Overhead-free I/O from enclaves

SysTEX'16
Trento, Italy

Meni Orenbach
Prof. Mark Silberstein

Research Statement:

Enclaves are accelerators
for secured execution

Accelerator system services and
Abstractions can be retrofitted
Inspire system services for enclaves

Case Example: GPU

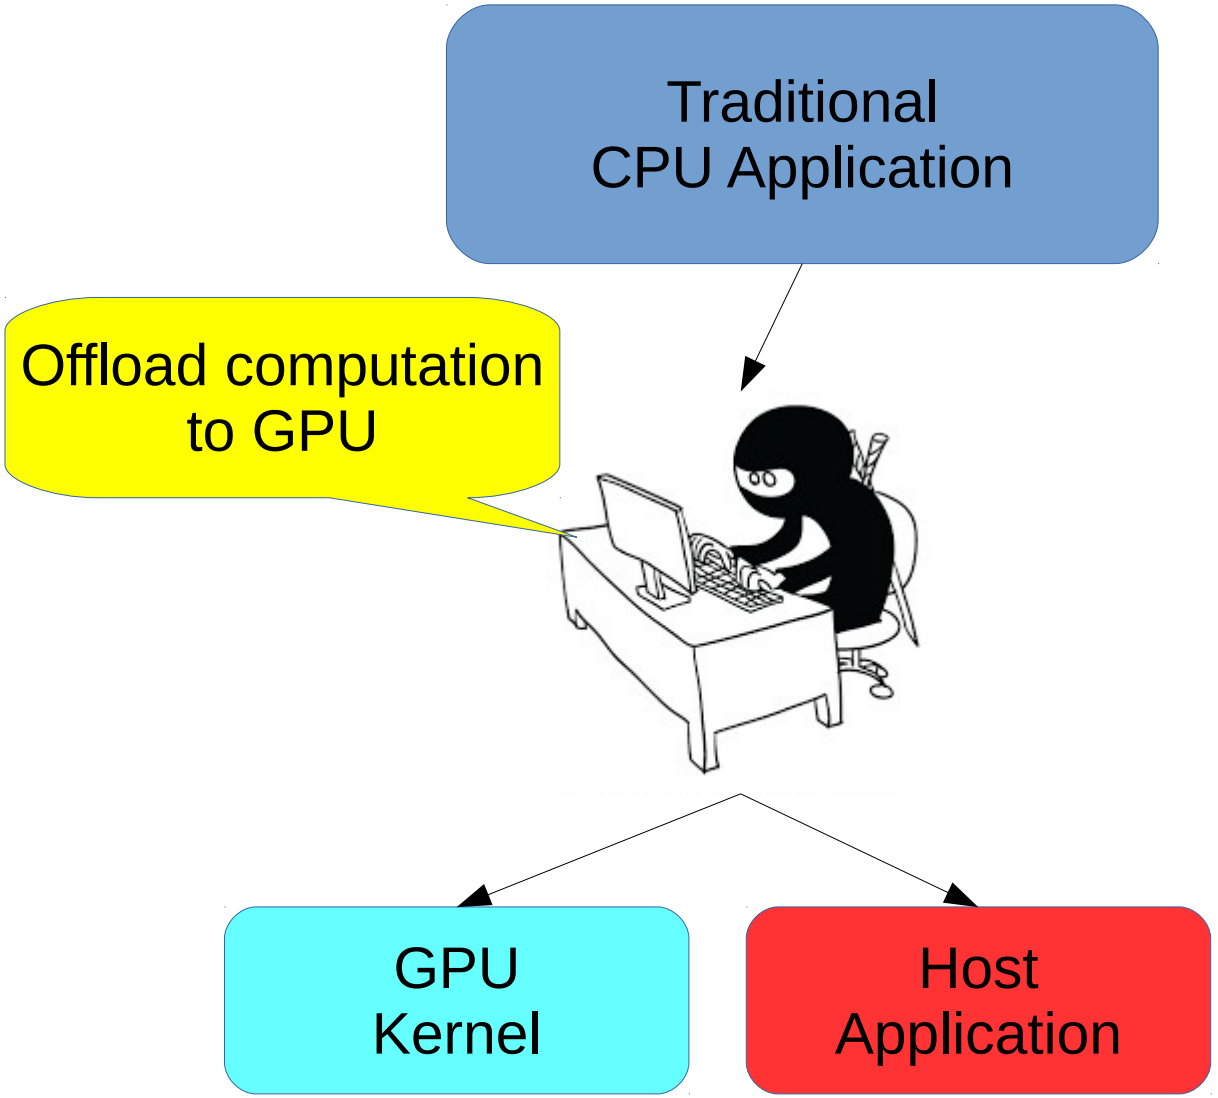

Partition: GPU and host

Background: GPU Kernels

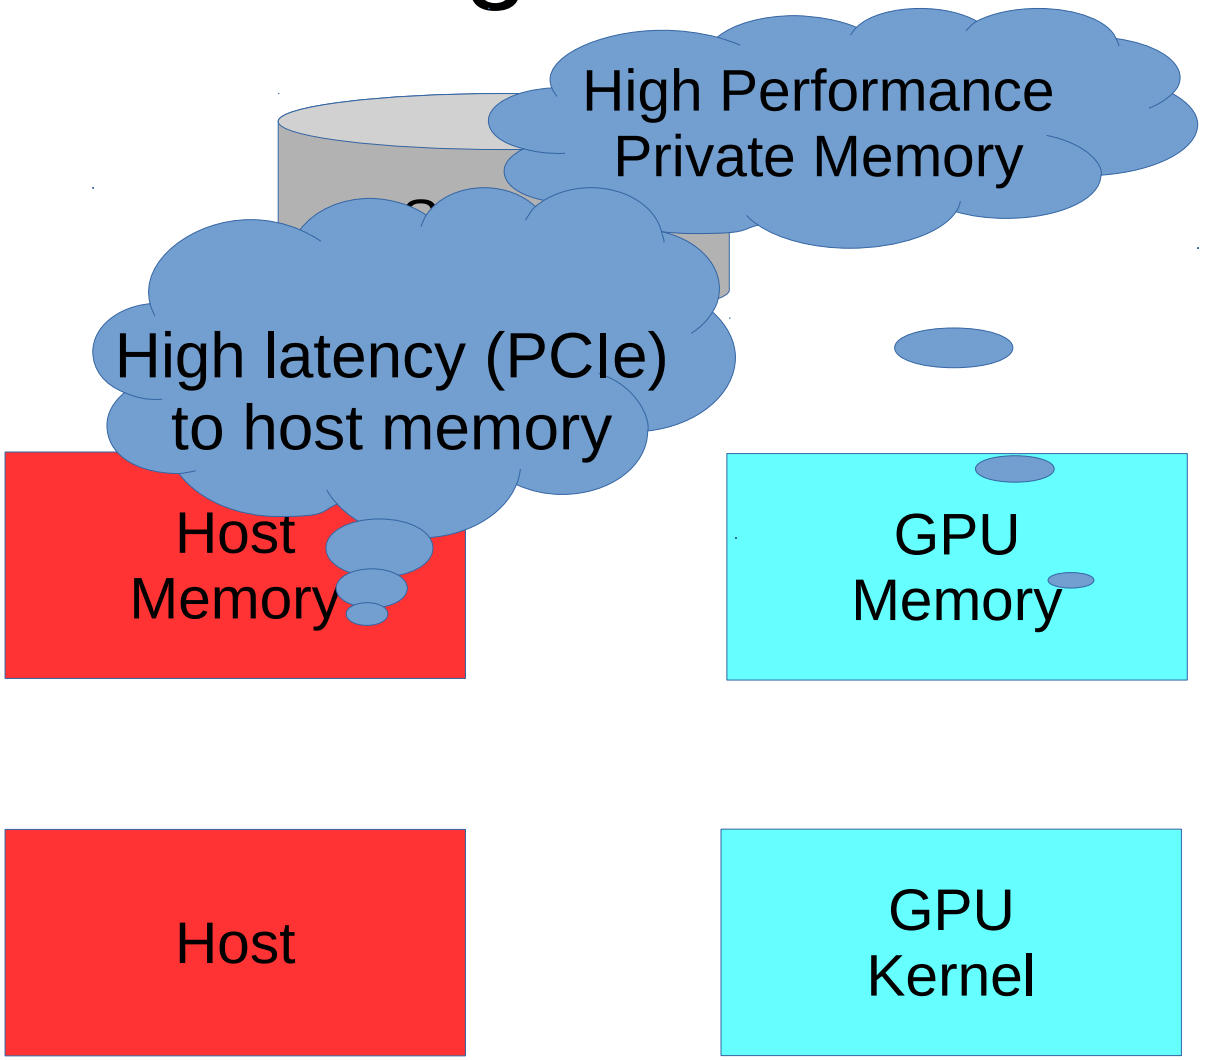

Partition: GPU and host

Separate GPU Memory

Background: GPU Kernels

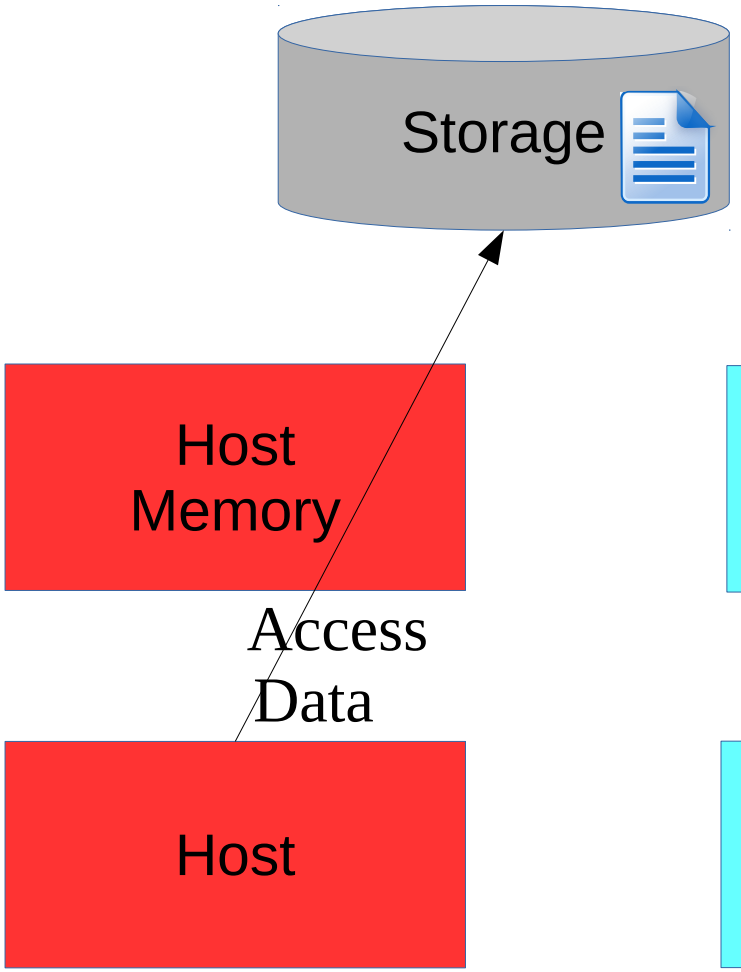

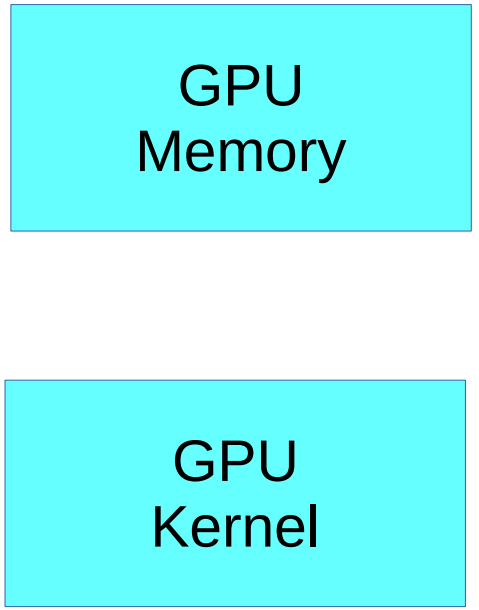

Partition: GPU and host

