Side-Channel Attacks on Optane Persistent Memory

Sihang Liu, **Suraaj Kanniwadi**, Martin Schwarzl, Andreas Kogler,
Daniel Gruss, Samira Khan

Usenix Security Symposium 2023

Overview: Motivation

Overview: Motivation

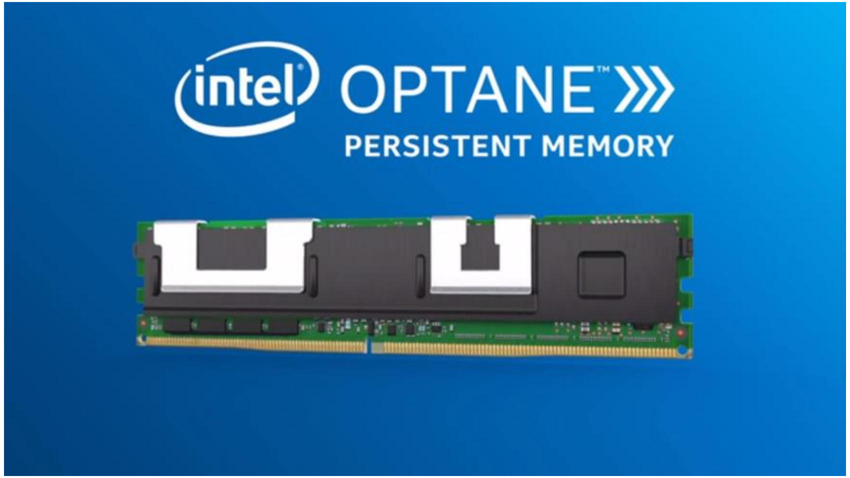

Overview: Motivation

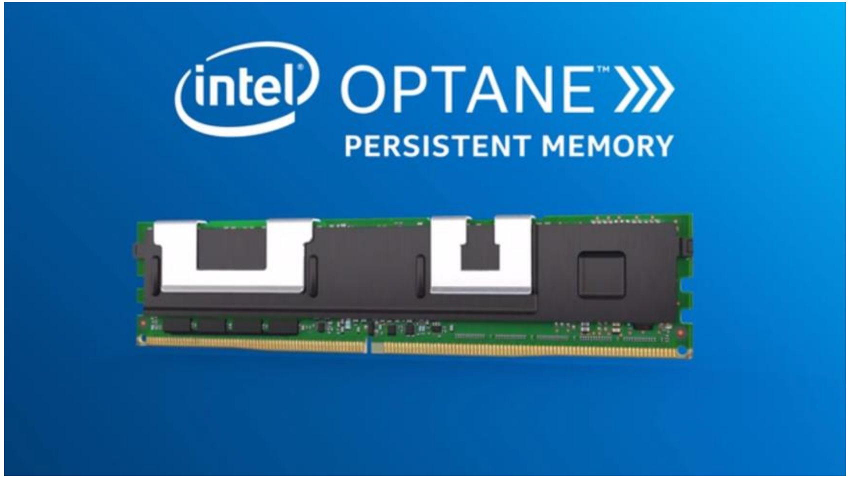

Overview: Contributions

Overview: Contributions

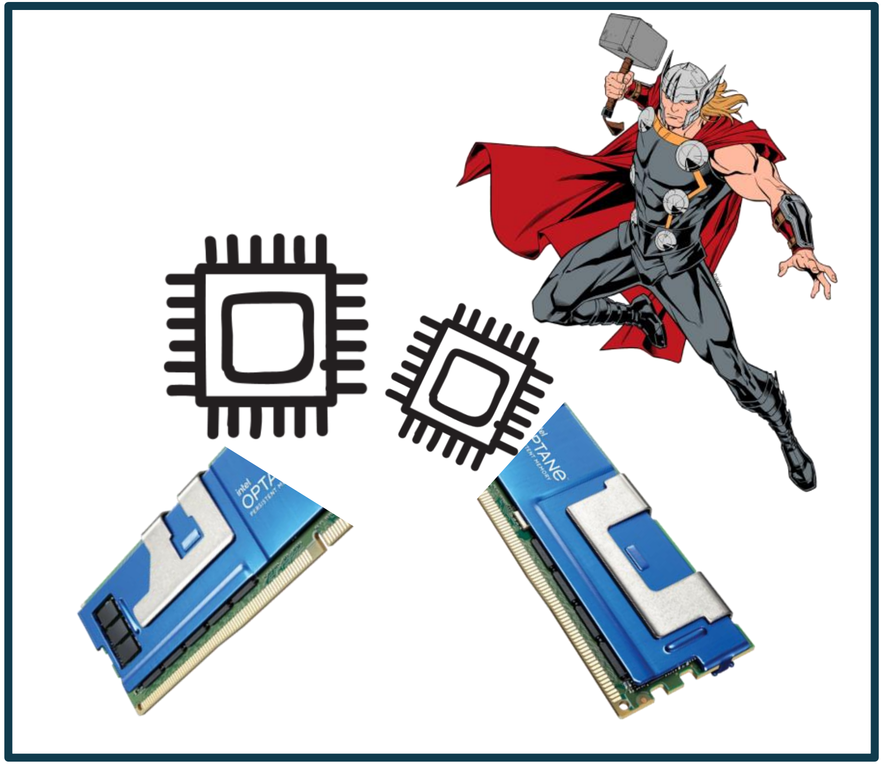

Overview: Contributions

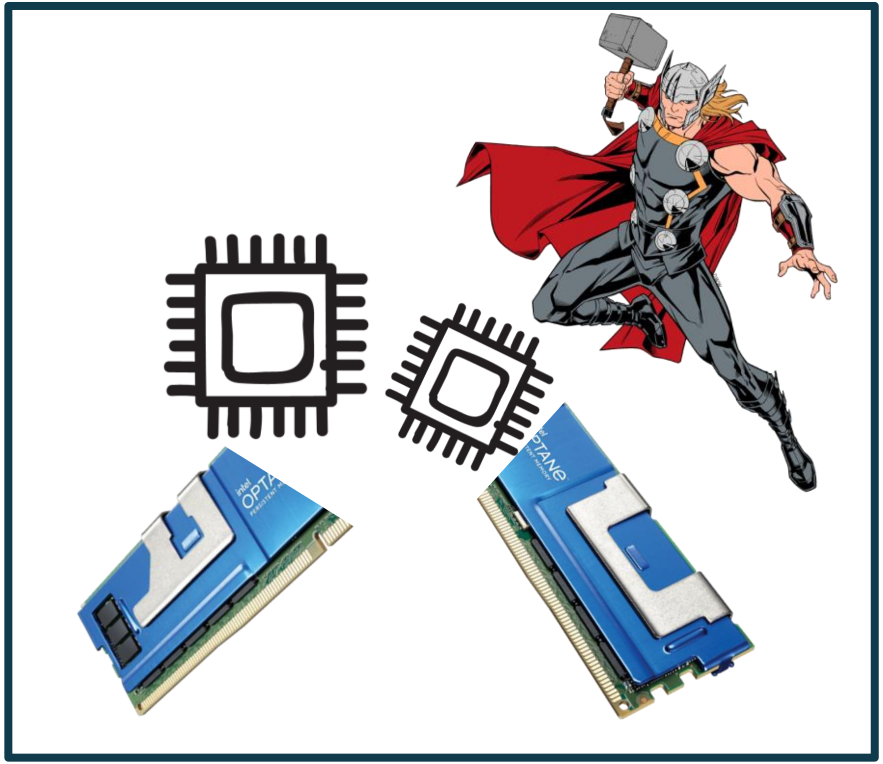

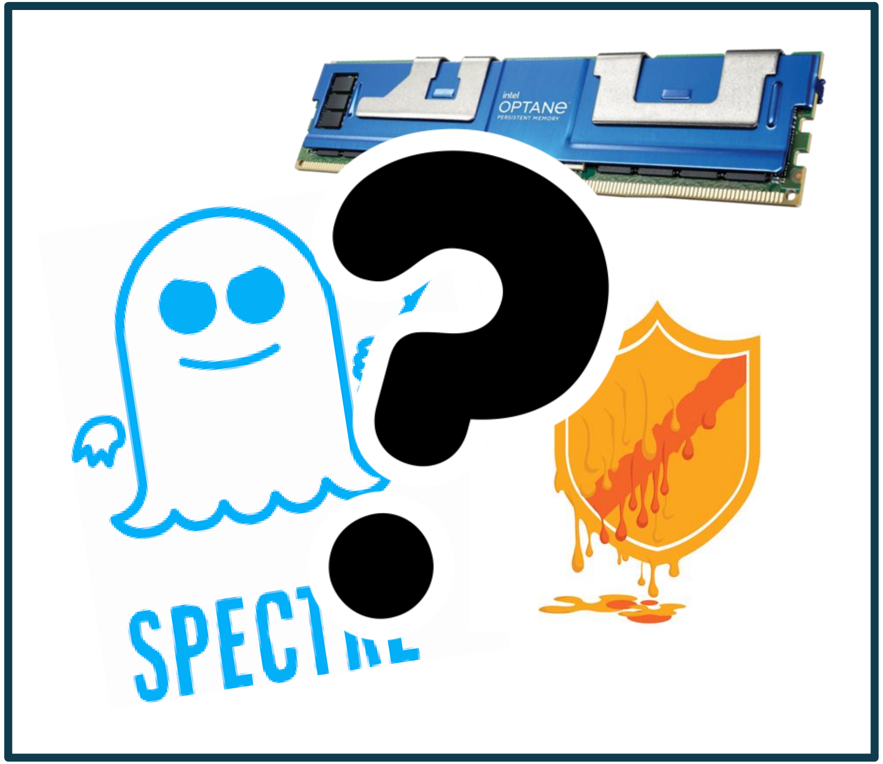

Background

Optane, Persistent Memory, and Side Channels

Optane Persistent Memory

Optane Persistent Memory

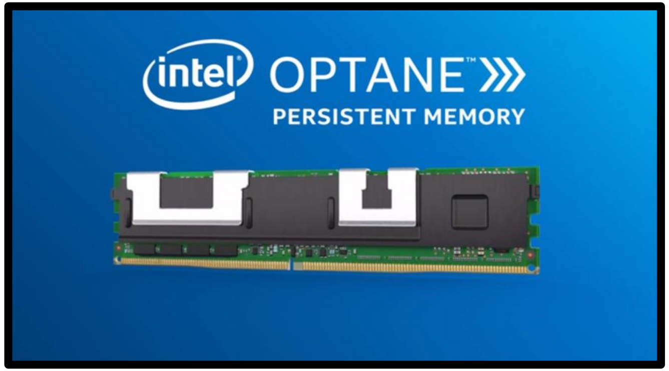

Optane Persistent Memory

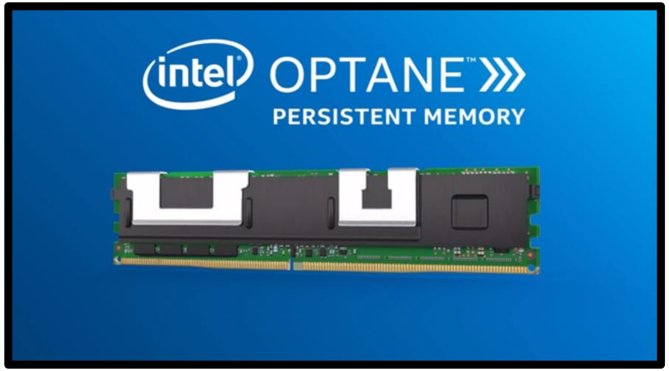

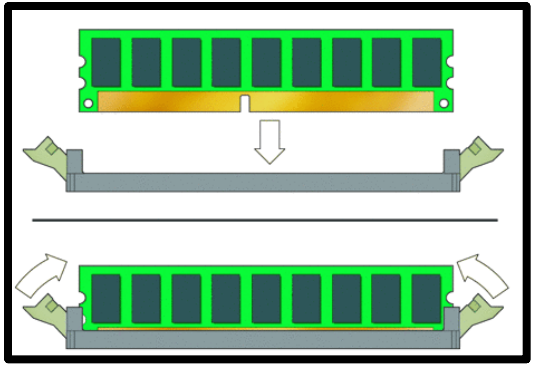

Optane Persistent Memory

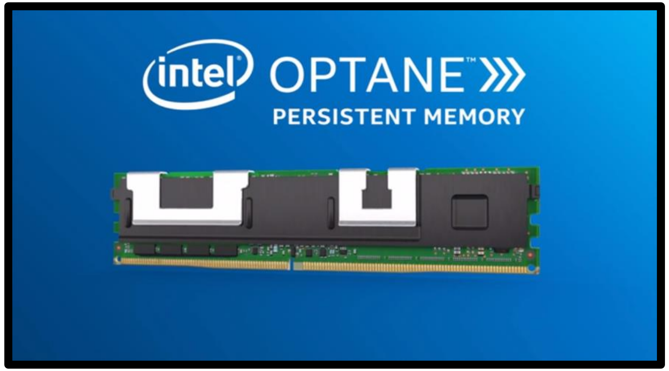

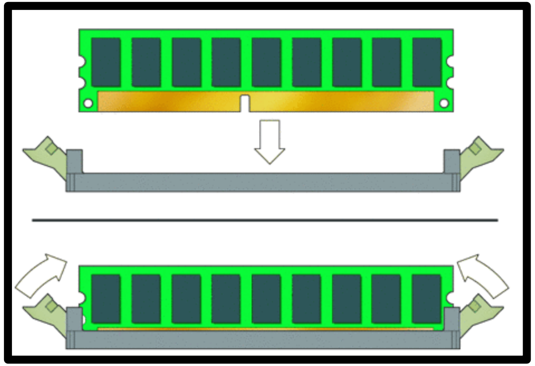

Optane Persistent Memory

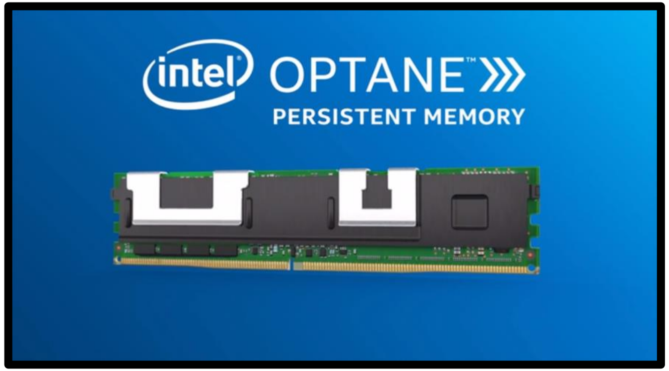

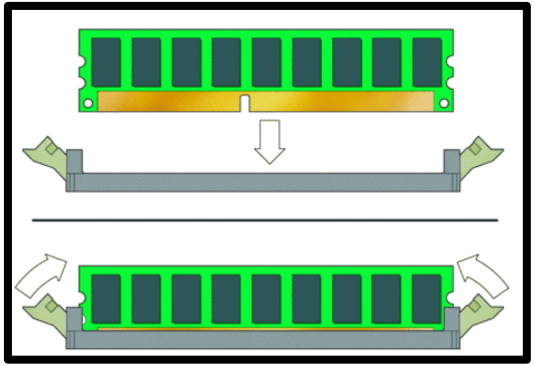

