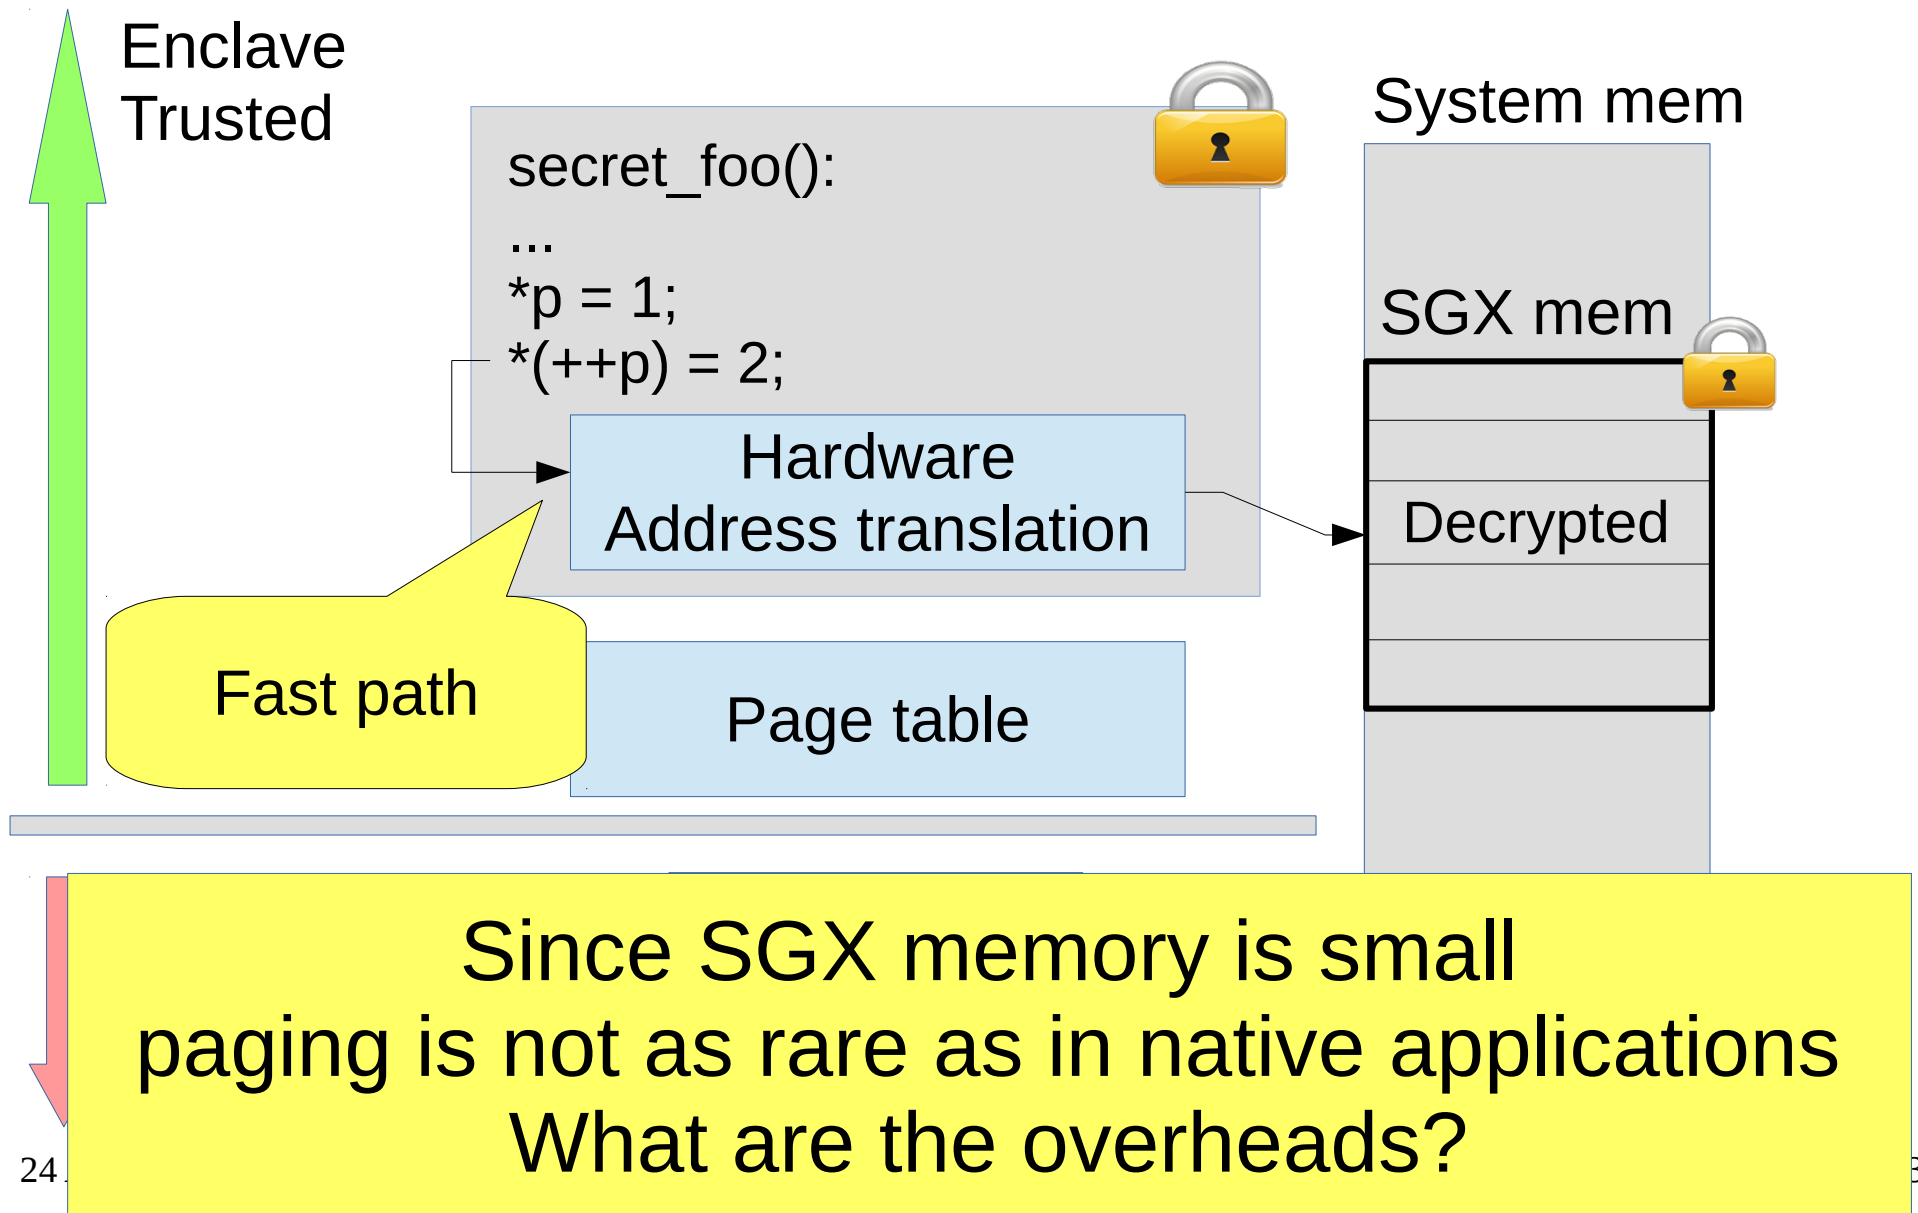
Background: SGX paging



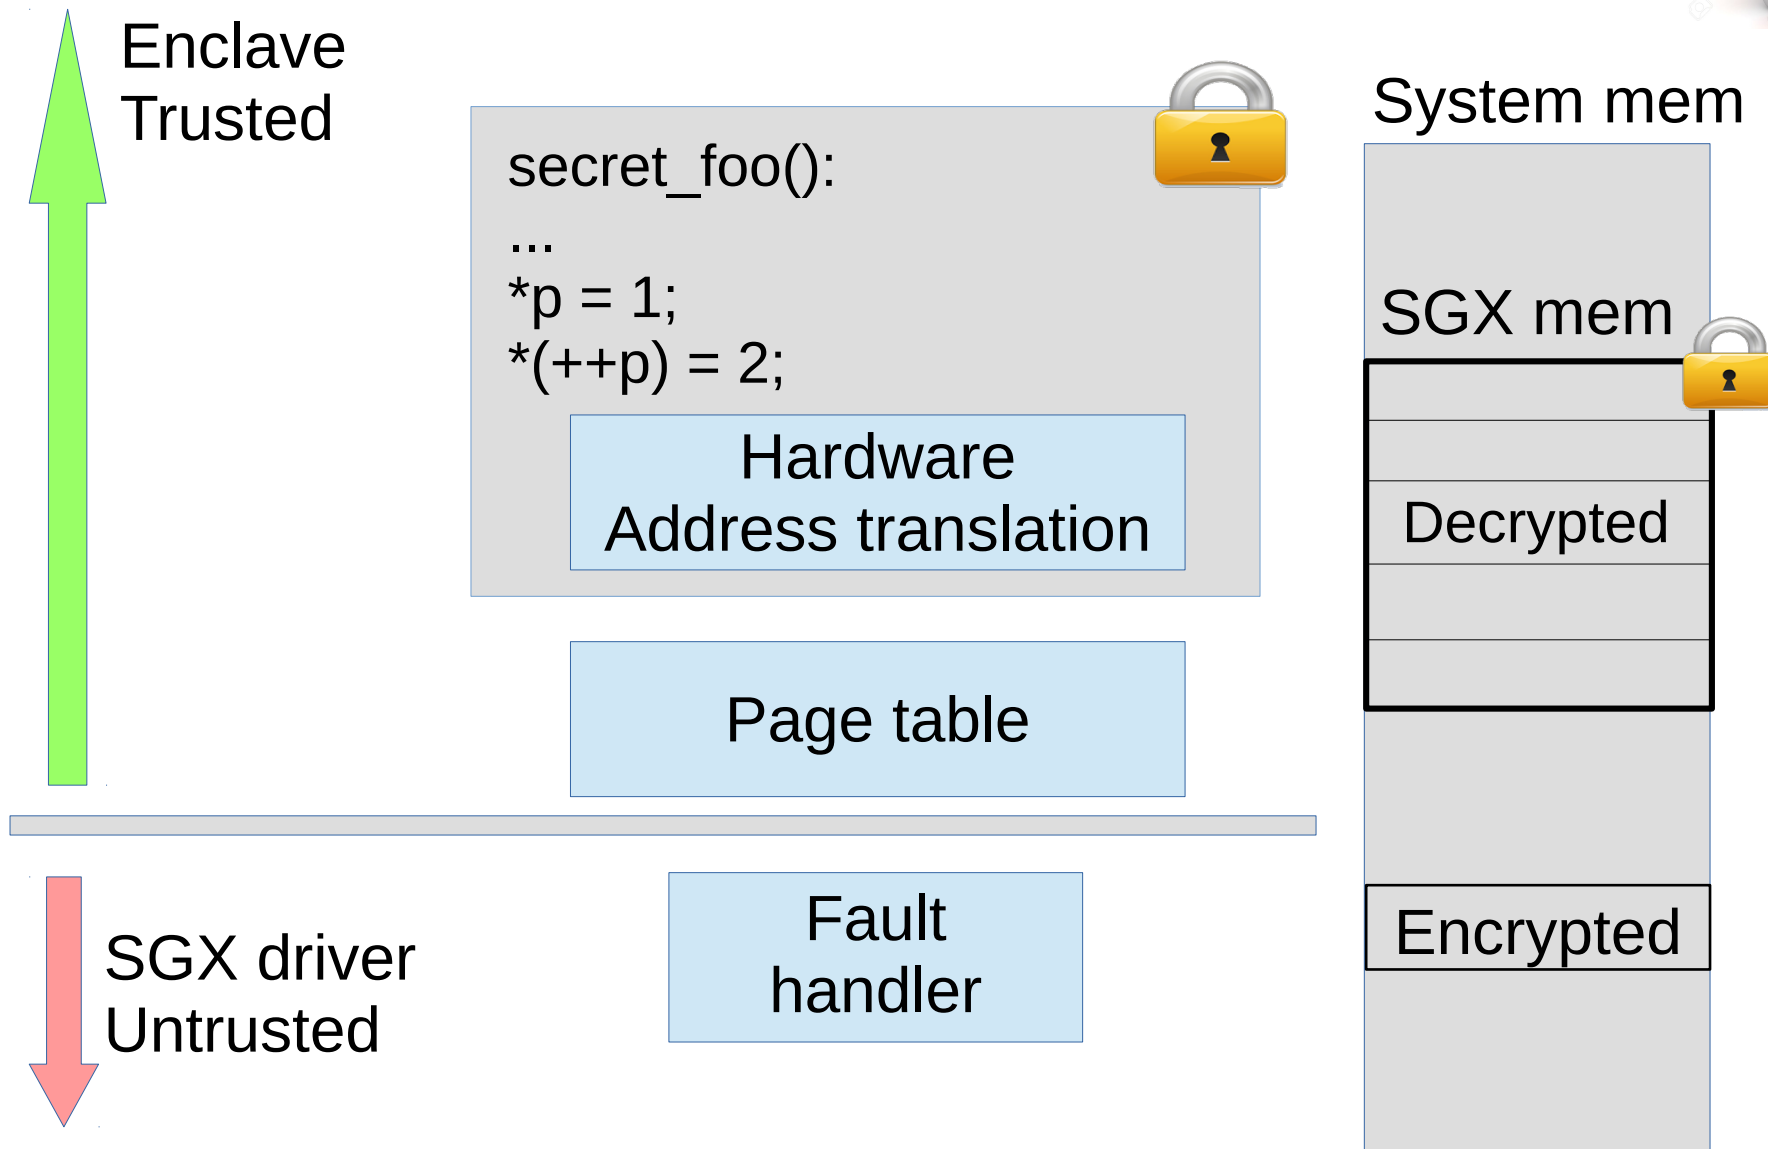
Background: SGX paging



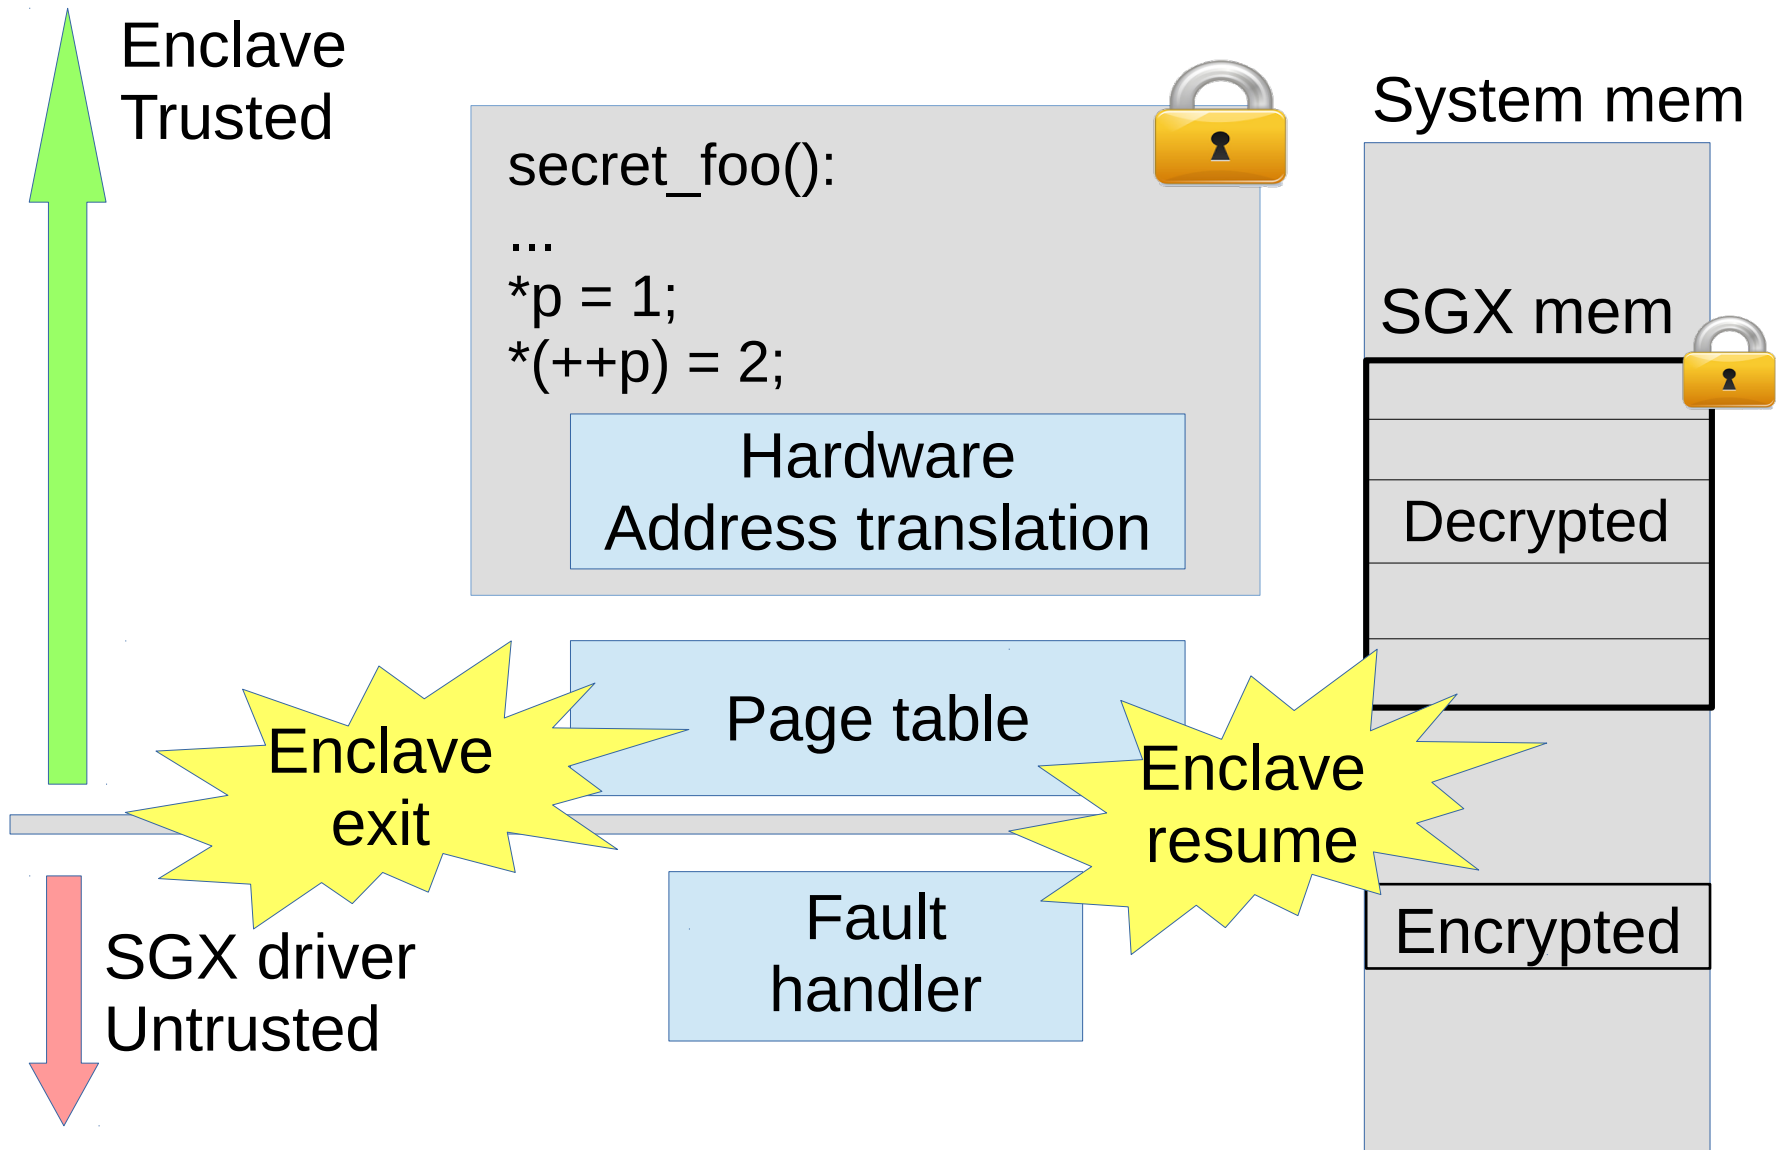
Background: SGX paging



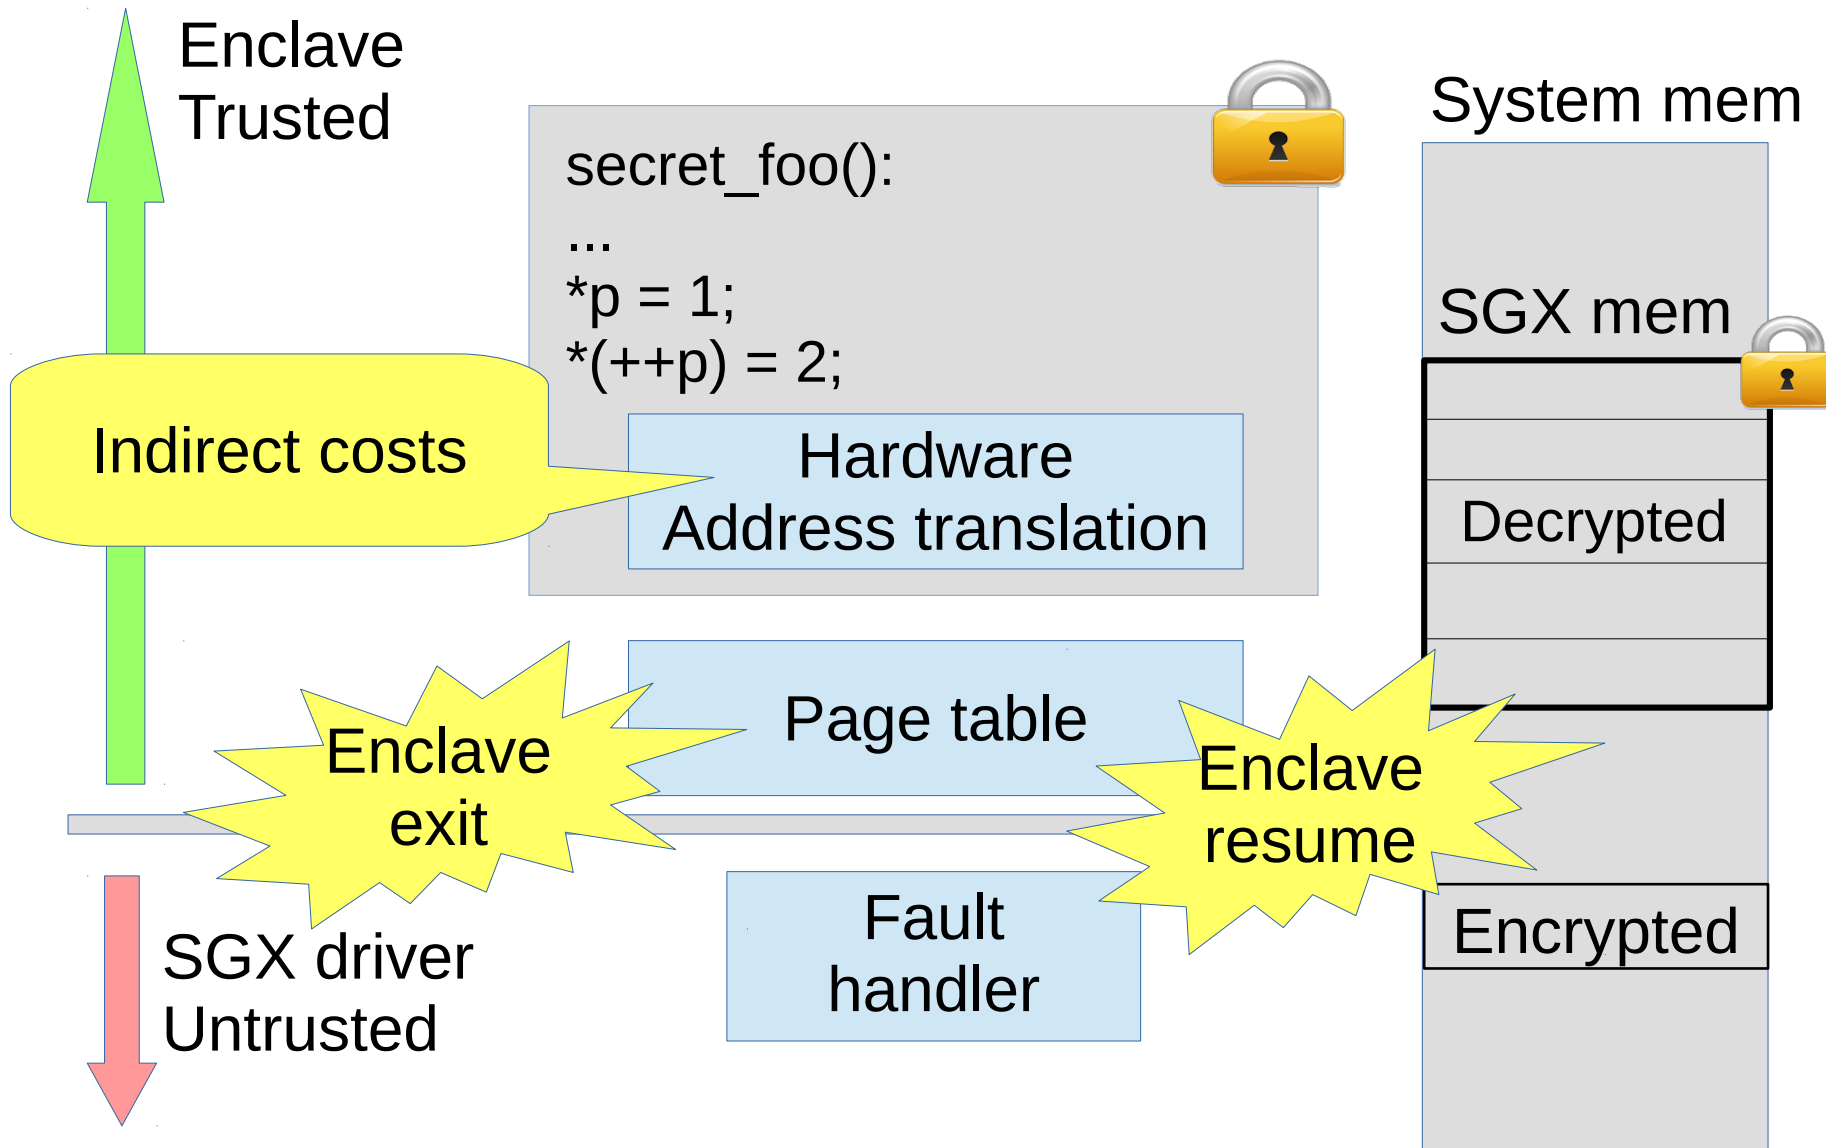
Background: SGX paging