Separate GPU Memory

Host manages OS services

GPU cannot invoke syscalls

Background: GPU Kernels

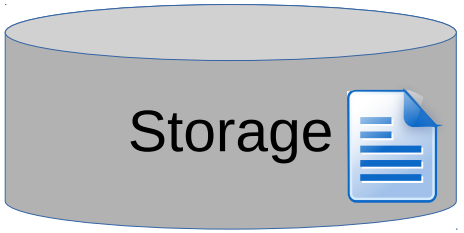

Host Memory

GPU Memory

Host

GPU Kernel

- Partition: GPU and host
- Separate GPU Memory
- Host manages OS services
- GPU cannot invoke syscalls
- Host operate on its memory

Background: GPU Kernels

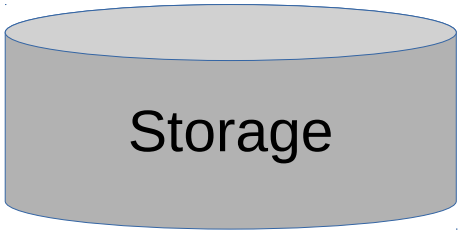

Host Memory

GPU Memory

Host

GPU Kernel

- Partition: GPU and host
- Separate GPU Memory
- Host manages OS services
- GPU cannot invoke syscalls
- Host operate on its memory
- Copy data to GPU memory