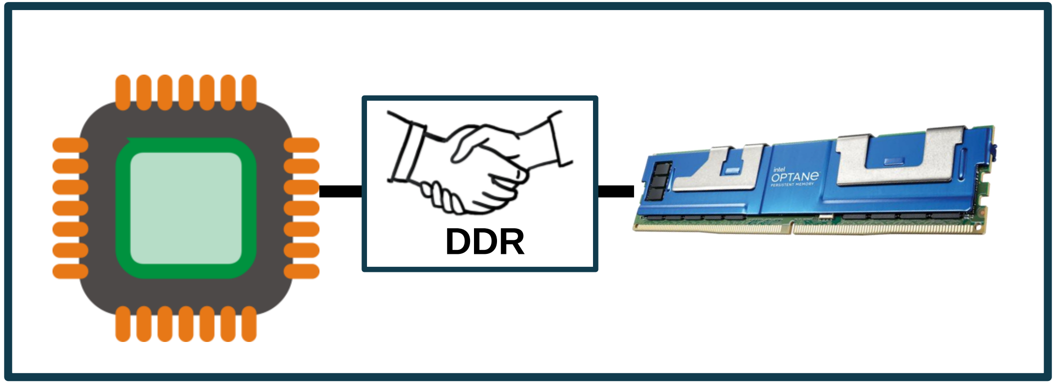

Optane Persistent Memory

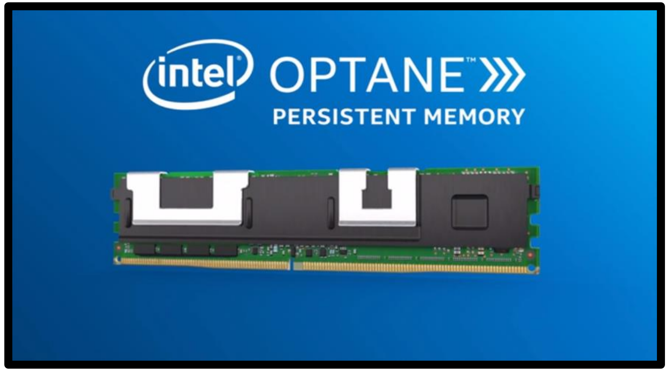

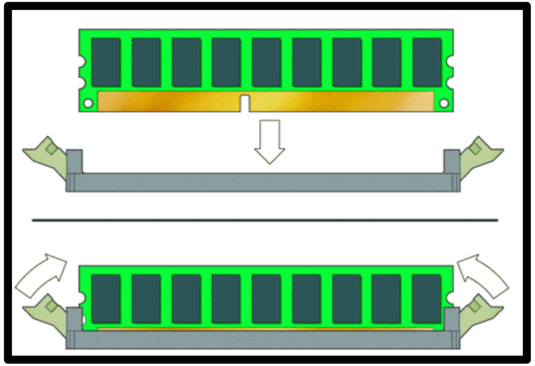

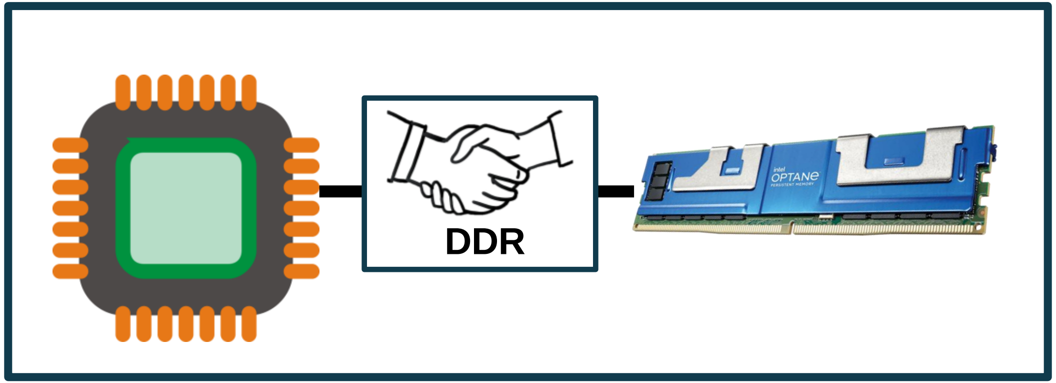

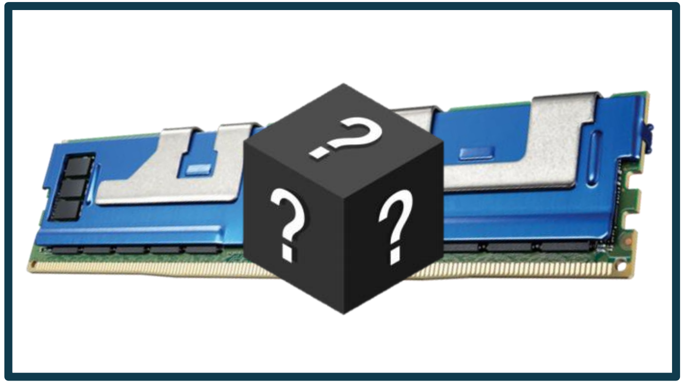

Stable storage with Direct Access

Stable storage with Direct Access

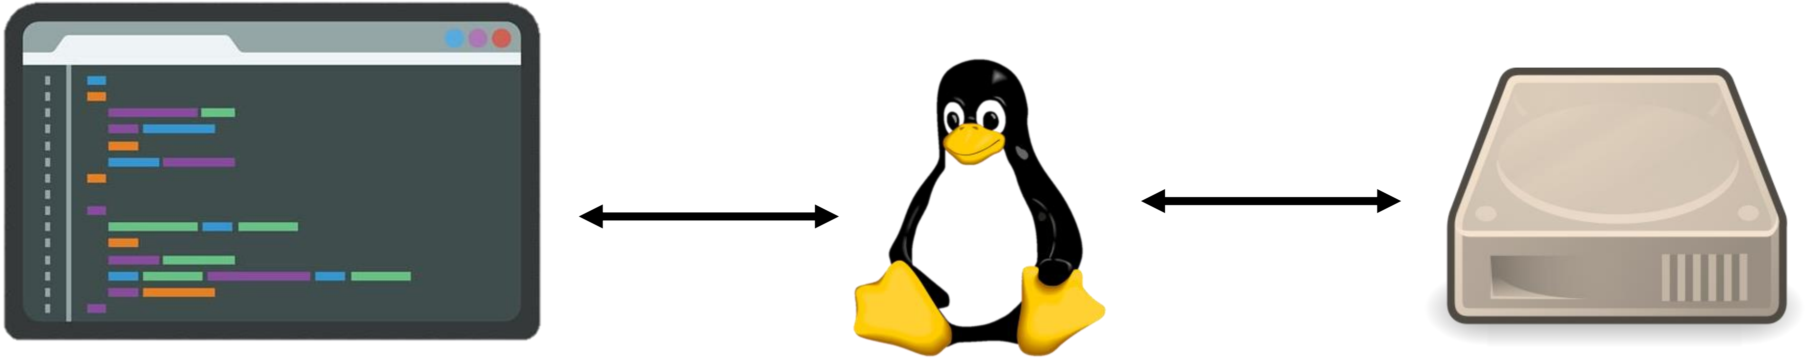

Stable storage with Direct Access

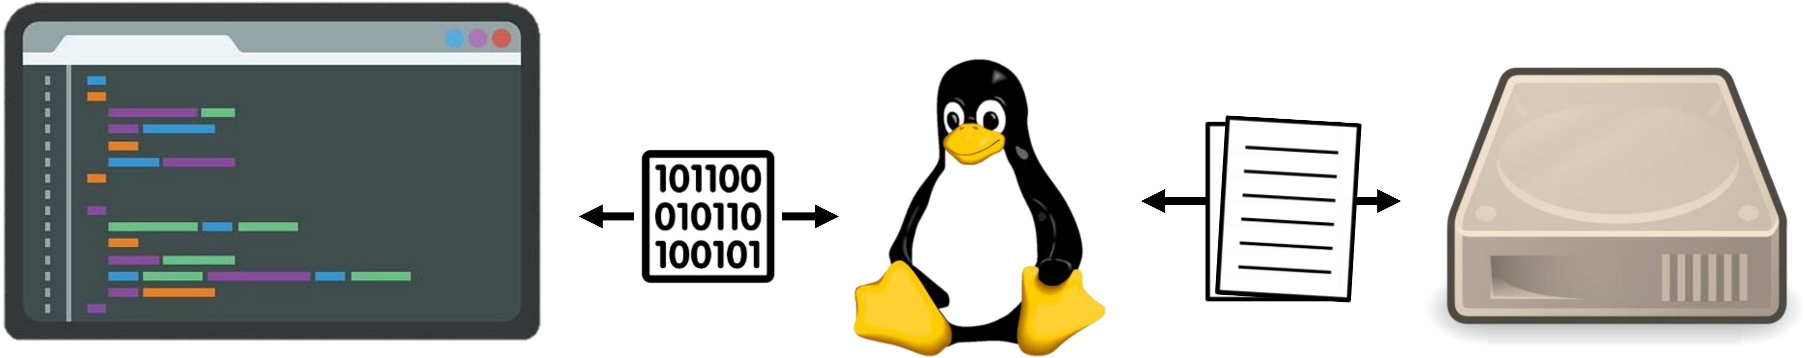

Stable storage with Direct Access

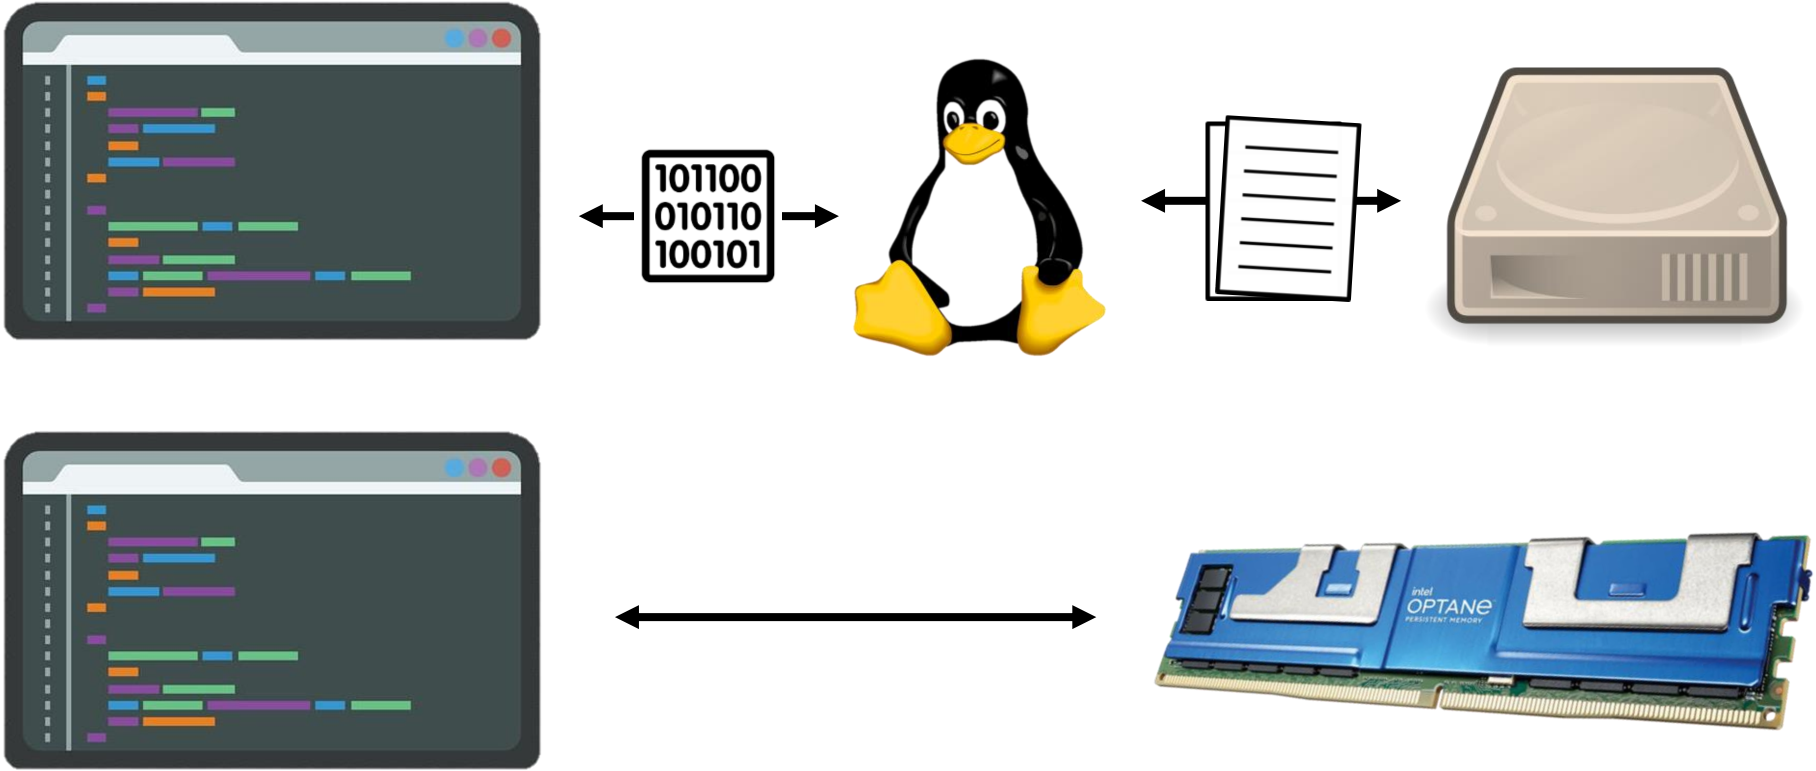

Stable storage with Direct Access

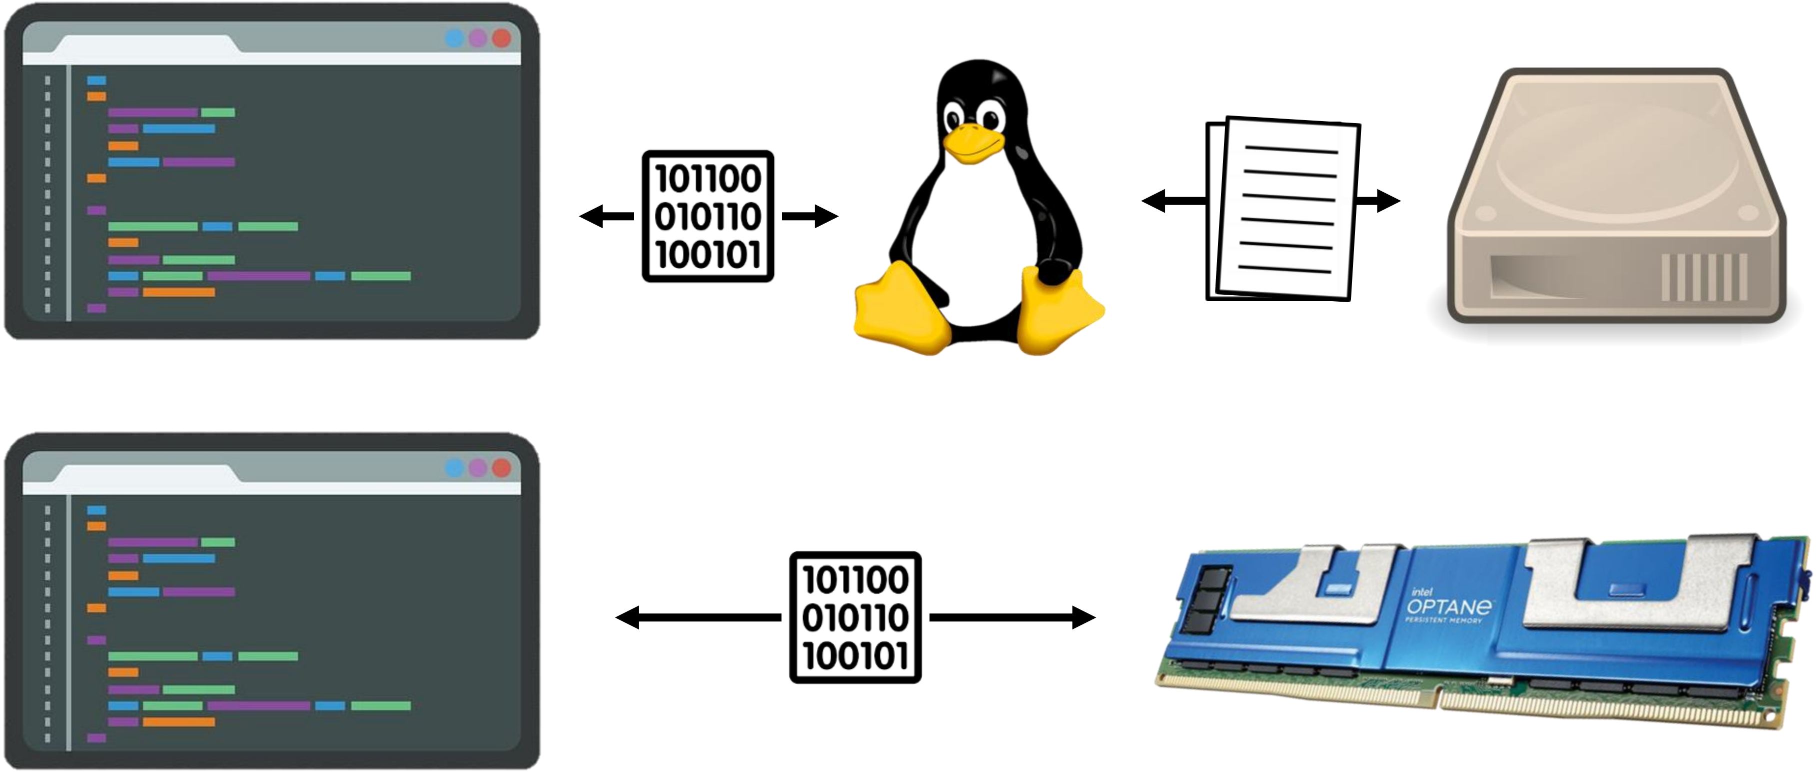

In the system heirarchy

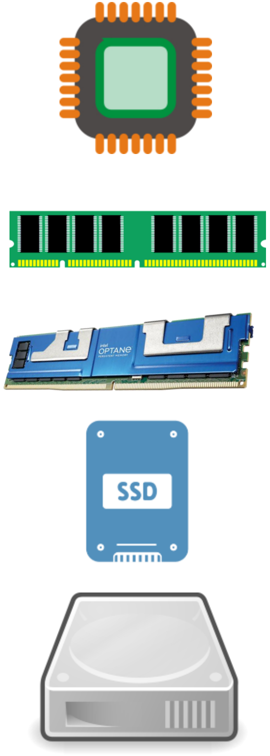

In the system heirarchy

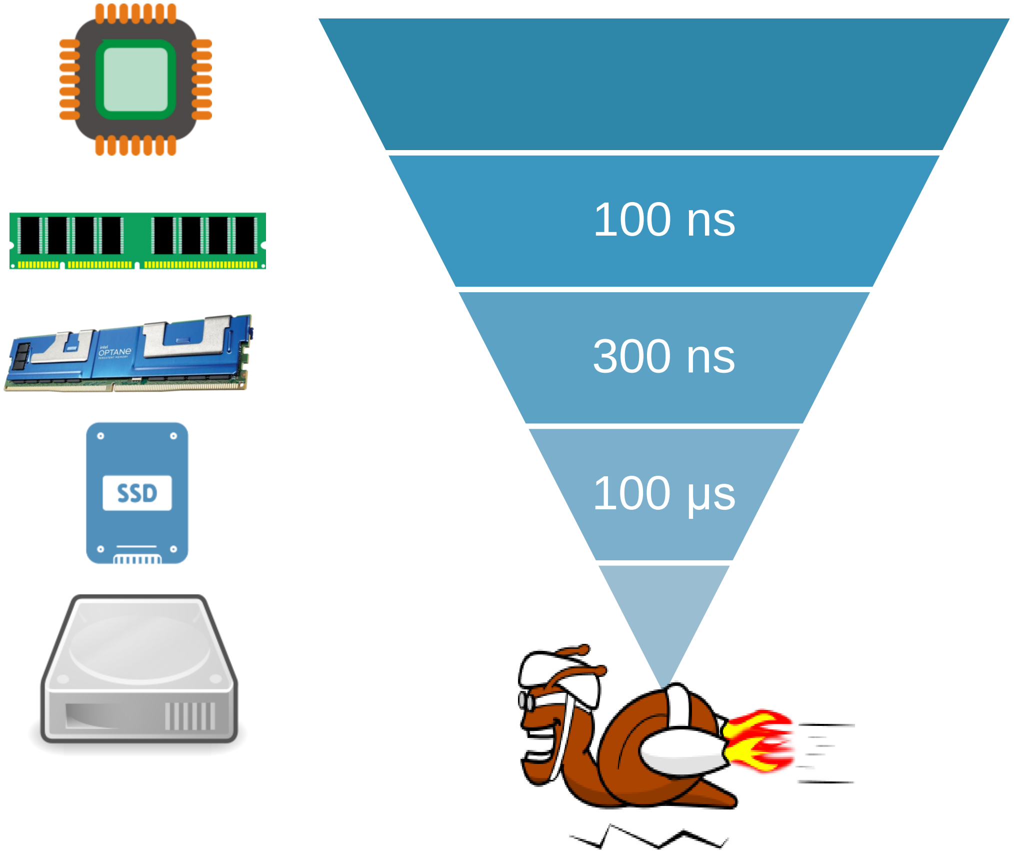

100 ns

300 ns

100 µs

In the system heirarchy

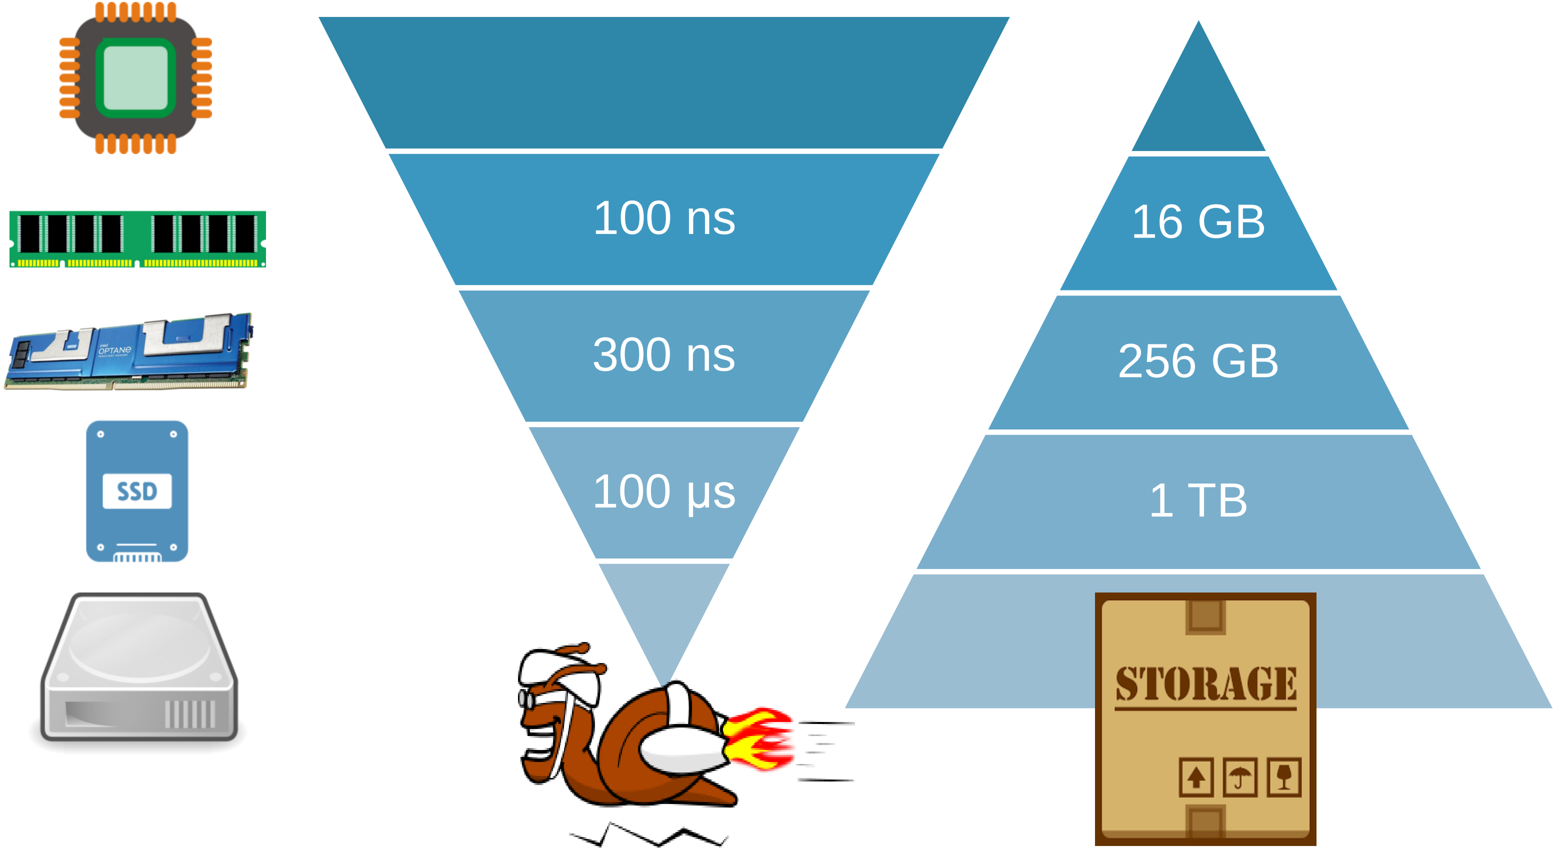

Side Channel Attack Recipe

Side Channel Attack Recipe

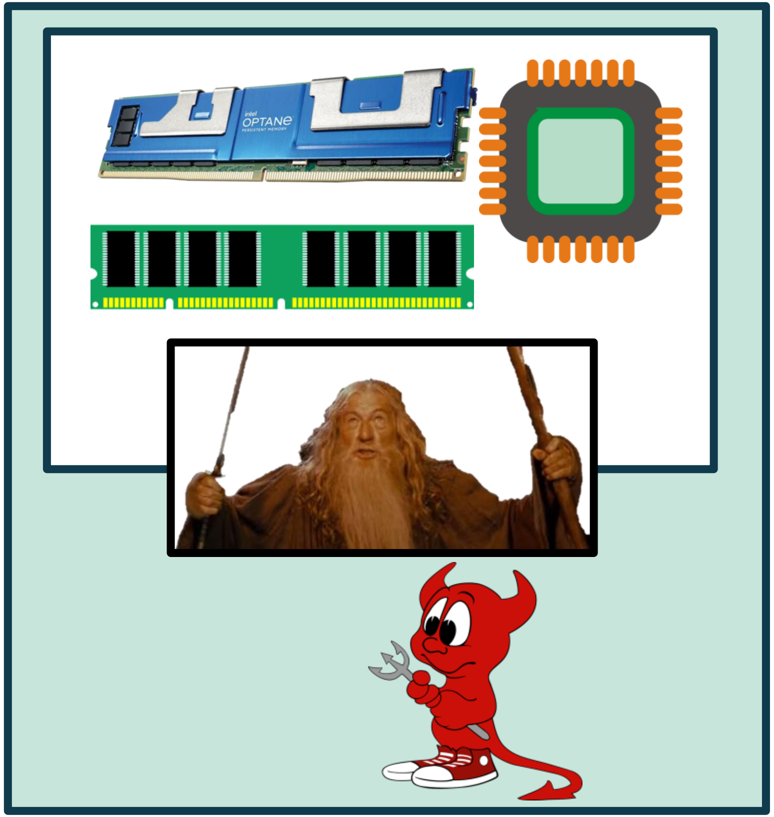

Side Channel Attack Recipe

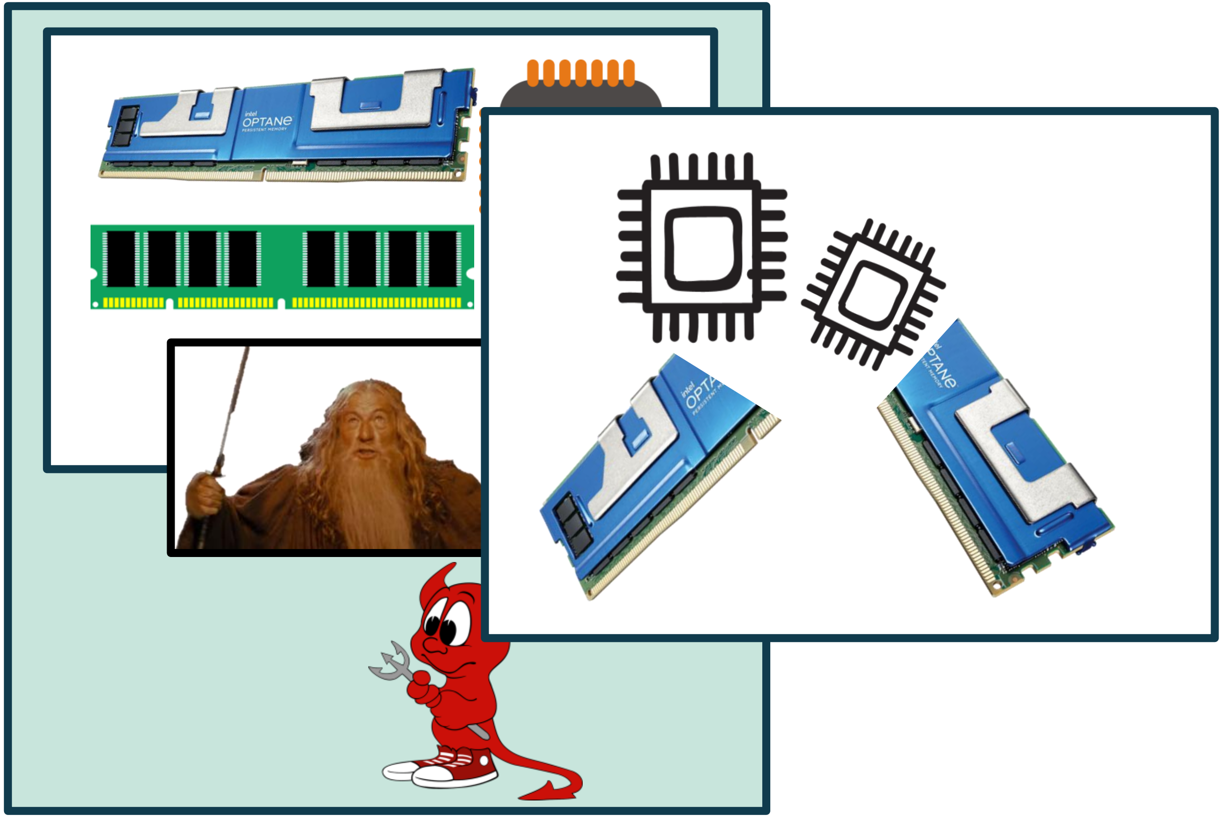

Side Channel Attack Recipe