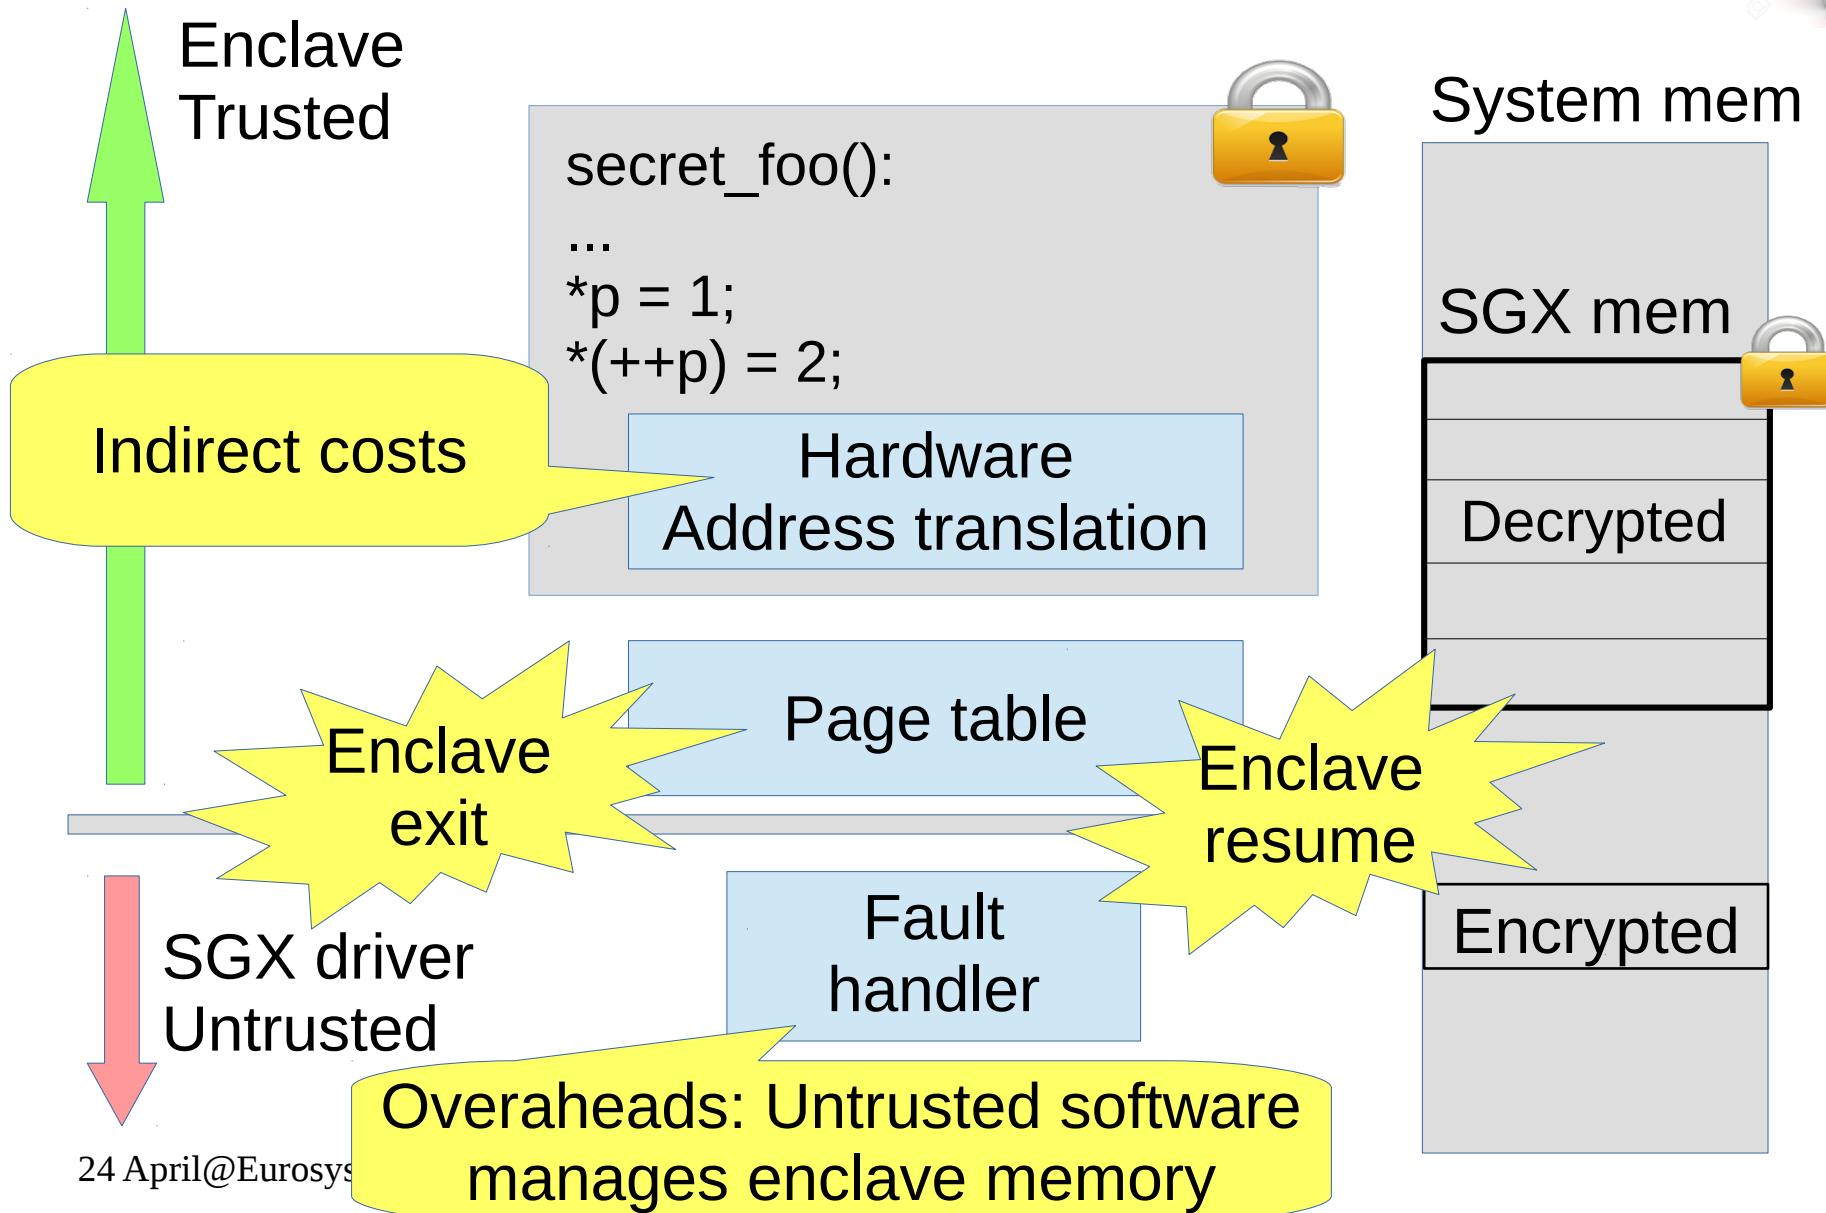

SGX paging overheads



SGX paging overheads



SGX paging overheads

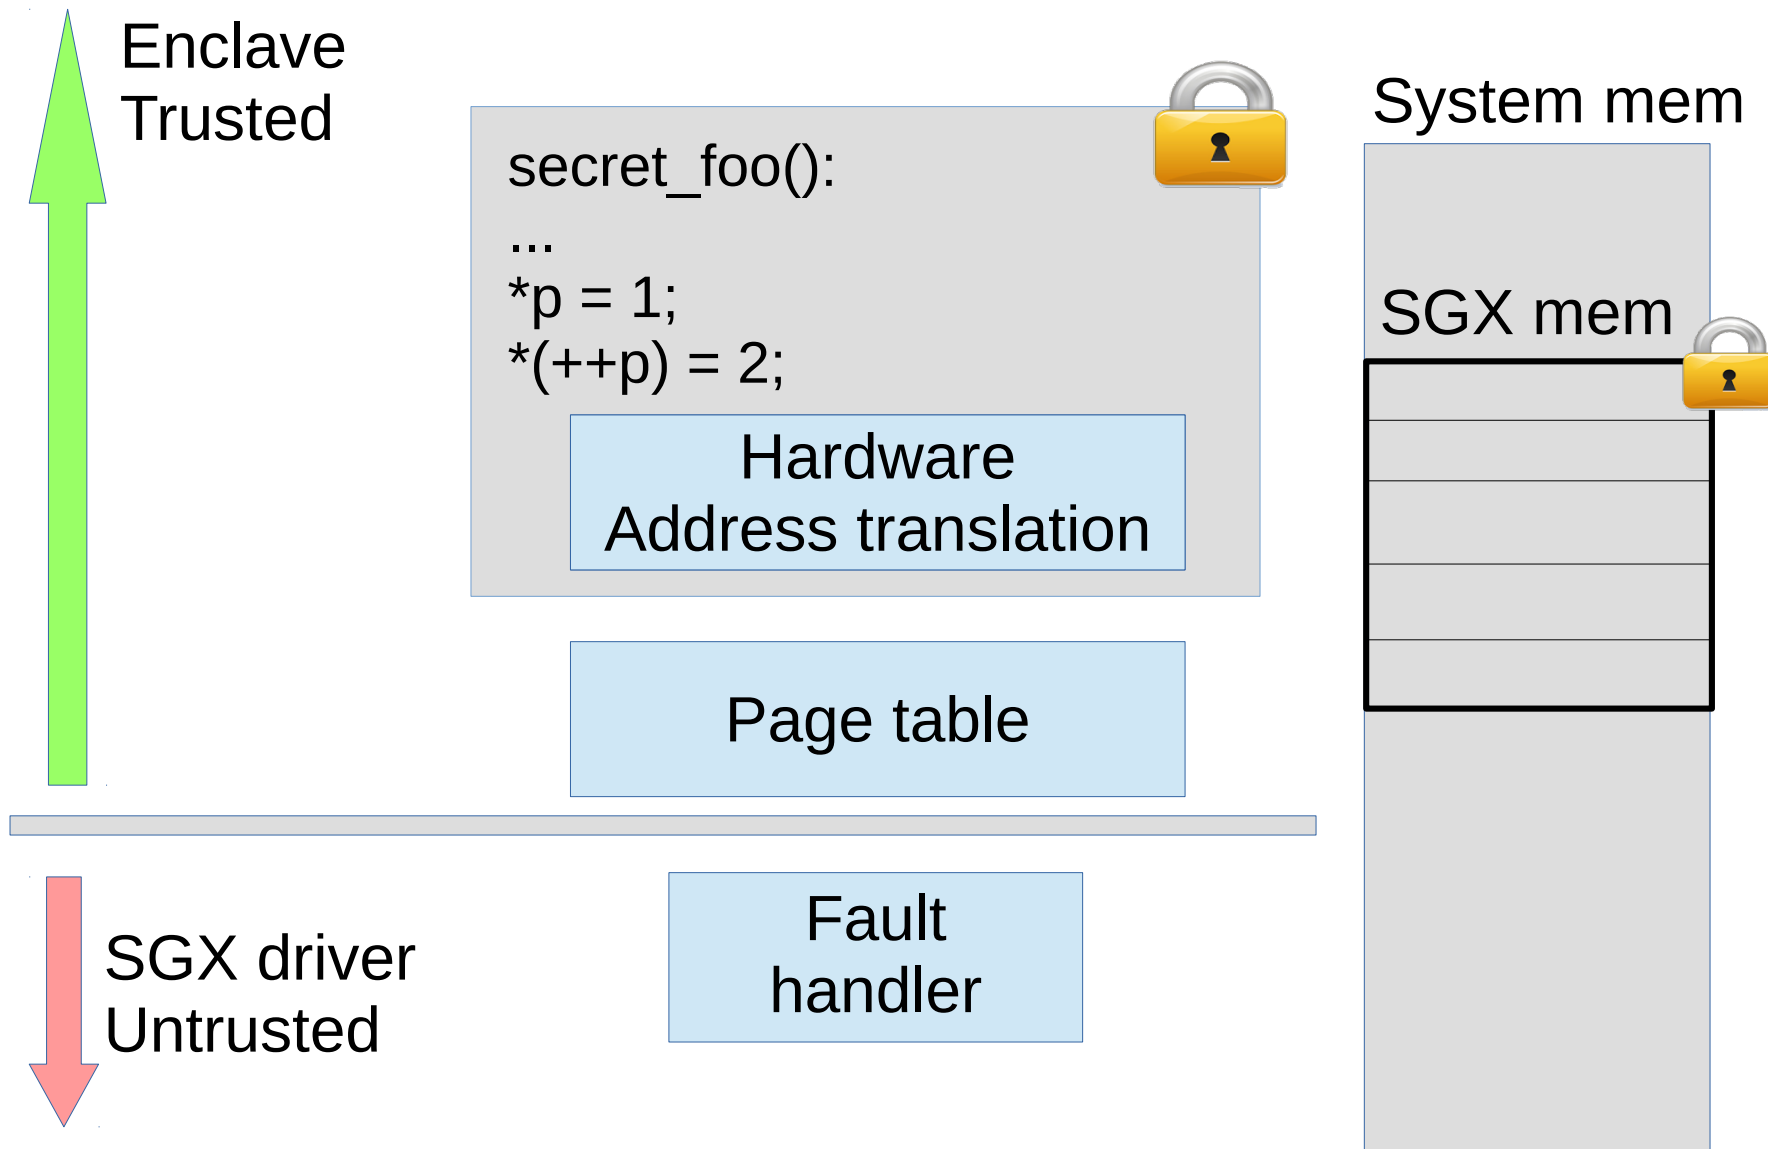


Wanted: In-enclave virtual memory management

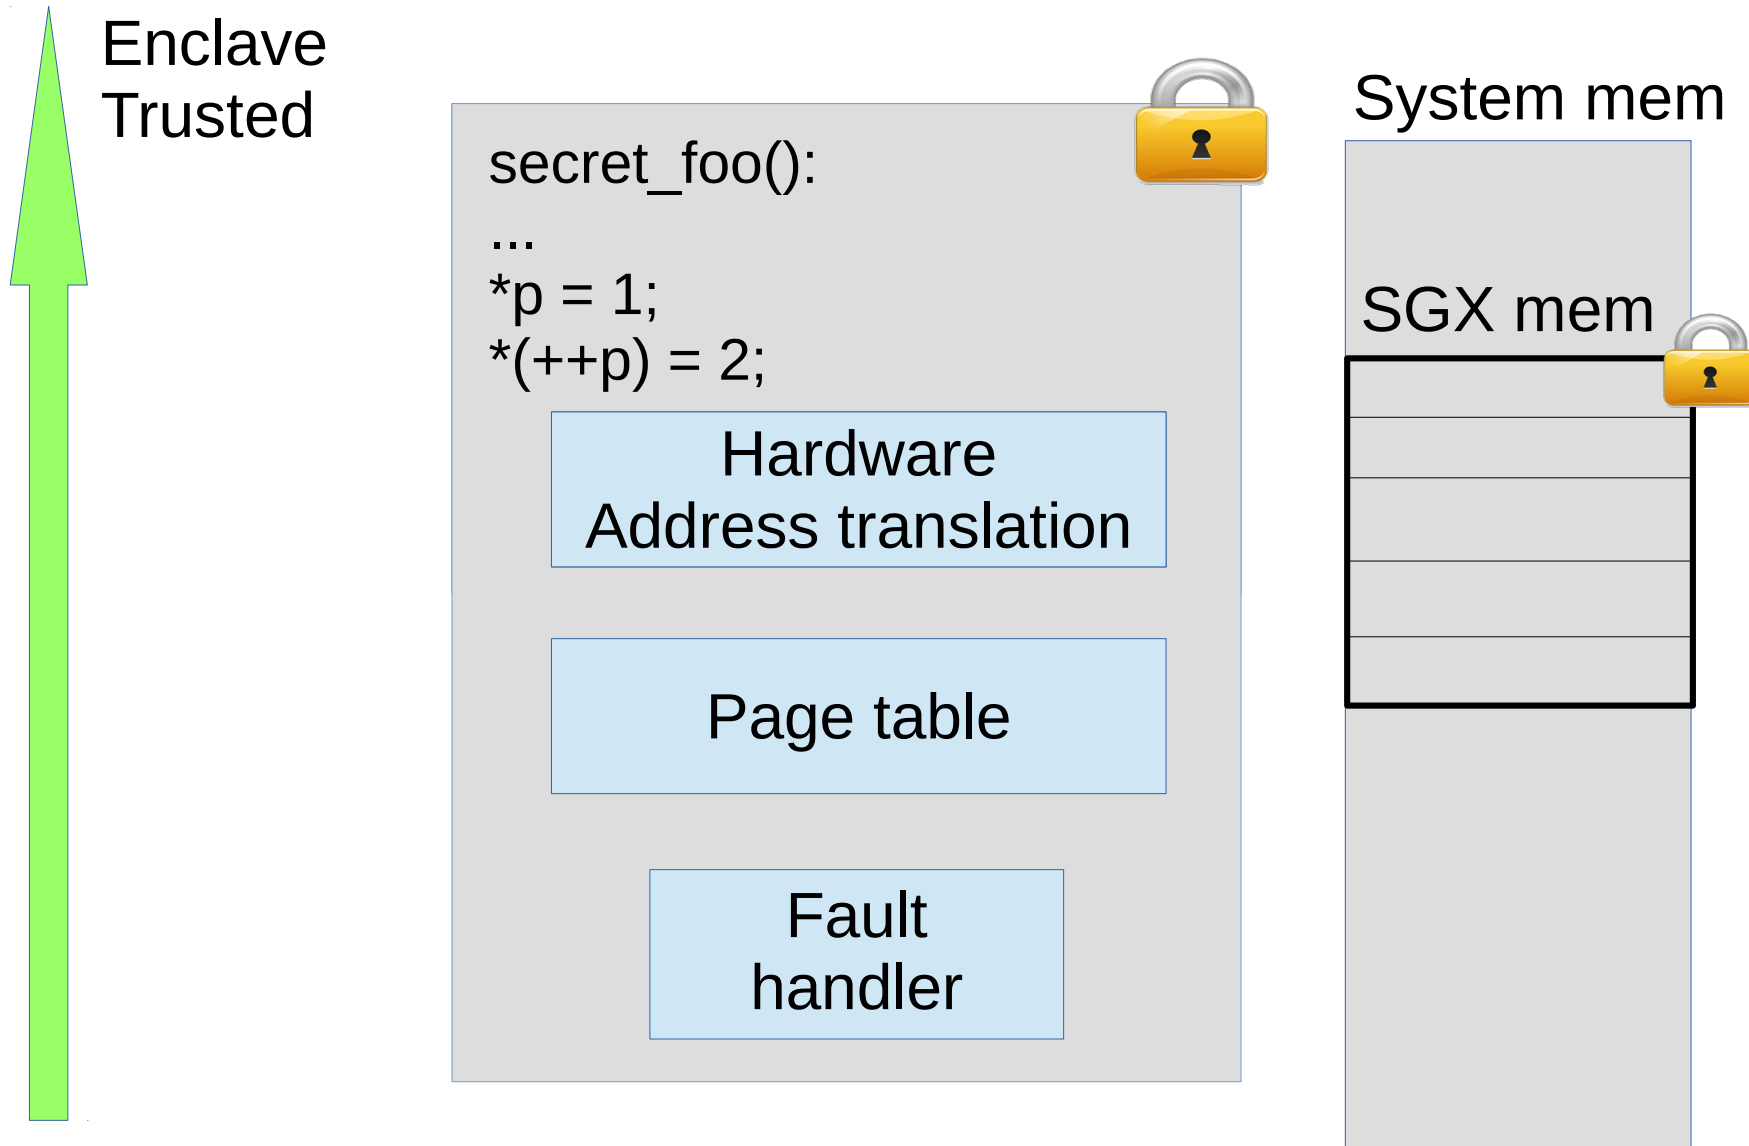


No more exits!