Background: GPU Kernels

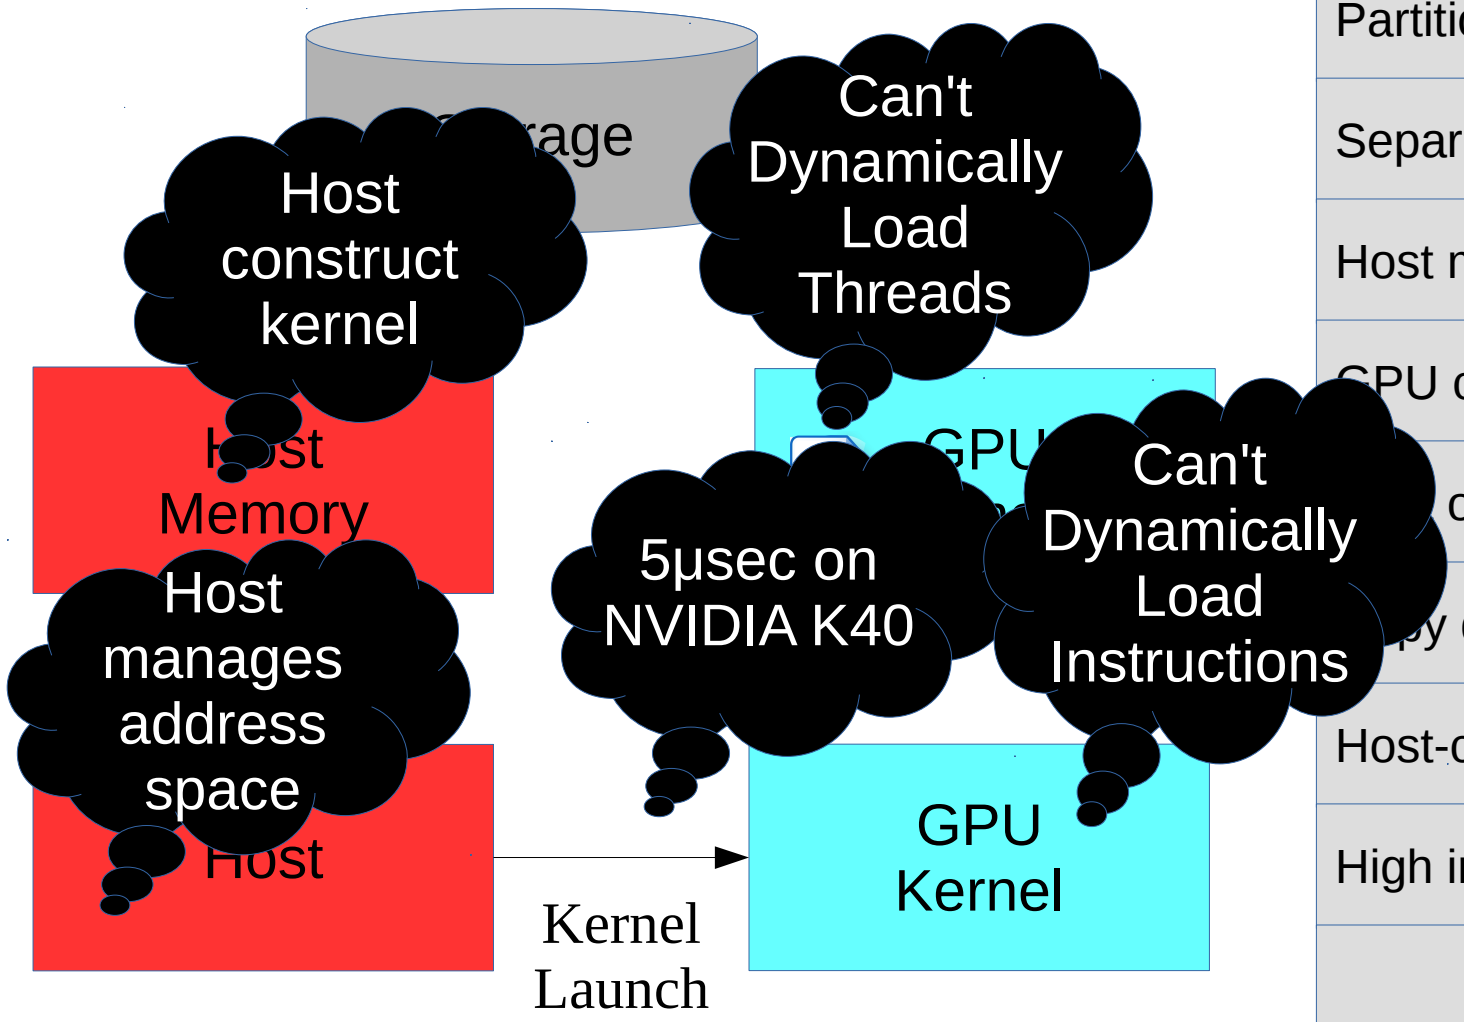

Partition: GPU and host
Separate GPU Memory
Host manages OS services
GPU cannot invoke syscalls
operate on its memory
y data to GPU memory
Host-centric management
High invocation costs

Background: GPU Kernels

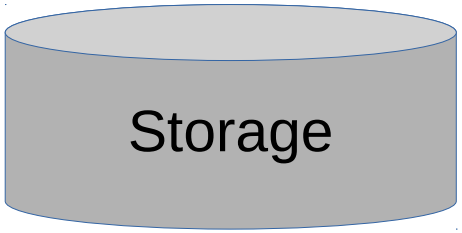

Host Memory

GPU Memory

Host

GPU Kernel

- Partition: GPU and host
- Separate GPU Memory
- Host manages OS services
- GPU cannot invoke syscalls
- Host operate on its memory
- Copy data to GPU memory
- Host-centric management
- High invocation costs
- GPU execute computation

Background: GPU Kernels

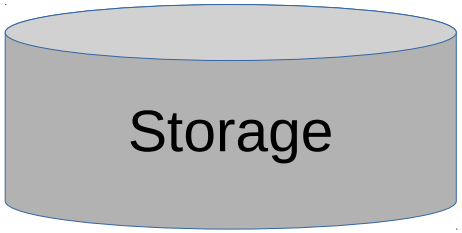

Host Memory

GPU Memory

Host

GPU Kernel

- Partition: GPU and host
- Separate GPU Memory
- Host manages OS services
- GPU cannot invoke syscalls
- Host operate on its memory
- Copy data to GPU memory
- Host-centric management
- High invocation costs
- GPU execute computation
- Copy back to host memory

What do GPU and enclave have in common?