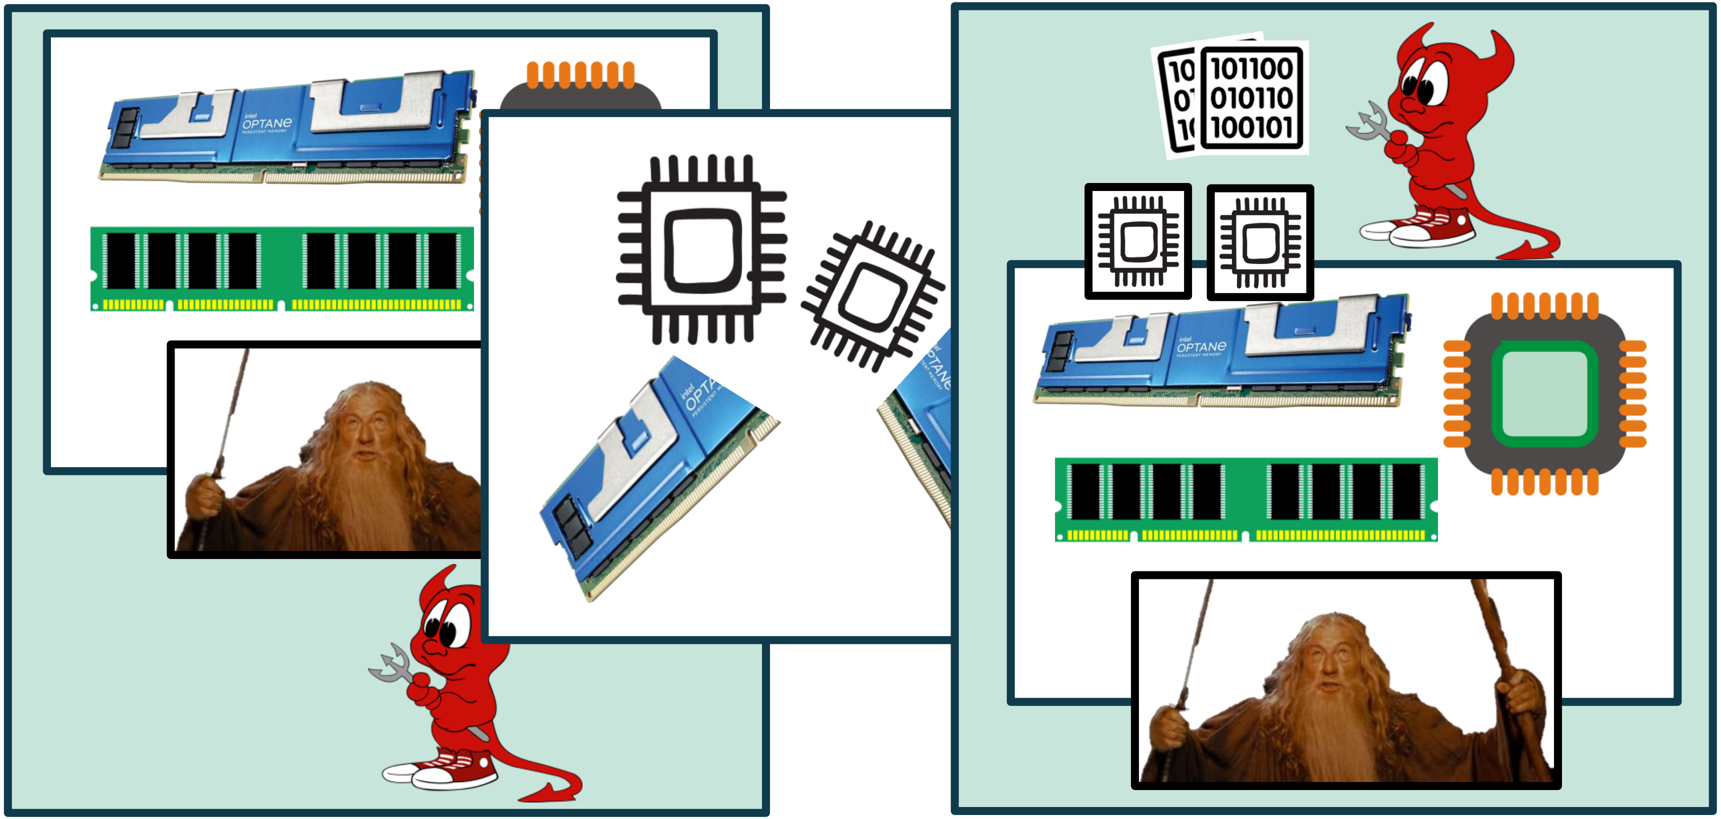

Reverse-Engineering of Optane

A glimpse into the Optane DIMM

Optane: Prior Work

Optane: Prior Work

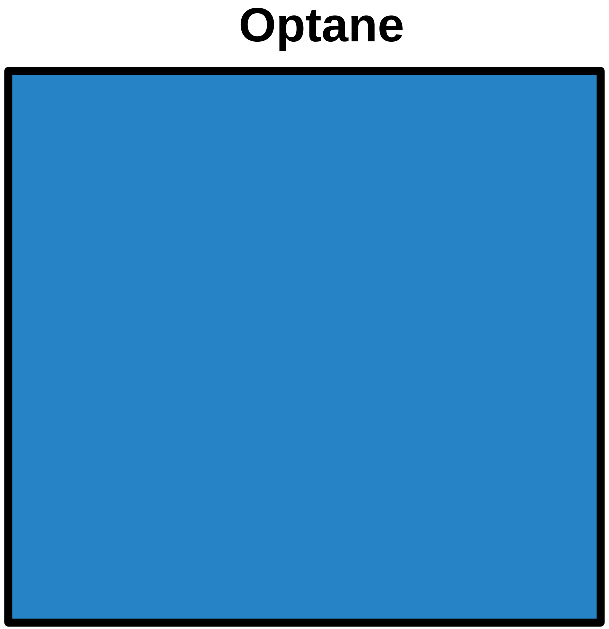

Optane: Prior Work

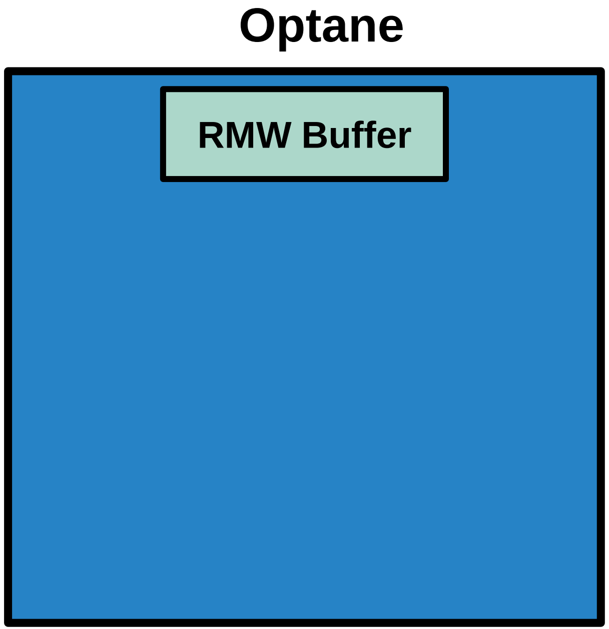

Optane: Prior Work

Optane

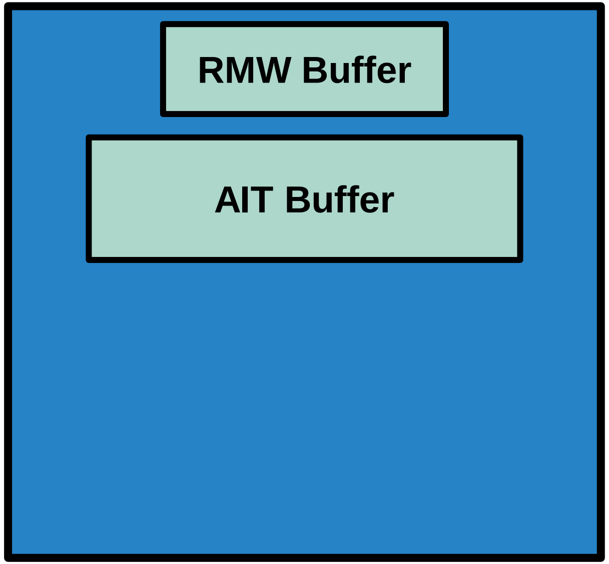

Optane: Prior Work

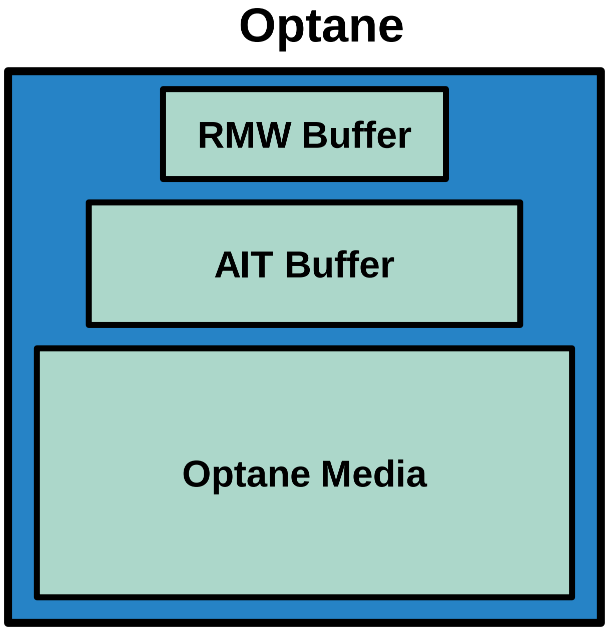

Optane: Prior Work

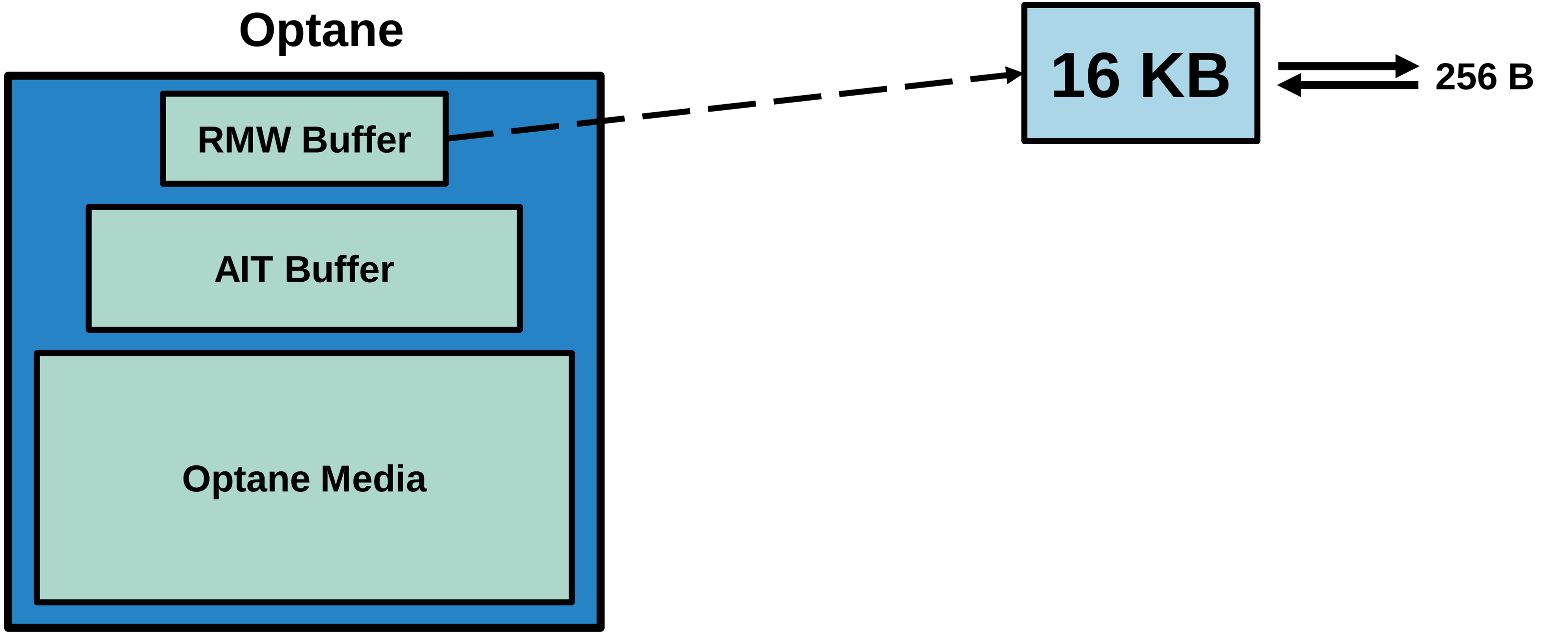

Optane: Prior Work

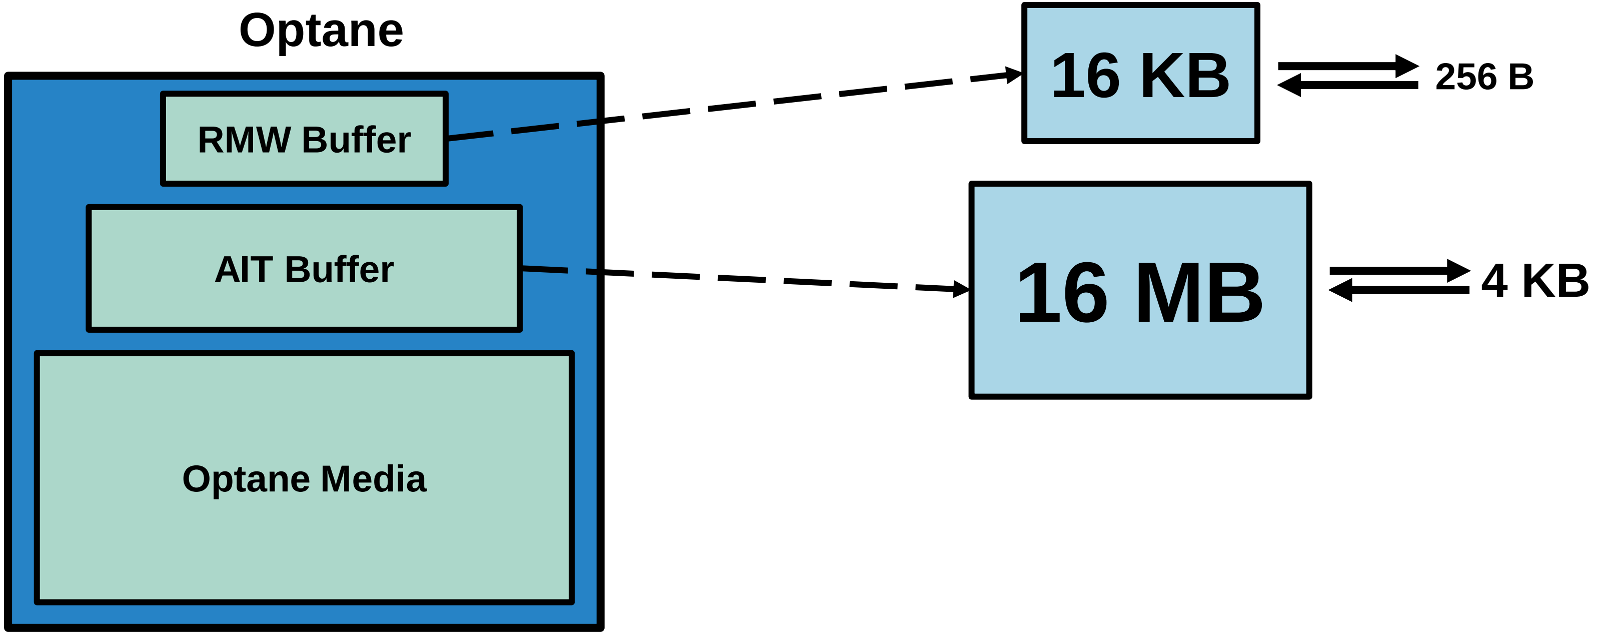

Optane: Prior Work

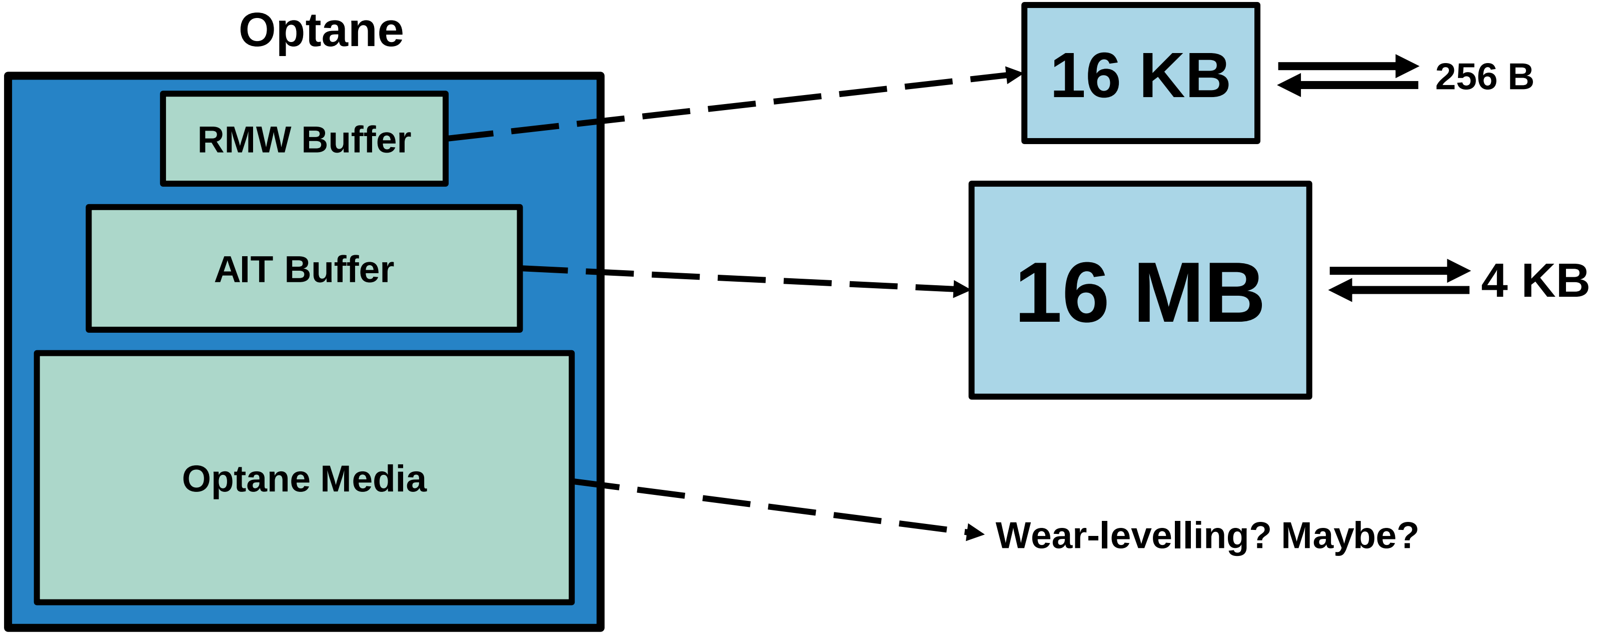

Optane: We have more!

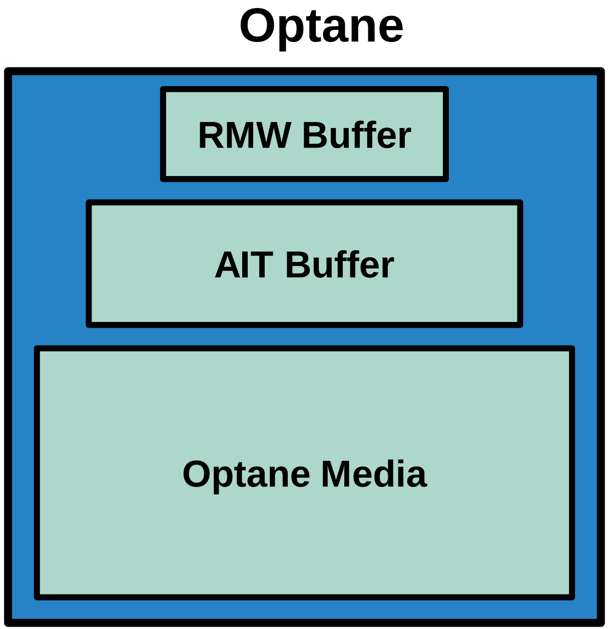

Optane: We have more!

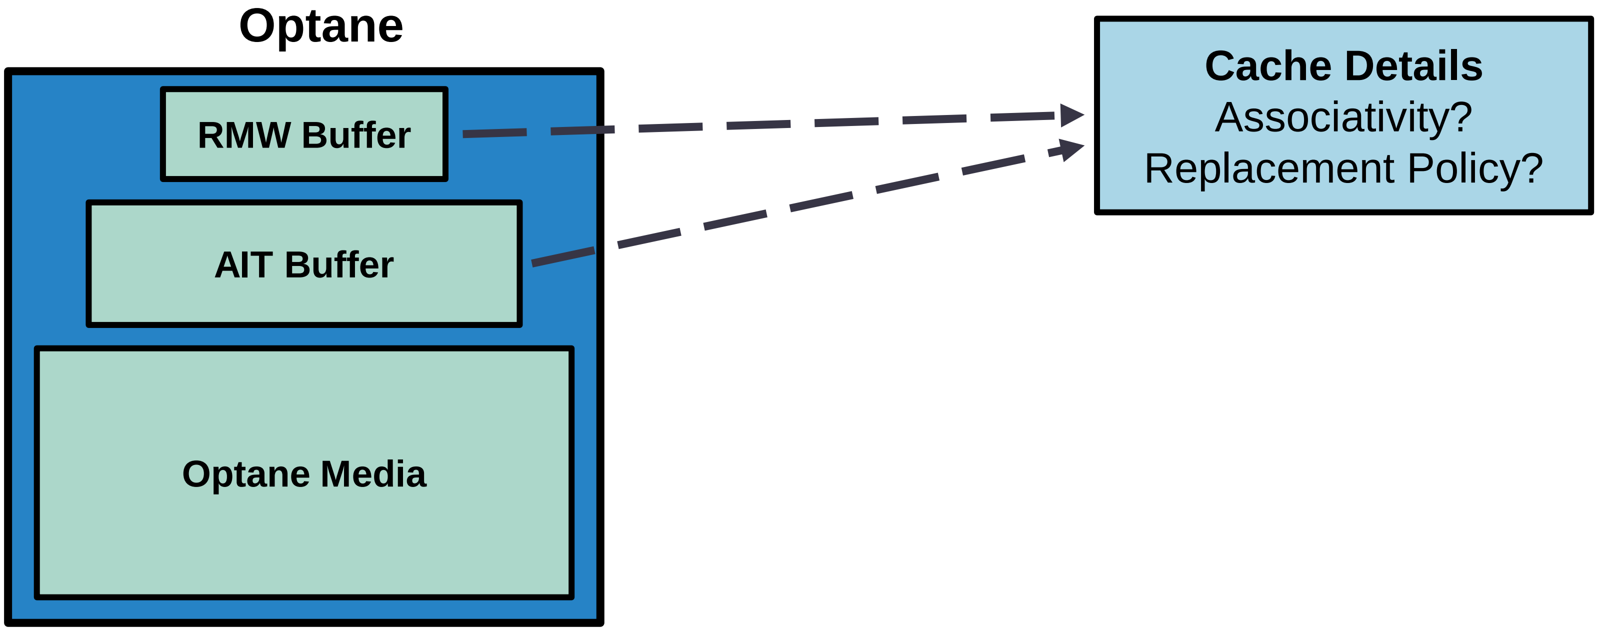

Optane: We have more!

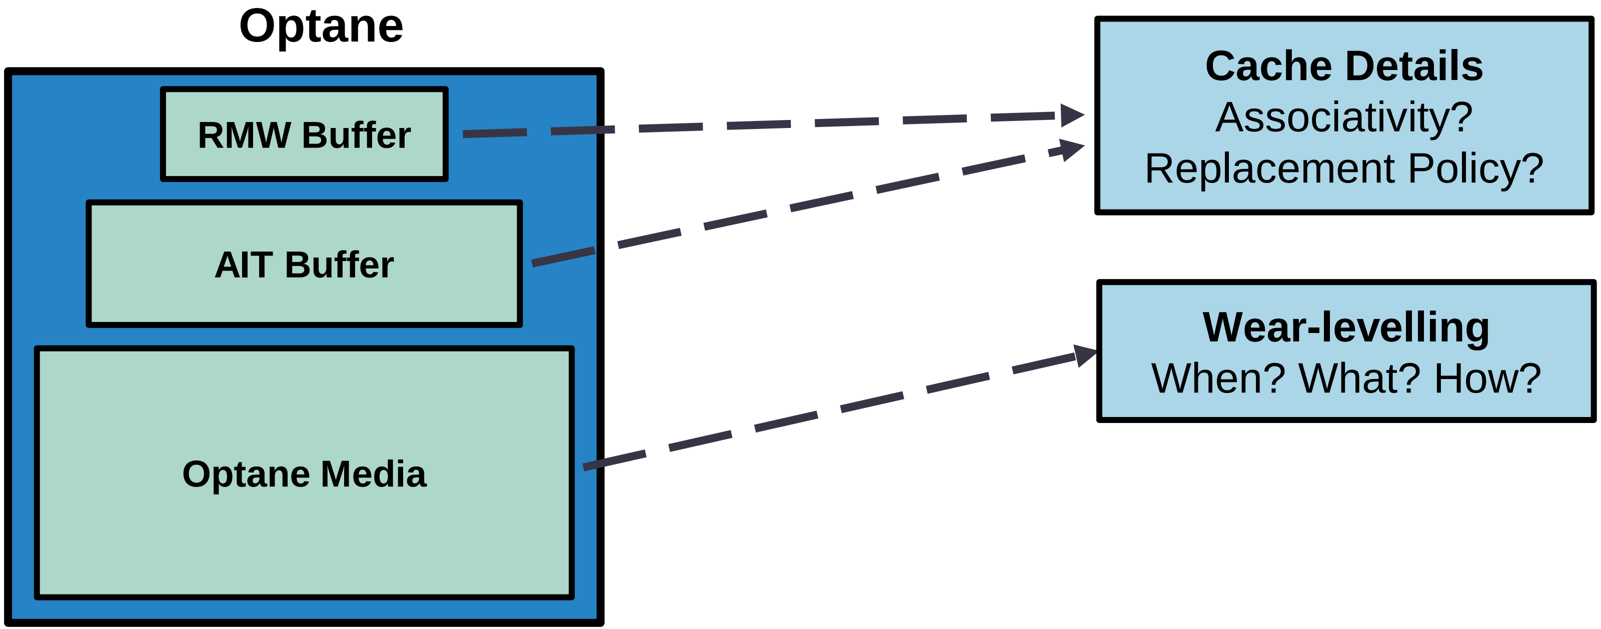

Cache Details
Associativity?
Replacement Policy?

Wear-levelling
When? What? How?

Optane: We have more!

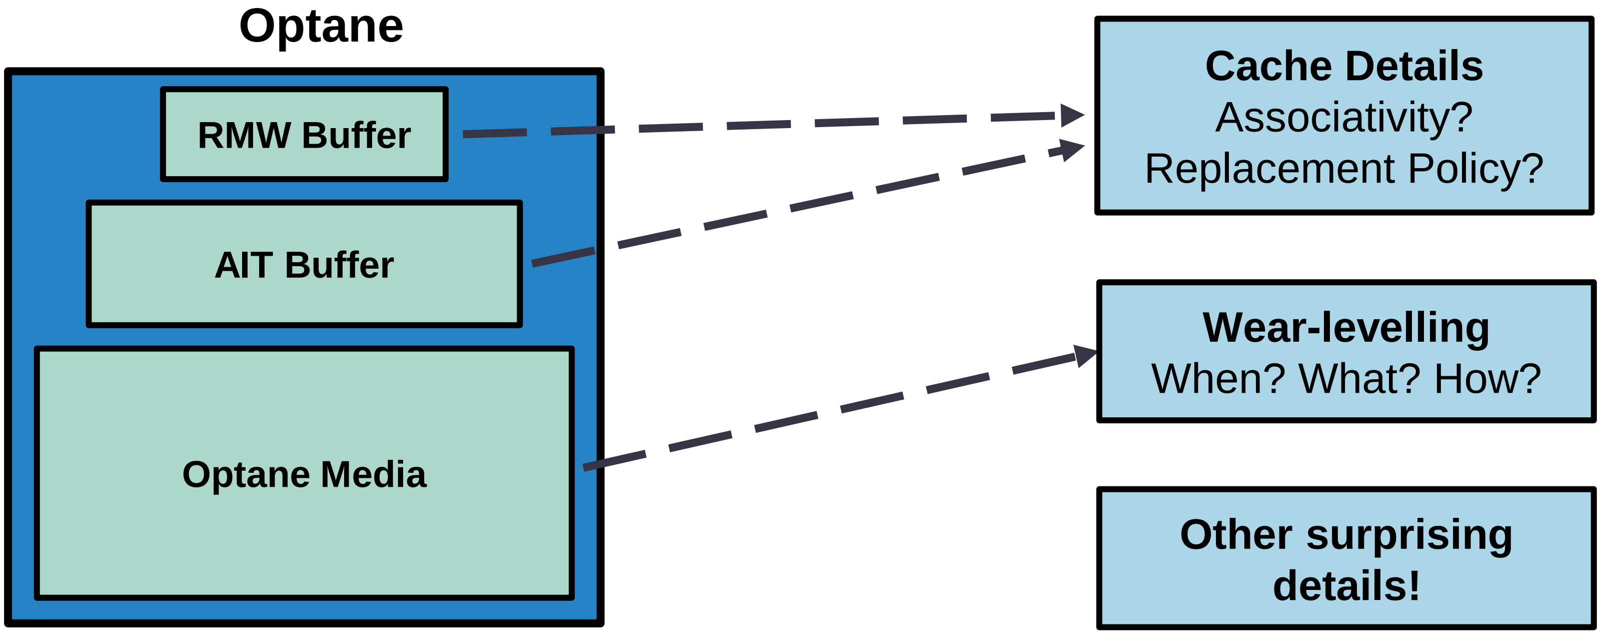

Cache Details
Associativity?
Replacement Policy?

Wear-levelling
When? What? How?

Other surprising details!