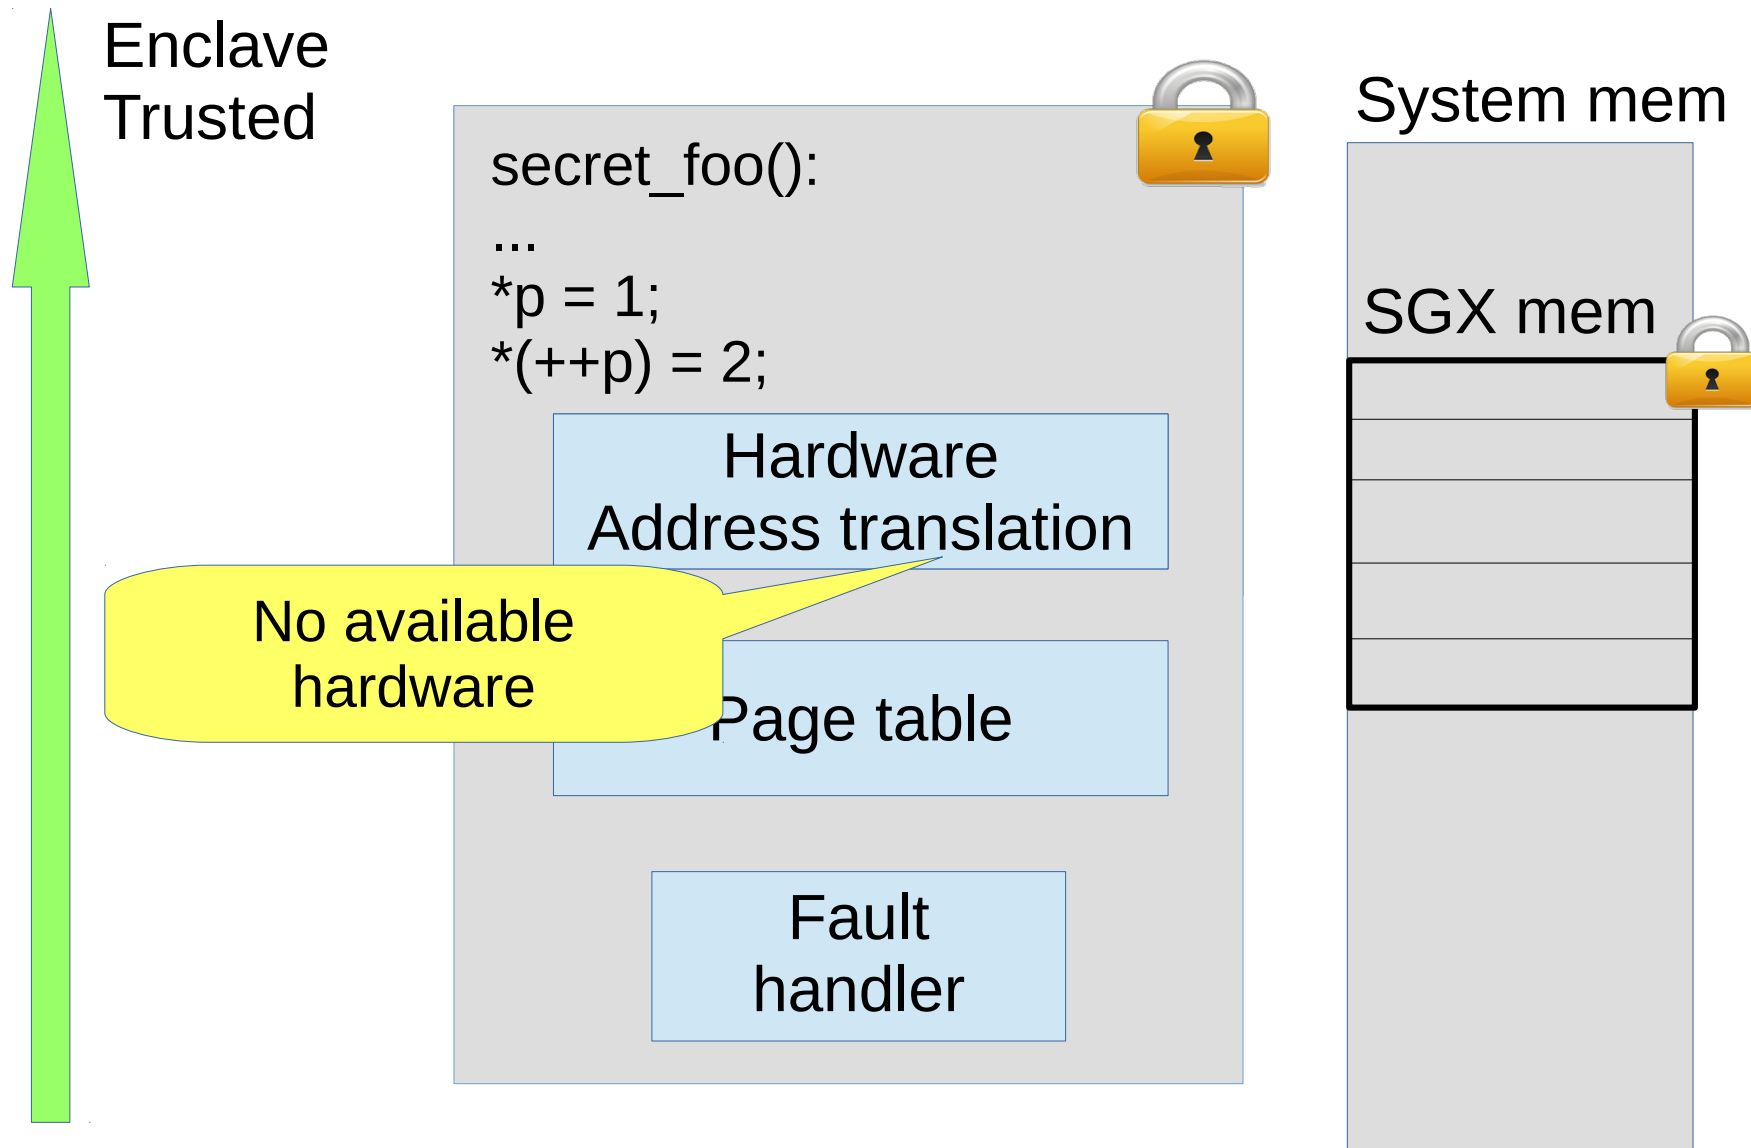
Ideal in-enclave VM management



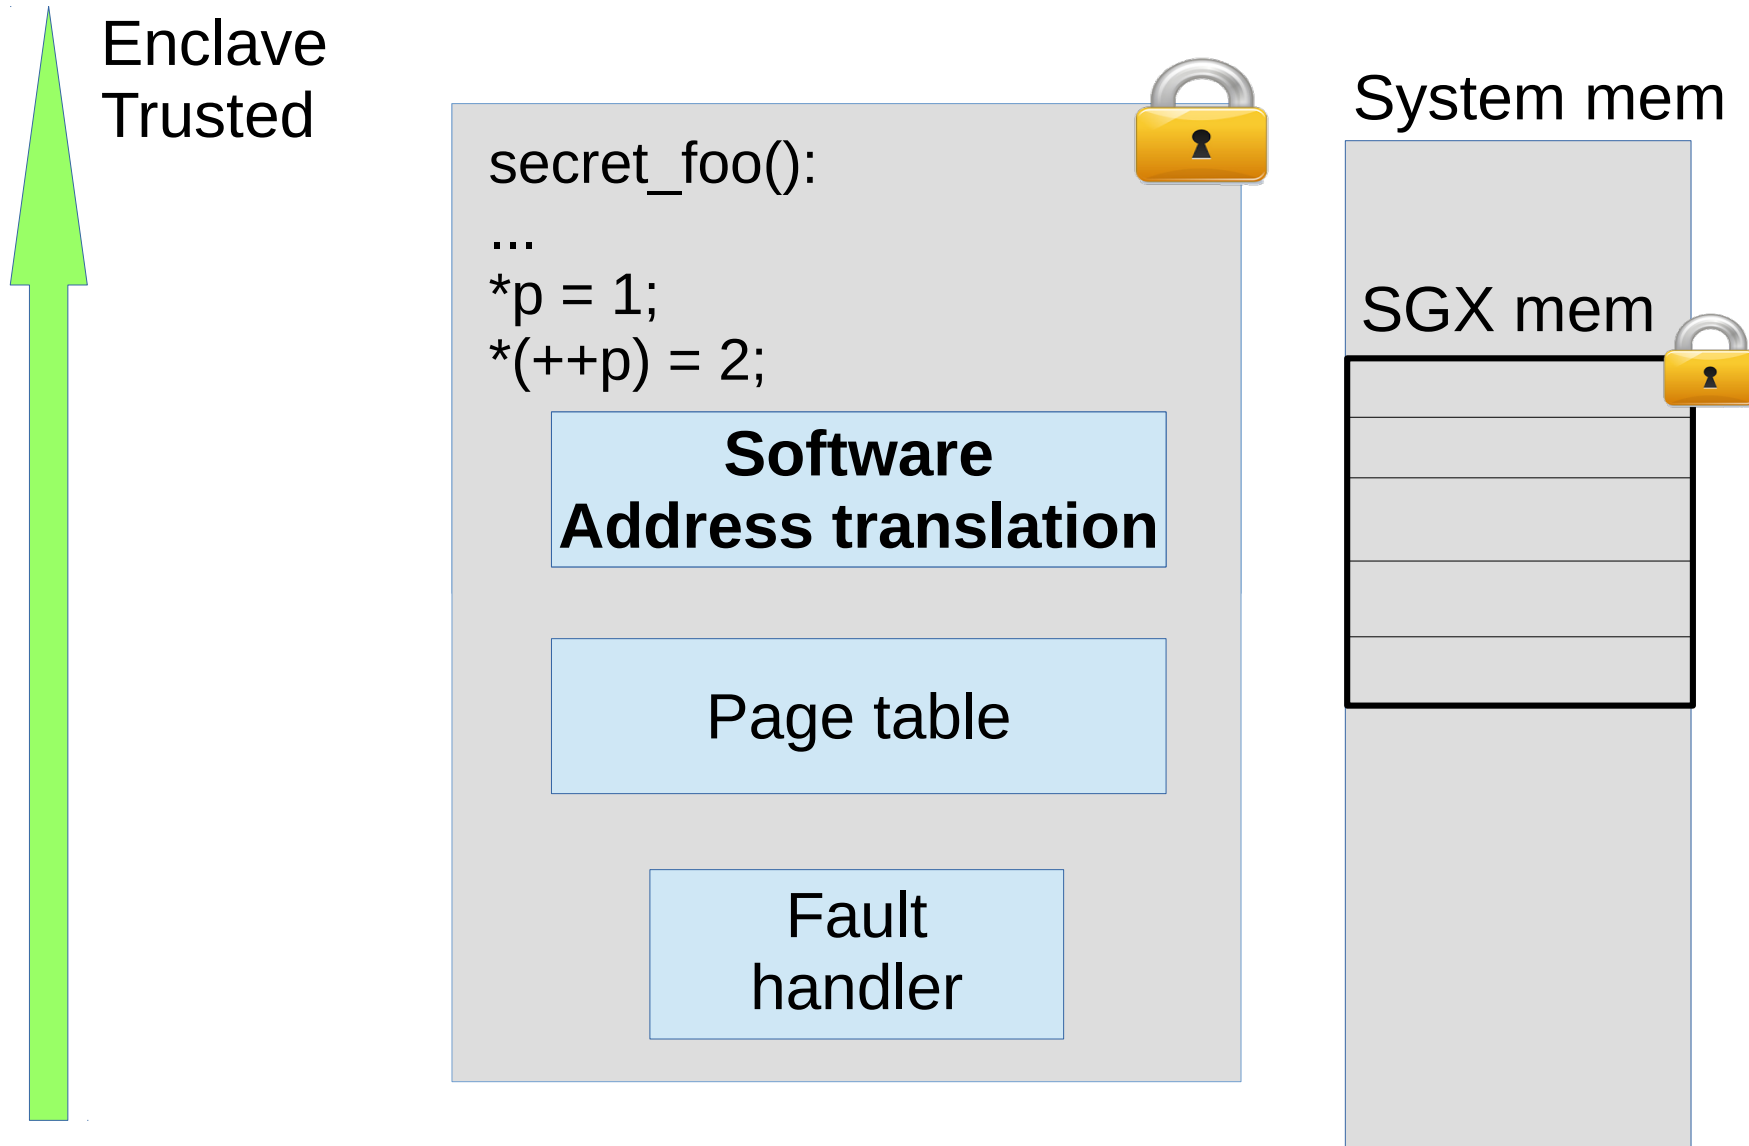
Ideal in-enclave VM management



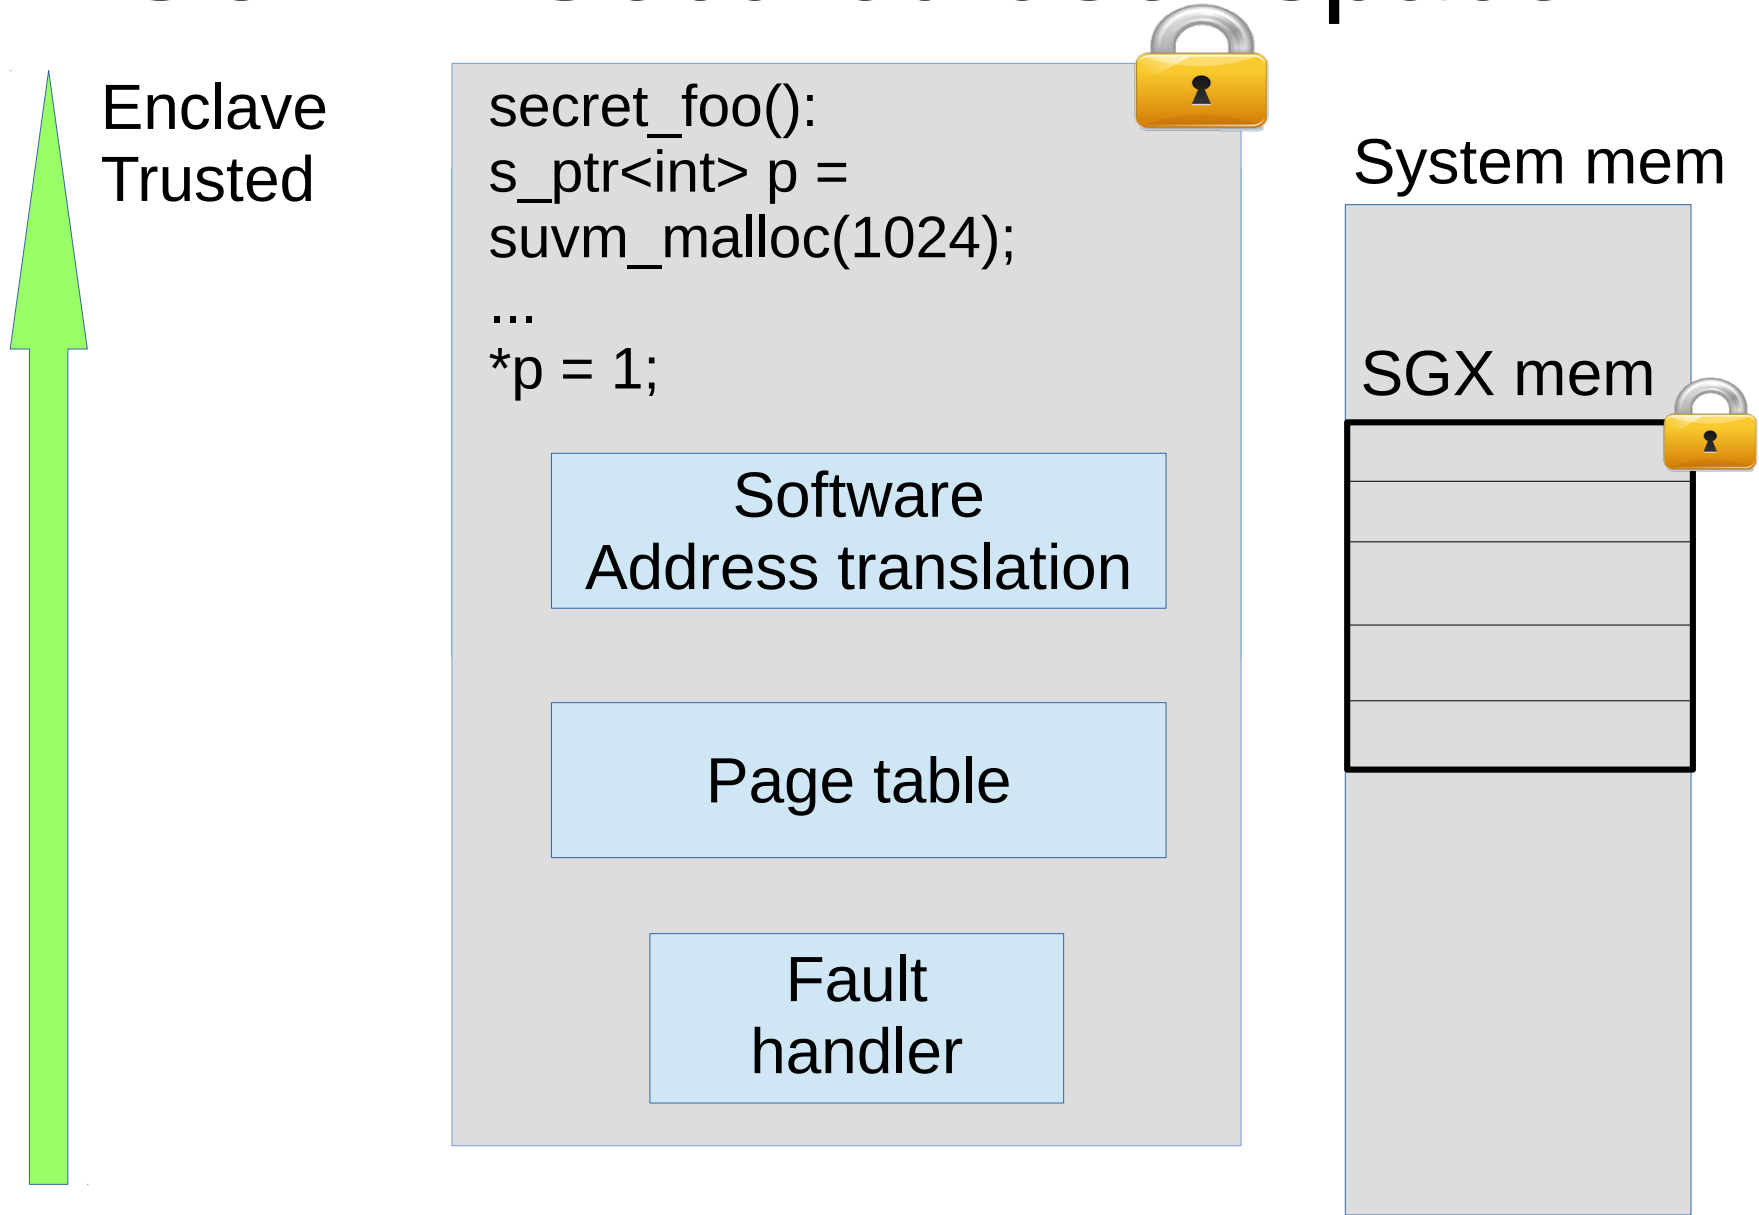
Ideal in-enclave VM management



Ideal in-enclave VM management



SUVM: Secured user-space VM



SUVM: Secured user-space VM

Enclave
Trusted

```
secret_foo():  
s_ptr<int> p =  
suvm_malloc(1024);  
...  
*p = 1;
```

Software
Address translation

Page table

Fault
handler



System mem

SGX mem



Template class:
SecuredPointer.

SUVM: Secured user-space VM

Enclave
Trusted

```
secret_foo():  
s_ptr<int> p =  
suvm_malloc(1024);  
...  
*p = 1;
```



Software
Address translation

Page table

Fault
handler

System mem

SGX mem



Encrypted

Template class:
SecuredPointer.

SUVM: Secured user-space VM

Enclave
Trusted

```
secret_foo():  
s_ptr<int> p =  
suvm_malloc(1024);  
...  
*p = 1;
```

Software
Address translation

Page table

Fault