Design an Enclave Application

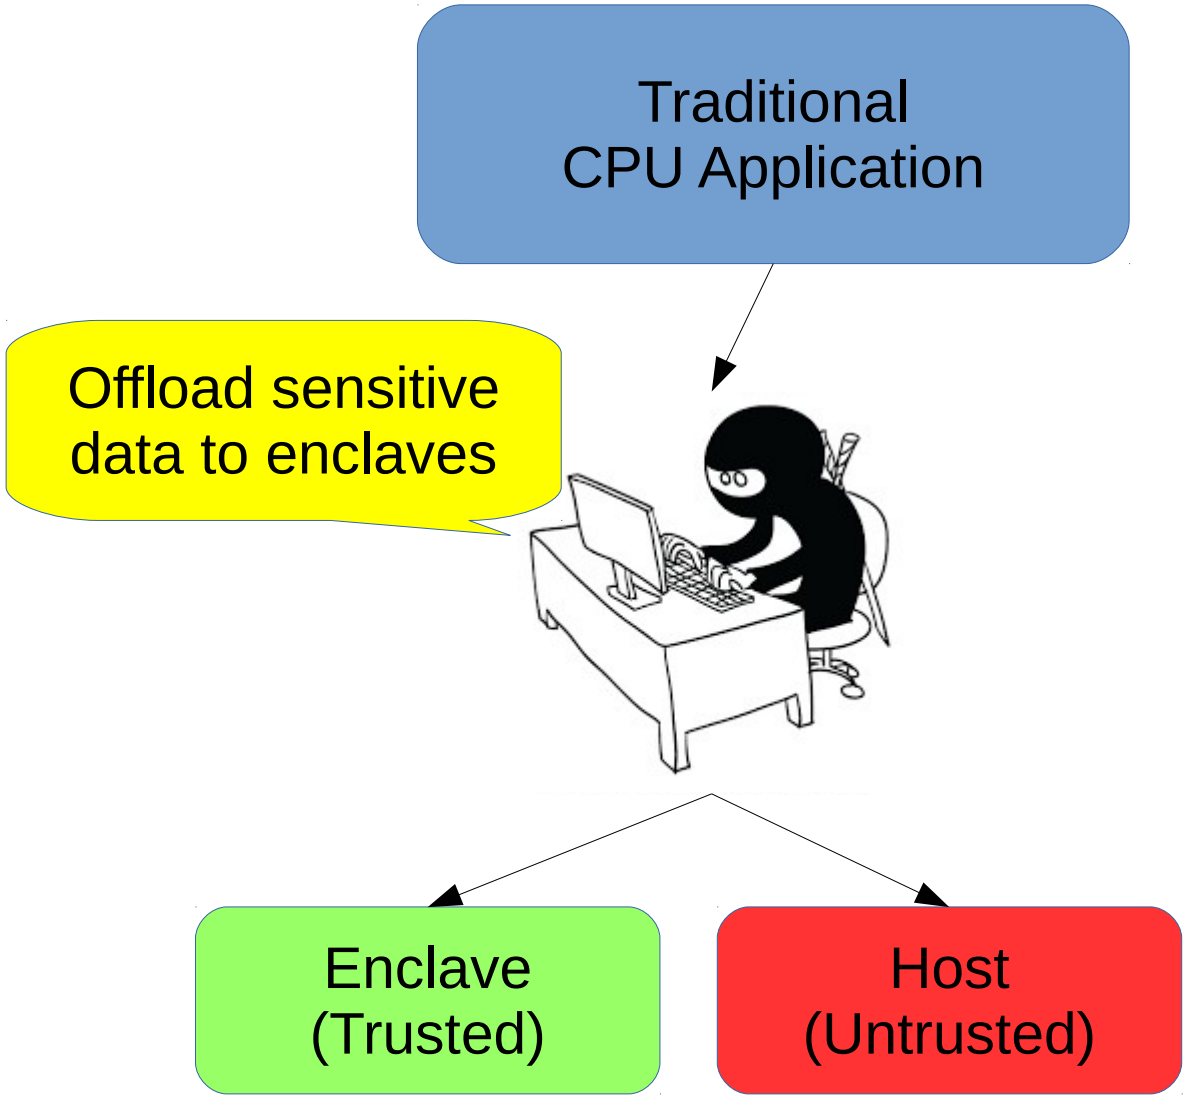

Partition:trusted and untrusted
Separate GPU Memory
Host manages OS services
GPU cannot invoke syscalls
Host operate on its memory
Copy data to GPU memory
Host-centric management
High invocation costs
GPU execute computation
Copy back to host memory

Private Reserved Memory

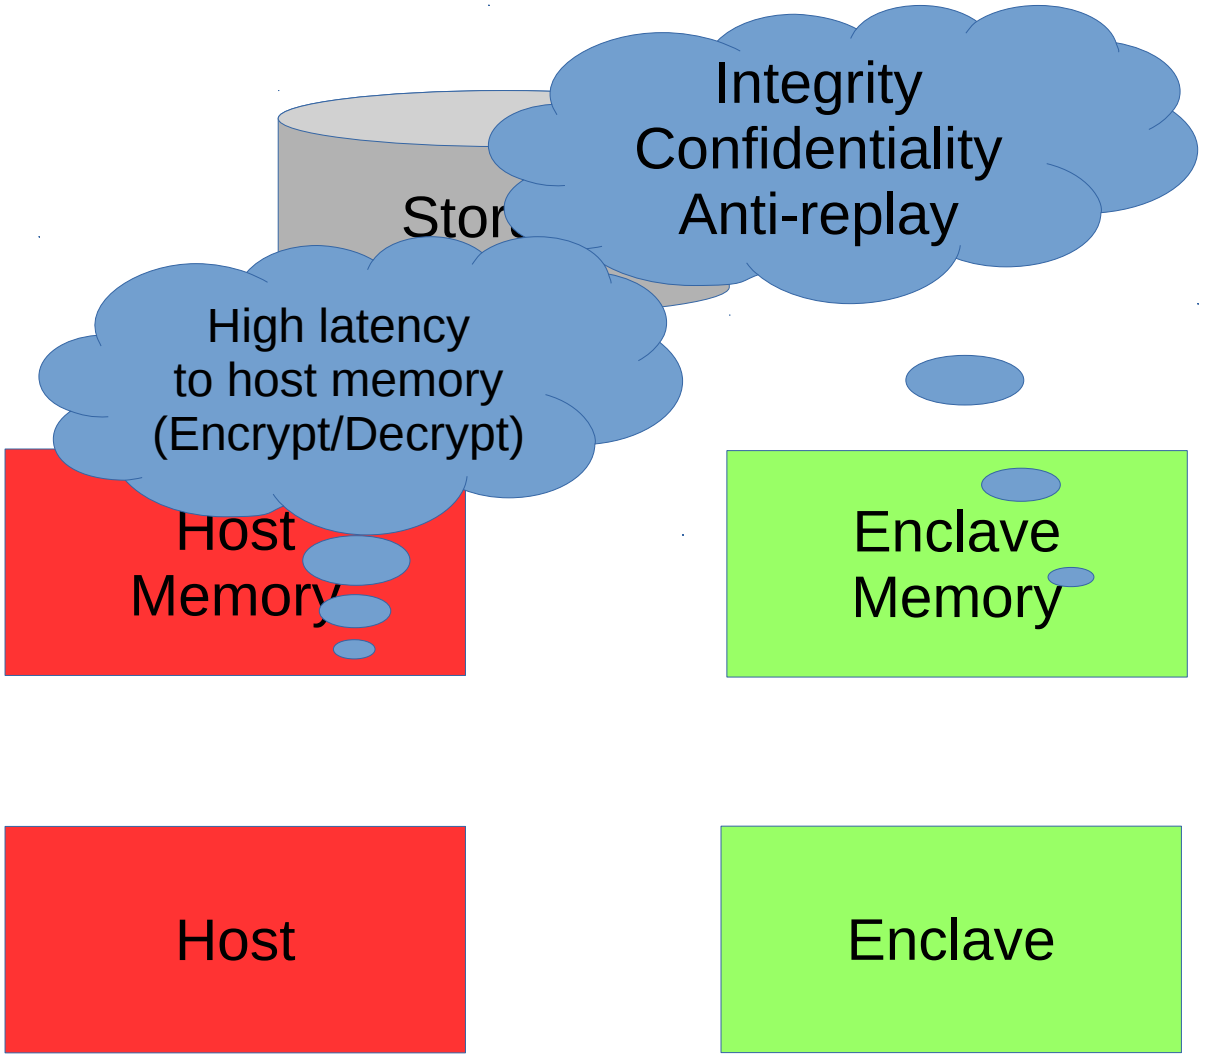

Partition:trusted and untrusted
Separate Enclave Memory
Host manages OS services
GPU cannot invoke syscalls
Host operate on its memory
Copy data to GPU memory
Host-centric management
High invocation costs
GPU execute computation
Copy back to host memory

The OS is untrusted

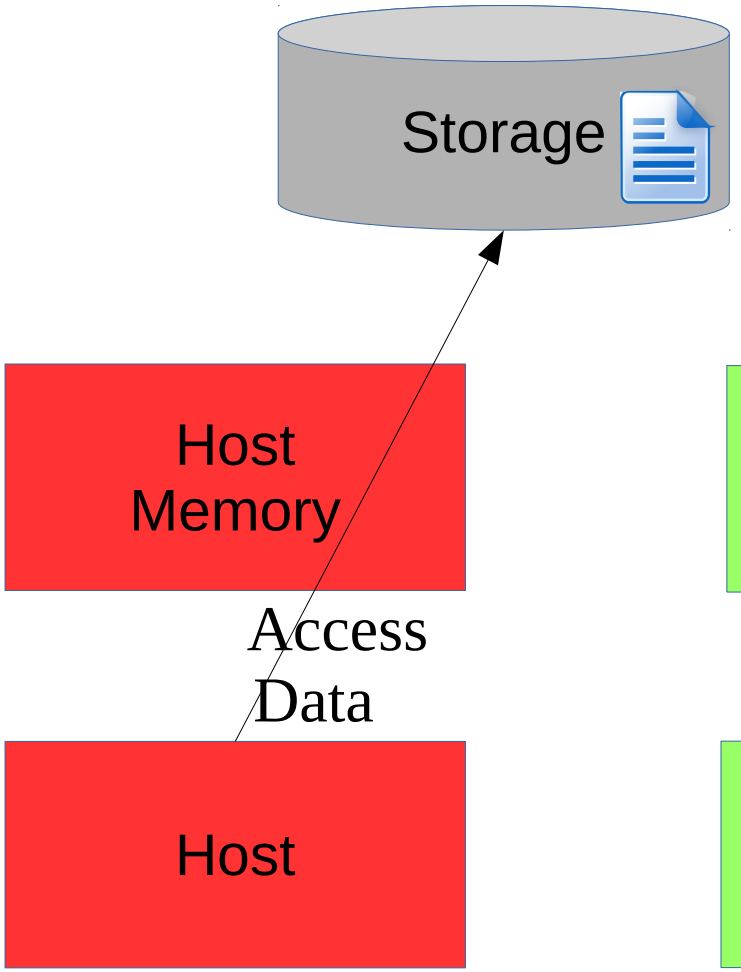

- Partition:trusted and untrusted
- Separate Enclave Memory
- Host manages OS services
- Enclave cannot invoke syscalls
- Host operate on its memory
- Copy data to GPU memory
- Host-centric management
- High invocation costs
- GPU execute computation
- Copy back to host memory