On-DIMM caches

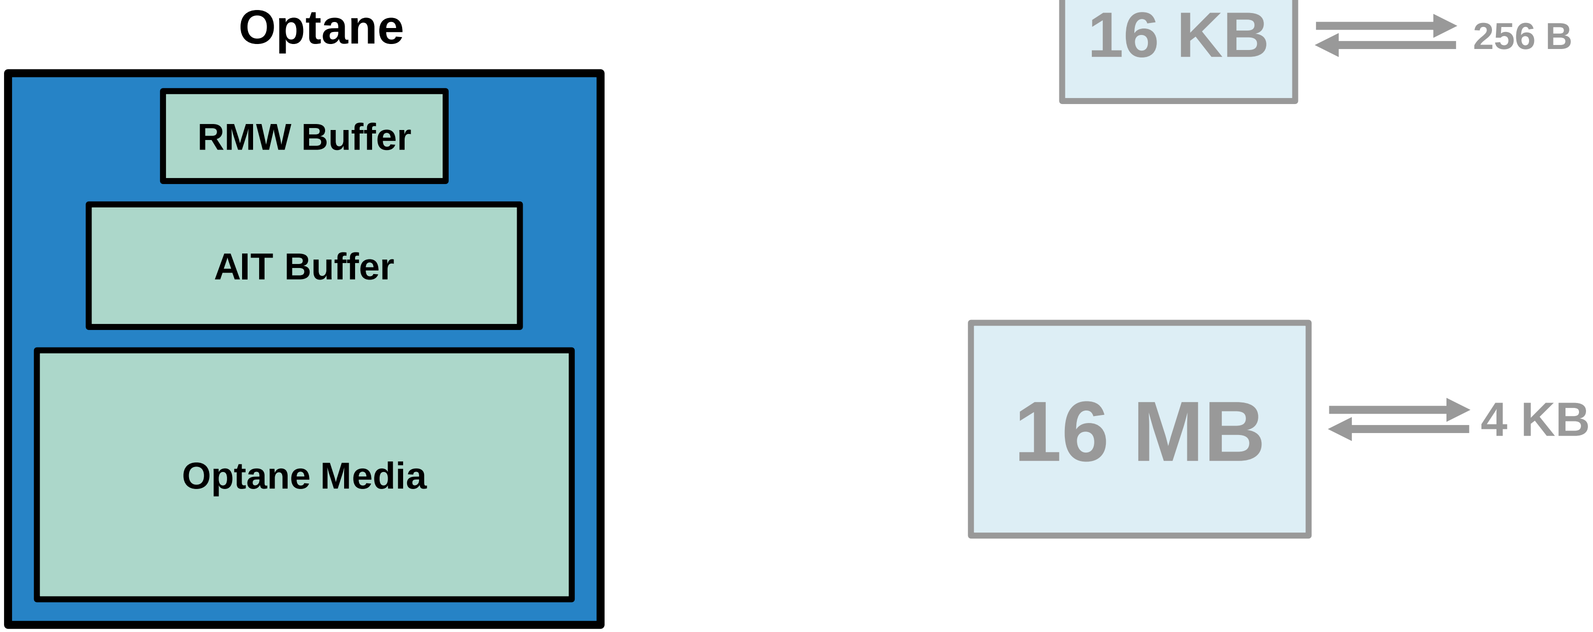

On-DIMM caches

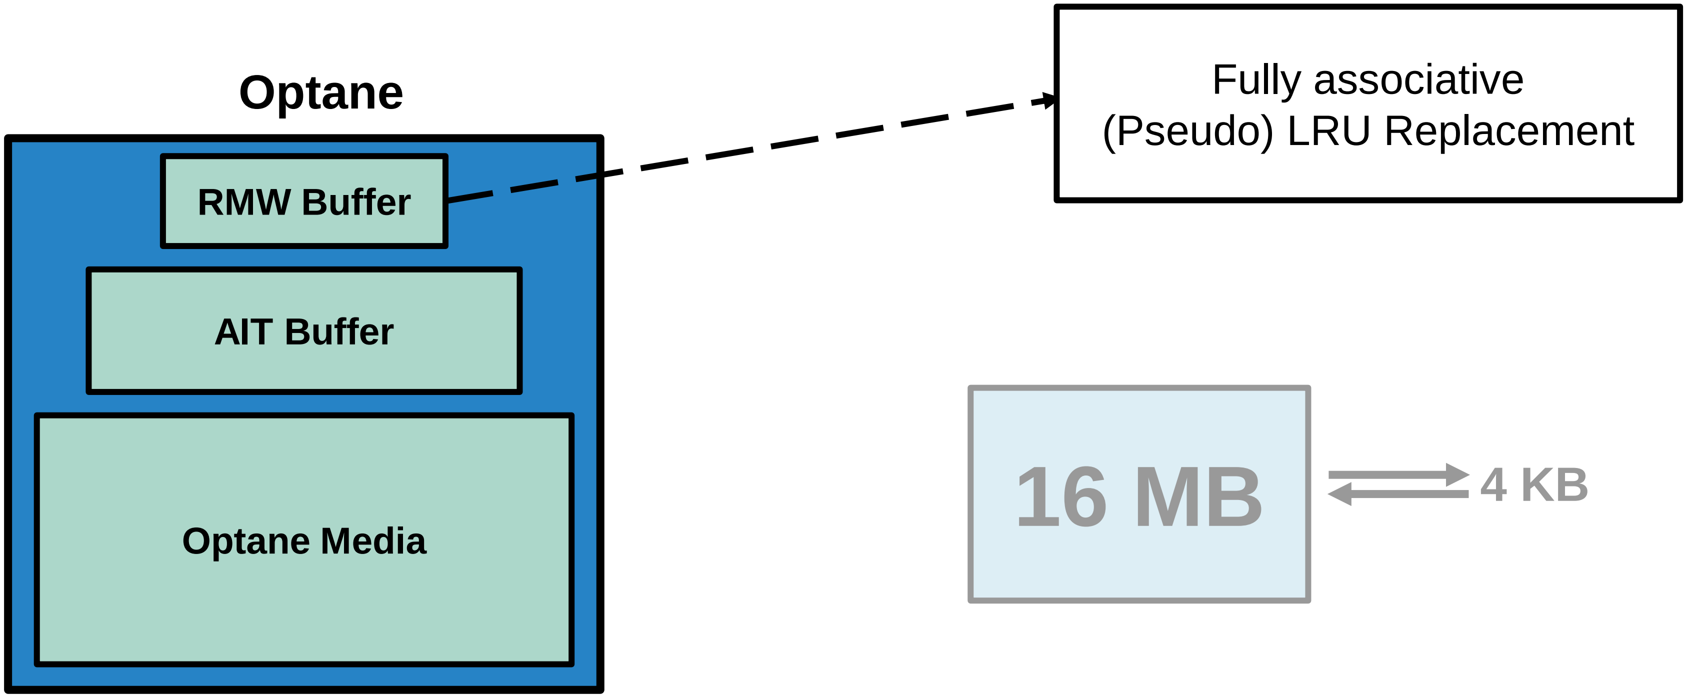

On-DIMM caches

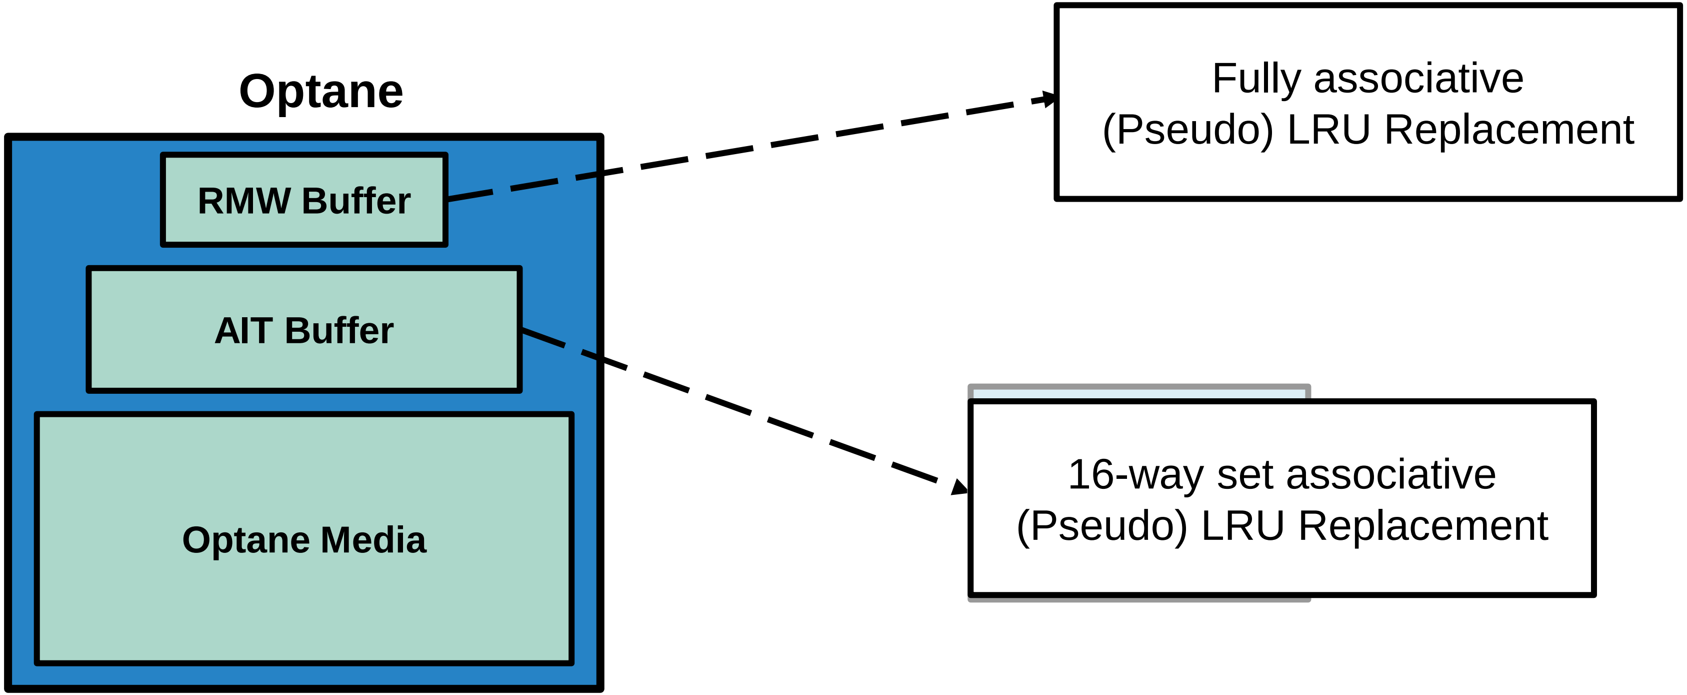

Wear-levelling in Memories

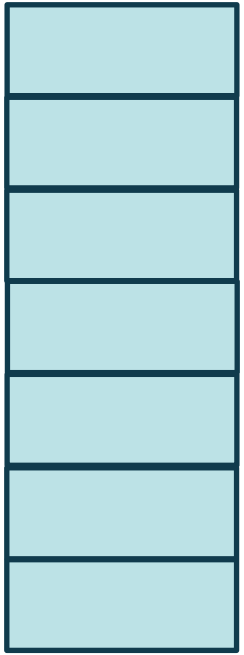

Memory Cells

Wear-levelling in Memories

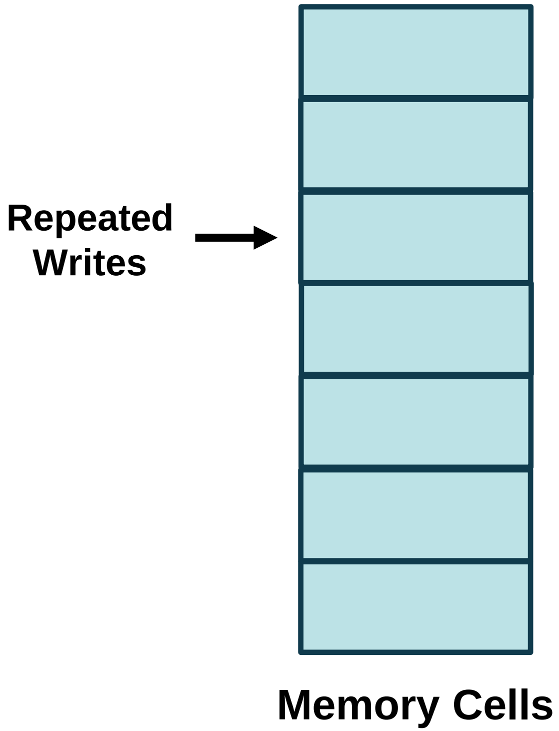

Wear-levelling in Memories

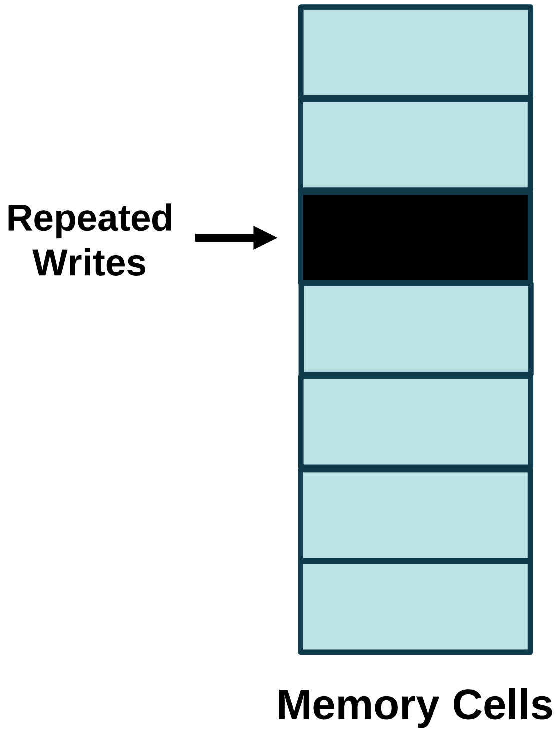

Wear-levelling in Memories

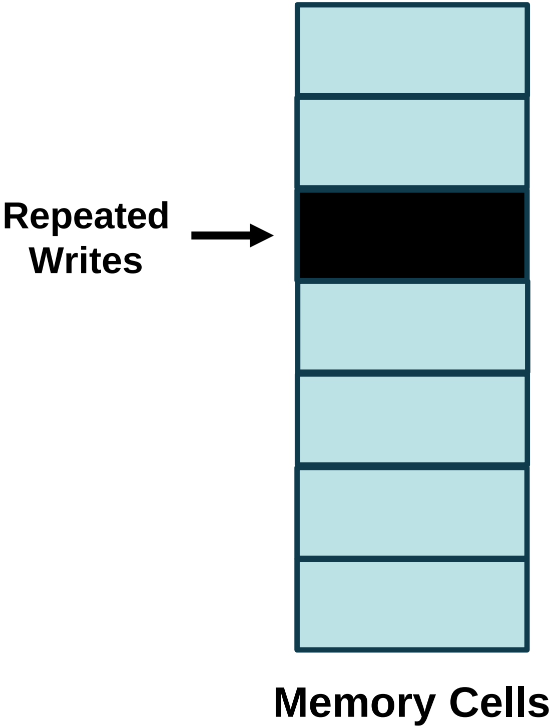

Wear-levelling in Memories

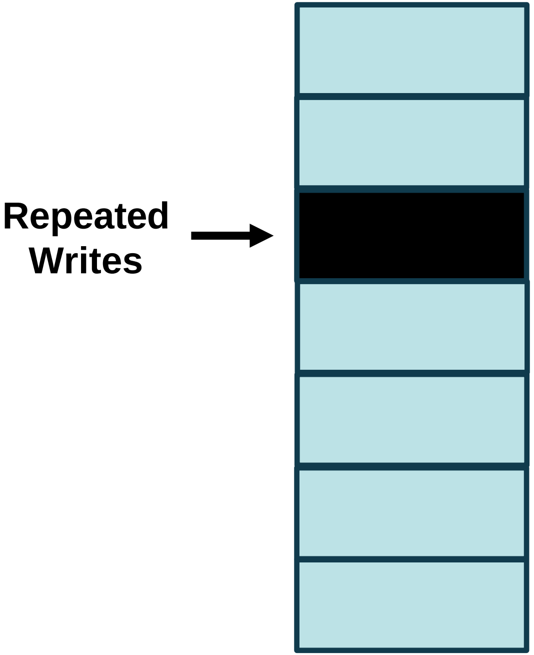

Repeated Writes

Memory Cells

Wear-levelling to the rescue!

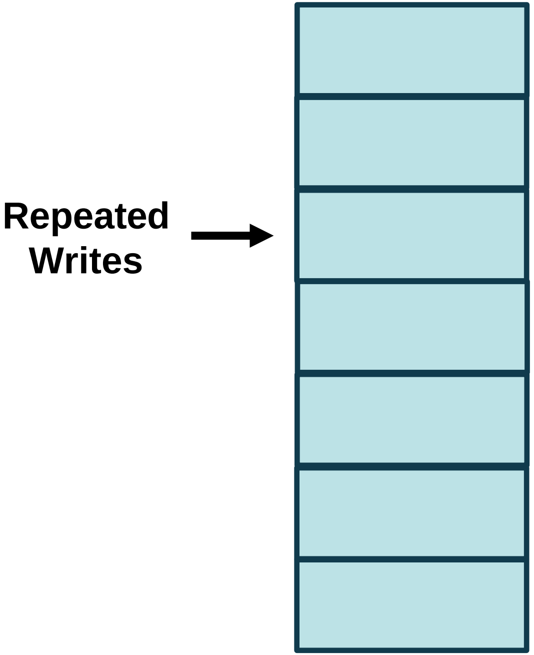

Repeated Writes

Memory Cells

Wear-levelling in Memories

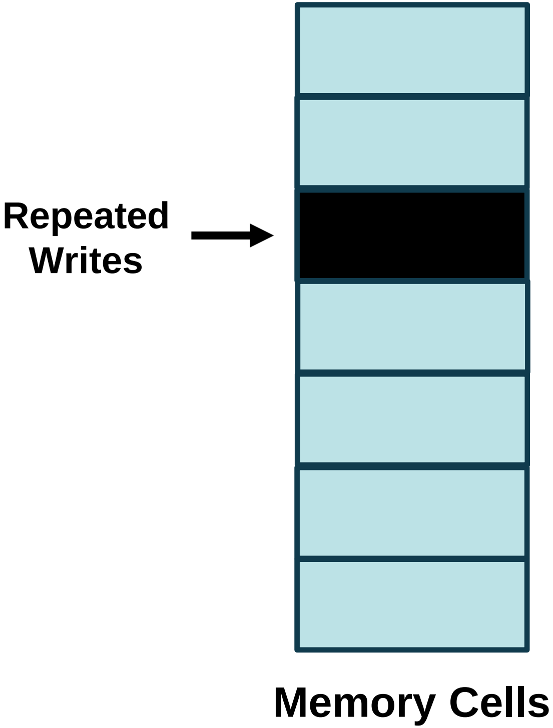

Wear-levelling to the rescue!

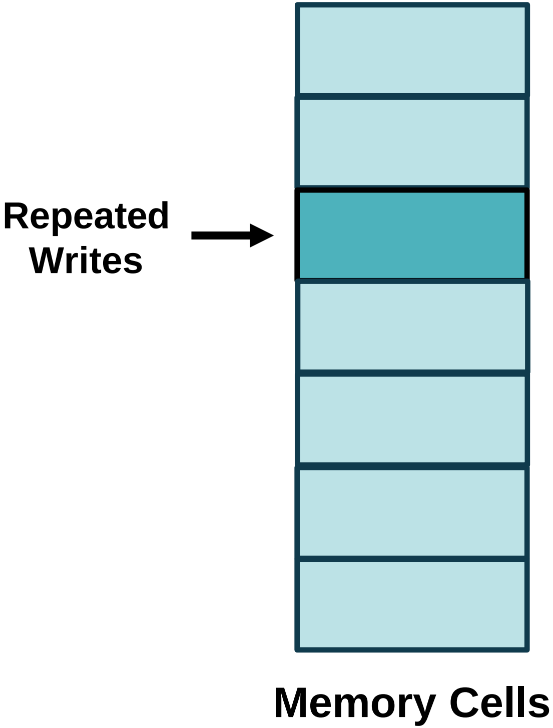

Wear-levelling in Memories

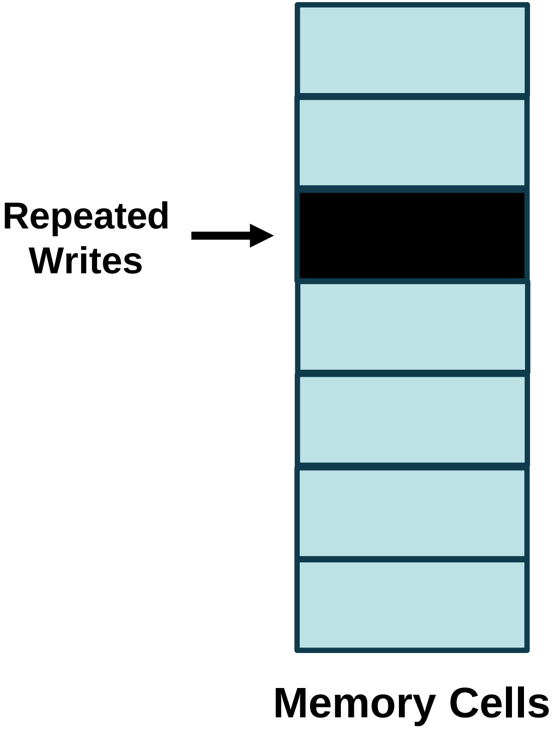

Wear-levelling to the rescue!

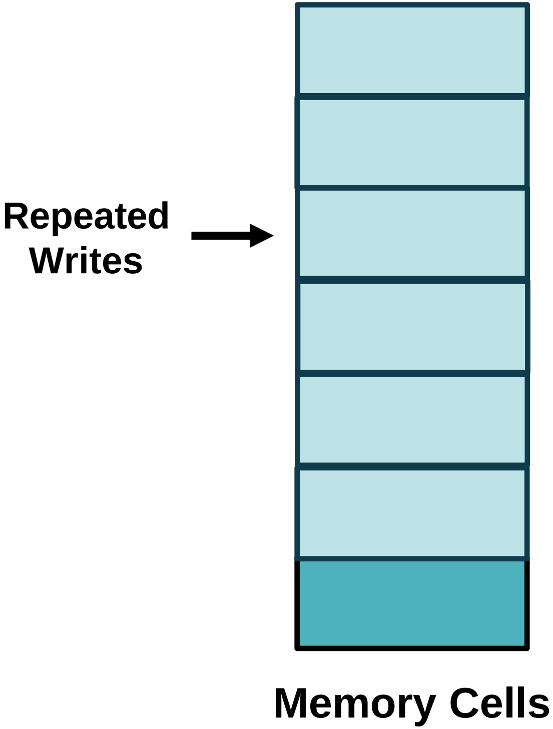

Wear-levelling in Memories

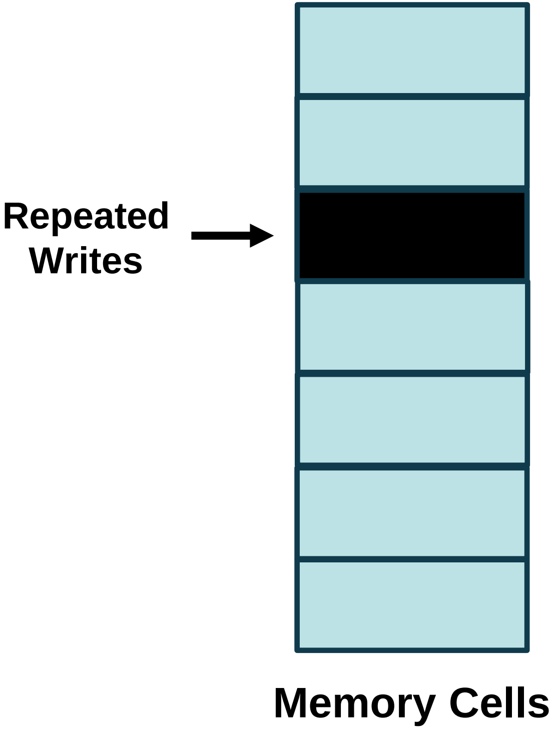

Wear-levelling to the rescue!

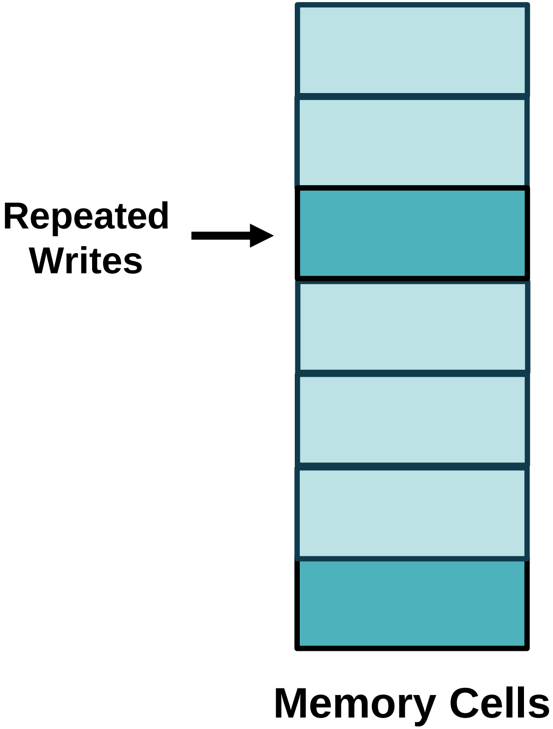

Wear-levelling in Memories

Repeated Writes

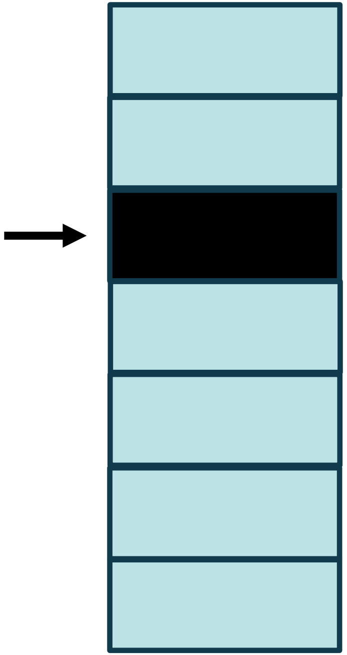

Memory Cells

Wear-levelling to the rescue!

Repeated Writes

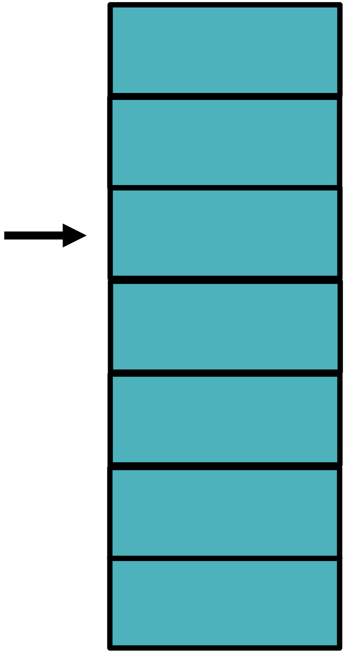

Memory Cells

Wear-levelling in Optane: When/What?

Wear-levelling in Optane: When/What?

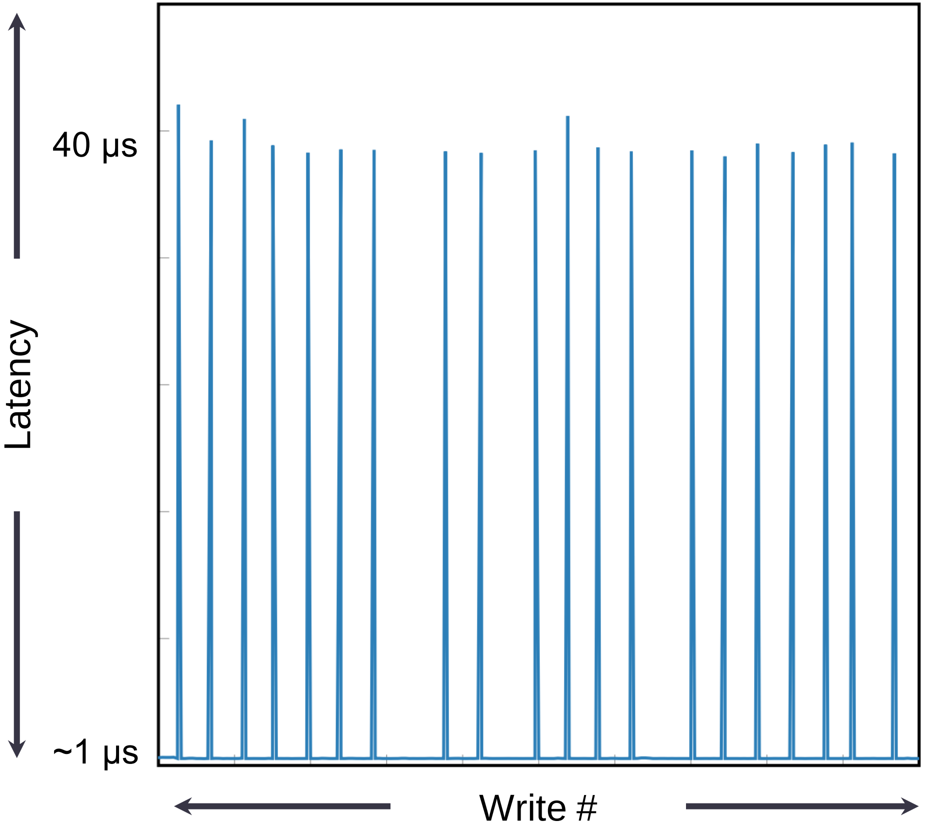

Wear-levelling in Optane: When/What?

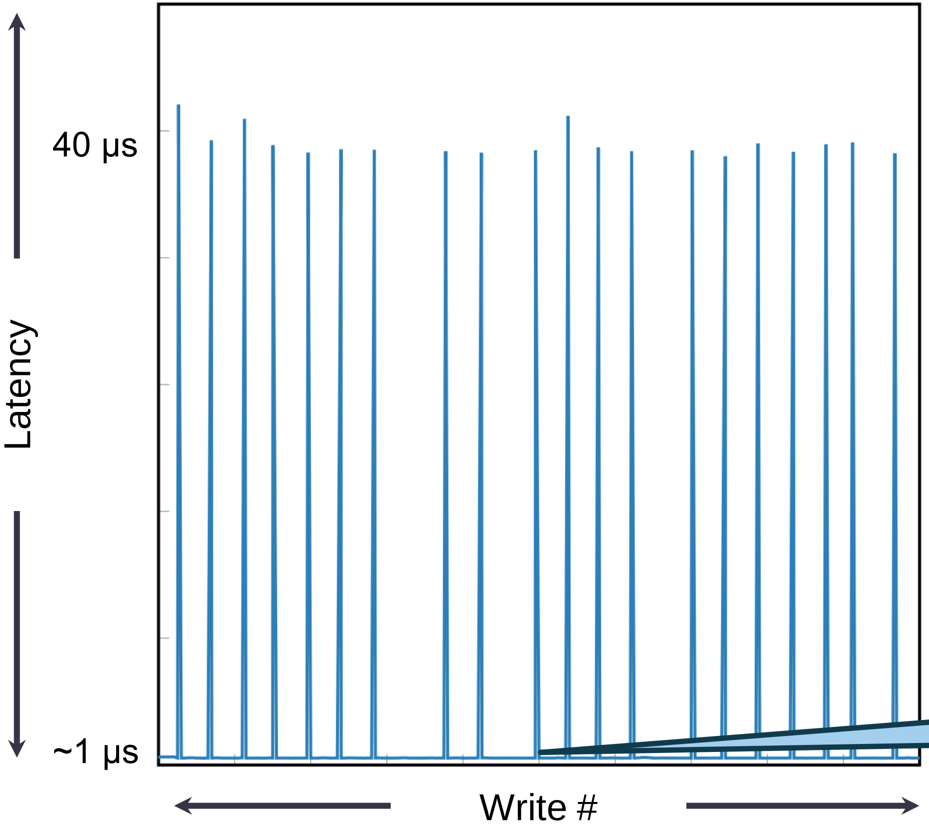

Typical Write Latency

Wear-levelling in Optane: When/What?