handler

System mem

SGX mem

Encrypted

Template class:
SecuredPointer.

Swapped-out

SUVM: Secured user-space VM

Enclave
Trusted

```
secret_foo():  
s_ptr<int> p =  
suvm_malloc(1024);  
...  
*p = 1;
```

Software
Address translation

Page table

Fault
handler

System mem

SGX mem

Encrypted

Template class:
SecuredPointer.

Swapped-out

SUVM: Secured user-space VM

Enclave
Trusted

```
secret_foo():  
s_ptr<int> p =  
suvm_malloc(1024);  
...  
*p = 1;
```

Software
Address translation

Page table

Fault
handler

System mem

SGX mem

Decrypted

Encrypted

Integrity
validation

Template class:
SecuredPointer.

Swapped-out

SUVM: Secured user-space VM

Enclave
Trusted

```
secret_foo():  
s_ptr<int> p =  
suvm_malloc(1024);  
...  
*p = 1;
```

Software
Address translation

Page table

Fault
handler

System mem

SGX mem

Decrypted

Encrypted

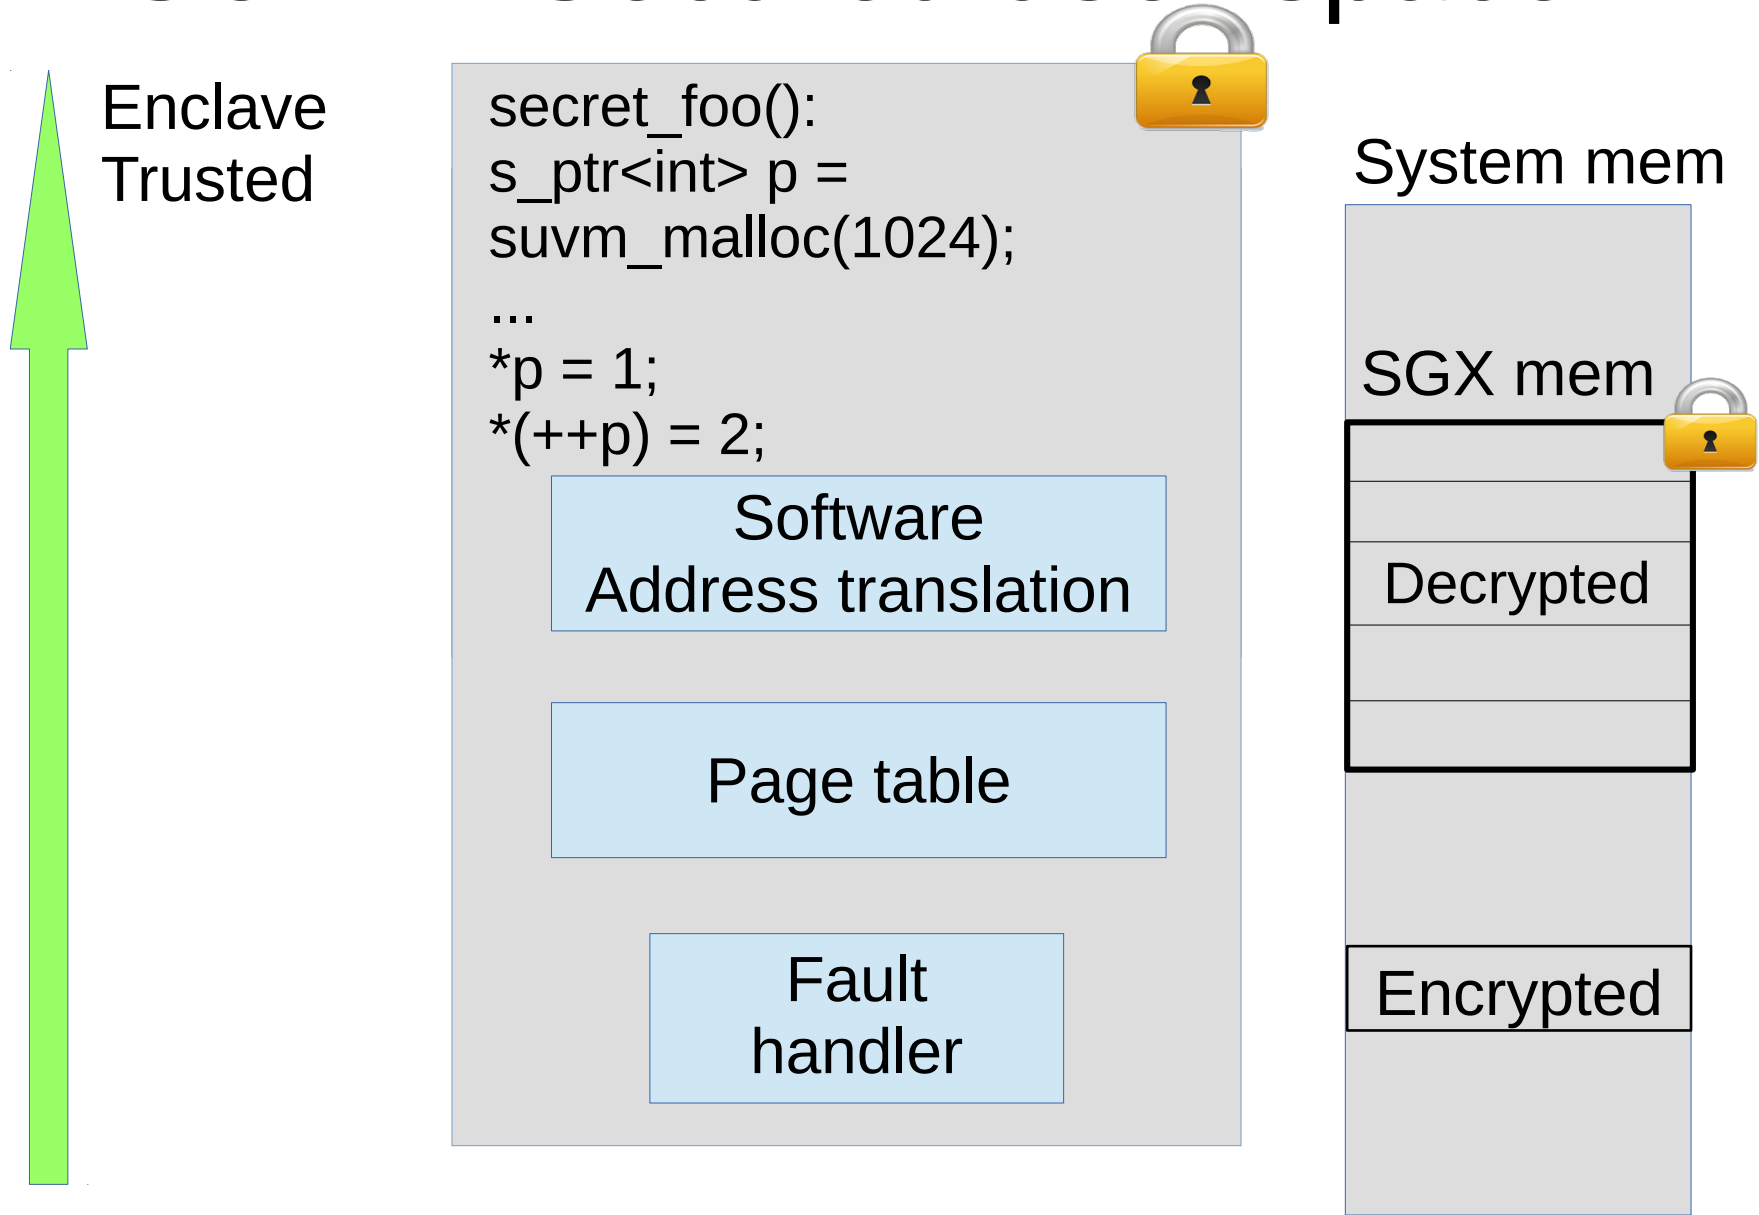
Integrity
validation

Template class:
SecuredPointer.