Untrusted code operates on untrusted memory

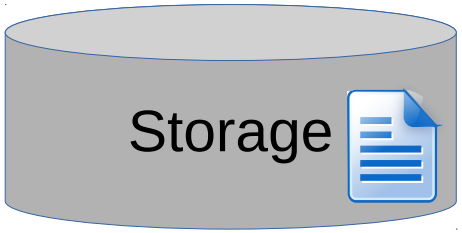

Host Memory

Enclave Memory

Host

Enclave

- Partition:trusted and untrusted
- Separate Enclave Memory
- Host manages OS services
- Enclave cannot invoke syscalls
- Host operate on its memory
- Copy data to GPU memory
- Host-centric management
- High invocation costs
- GPU execute computation
- Copy back to host memory

Trusted code operates on Trusted memory

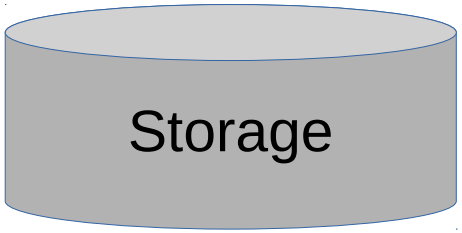

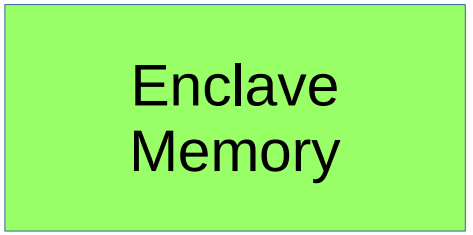

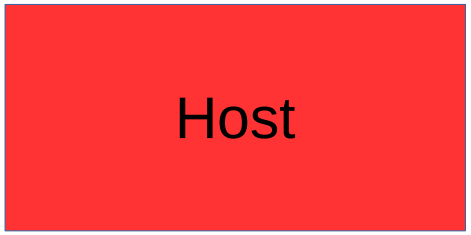

- Partition:trusted and untrusted
- Separate Enclave Memory
- Host manages OS services
- Enclave cannot invoke syscalls
- Host operate on its memory
- Copy data to enclave memory
- Host-centric management
- High invocation costs
- GPU execute computation
- Copy back to host memory

Host-centric management

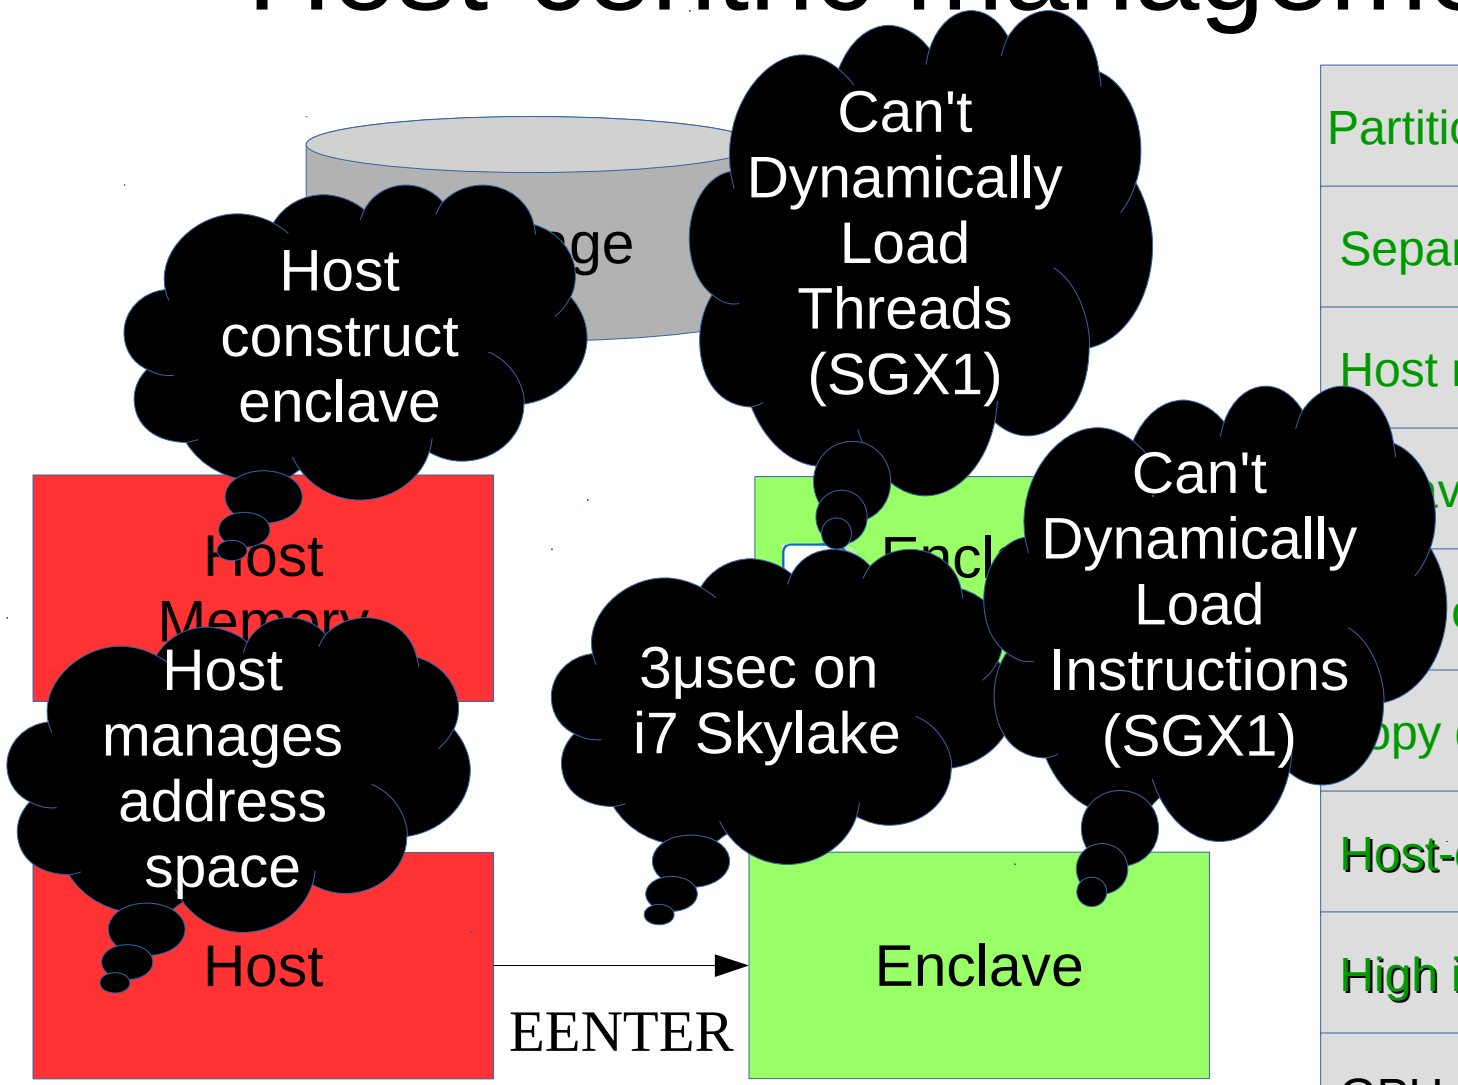

- Partition:trusted and untrusted
- Separate Enclave Memory
- Host manages OS services
- ve cannot invoke syscalls
- operate on its memory
- py data to enclave memory
- ~~Host-centric management~~
- ~~High invocation costs~~
- GPU execute computation
- Copy back to host memory