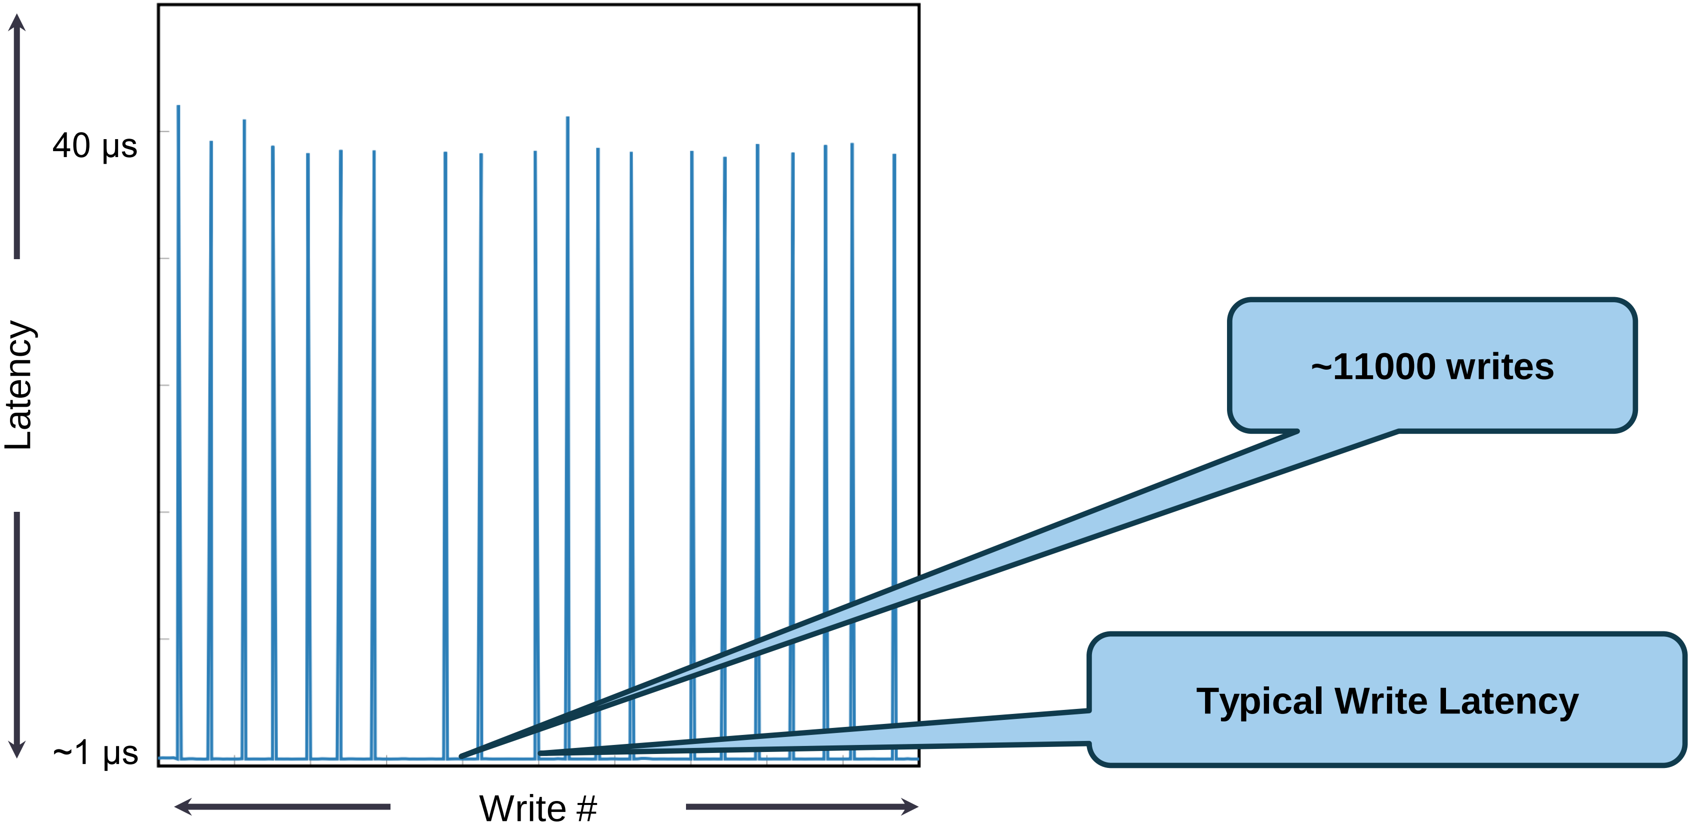

~11000 writes

Typical Write Latency

Wear-levelling in Optane: When/What?

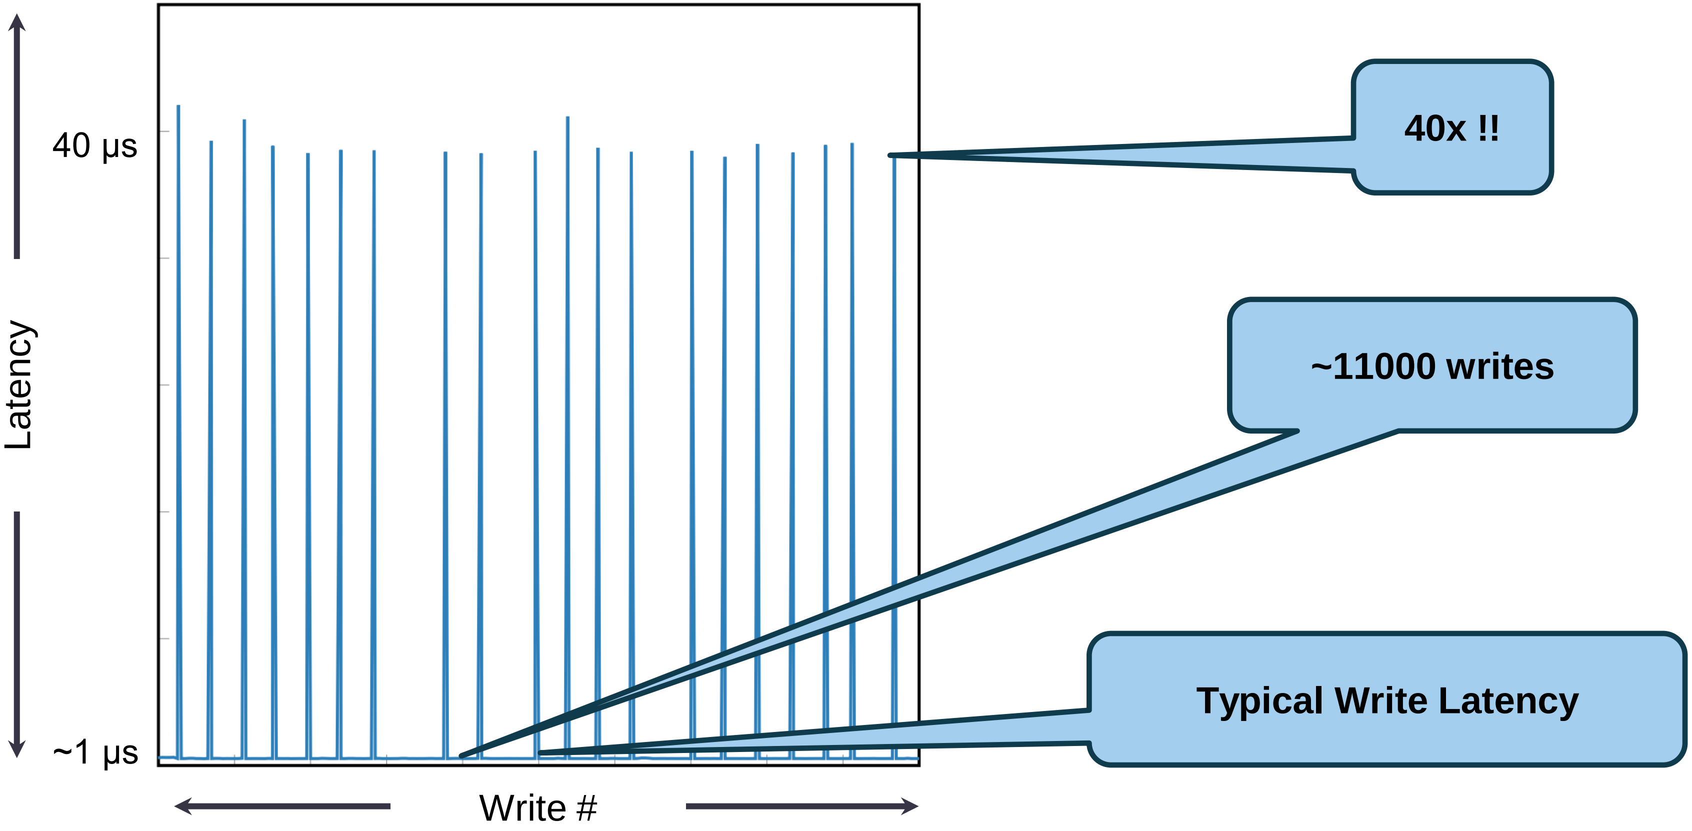

40x !!

~11000 writes

Typical Write Latency

Wear-levelling in Optane: When/What?

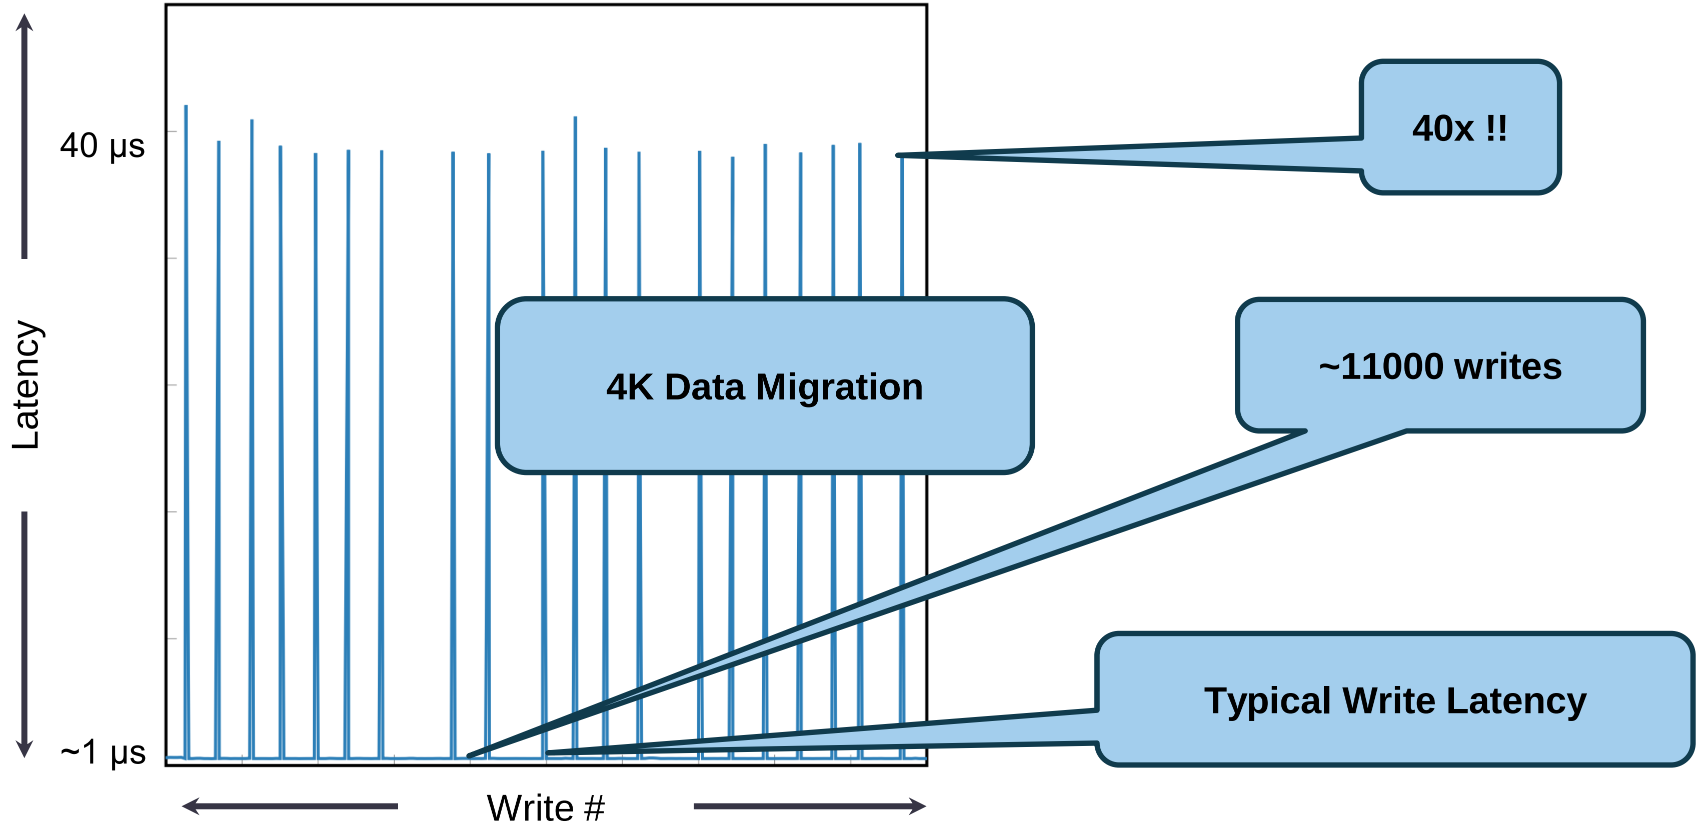

Wear-levelling in Optane: How?

Expectations	Reality

Wear-levelling in Optane: How?

Expectations

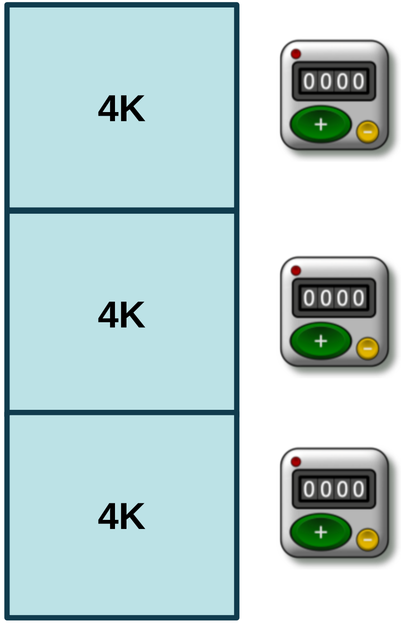

Reality

Wear-levelling in Optane: How?

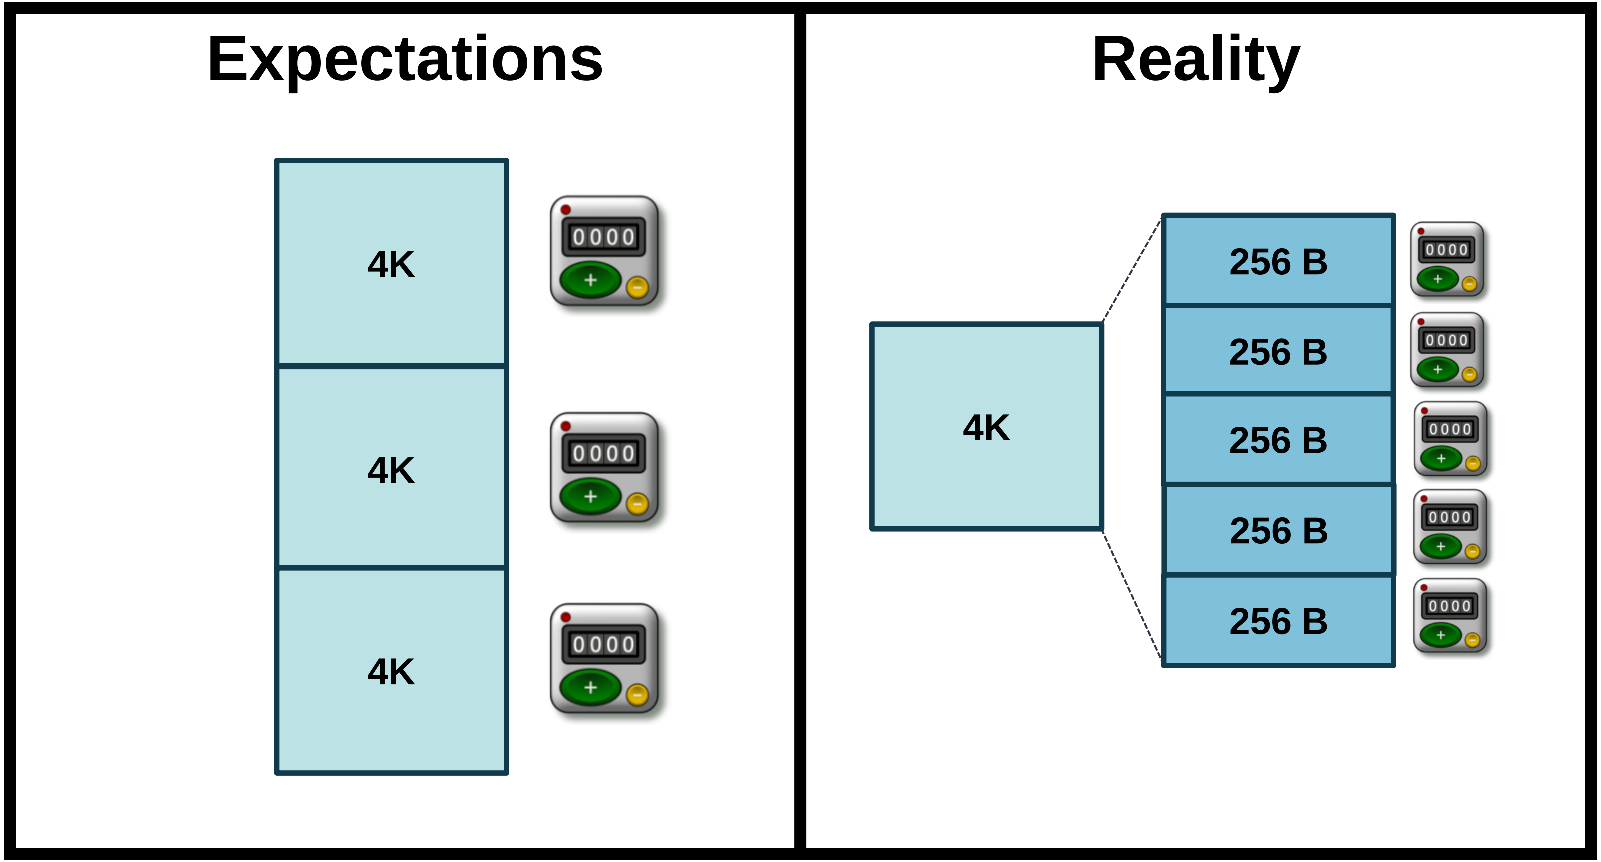

An Optane Curveball: *clflush*

Expectations	Reality

An Optane Curveball: *clflush*

Expectations	Reality
cl == cache line	

An Optane Curveball: *clflush*

Expectations

cl == cache line

== CPU cache line

Reality

An Optane Curveball: *clflush*

Expectations

cl == cache line

== CPU cache line

“*clflush* flushes only CPU caches”

Reality

An Optane Curveball: *clflush*

Expectations

cl == cache line

== CPU cache line

“*clflush* flushes only CPU caches”

Reality

clflush reaches Optane!

An Optane Curveball: *clflush*

Expectations

cl == cache line

== CPU cache line

“*clflush* flushes only CPU caches”

Reality

clflush reaches Optane!

Flushes Optane Buffer!

An Optane Curveball: *clflush*

Expectations

cl == cache line

== CPU cache line

“*clflush* flushes only CPU caches”

Reality

clflush reaches Optane!

Flushes Optane Buffer!

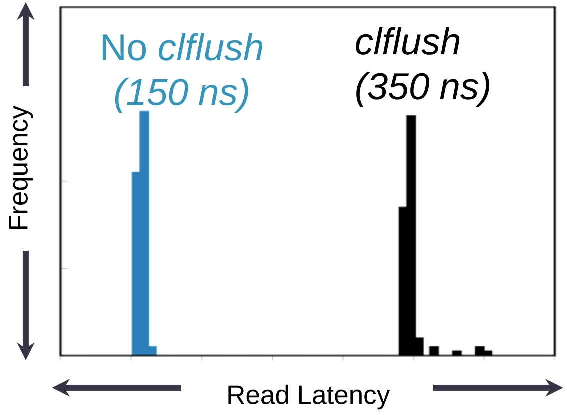

An Optane Curveball: R/W Contention

An Optane Curveball: R/W Contention

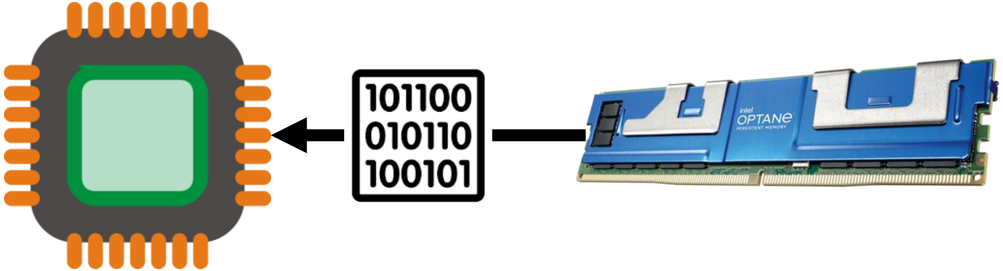

An Optane Curveball: R/W Contention

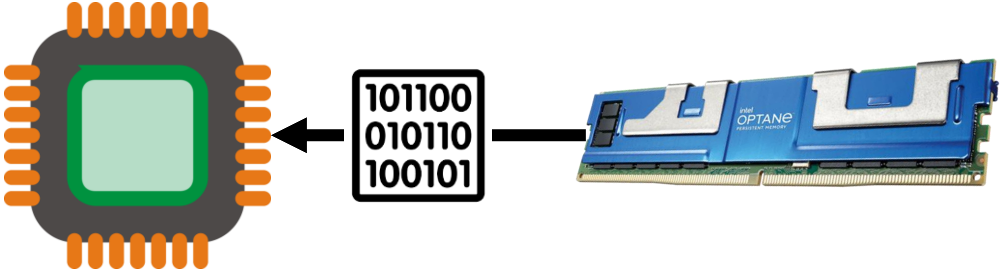

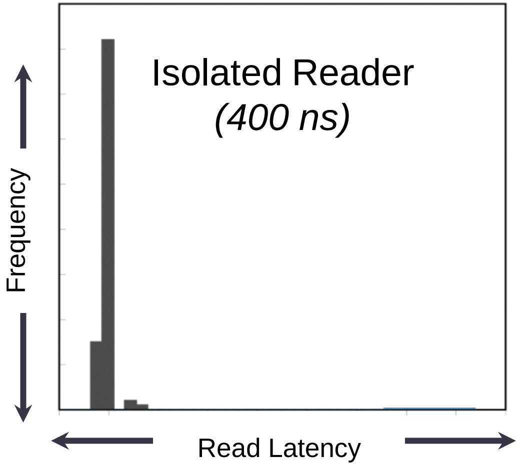

An Optane Curveball: R/W Contention

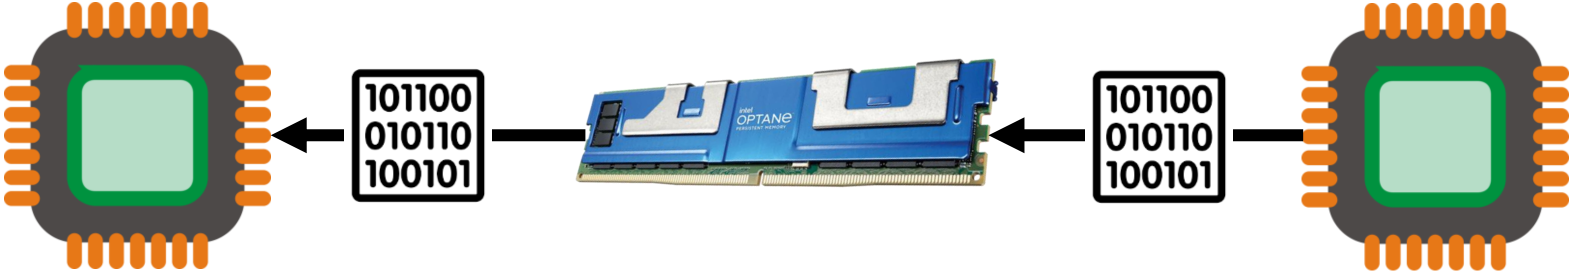

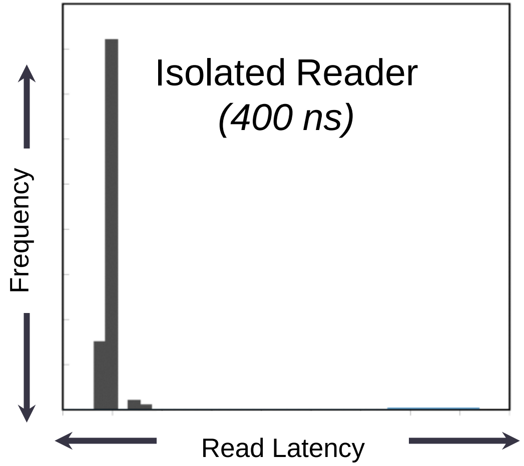

Isolated Reader
(400 ns)