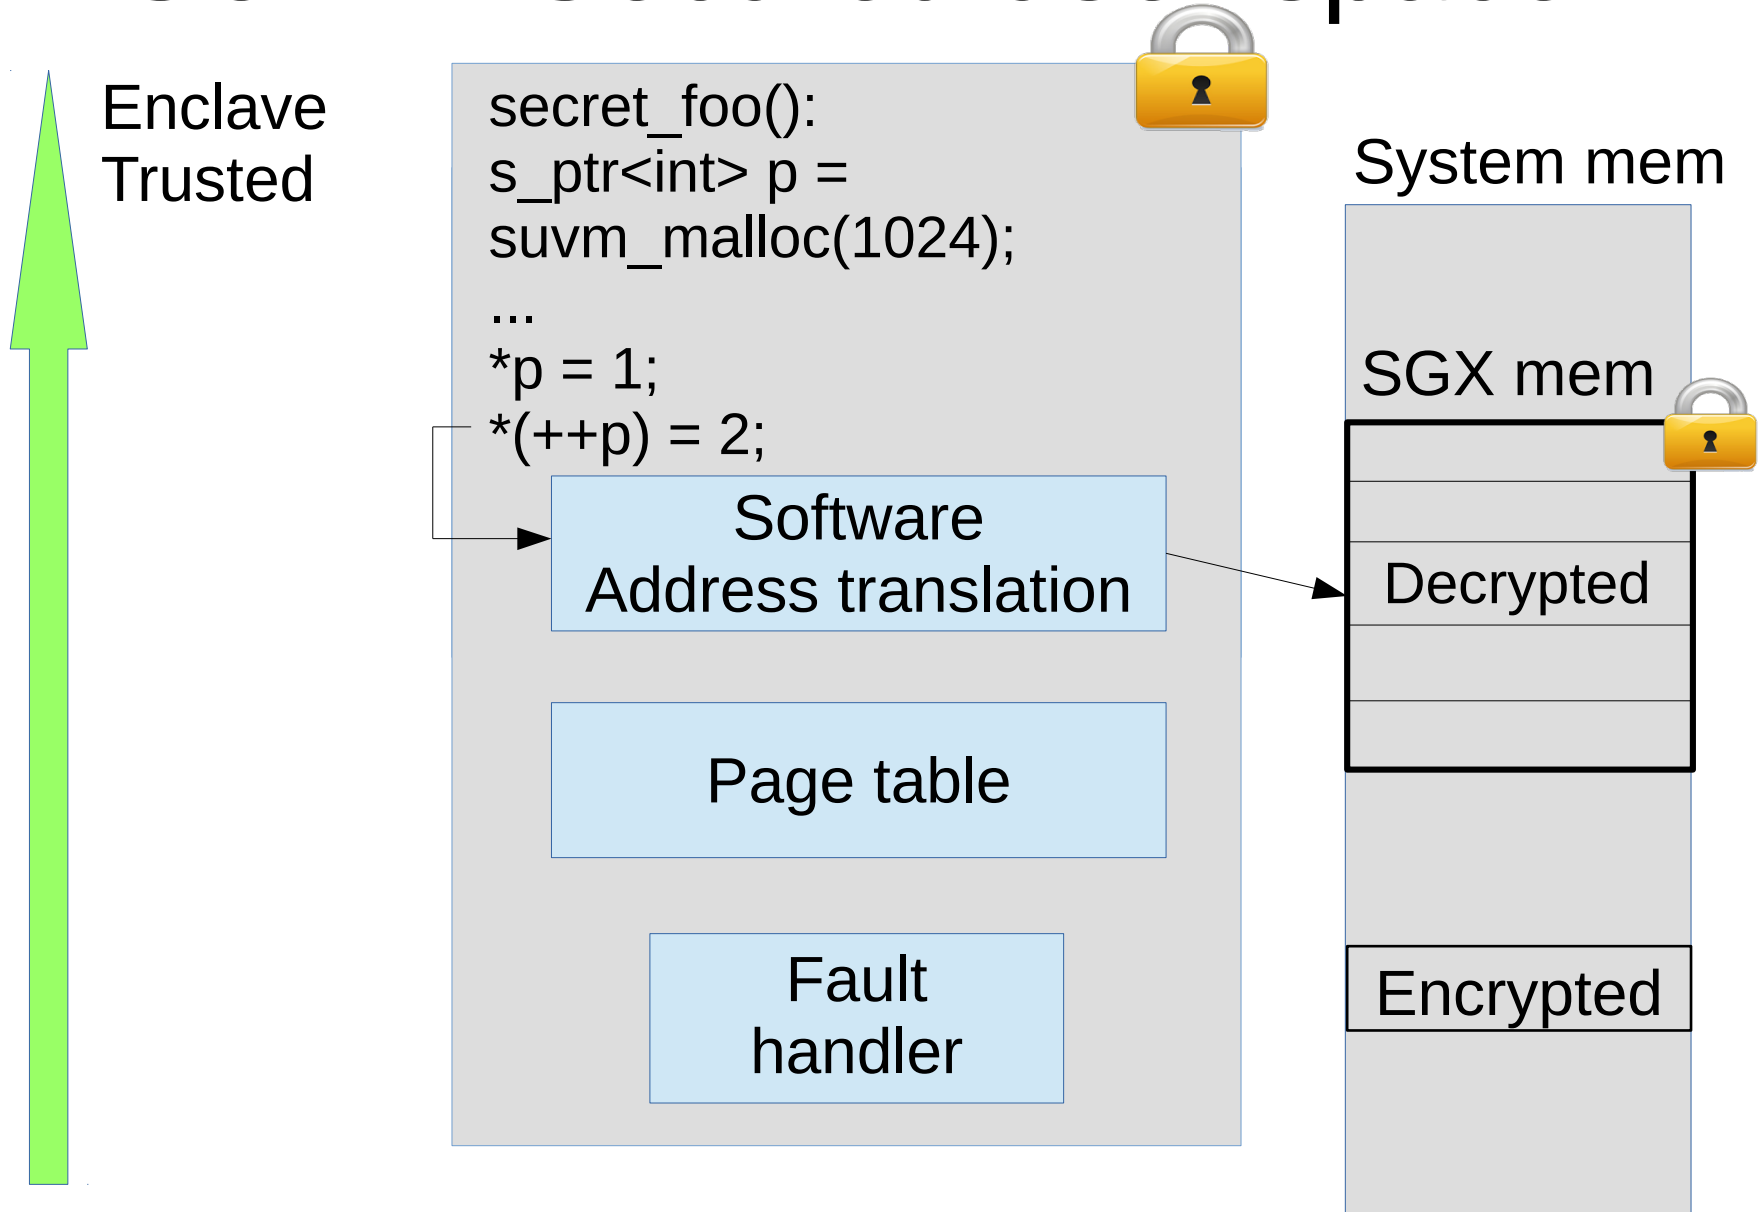
Control path
in-enclave

Swapped-out

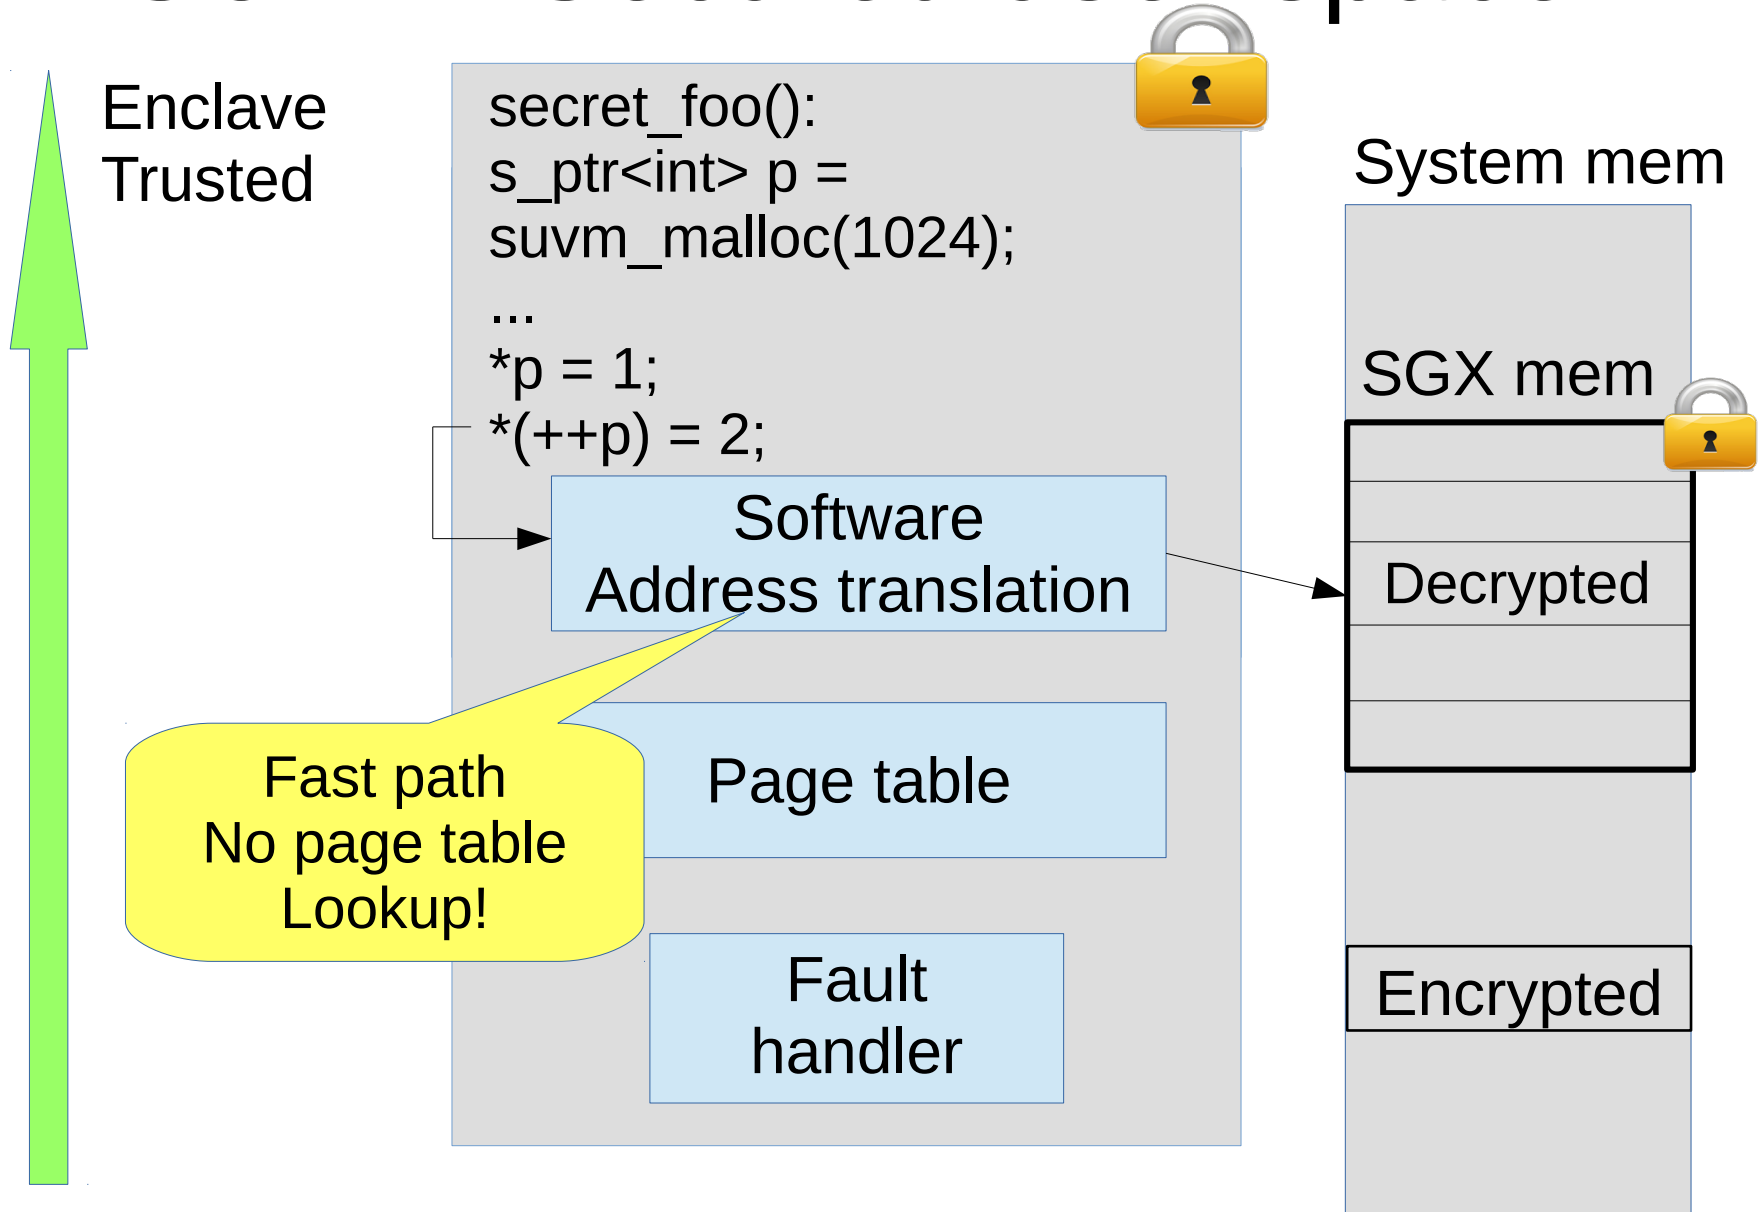
SUVM: Secured user-space VM



SUVM: Secured user-space VM



SUVM: Secured user-space VM



Wait...Software based VM management?



Based on software address translation
on GPUs, ActivePointers [ISCA'2016]

SUVM key contributions

- Multi-threaded

Compared to SGX:

Fast path: up to 20% overheads

Slow path: Eliminates costs of exits

	1 Thread	4 Threads
READ	5.5x	7x
WRITE	3.5x	5.9x

Throughput speedup

Software address translation offers new optimizations

- Customized page size
- Customized eviction policy
- Multi-enclave memory coordination
- Write-back only dirty pages
- Sub-page direct access to backing store

Software address translation offers new optimizations

- Customized page size
- Customized eviction policy
- Multi-enclave memory coordination
- Write-back only dirty pages
- Sub-page direct access to backing store



Virtual Machine
ballooning

Software address translation offers new optimizations

- Customized page size
- Customized eviction policy
- Multi-enclave memory coordination
- Write-back only dirty pages
- Sub-page direct access to backing store

Virtual Machine
ballooning



- Background
- Motivation
- Overhead analysis
- Eleos design
- **Evaluation**



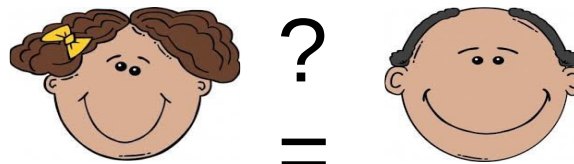
Biometric Identity checking server

Workload generator

Face verification server



+
ID

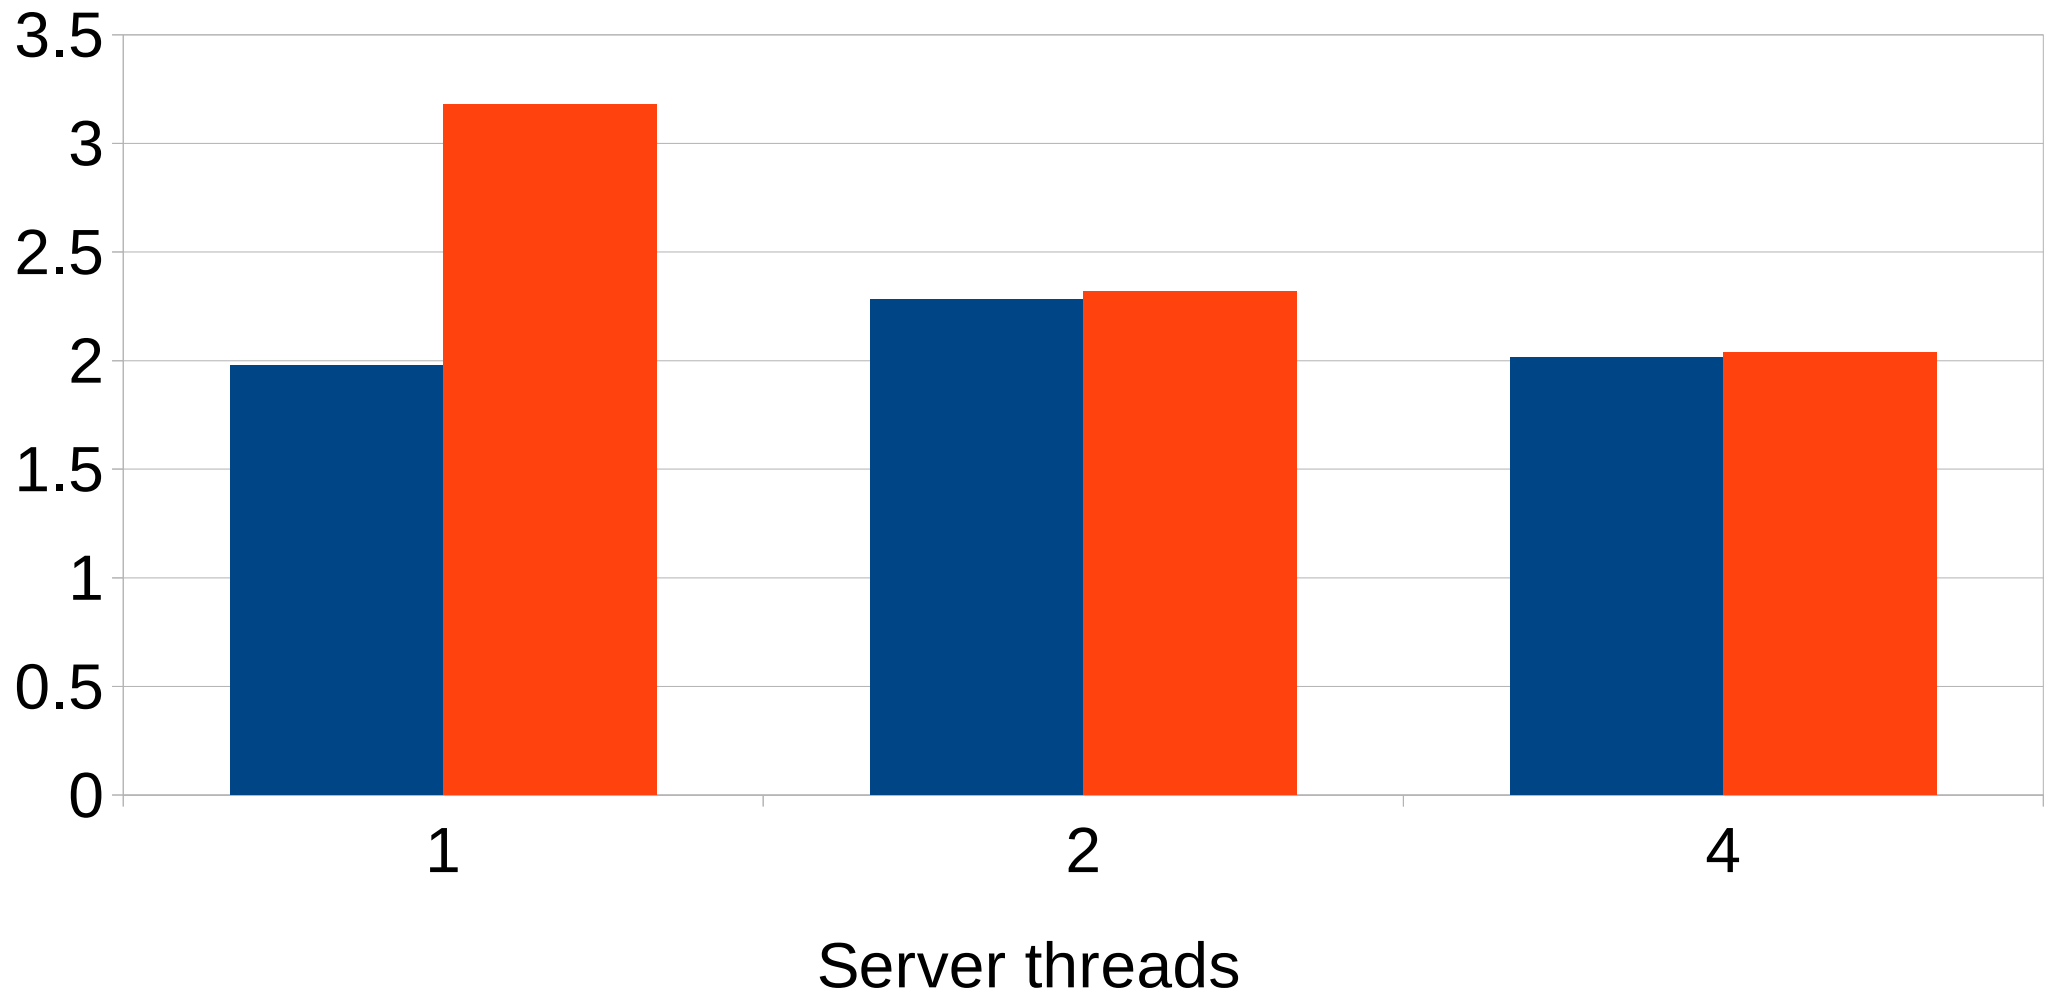


450MB DB
(5X SGX mem)