Isolated execution

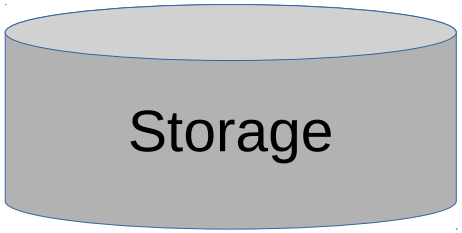

Host Memory


Enclave Memory

Host

Enclave

- Partition:trusted and untrusted
- Separate Enclave Memory
- Host manages OS services
- Enclave cannot invoke syscalls
- Host operate on its memory
- Copy data to enclave memory
- Host-centric management
- High invocation costs
- Enclave execute computation
- Copy back to host memory

Communication through untrusted memory

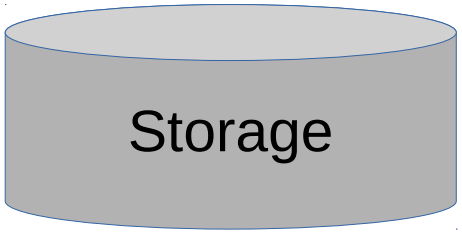

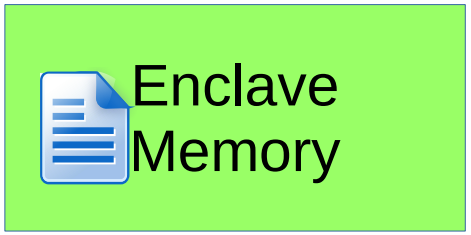

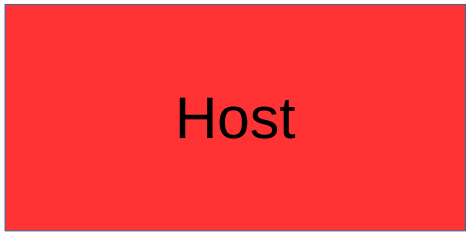

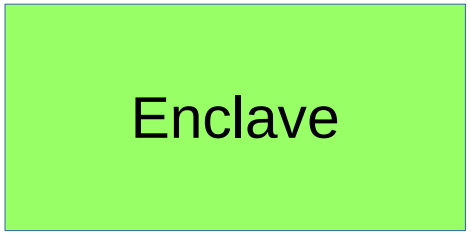

- Partition:trusted and untrusted
- Separate Enclave Memory
- Host manages OS services
- Enclave cannot invoke syscalls
- Host operate on its memory
- Copy data to enclave memory
- Host-centric management
- High invocation costs
- Enclave execute computation
- Copy back to host memory

The reason is... Isolation by design

Enclaves use strong isolation to provide strong security

Accelerators run on different hardware
Accelerators are isolated by necessity

Effect on processes' runtime

- Simplified parameter server in and out of enclave
  - Network server
  - Private model & data
  - Store model in hash table
  - Clients send 100k random requests to update items
  - Server issues recv() to get requests and update
  - Enclave encapsulate recv() in OCALL.