An Optane Curveball: R/W Contention

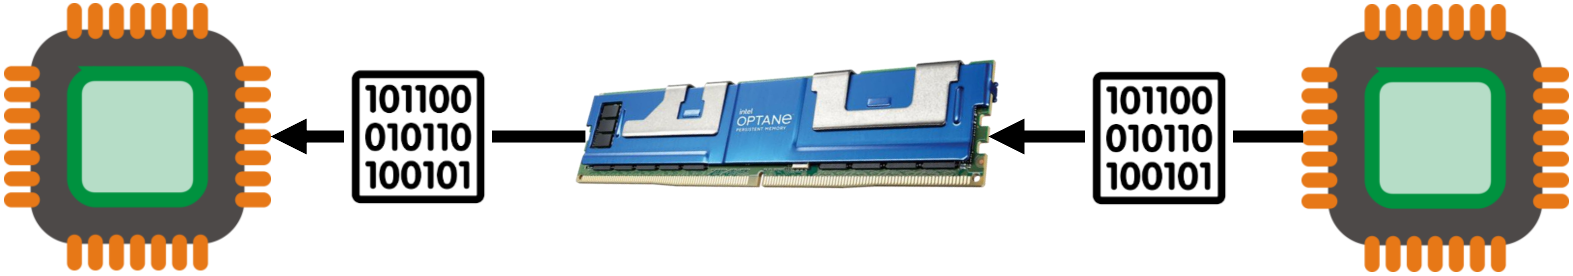

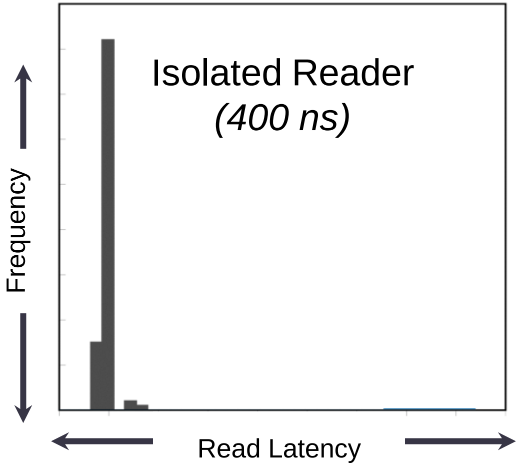

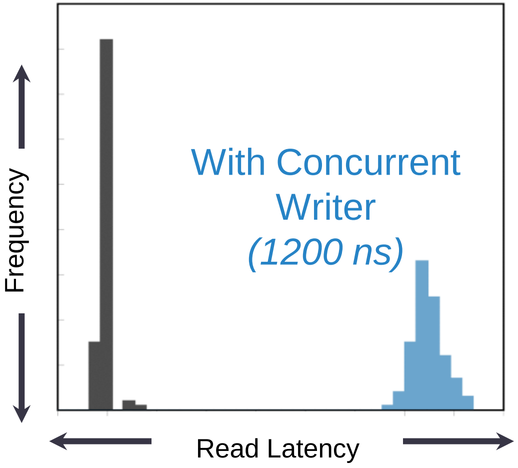

The Attacks

Exploring the security implications of our new attack primitives

A Bird’s Eye view

A Bird’s Eye view

Optane

A Bird’s Eye view

Optane

Attack Primitives

A Bird’s Eye view

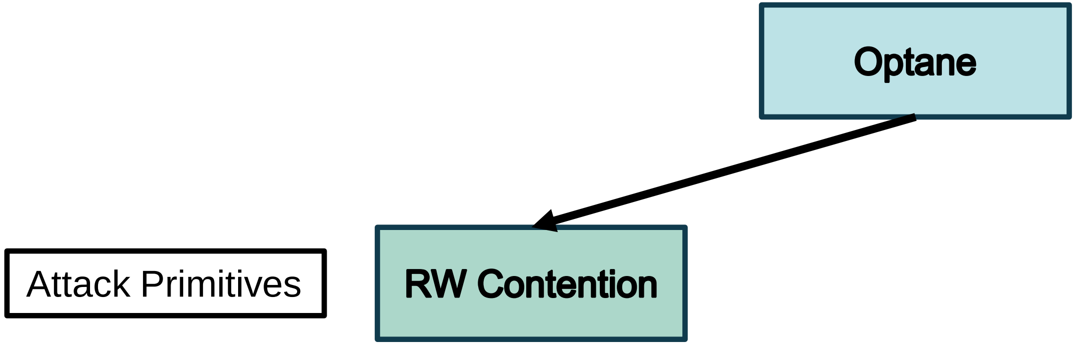

A Bird's Eye view

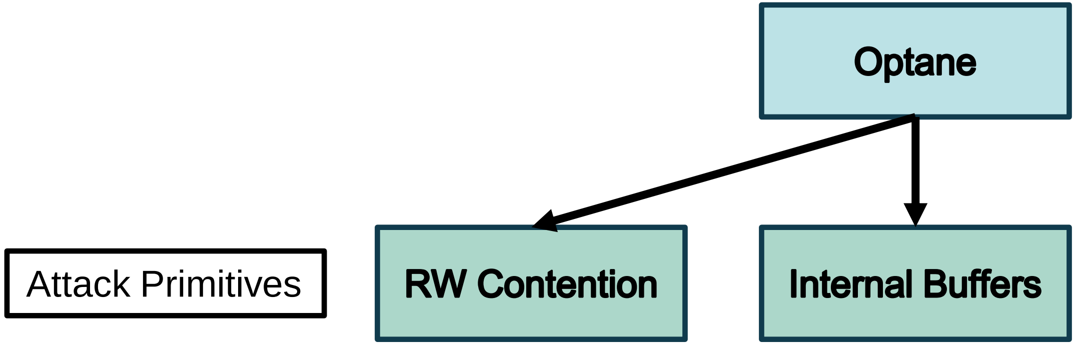

Optane

Attack Primitives

RW Contention

Internal Buffers

A Bird’s Eye view

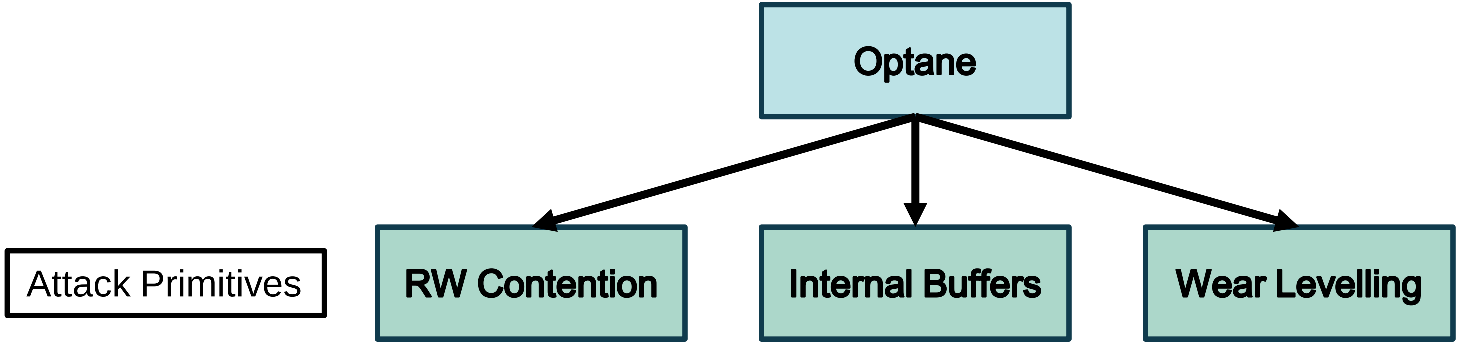

A Bird’s Eye view

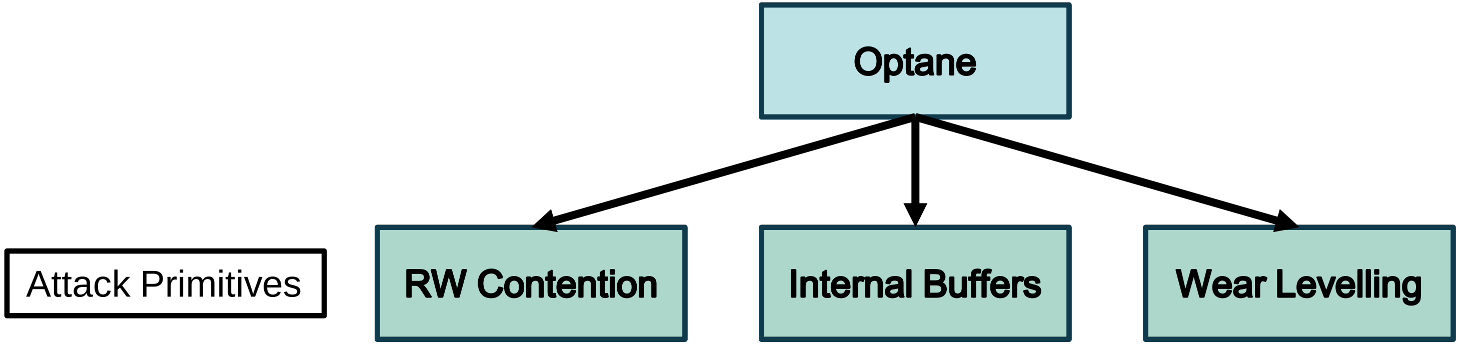

Our Attacks

A Bird’s Eye view

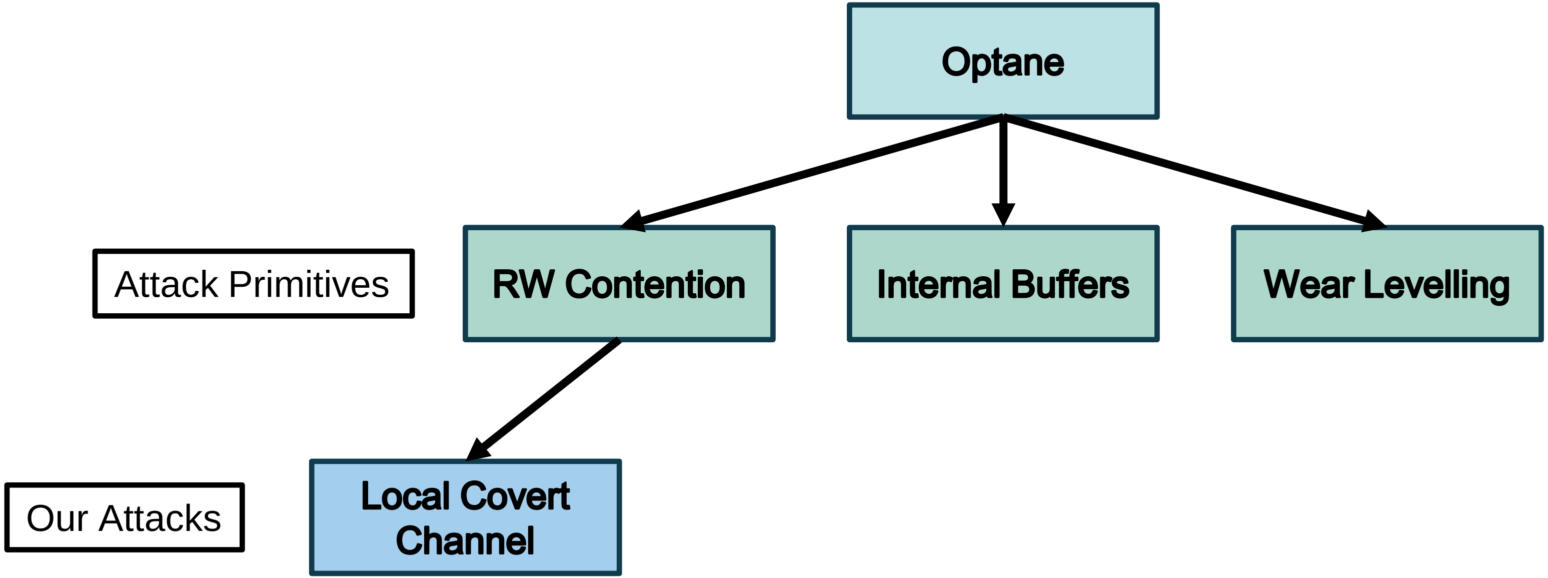

Optane

Attack Primitives

RW Contention

Internal Buffers

Wear Levelling

Our Attacks

Local Covert Channel

A Bird’s Eye view

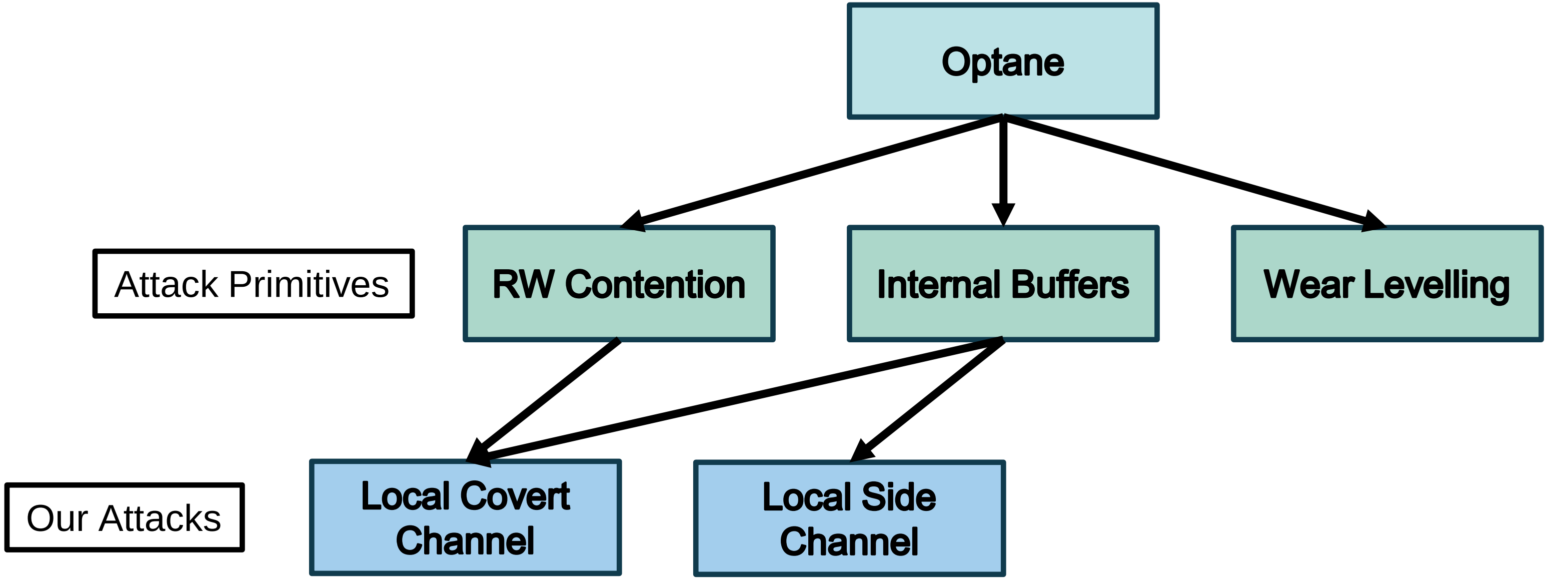

A Bird's Eye view

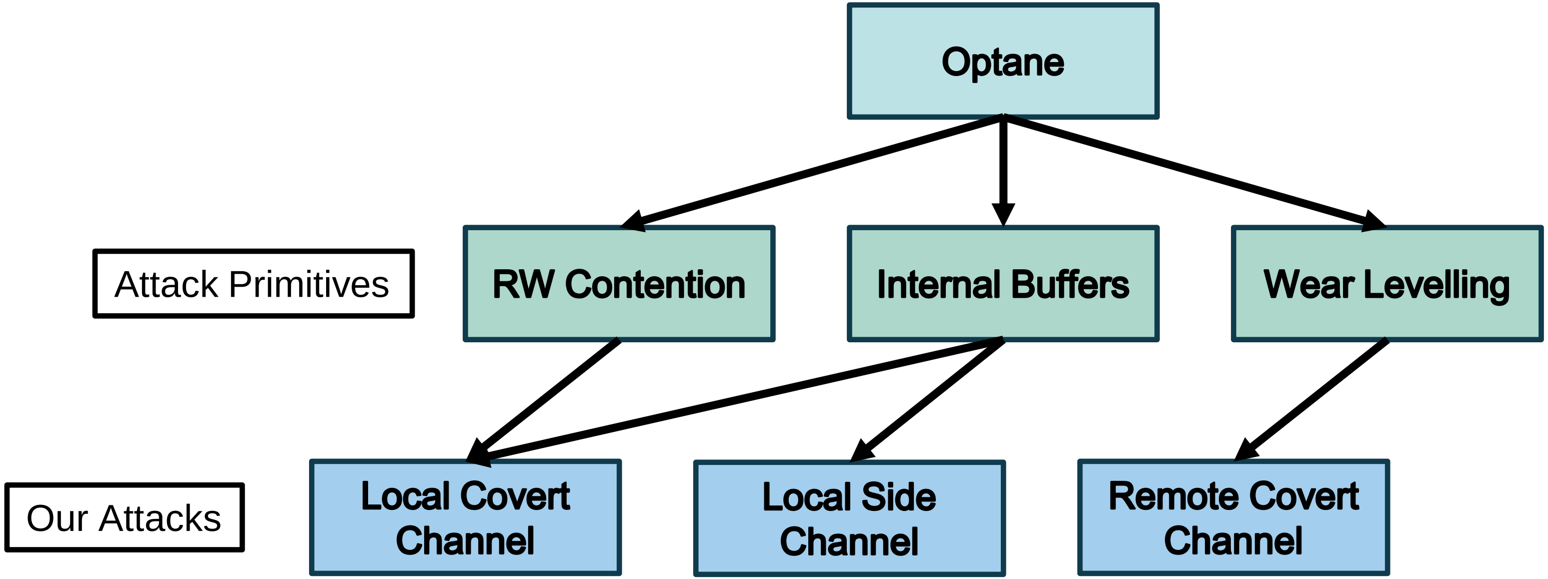

A Bird's Eye view

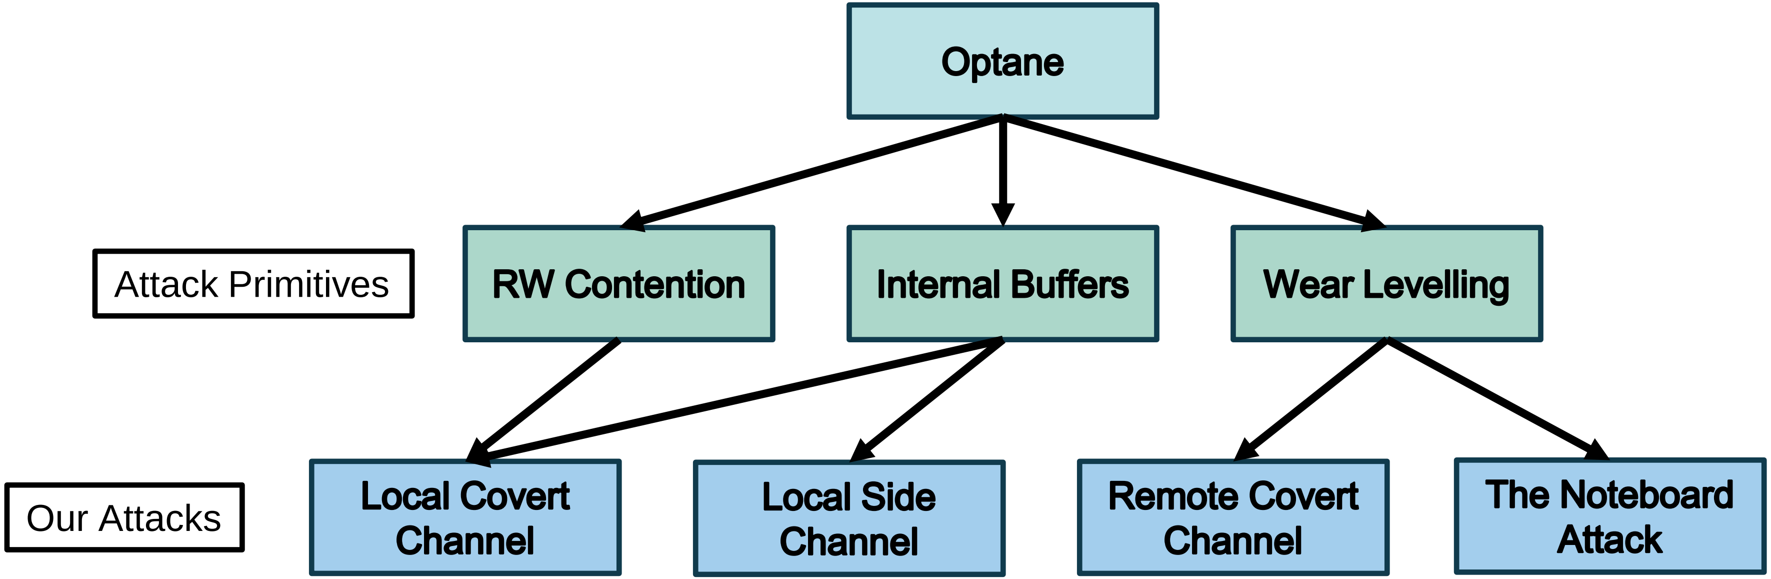

A Bird’s Eye view

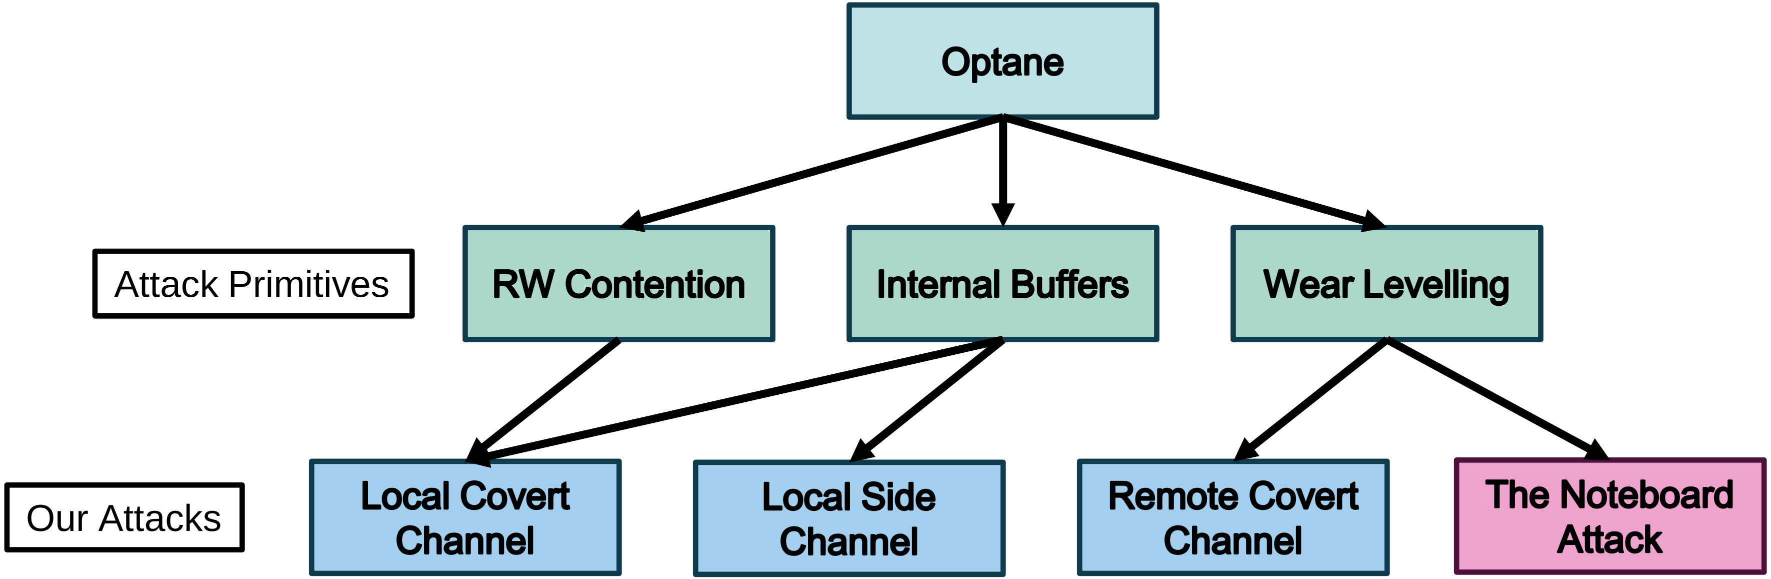

Attack: Noteboard Covert Channel

Encoding secret messages on Optane’s wear-levelling metadata

The Idea

The Idea

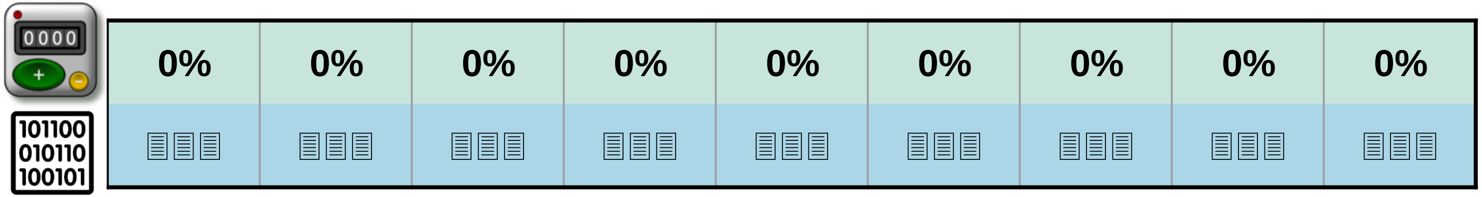

0000	0%	0%	0%	0%	0%	0%	0%	0%	0%
101100 010110 100101	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

The Idea

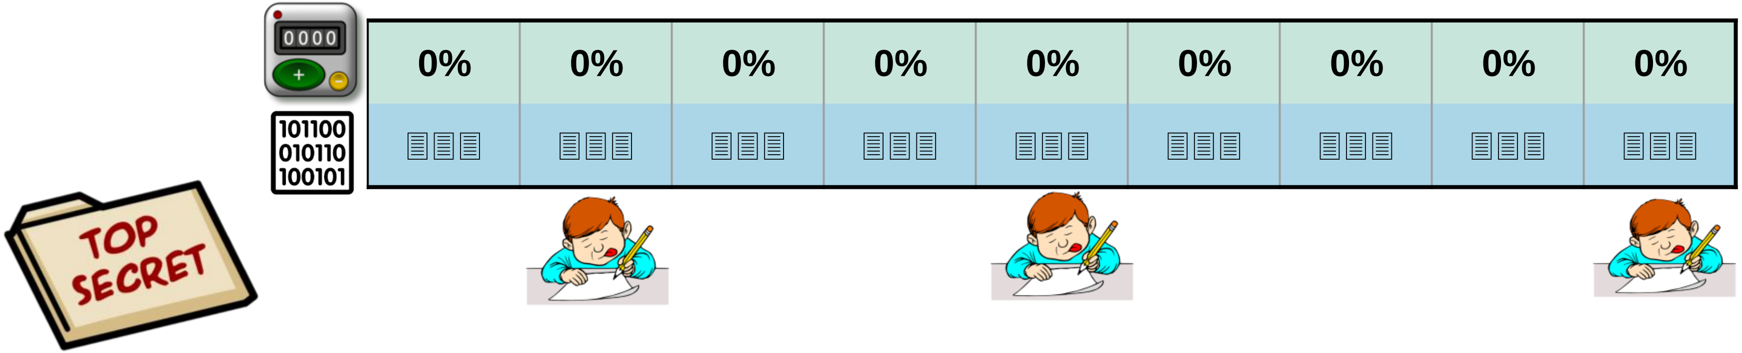

0%	0%	0%	0%	0%	0%	0%	0%	0%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

The Idea

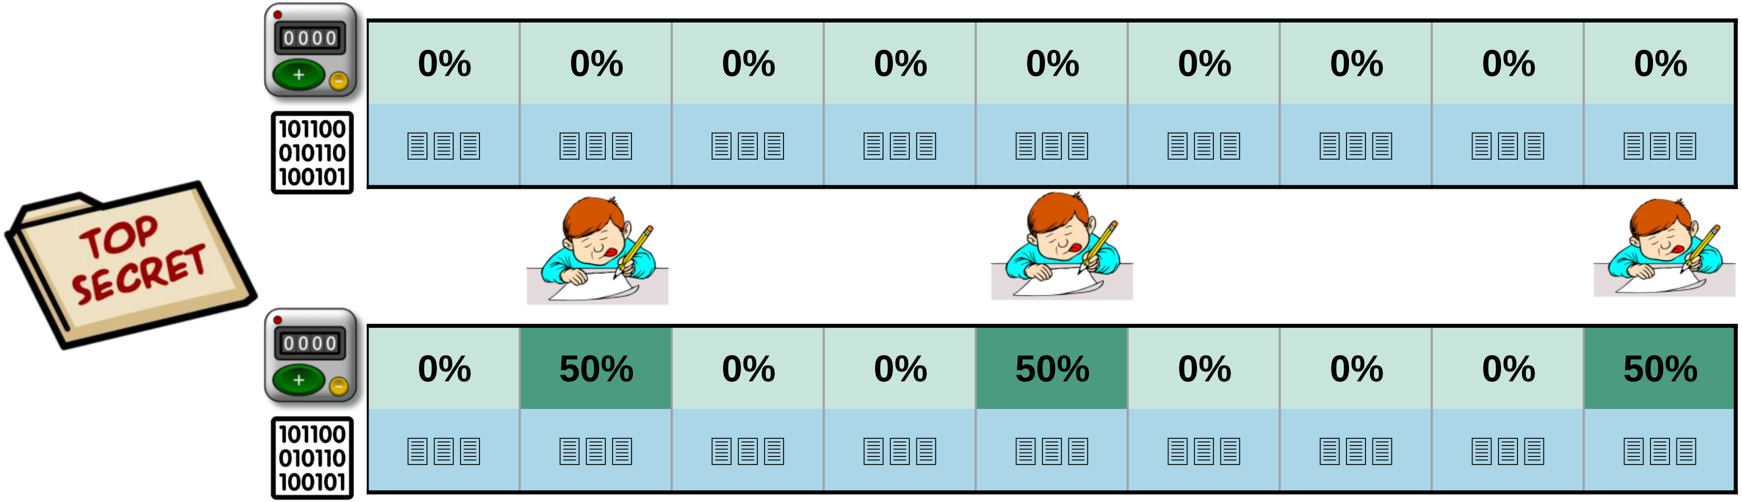

0%	0%	0%	0%	0%	0%	0%	0%	0%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

0%	50%	0%	0%	50%	0%	0%	0%	50%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