Biometric Identity validating server

Speedup compared to vanilla SGX

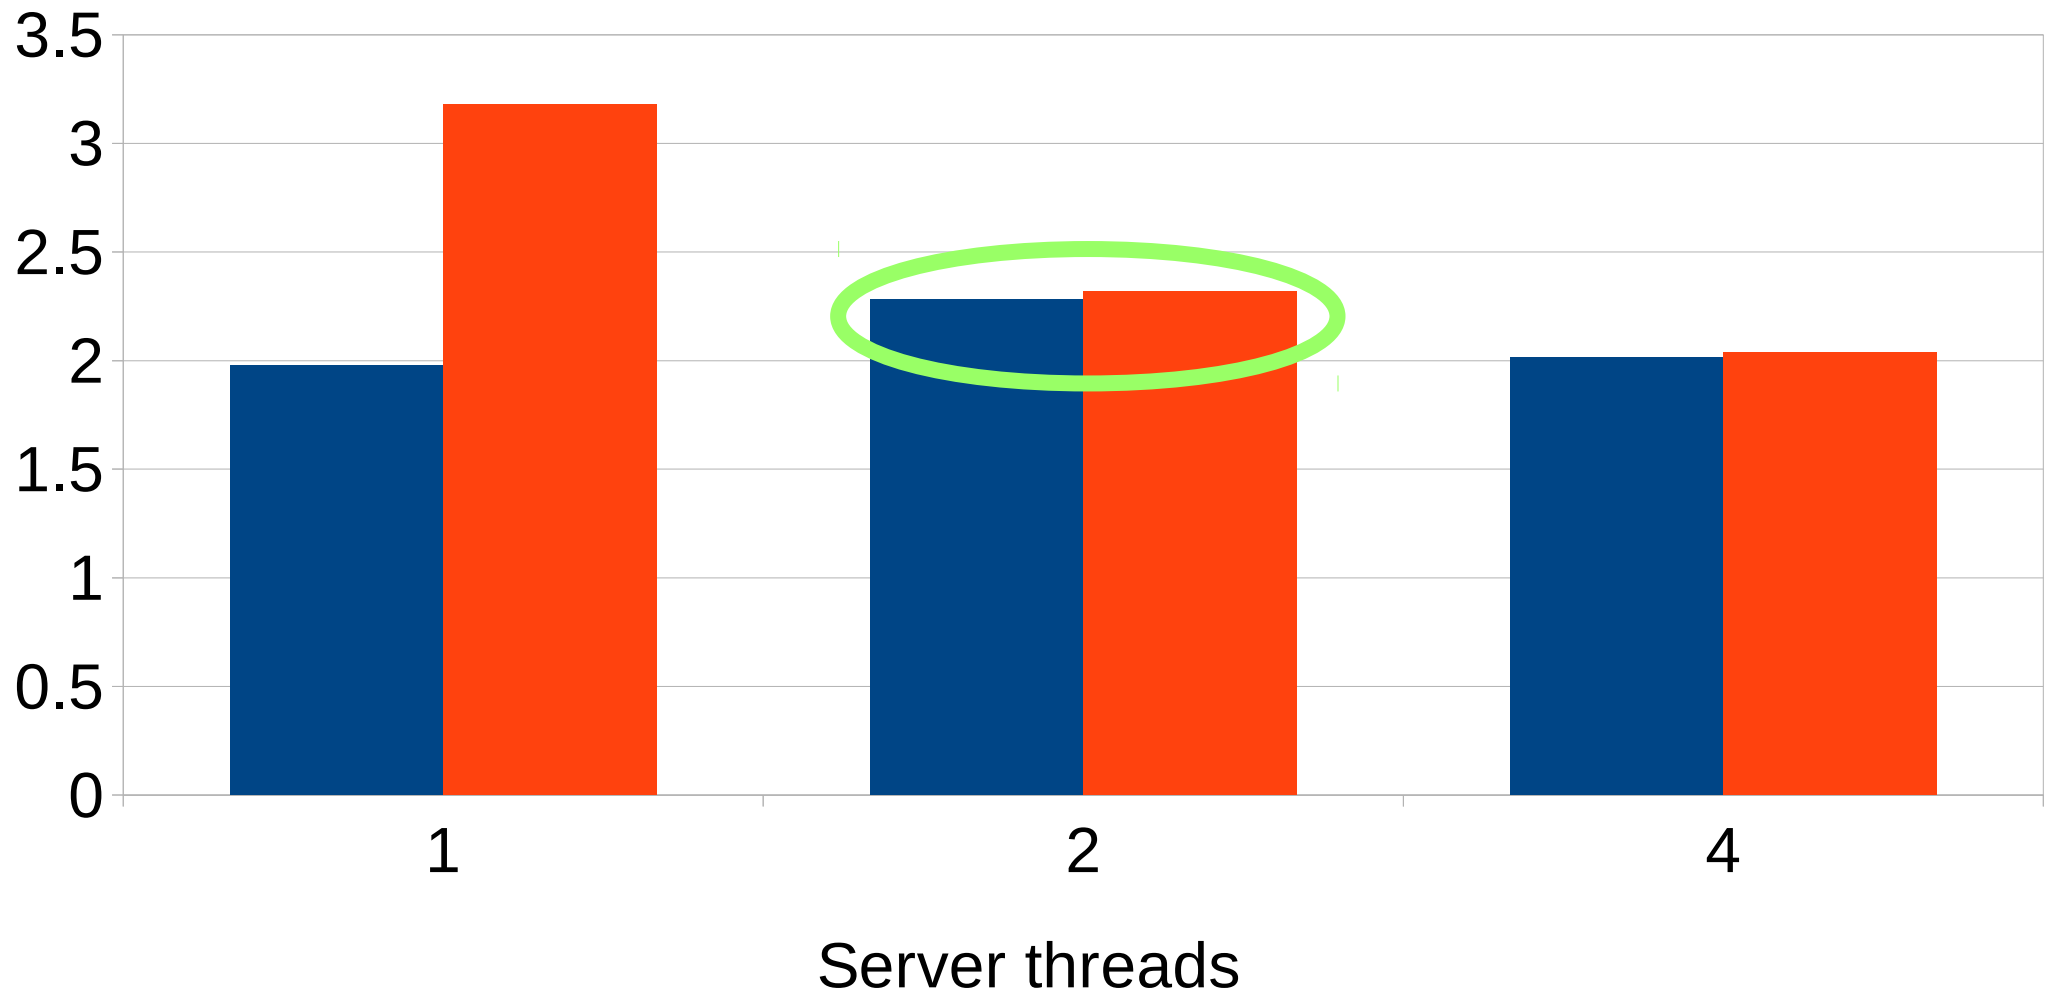
■ Eleos ■ Native



Biometric Identity validating server

Speedup compared to vanilla SGX

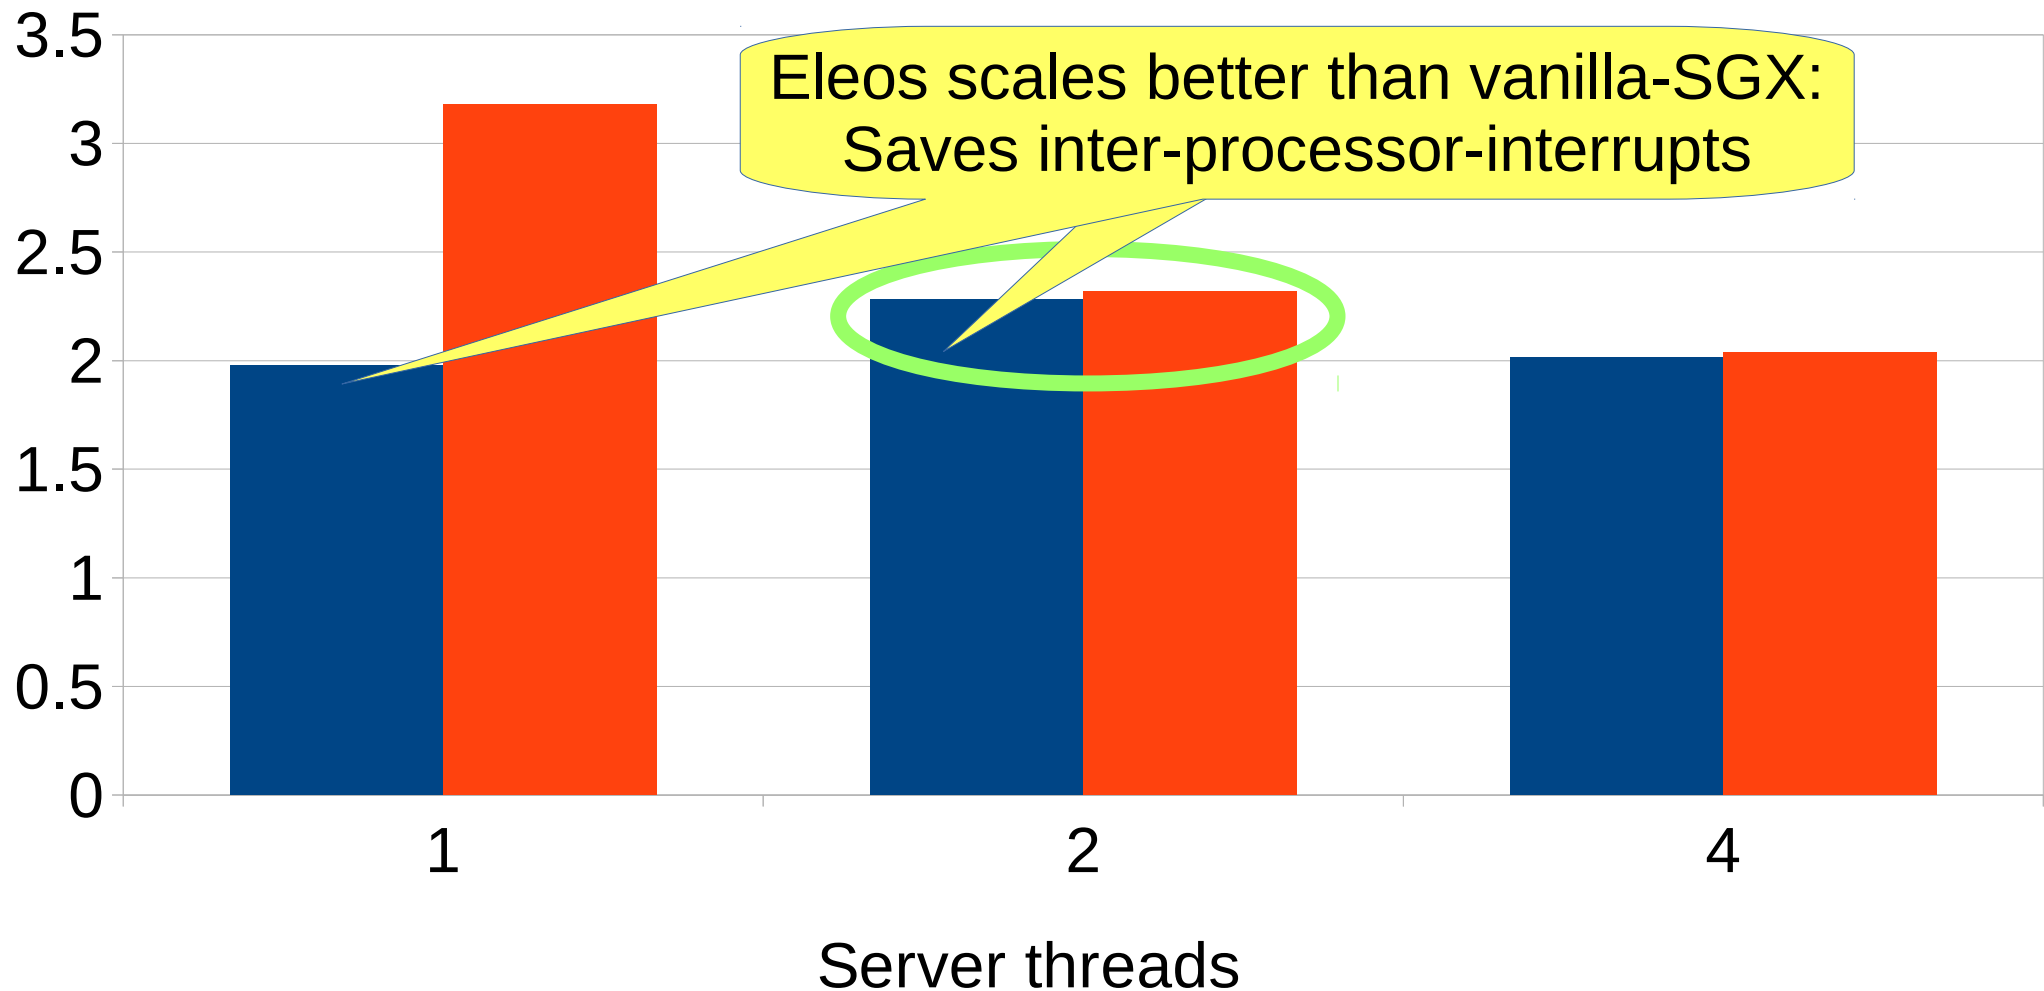
■ Eleos ■ Native



Biometric Identity validating server

Speedup compared to vanilla SGX

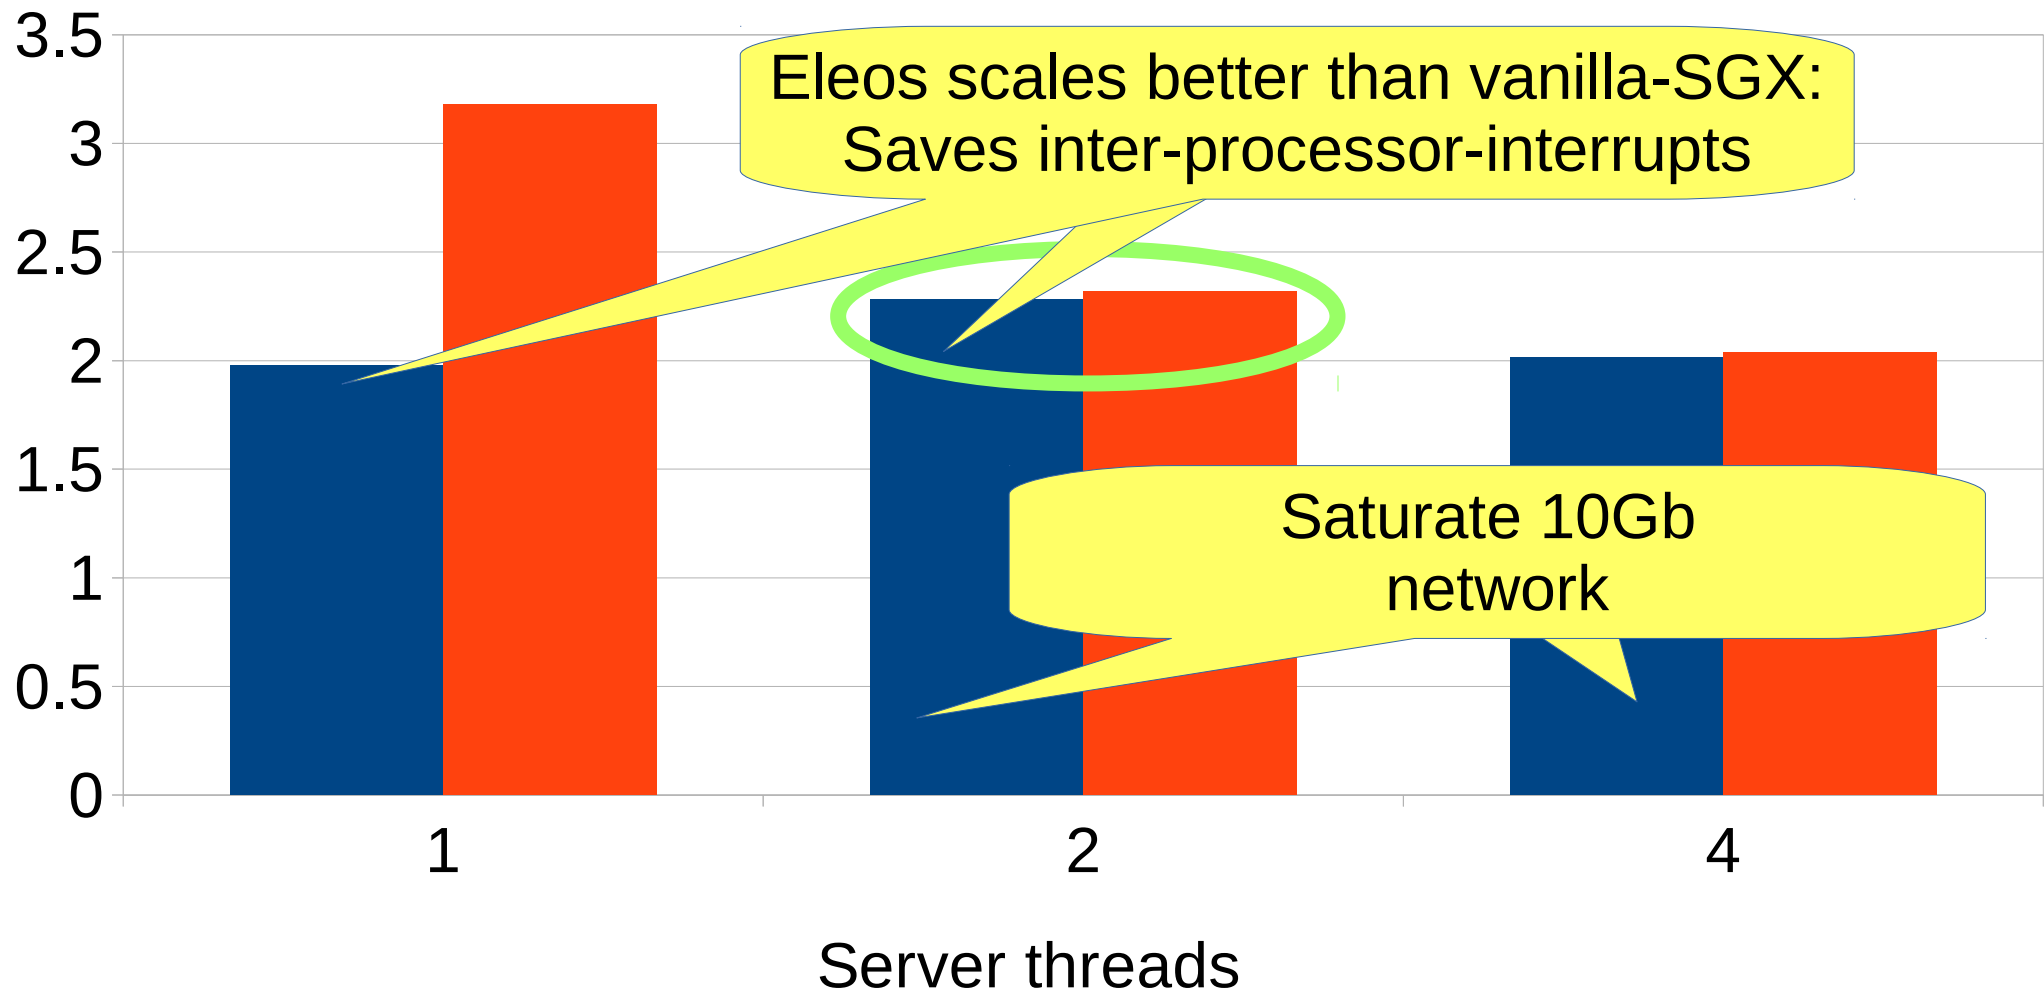
■ Eleos ■ Native



Biometric Identity validating server

Speedup compared to vanilla SGX

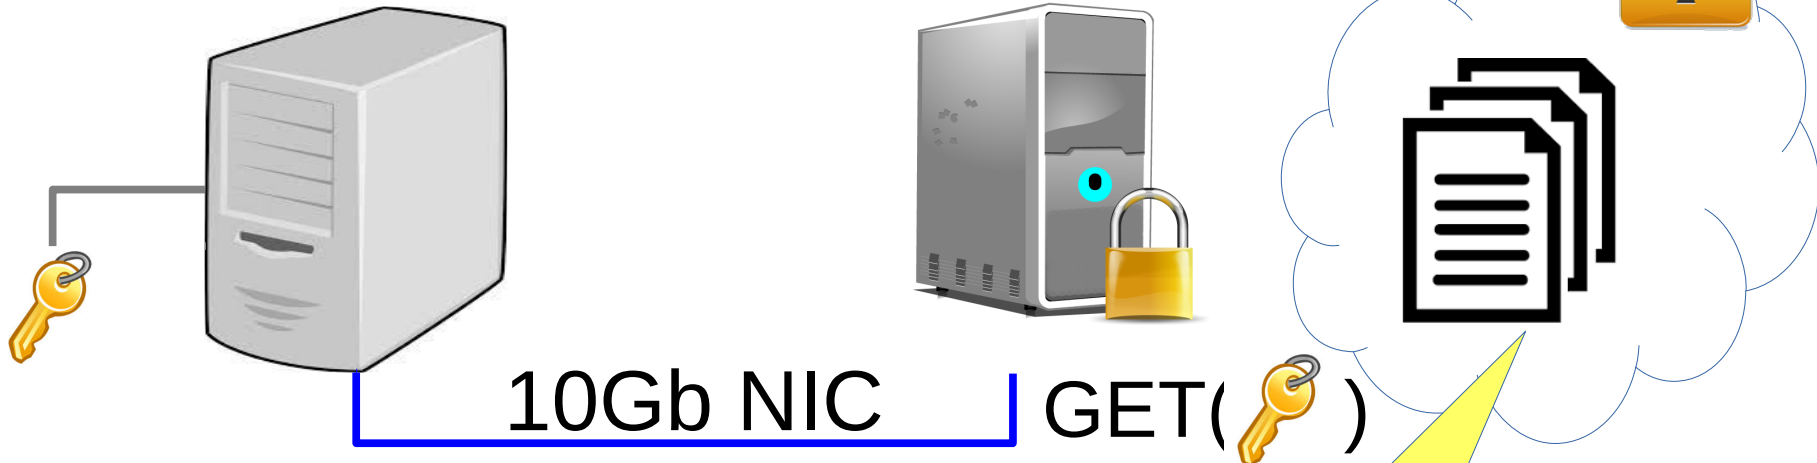
■ Eleos ■ Native



Memcached

Workload
Generator
(memaslap)

Memcached
Graphene LibOS
[Eurosys'2014]