Simplified parameter server

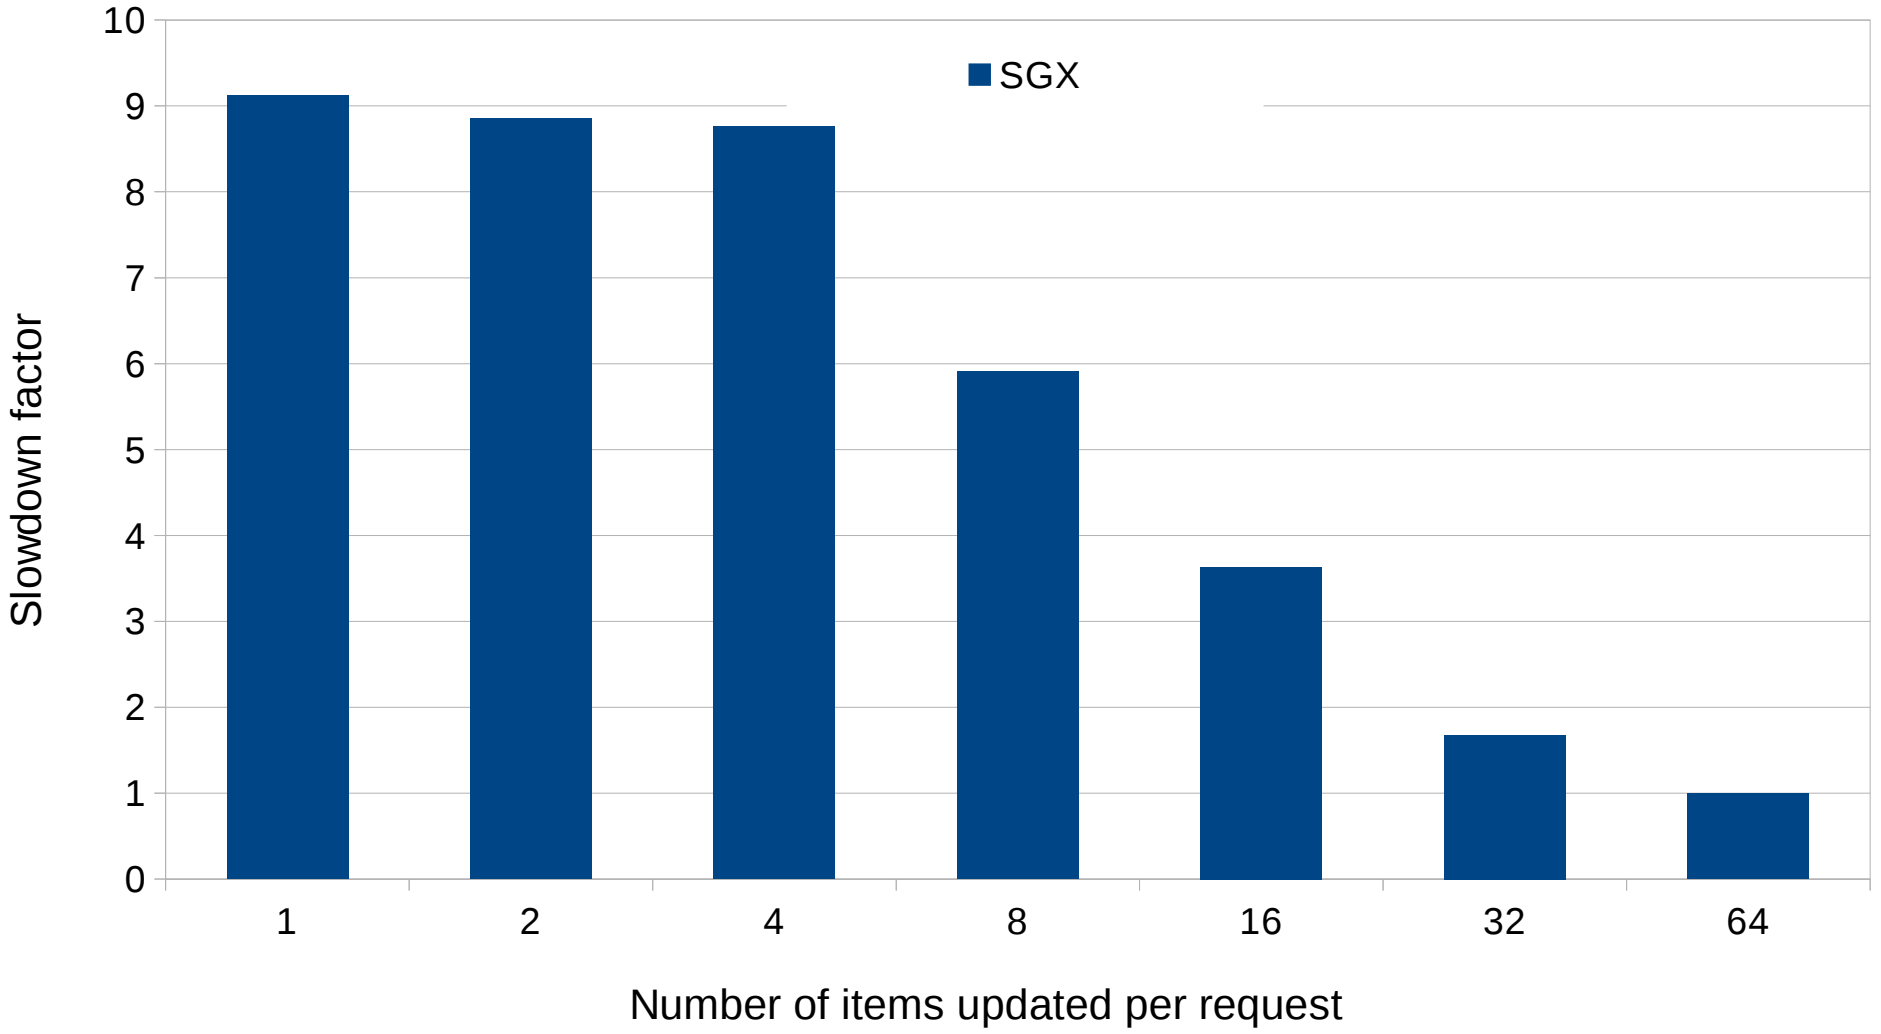

What can we learn from GPUs?

SGXIO: Overhead-free I/O from enclaves

Based on GPUfs [ASPLOS'2013]