The Idea

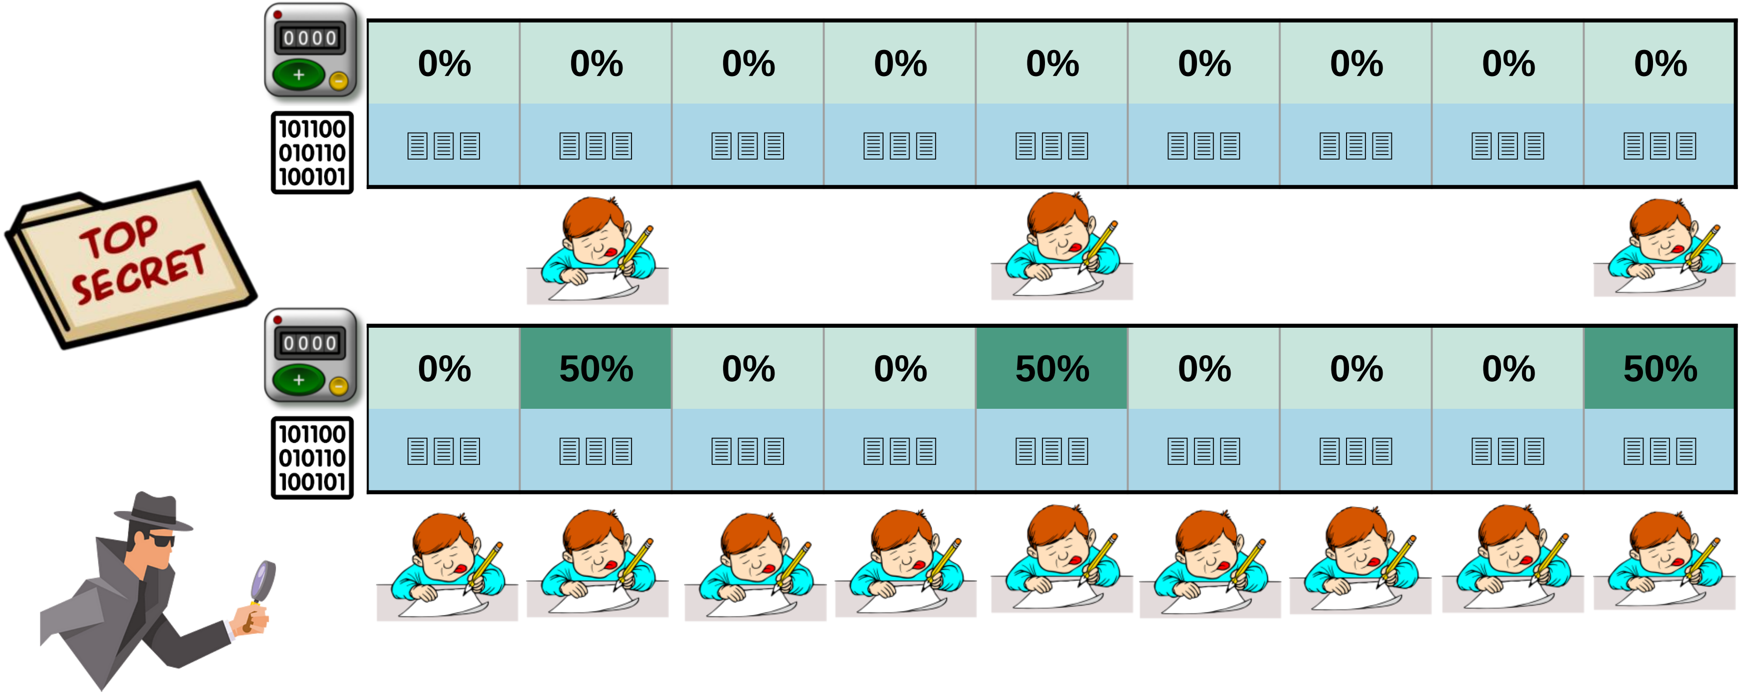

0%	0%	0%	0%	0%	0%	0%	0%	0%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

0%	50%	0%	0%	50%	0%	0%	0%	50%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

The Idea

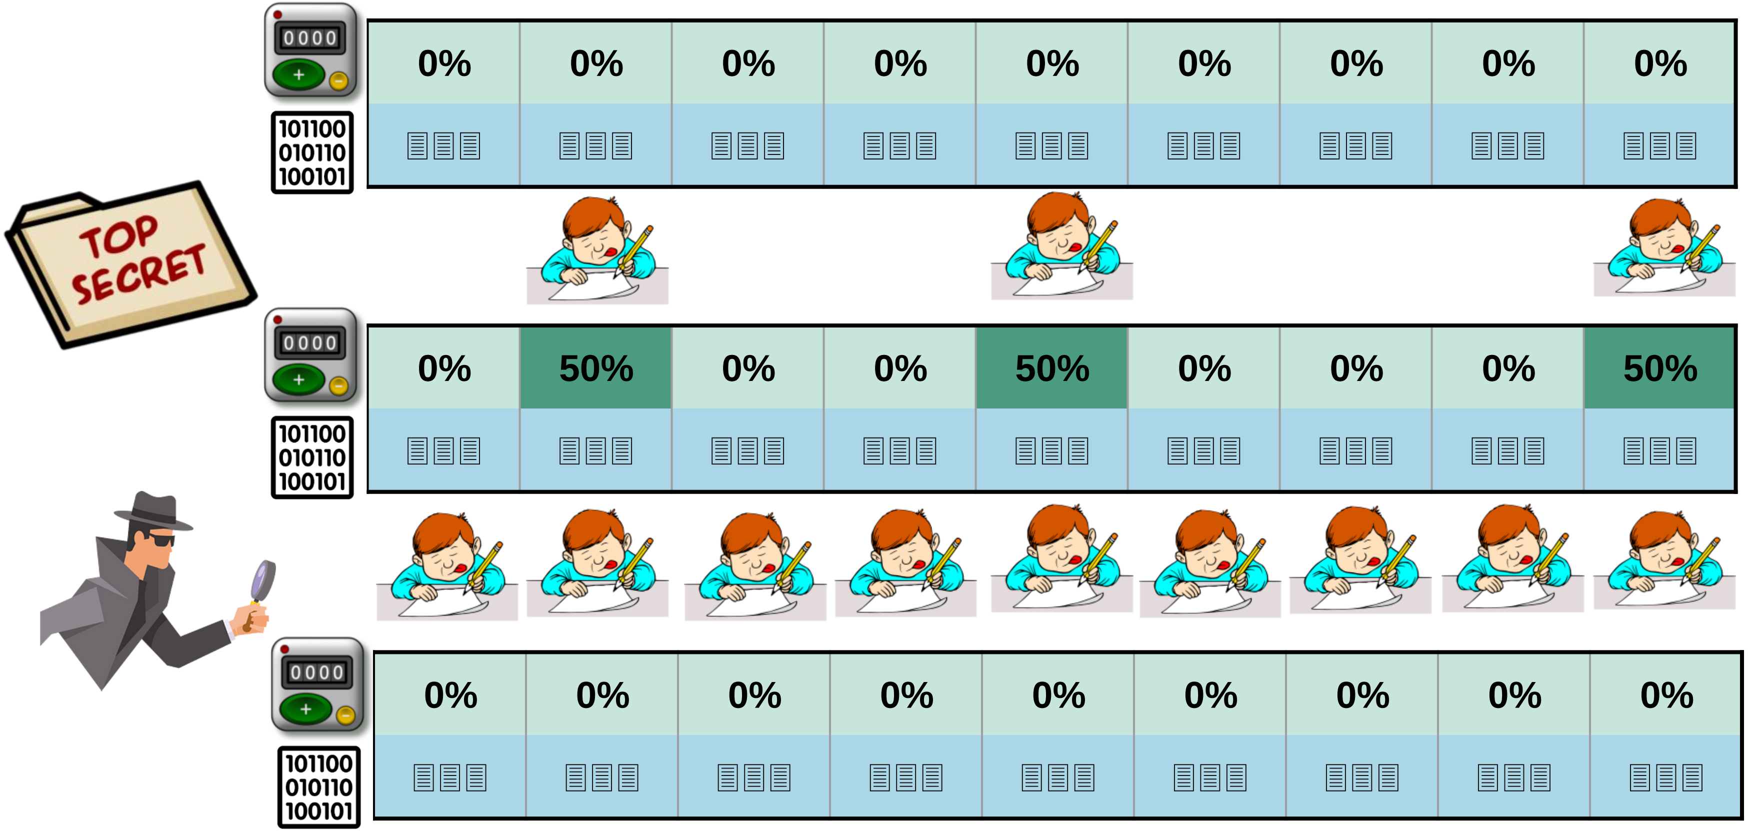

0%	0%	0%	0%	0%	0%	0%	0%	0%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

0%	50%	0%	0%	50%	0%	0%	0%	50%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

0%	0%	0%	0%	0%	0%	0%	0%	0%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

The Idea

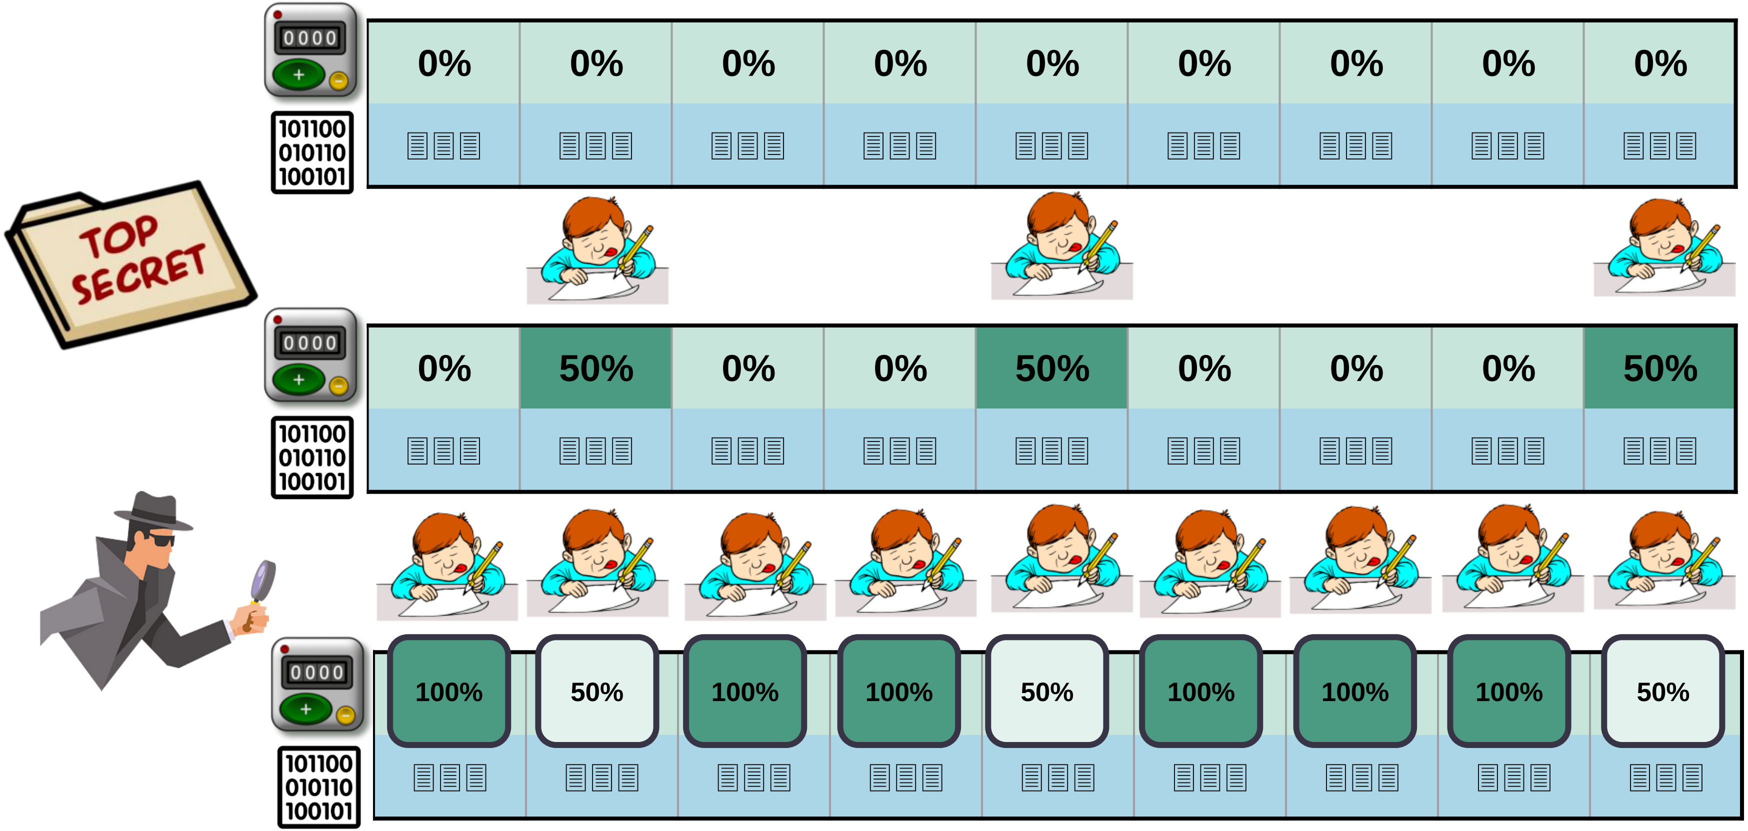

0%	0%	0%	0%	0%	0%	0%	0%	0%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

0%	50%	0%	0%	50%	0%	0%	0%	50%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

100%	50%	100%	100%	50%	100%	100%	100%	50%
▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤	▤▤▤

The Idea

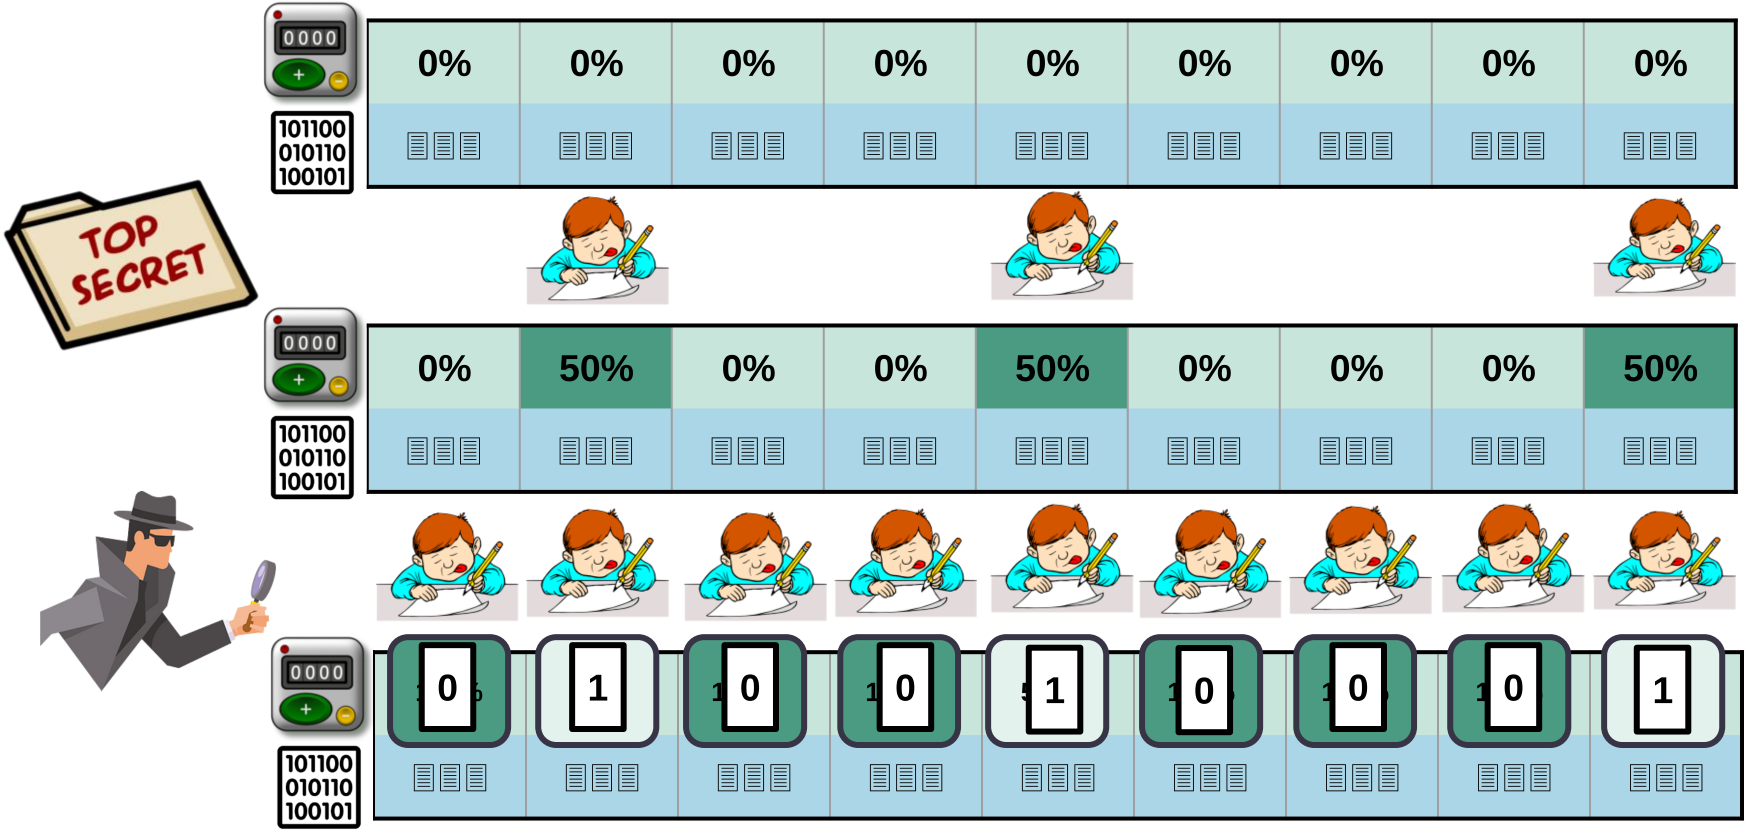

0%	0%	0%	0%	0%	0%	0%	0%	0%

0%	50%	0%	0%	50%	0%	0%	0%	50%

0	1	0	0	1	0	0	0	1

A Realization

A Realization

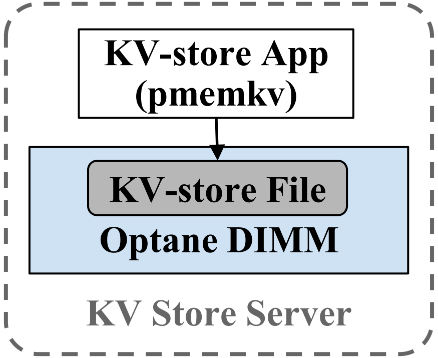

A Realization

Remote Sender

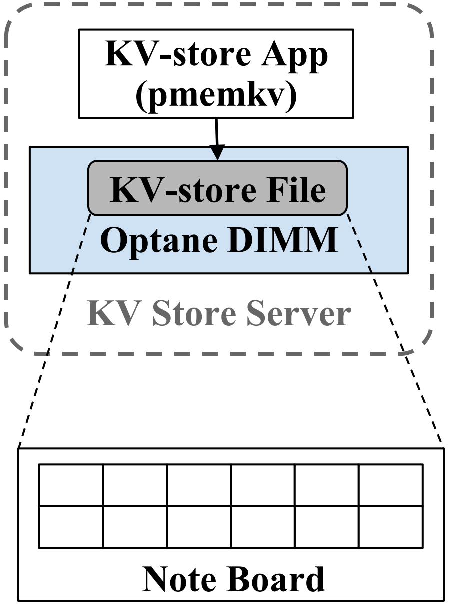

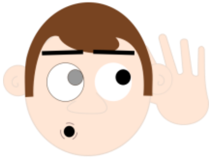

Remote Receiver

A Realization

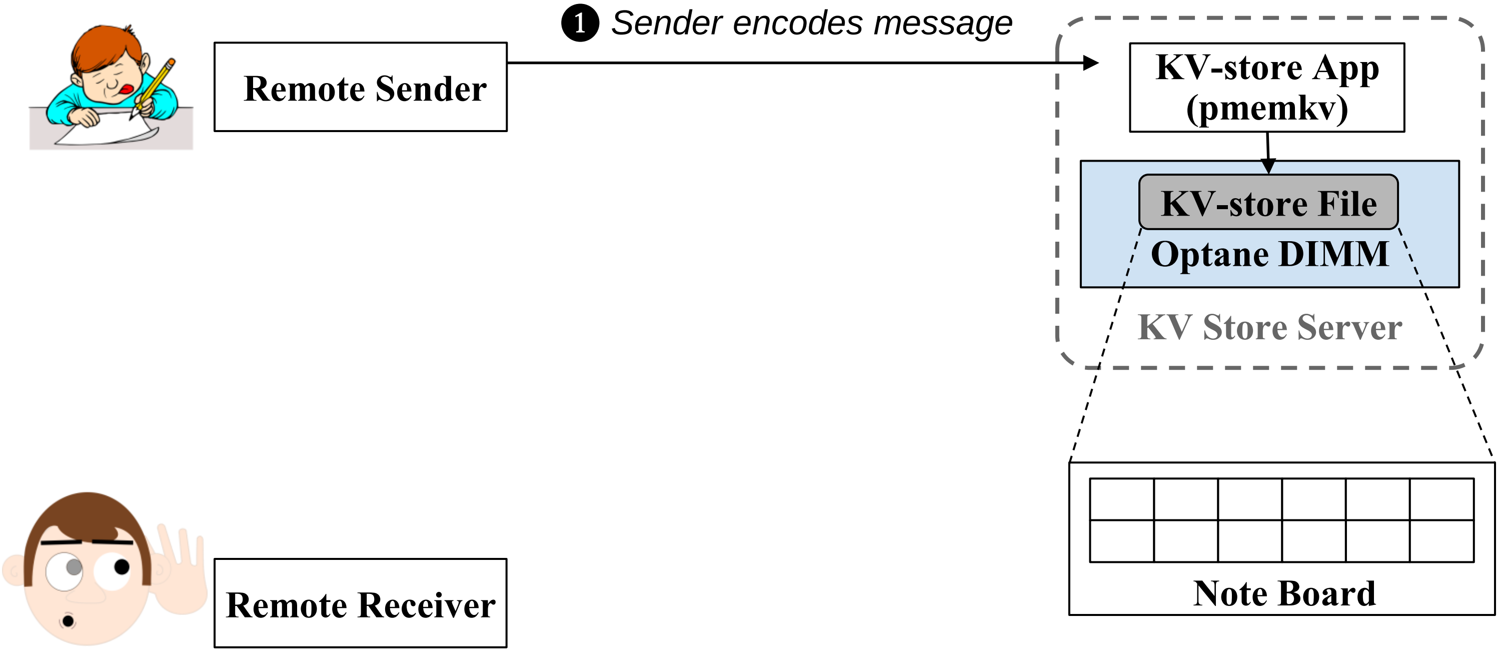

❶ *Sender encodes message*

Remote Sender

KV-store App (pmemkv)

KV-store File

Optane DIMM

KV Store Server

Note Board

Remote Receiver

A Realization

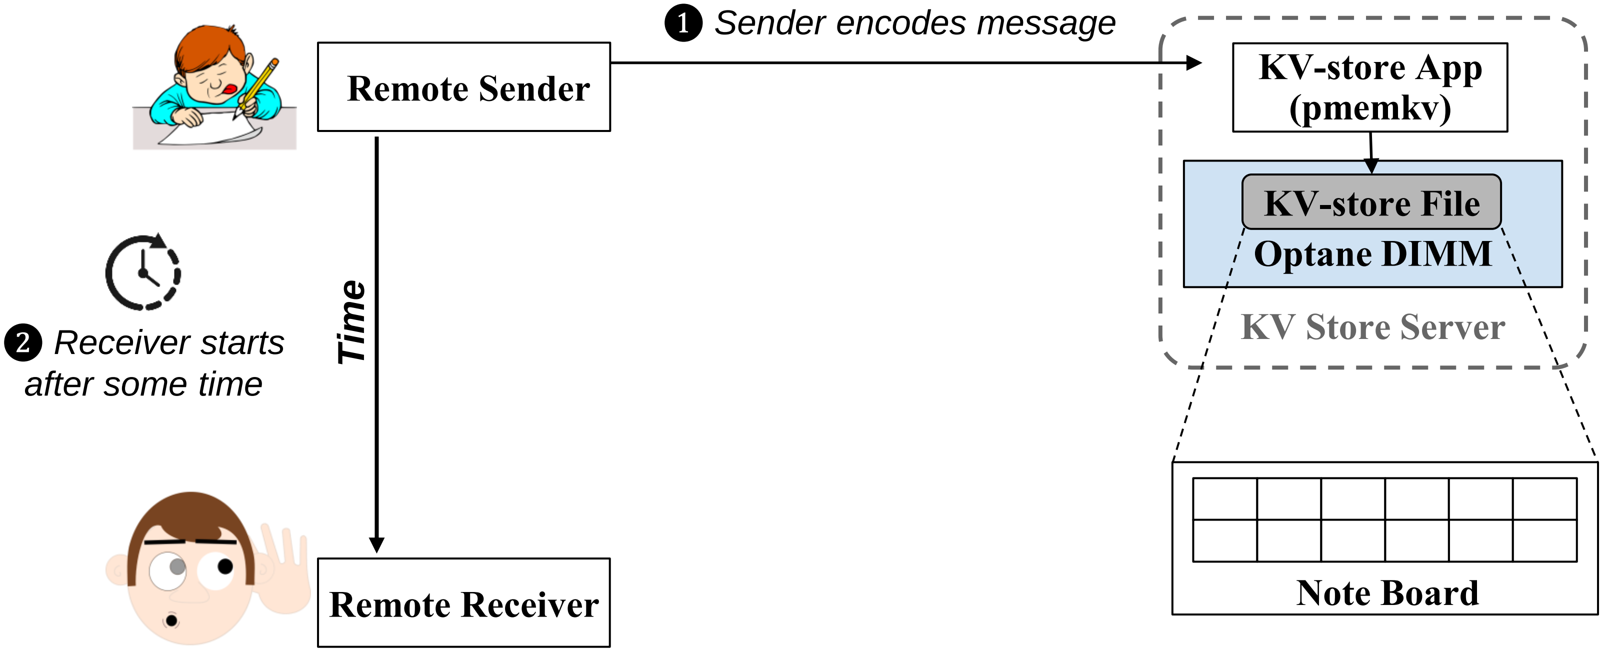

❶ *Sender encodes message*

Remote Sender

KV-store App (pmemkv)