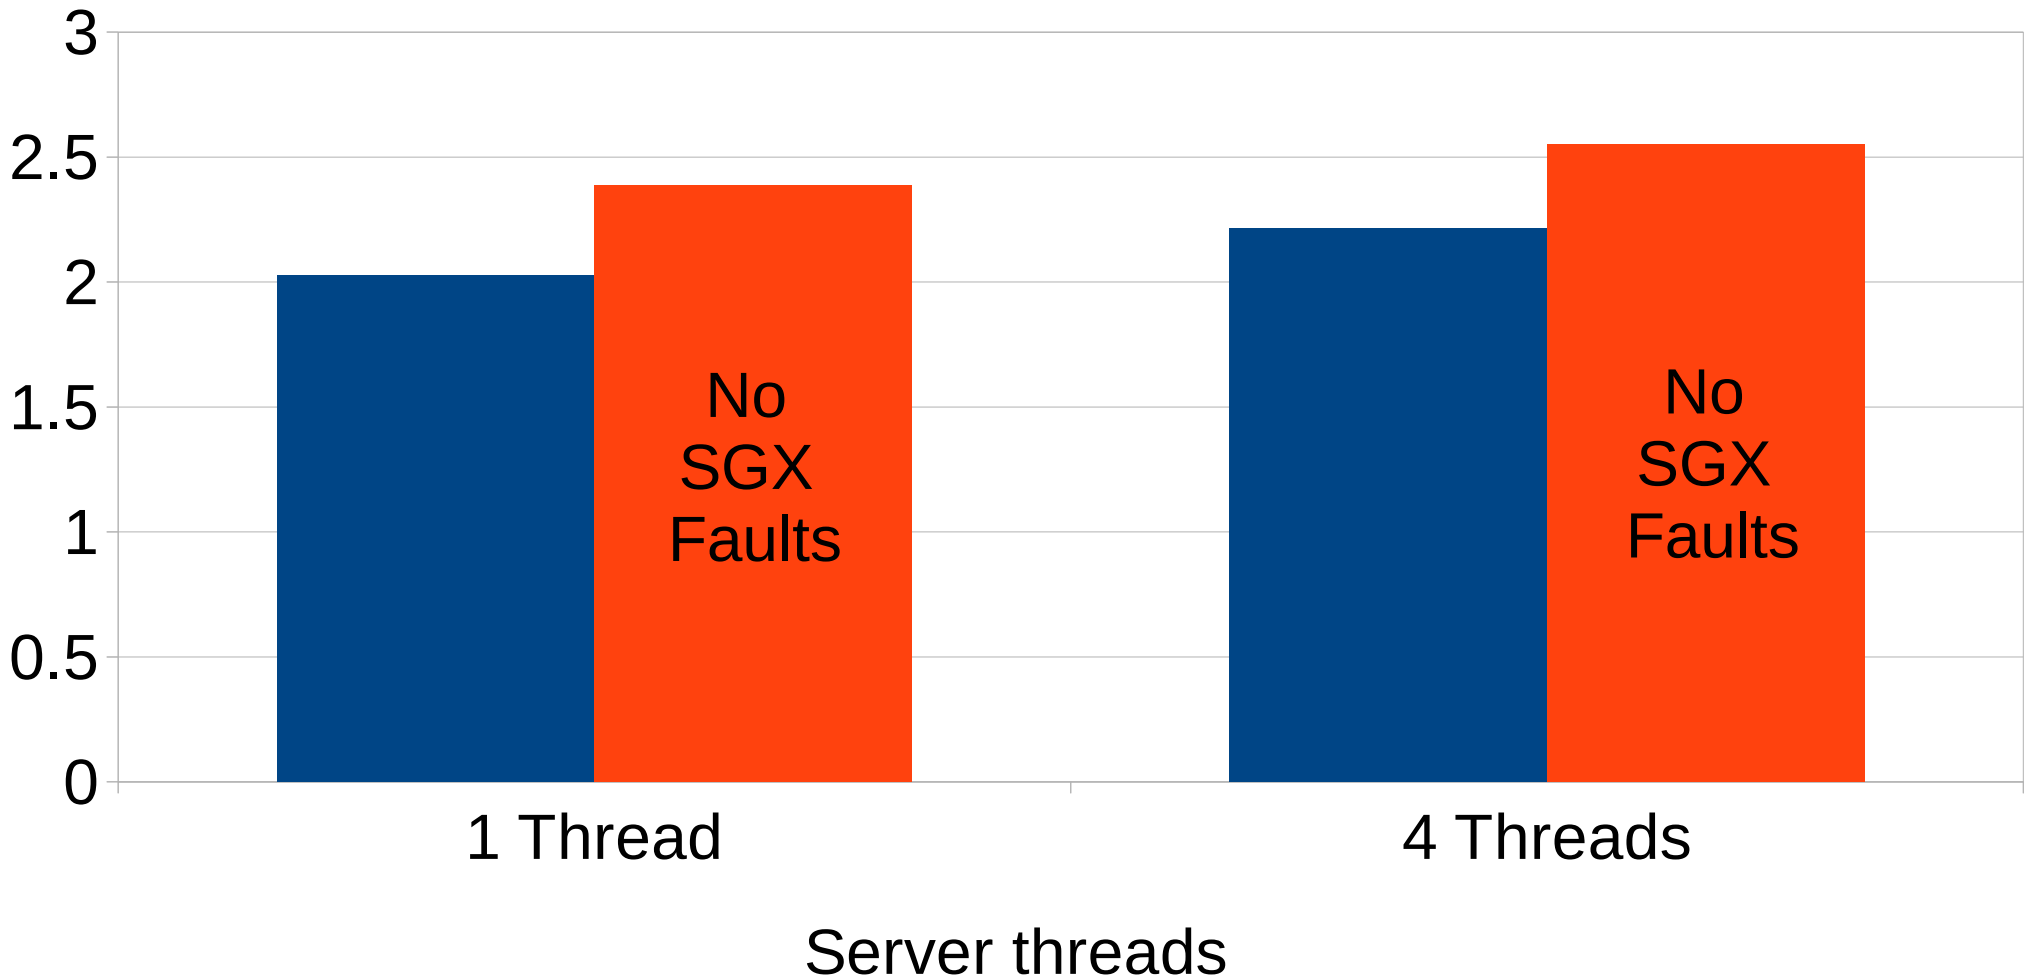
~75 LOC
modification
for SUVM

500MB DB
(5.5X SGX mem)

Memcached

Speedup compared to vanilla SGX (500 MB)

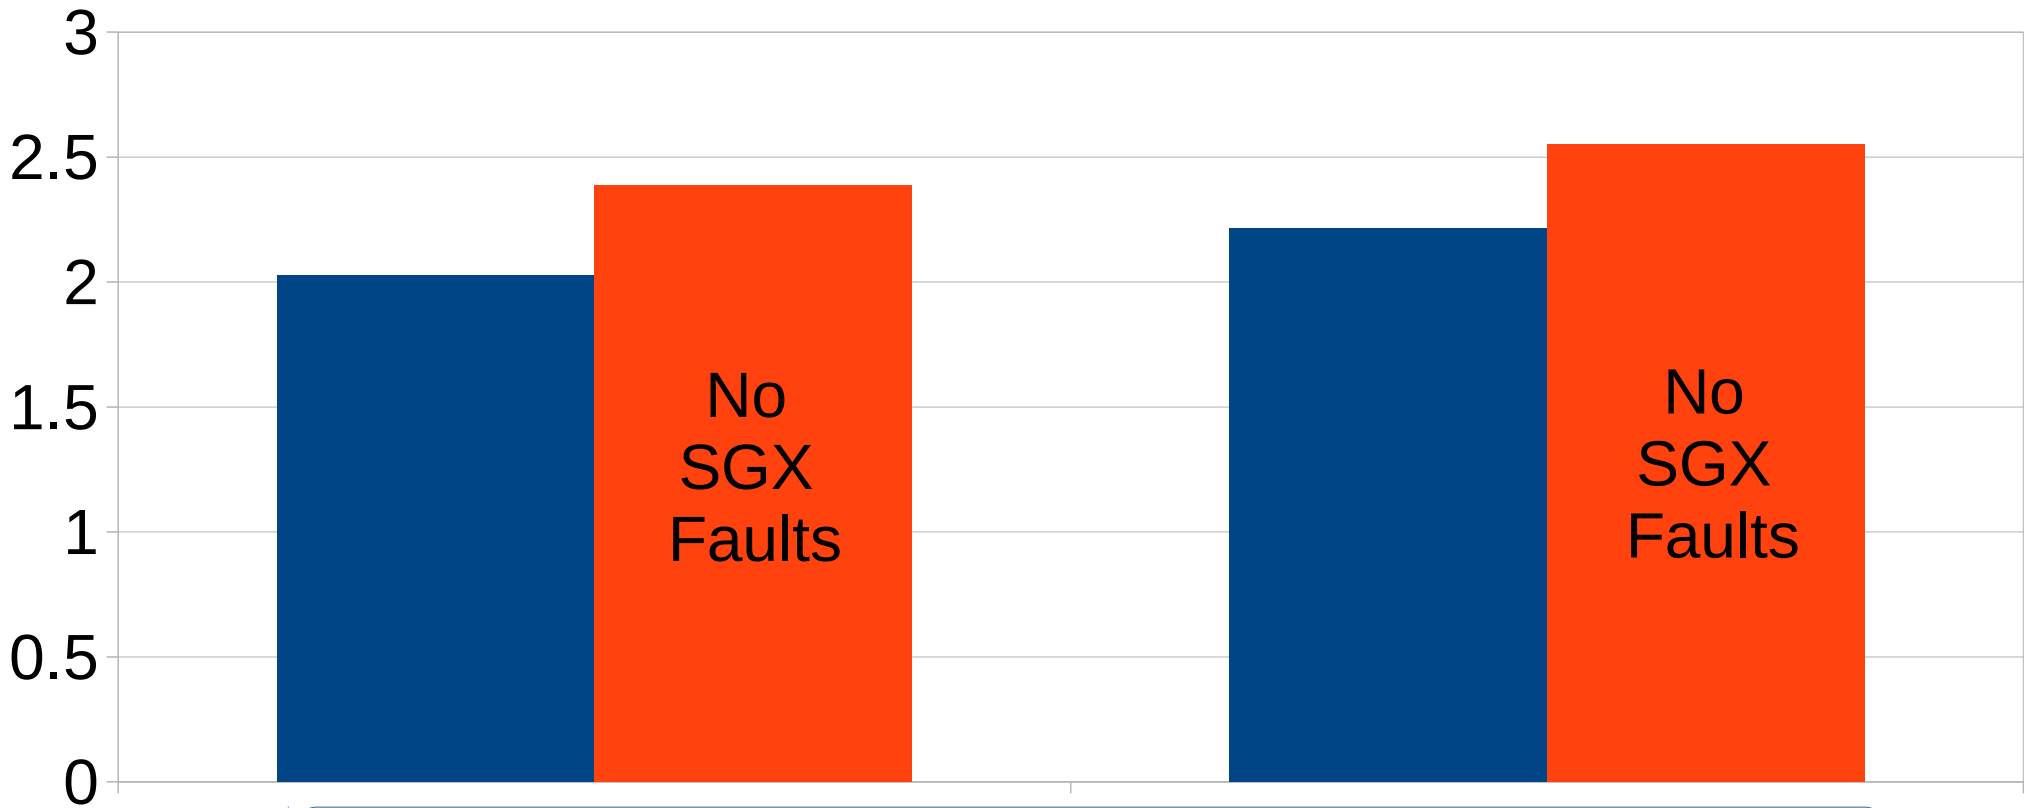
■ Eleos (500MB DB) ■ vanilla SGX (20MB DB)



Memcached

Speedup compared to vanilla SGX (500 MB)

■ Eleos (500MB DB) ■ vanilla SGX (20MB DB)

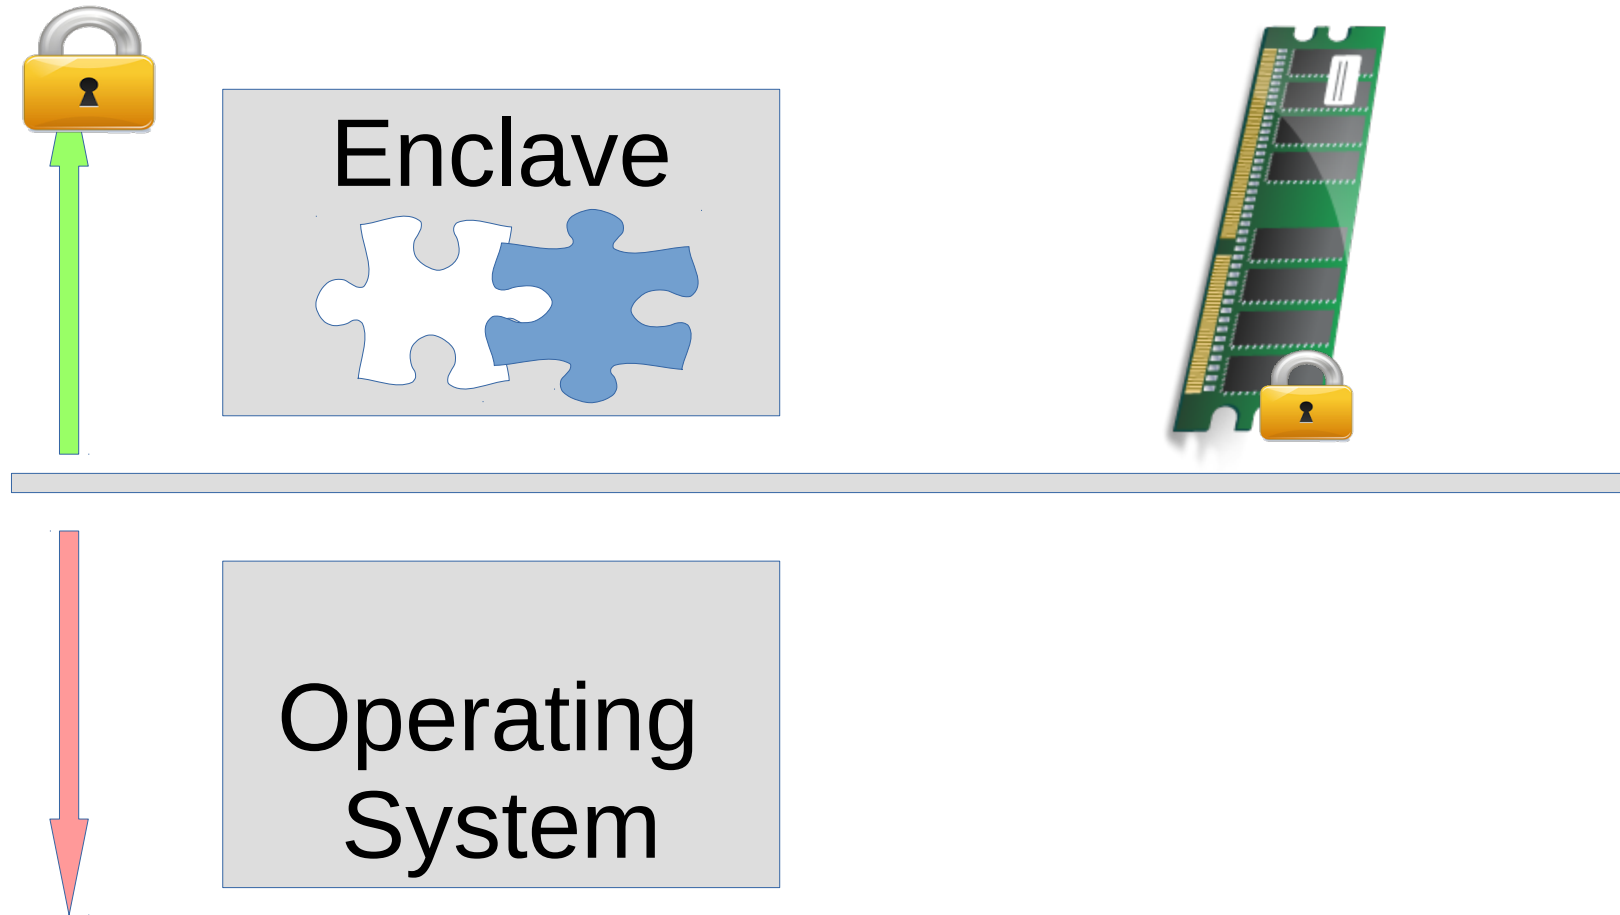


Disclaimer: Eleos+Graphene is 3x slower than native

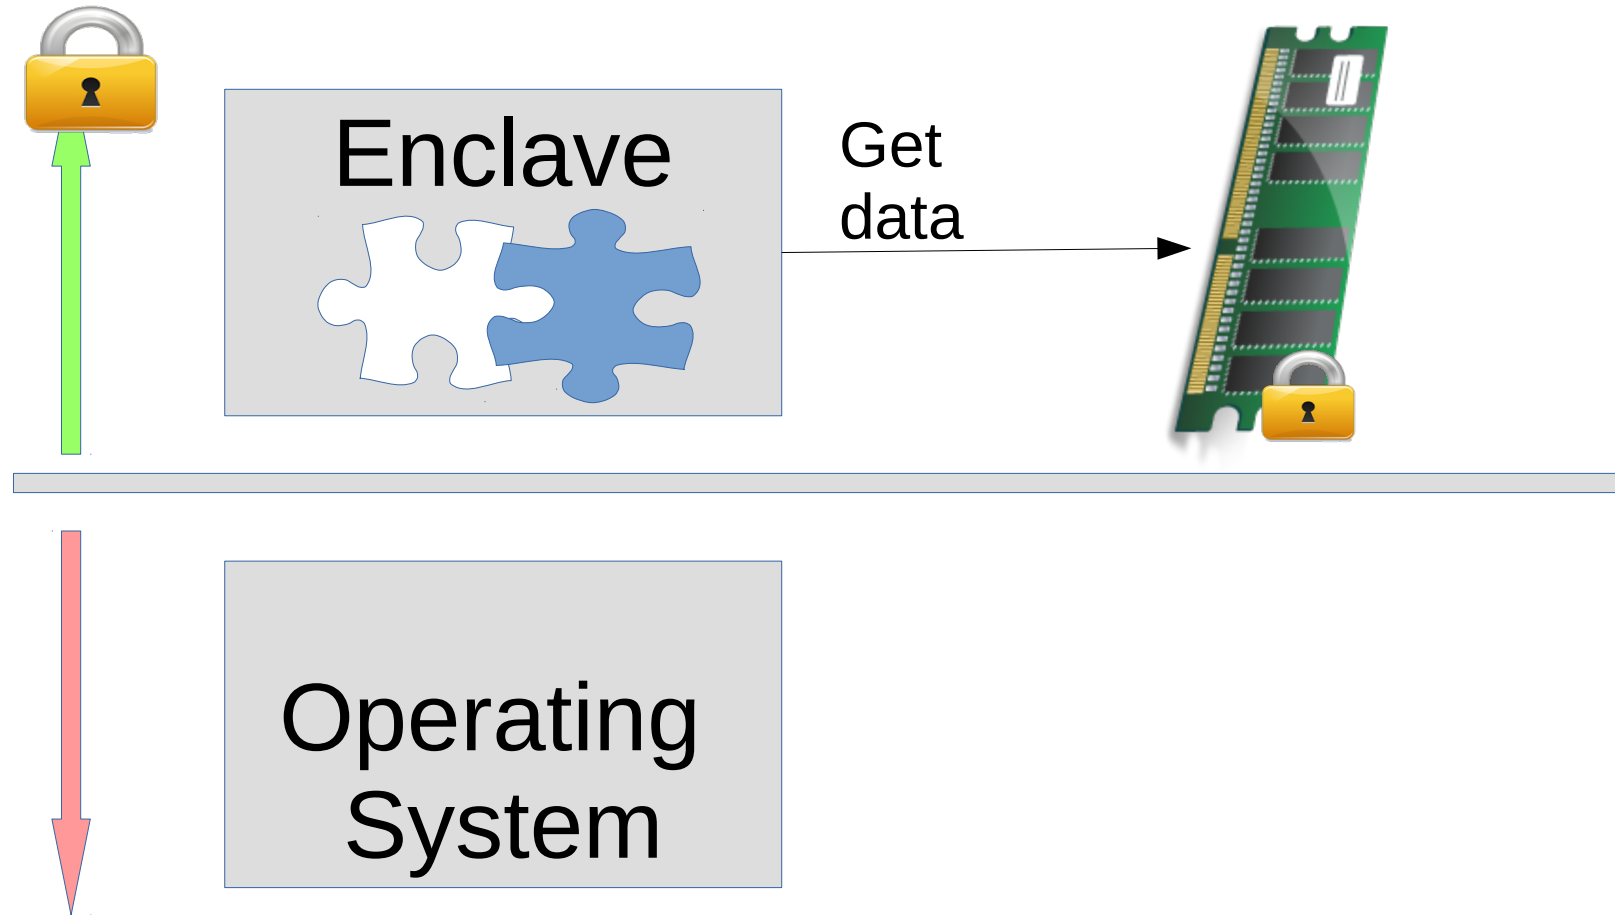
Take aways

- Eleos eliminates enclave exits costs
- Eleos available for Windows and Linux
 - Makes memory demanding applications available on Windows today
- Eleos takes a modularize approach
 - Memory demanding app? Link to SUVM
 - I/O intensive app? Link to RPC
 - Maintaining small TCB