- RPC communication infrastructure

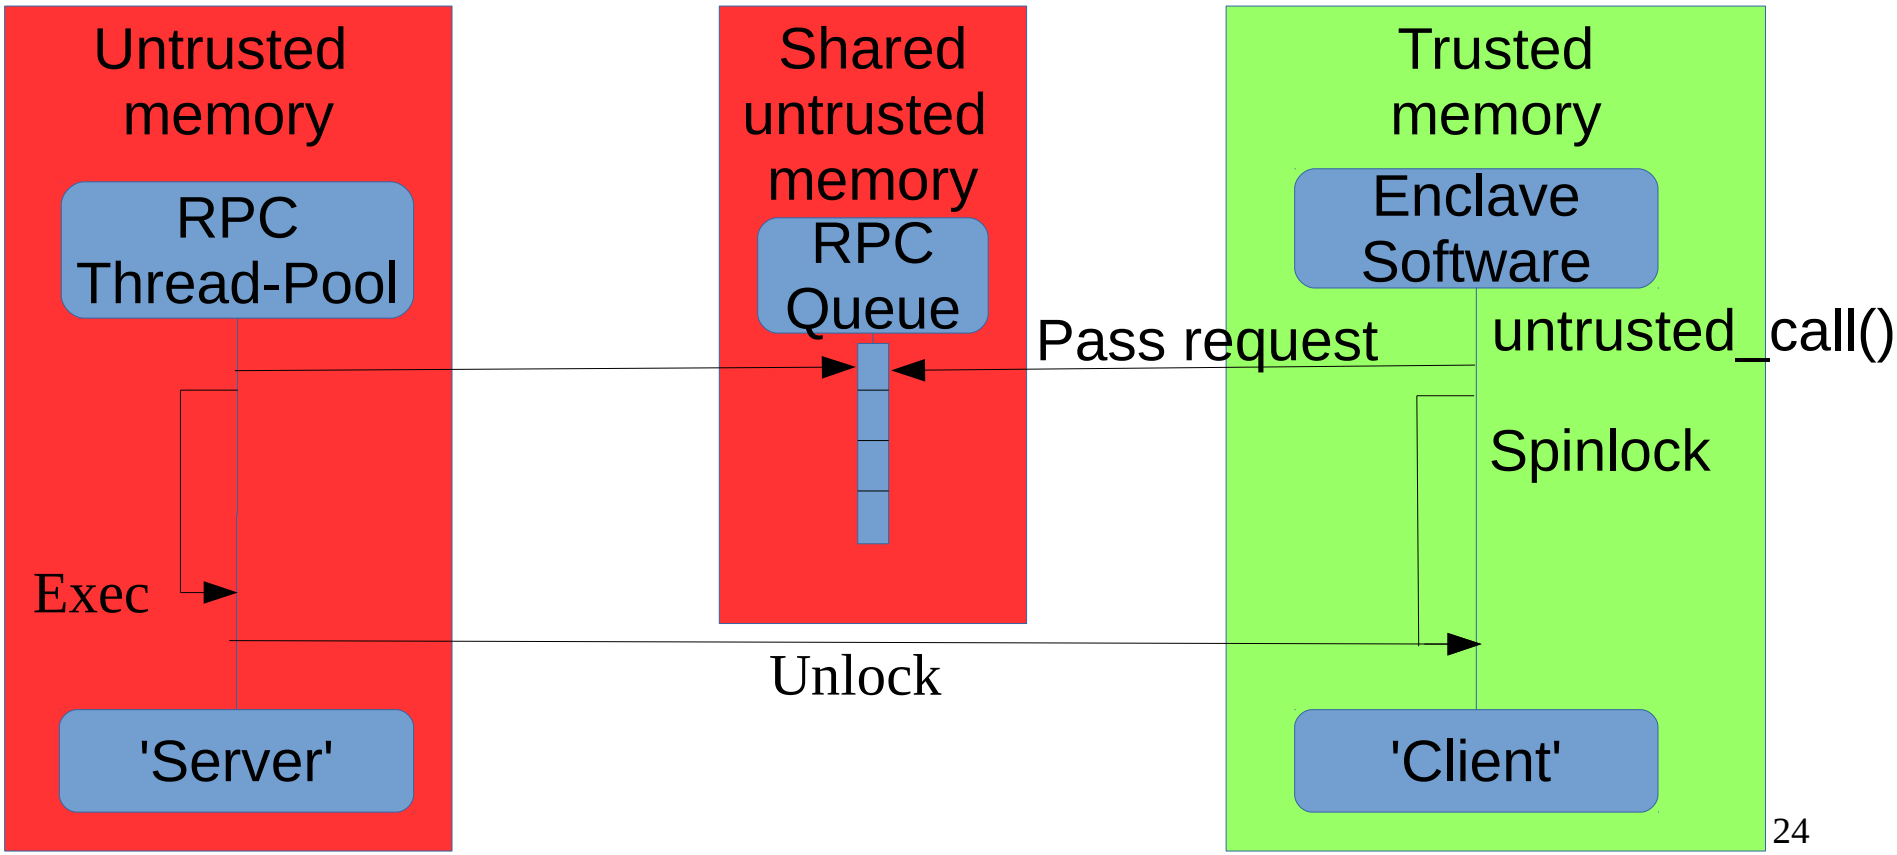

Simplified parameter server

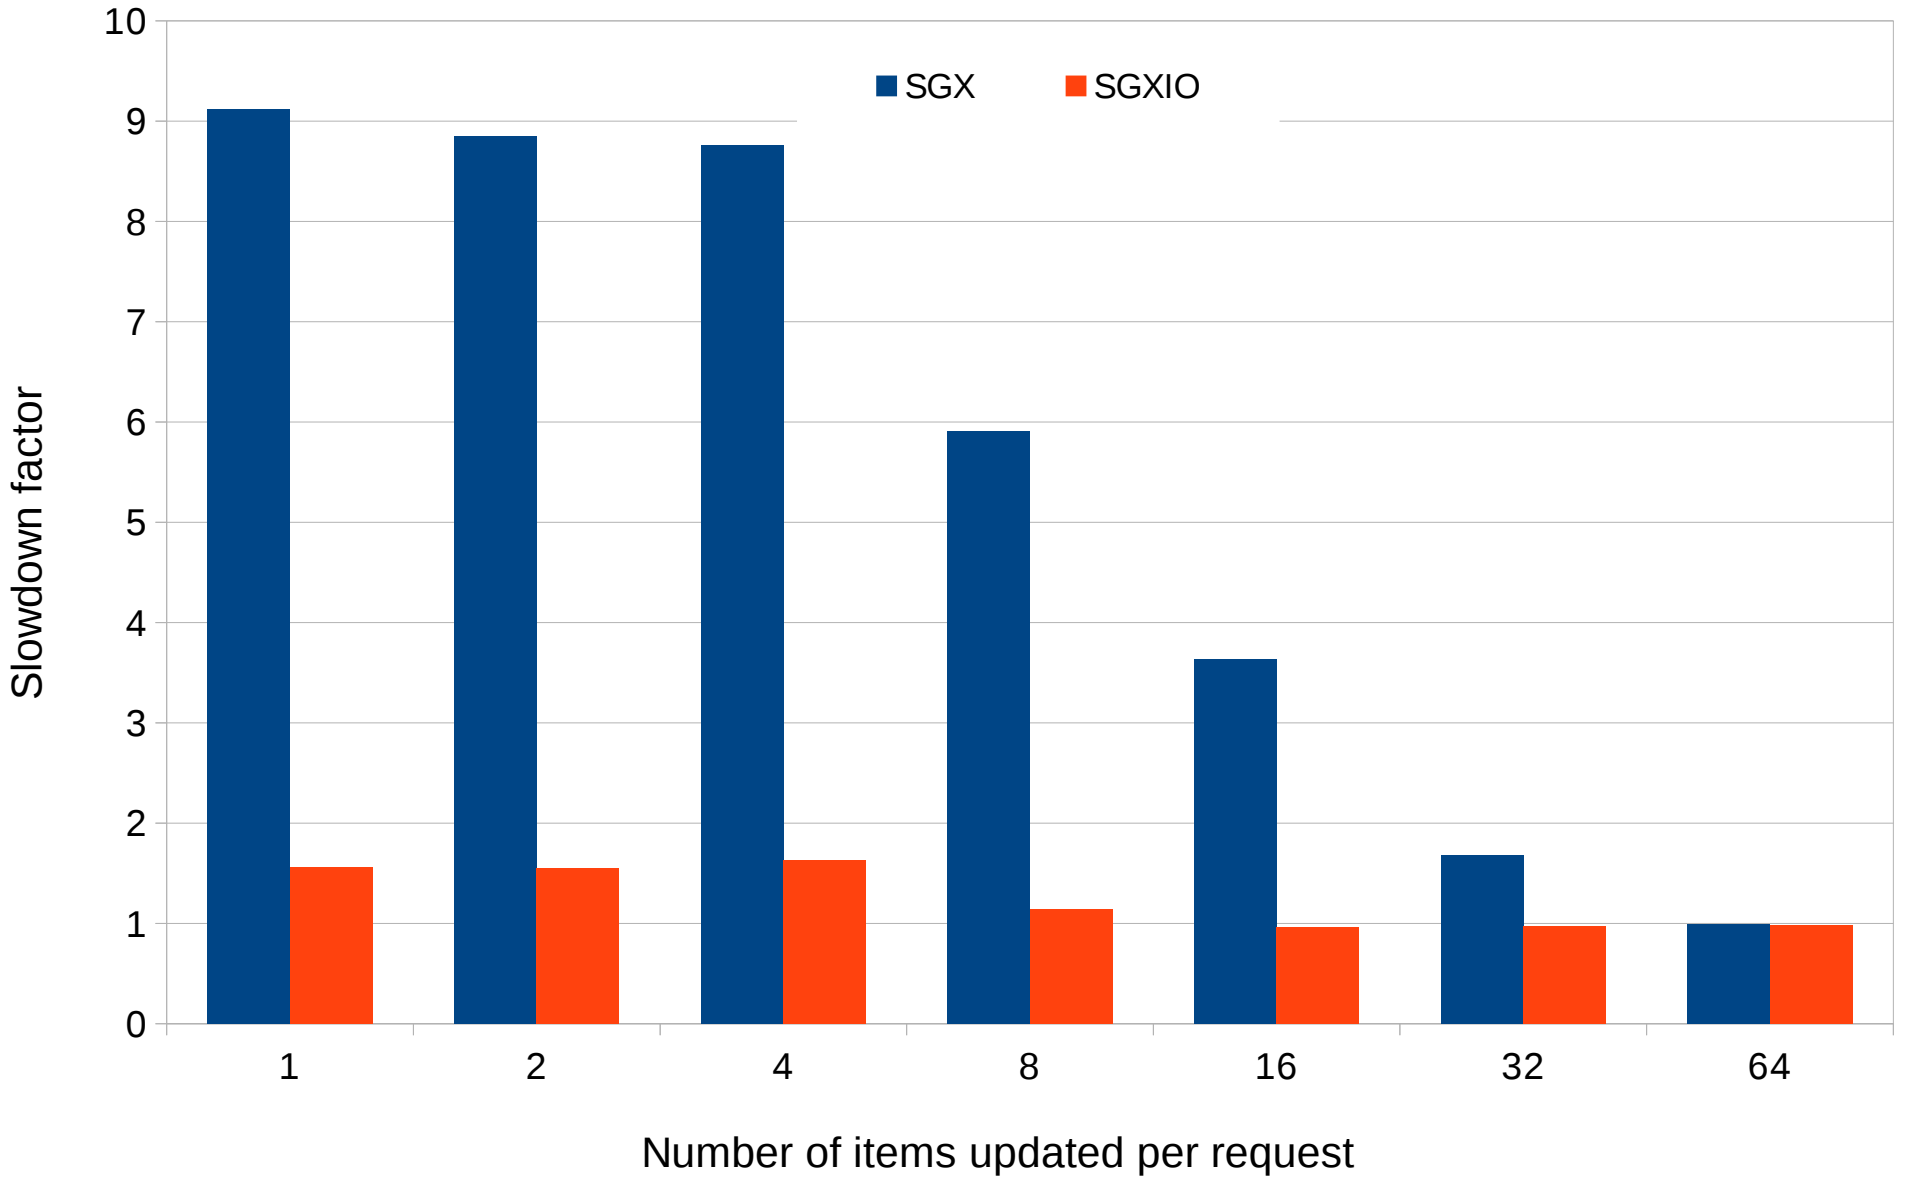

Same, Same but different

- Enclaves **are not** traditional accelerators
  - Latency to host memory
    - MMU vs PCIe
  - Atomic instructions shared with the host
  - Internal management
    - E.g., Enclave Thread-scheduler

Enclaves bring new possibilities

Retrofitting accelerators' ideas for enclaves

- SGXIO: OS services for enclaves
- Asynchronous DMA host copies
- Non-blocking enclave launches
- In-enclave virtual memory management

Thank you!

Questions?

shmeni@tx.technion.ac.il
mark@ee.technion.ac.il