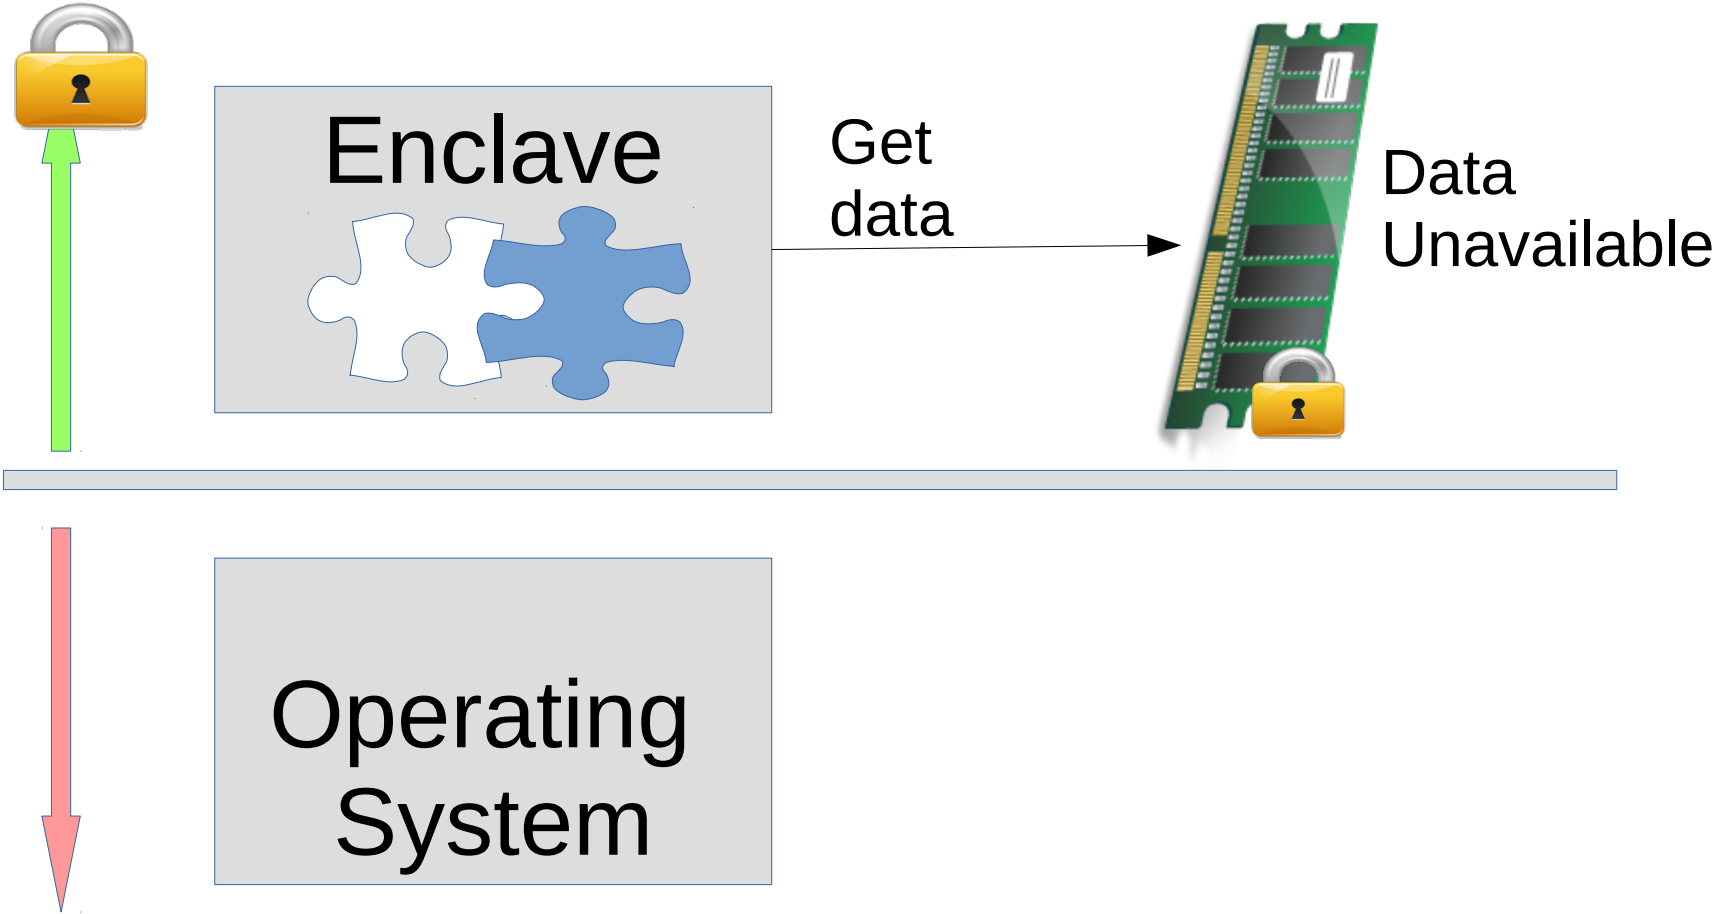
Traditional SGX: Host-centric OS services



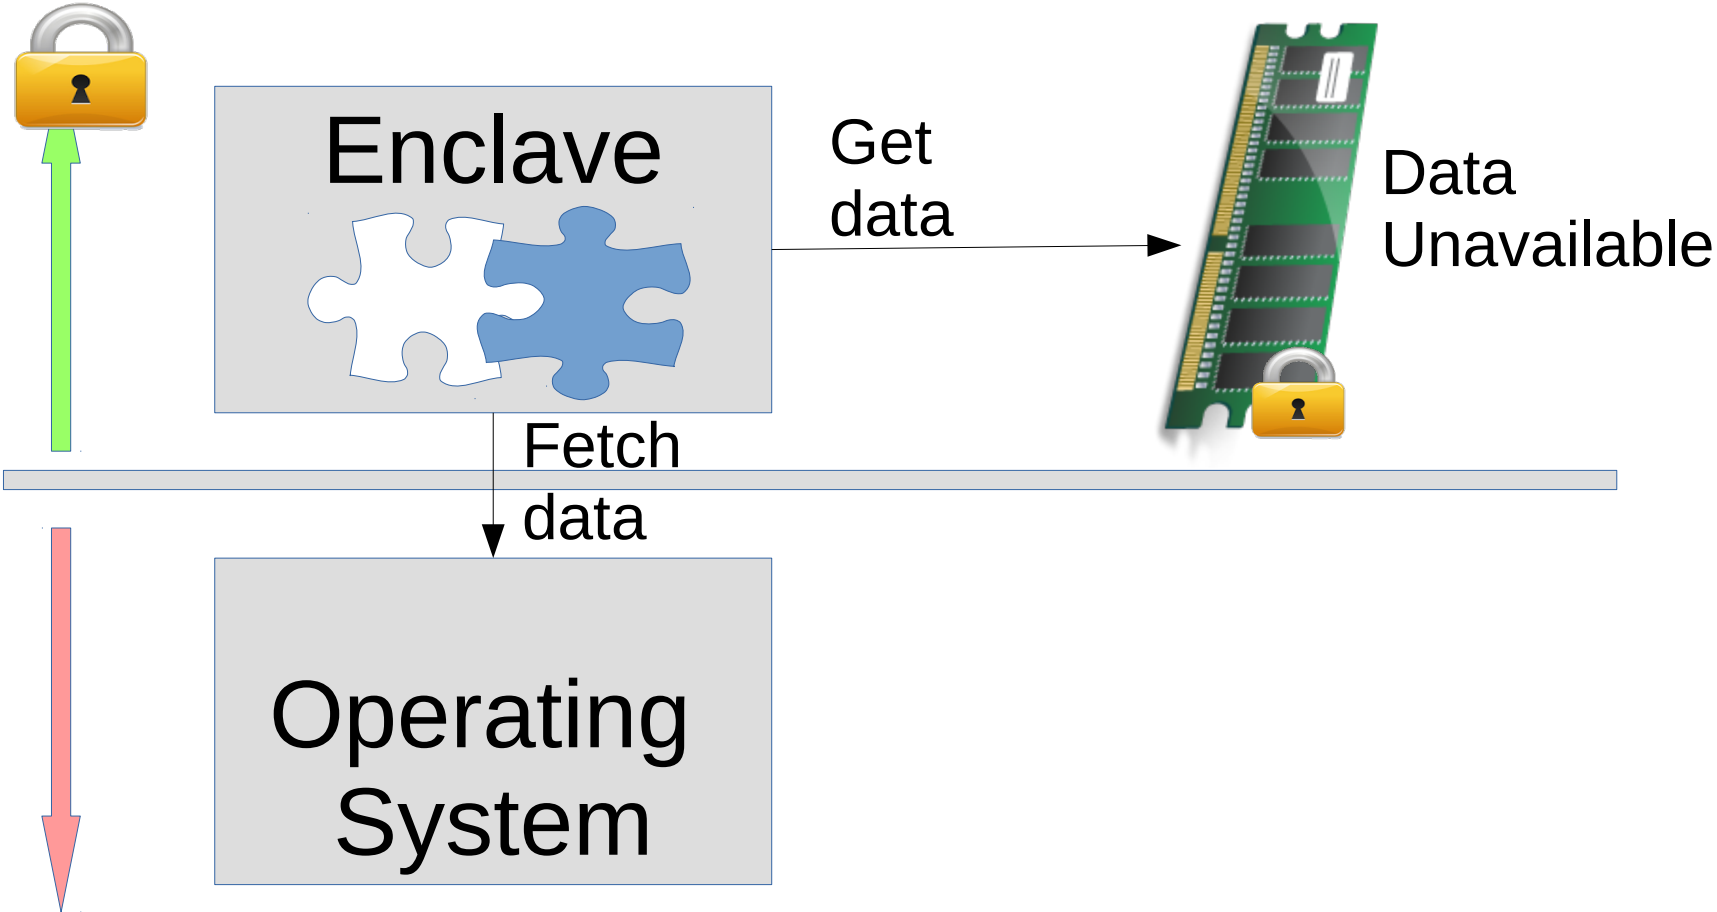
Traditional SGX: Host-centric OS services



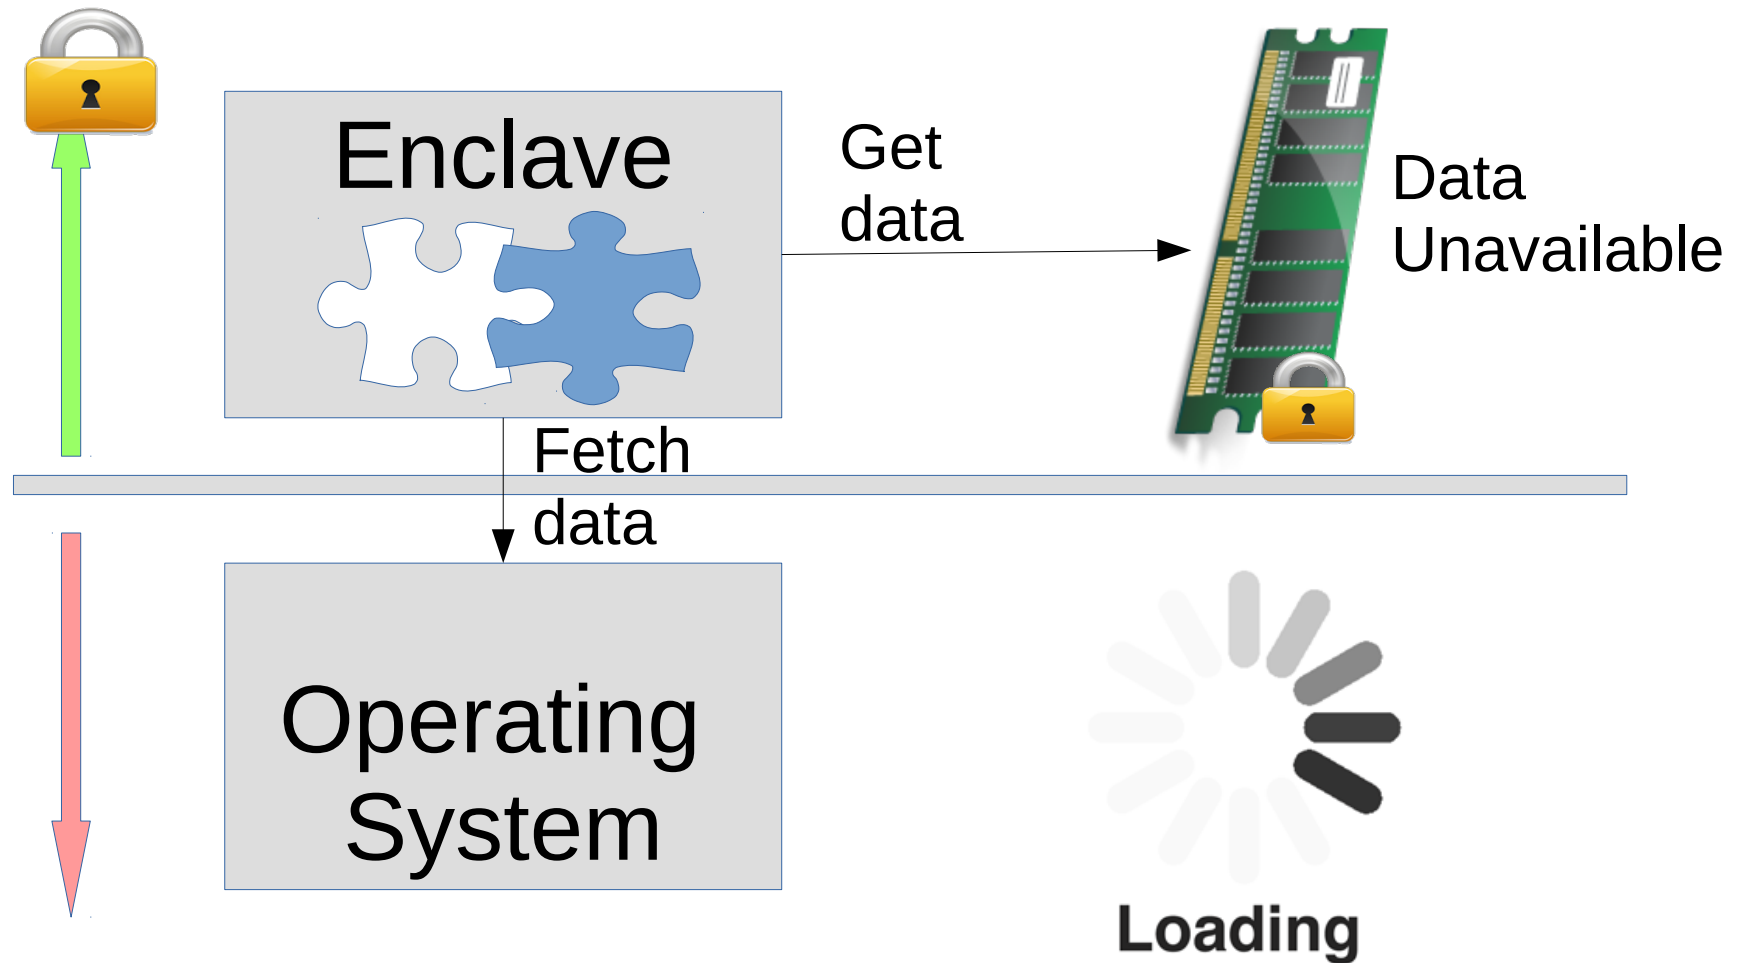
Traditional SGX: Host-centric OS services



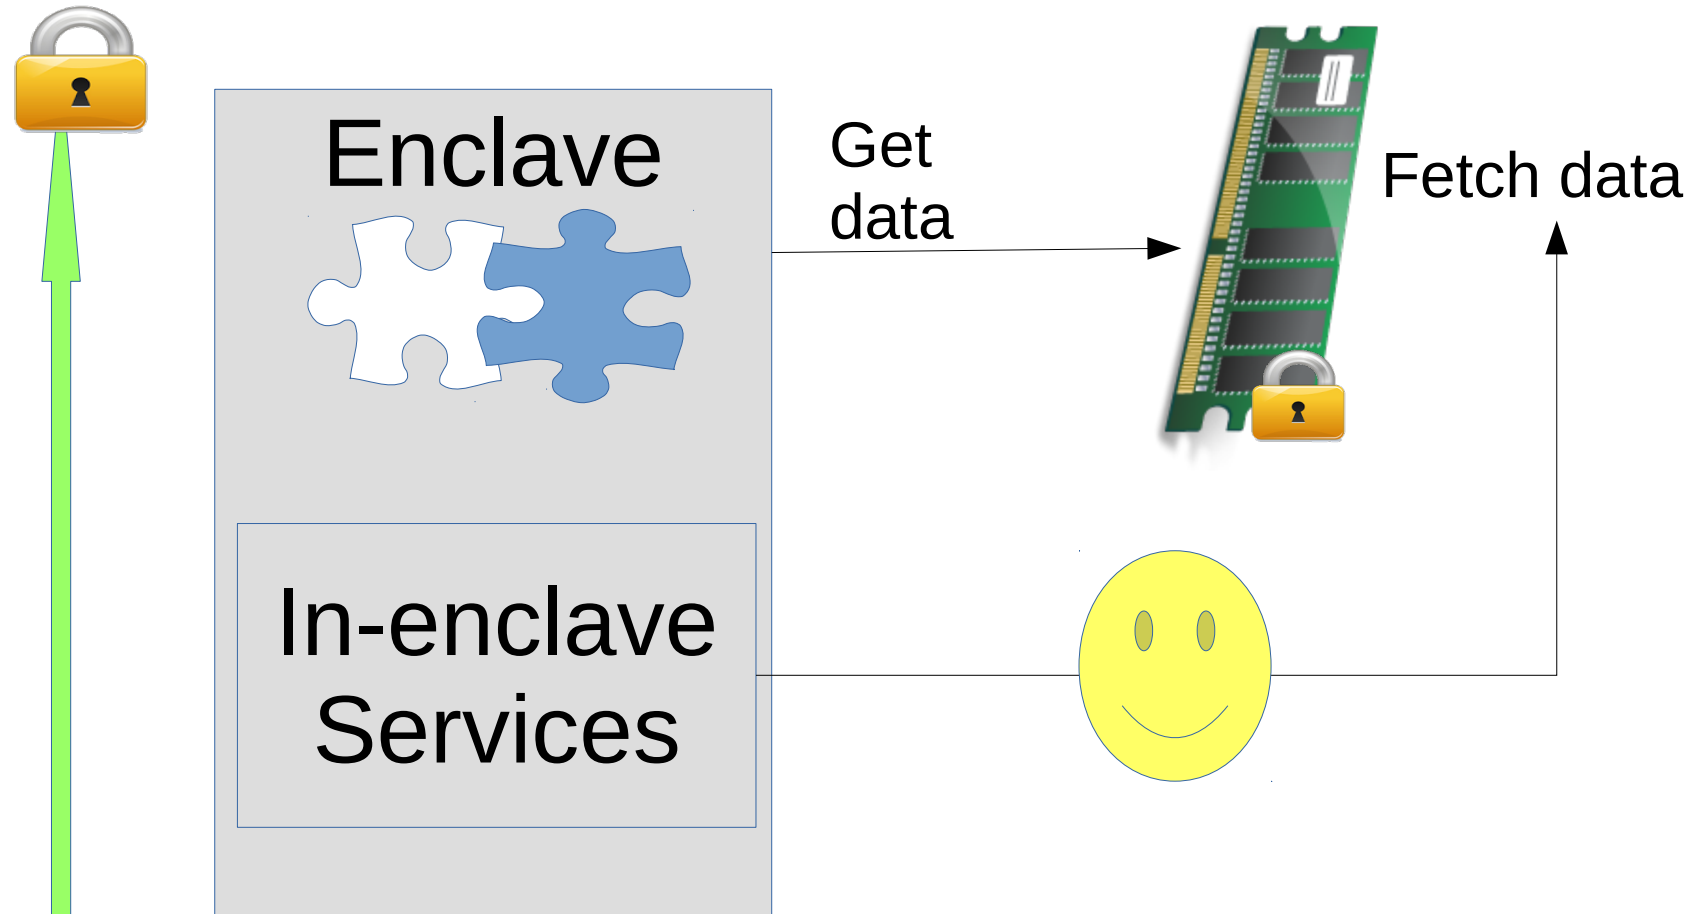
Traditional SGX: Host-centric OS services



Traditional SGX: Host-centric OS services



Eleos Insight: Enclave-centric OS services



Take aways (2)

- Eleos adapts 'accelerator-centric management'
 - System calls: GPUfs [ASPLOS'13], GPUnet [OSDI'14]
 - Virtual memory: ActivePointers [ISCA'16]
- We can do more!
 - Asynchronous DMA host copies
 - Non-blocking enclave launches

More information at:

“SGX Enclaves as Accelerators” [Systex'16]

Thank you



Code is available at:

<https://github.com/acsl-technion/eleos>



shmeni@tx.technion.ac.il

Backup slides