KV-store File

Optane DIMM

KV Store Server

Time

❷ *Receiver starts after some time*

Remote Receiver

Note Board

A Realization

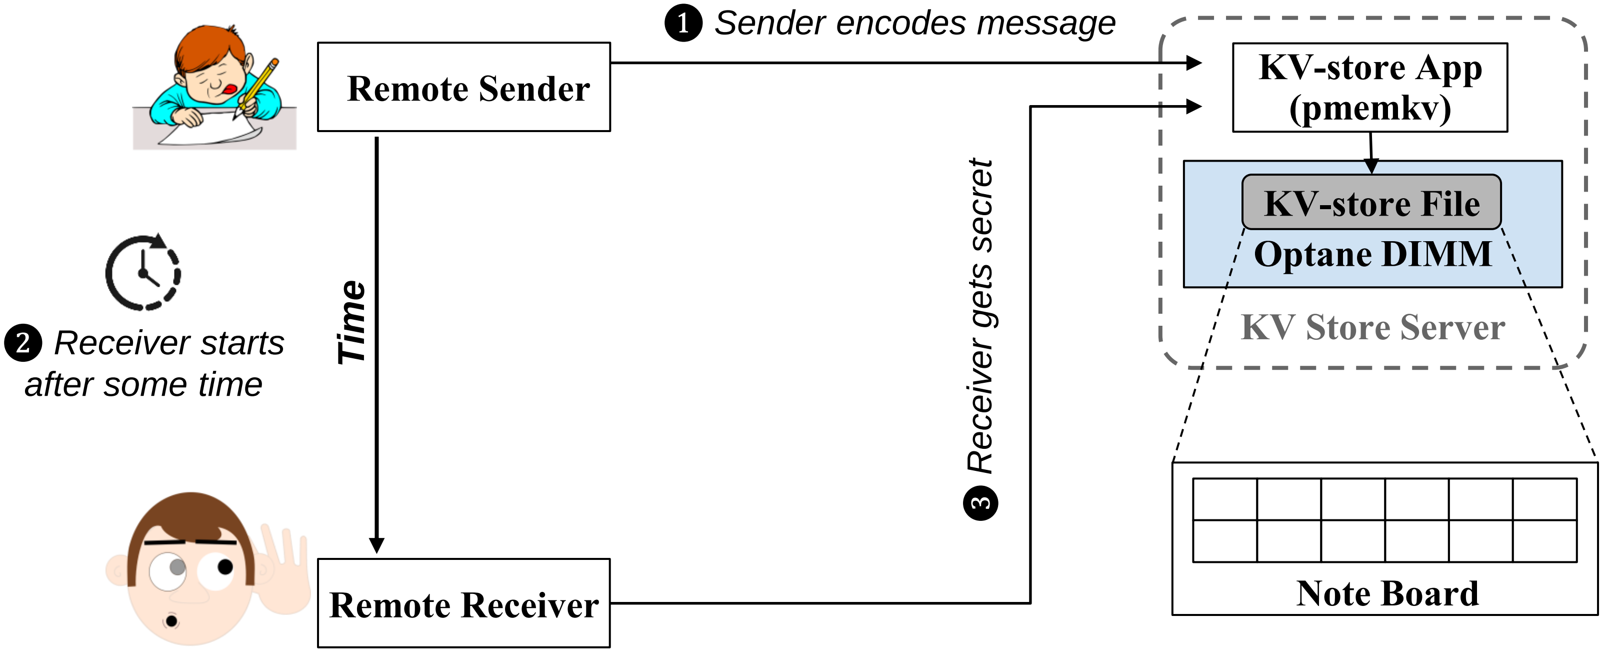

A Realization

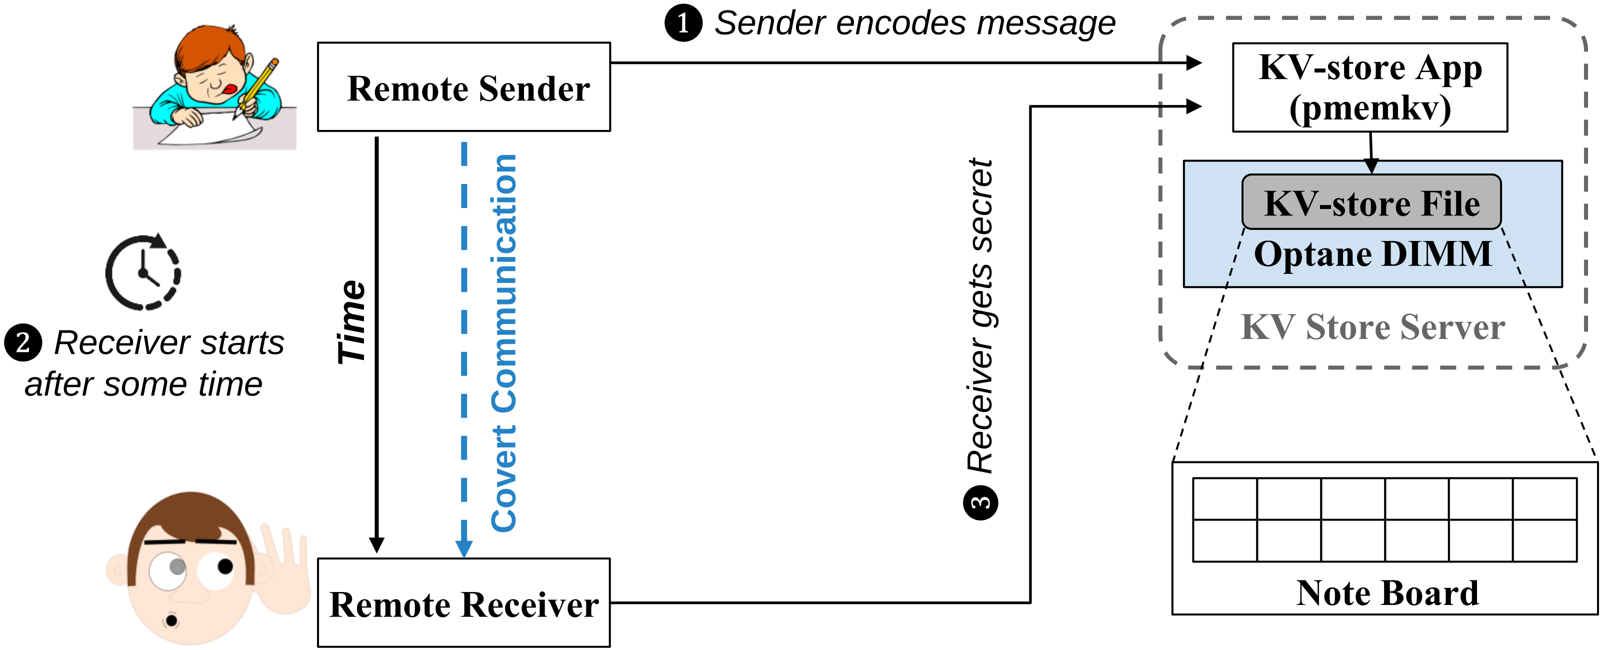

A Realization

Result

A Realization

Result

A Realization

Result

A Realization

Result

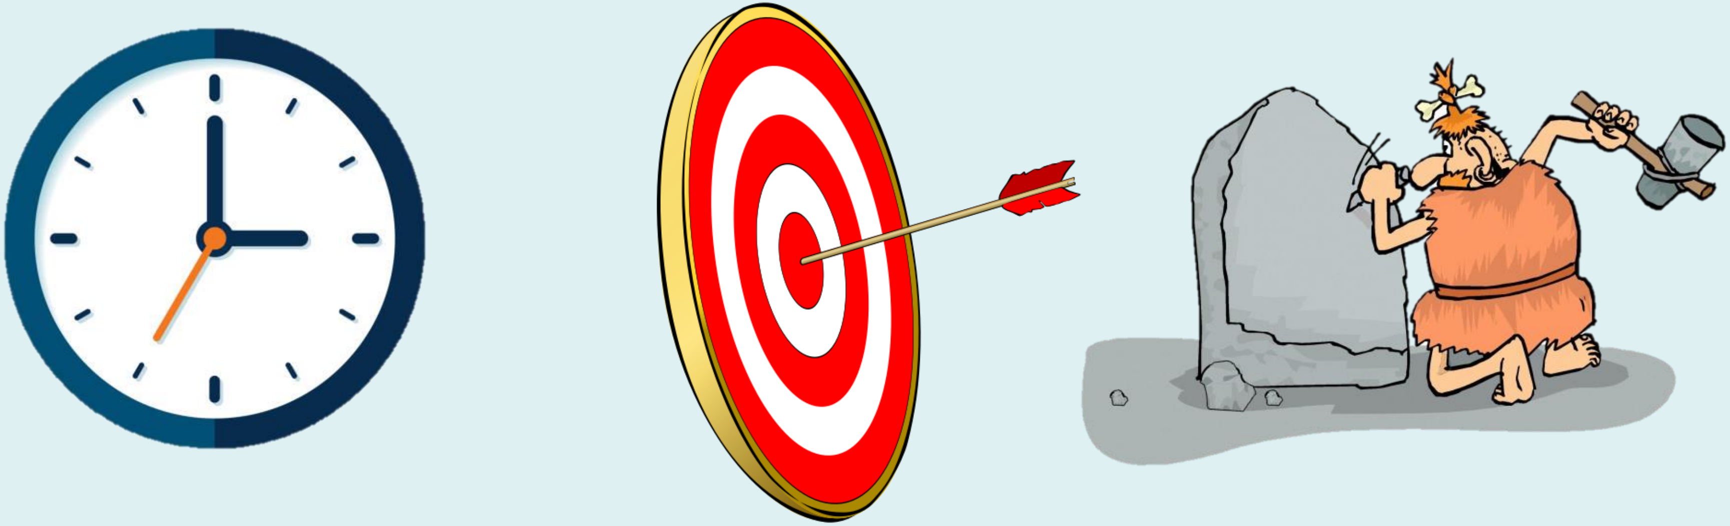

Looking at the Future

Looking at the Future

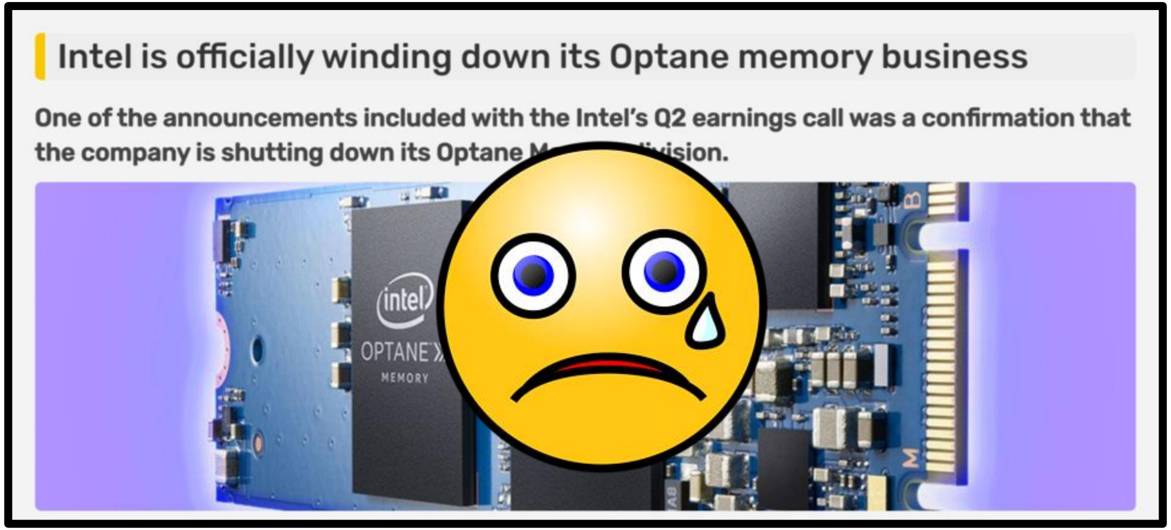

Intel is officially winding down its Optane memory business

One of the announcements included with the Intel's Q2 earnings call was a confirmation that the company is shutting down its Optane M[illegible]ivision.

Looking at the Future

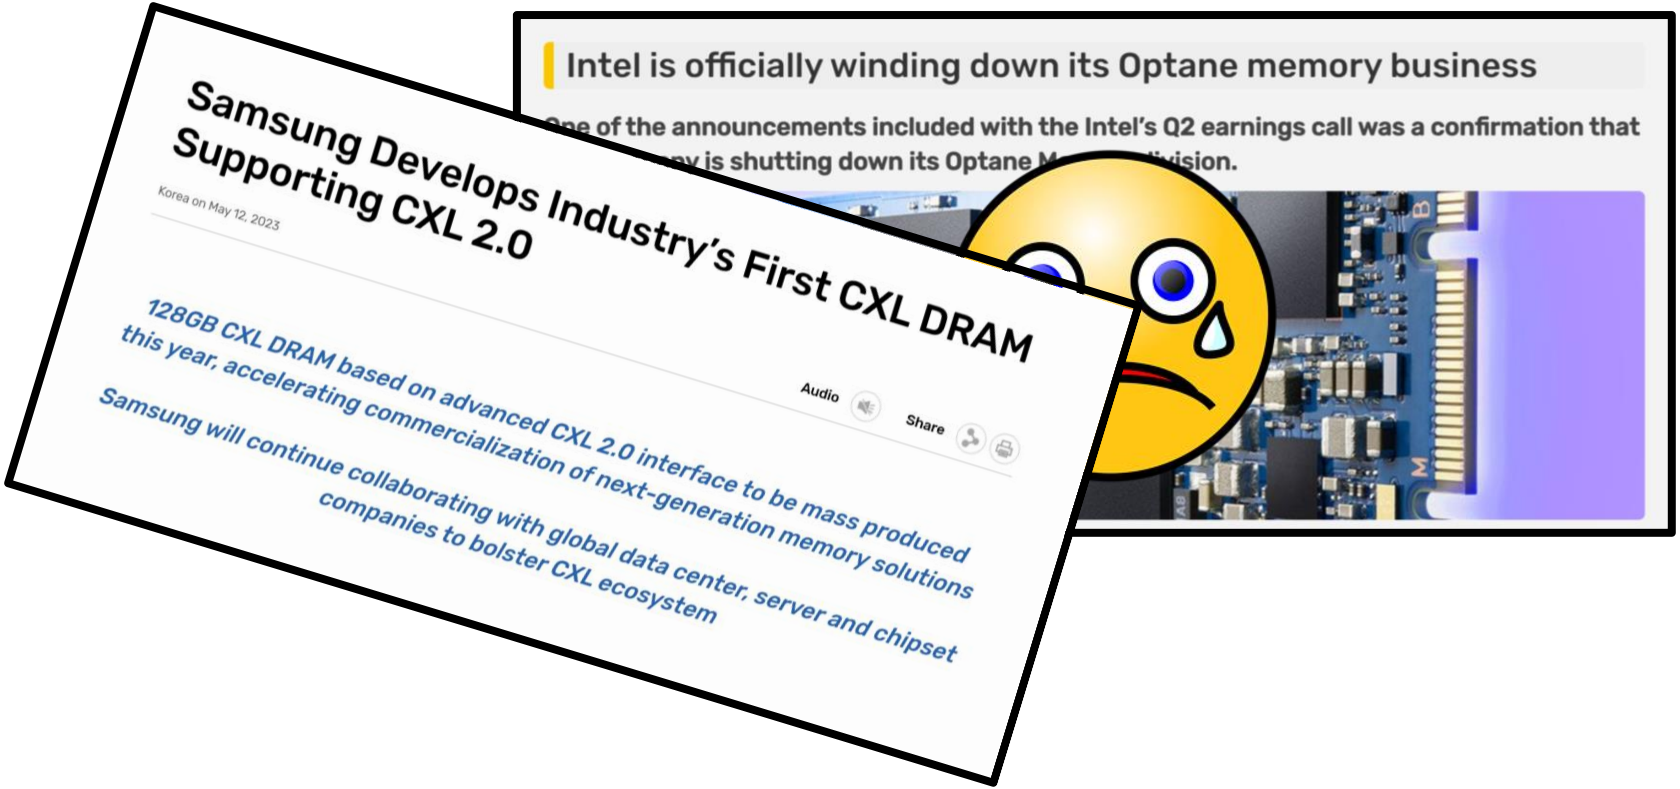

Samsung Develops Industry's First CXL DRAM Supporting CXL 2.0

Korea on May 12, 2023

Audio

Share

128GB CXL DRAM based on advanced CXL 2.0 interface to be mass produced this year, accelerating commercialization of next-generation memory solutions

Samsung will continue collaborating with global data center, server and chipset companies to bolster CXL ecosystem

Intel is officially winding down its Optane memory business

One of the announcements included with the Intel's Q2 earnings call was a confirmation that ny is shutting down its Optane M ivision.

Looking at the Future

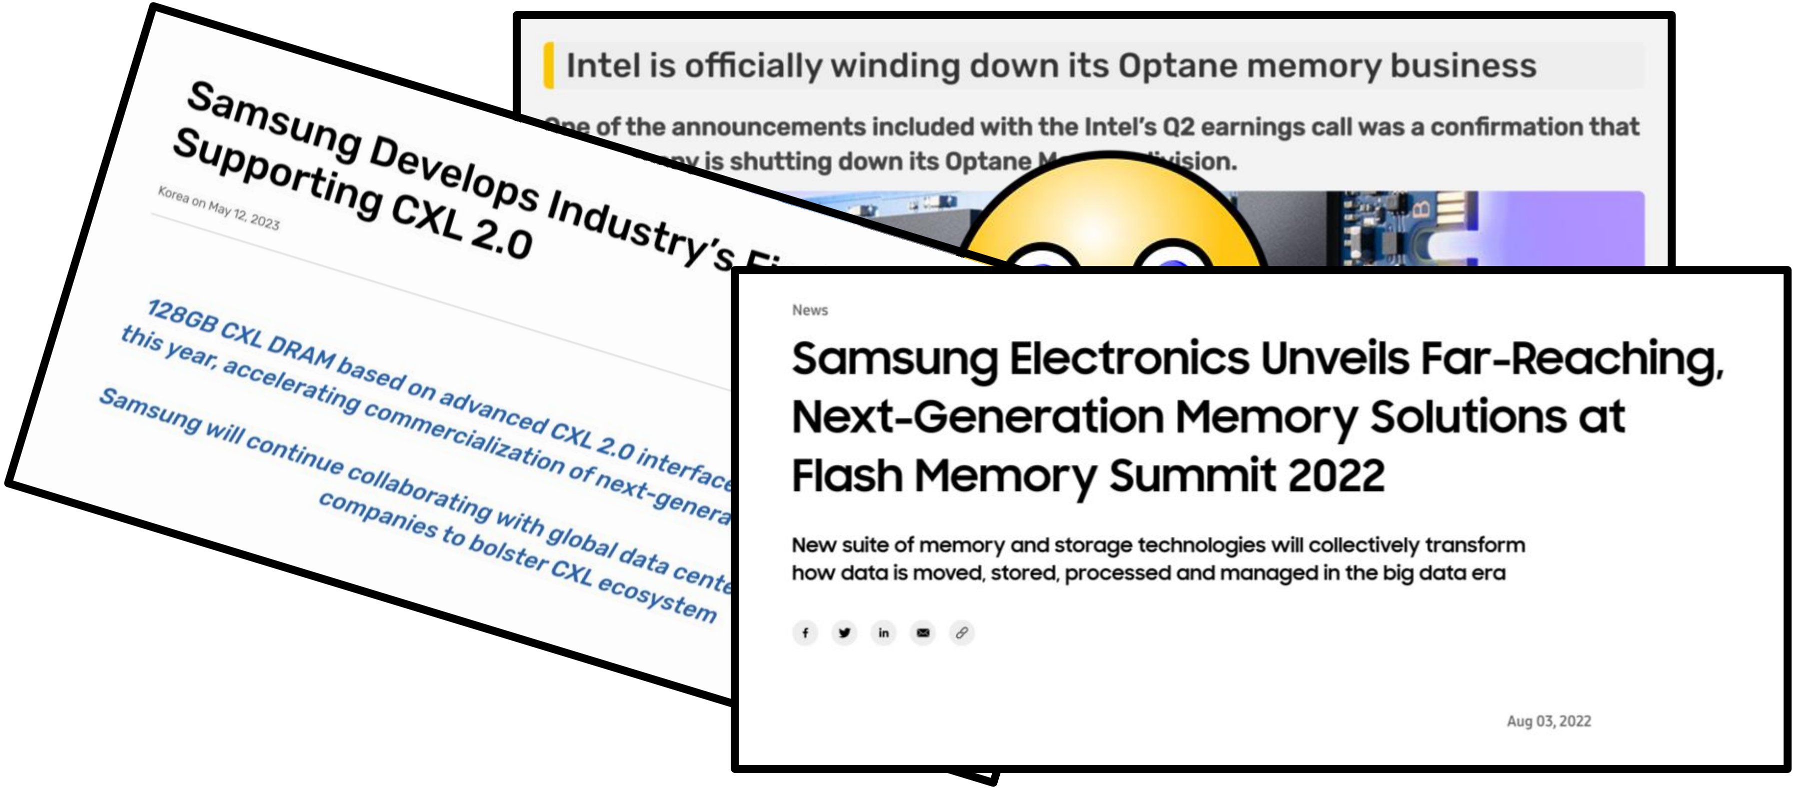

Intel is officially winding down its Optane memory business

ne of the announcements included with the Intel's Q2 earnings call was a confirmation that
pany is shutting down its Optane M ivision.

Samsung Develops Industry's Fi
Supporting CXL 2.0

Korea on May 12, 2023

128GB CXL DRAM based on advanced CXL 2.0 interface
this year, accelerating commercialization of next-genera
Samsung will continue collaborating with global data cente
companies to bolster CXL ecosystem

News

Samsung Electronics Unveils Far-Reaching, Next-Generation Memory Solutions at Flash Memory Summit 2022

New suite of memory and storage technologies will collectively transform how data is moved, stored, processed and managed in the big data era

Aug 03, 2022

Looking at the Future

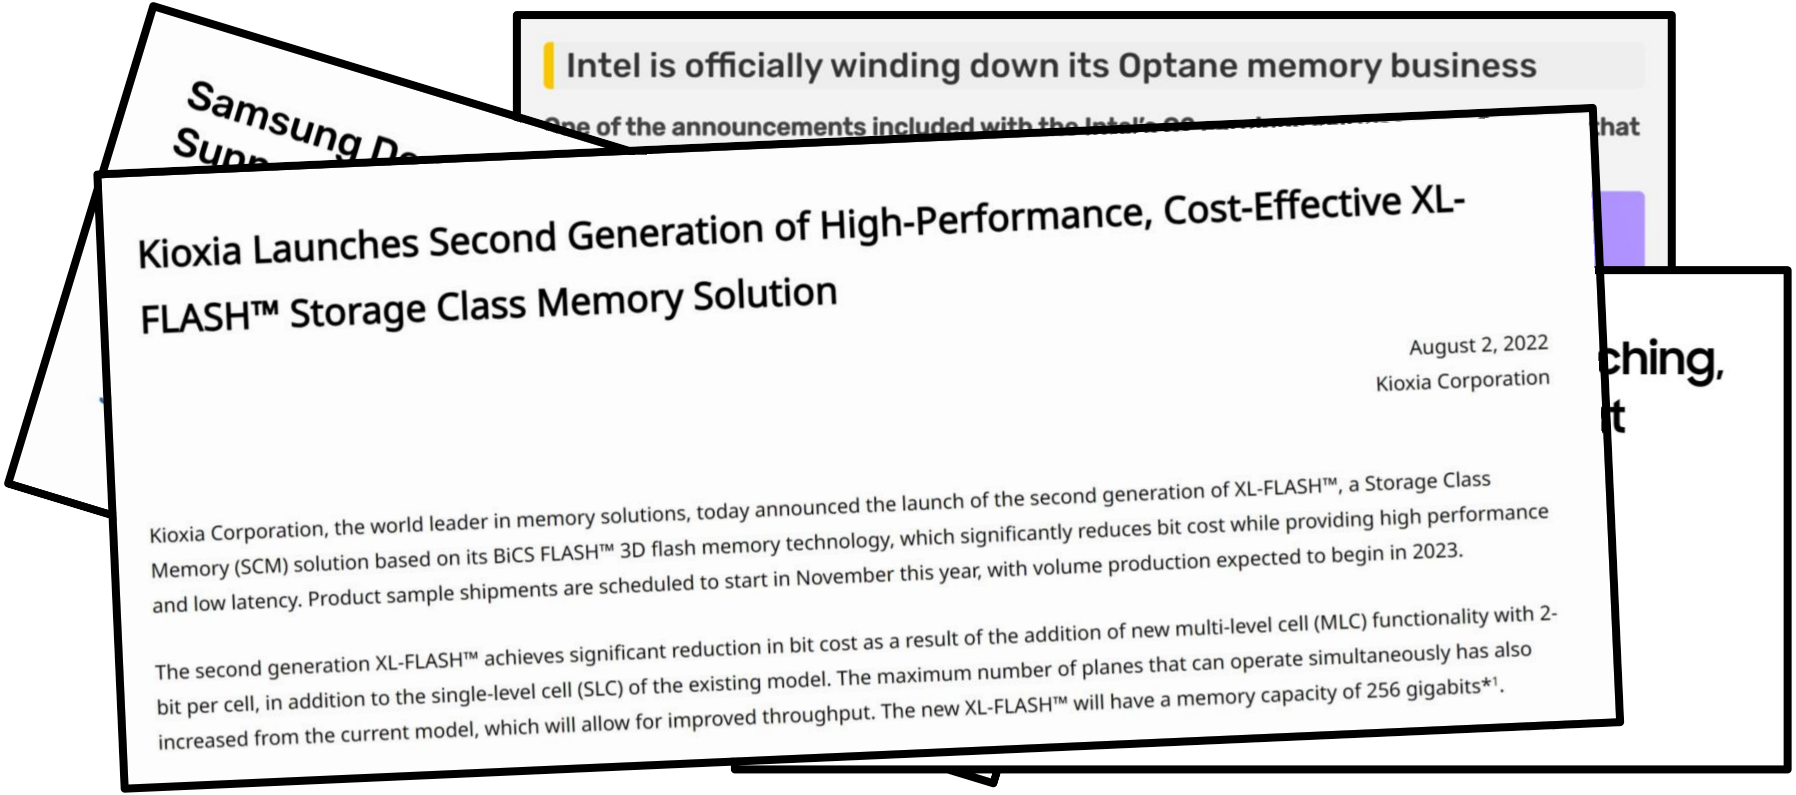

Samsung D

Sup

Intel is officially winding down its Optane memory business

One of the announcements included with the Intel's Q2 that

ching,

t

Kioxia Launches Second Generation of High-Performance, Cost-Effective XL-FLASH™ Storage Class Memory Solution

August 2, 2022

Kioxia Corporation

Kioxia Corporation, the world leader in memory solutions, today announced the launch of the second generation of XL-FLASH™, a Storage Class Memory (SCM) solution based on its BiCS FLASH™ 3D flash memory technology, which significantly reduces bit cost while providing high performance and low latency. Product sample shipments are scheduled to start in November this year, with volume production expected to begin in 2023.

The second generation XL-FLASH™ achieves significant reduction in bit cost as a result of the addition of new multi-level cell (MLC) functionality with 2-bit per cell, in addition to the single-level cell (SLC) of the existing model. The maximum number of planes that can operate simultaneously has also increased from the current model, which will allow for improved throughput. The new XL-FLASH™ will have a memory capacity of 256 gigabits*[1].

Summary

Summary

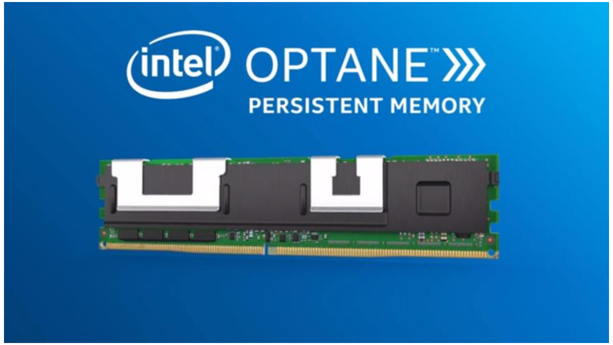

Summary

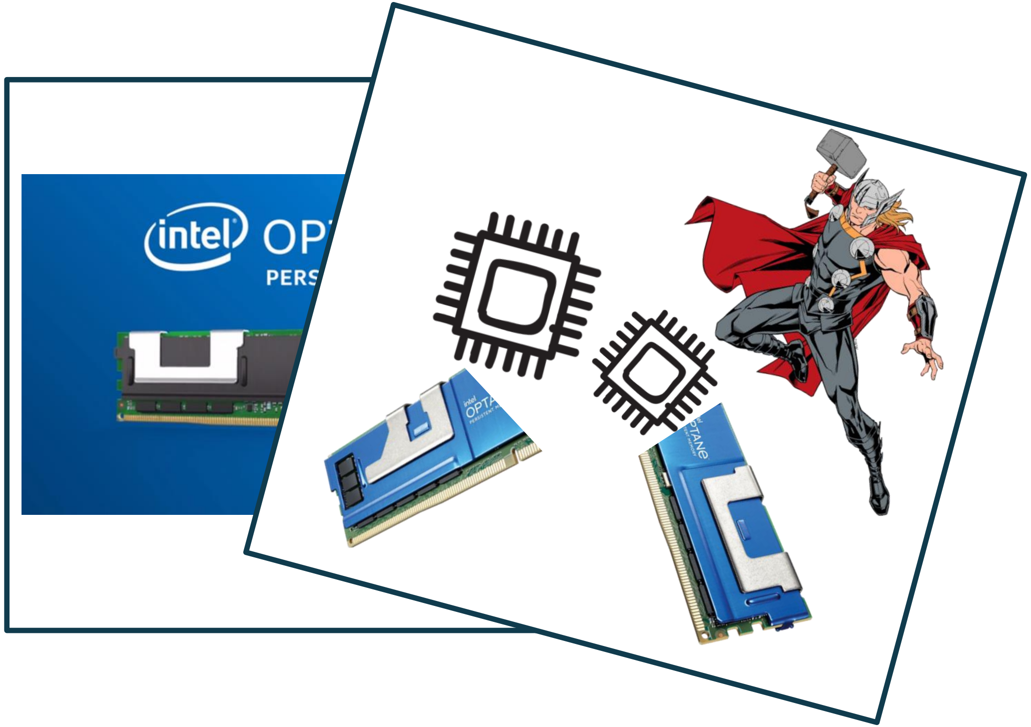

Summary

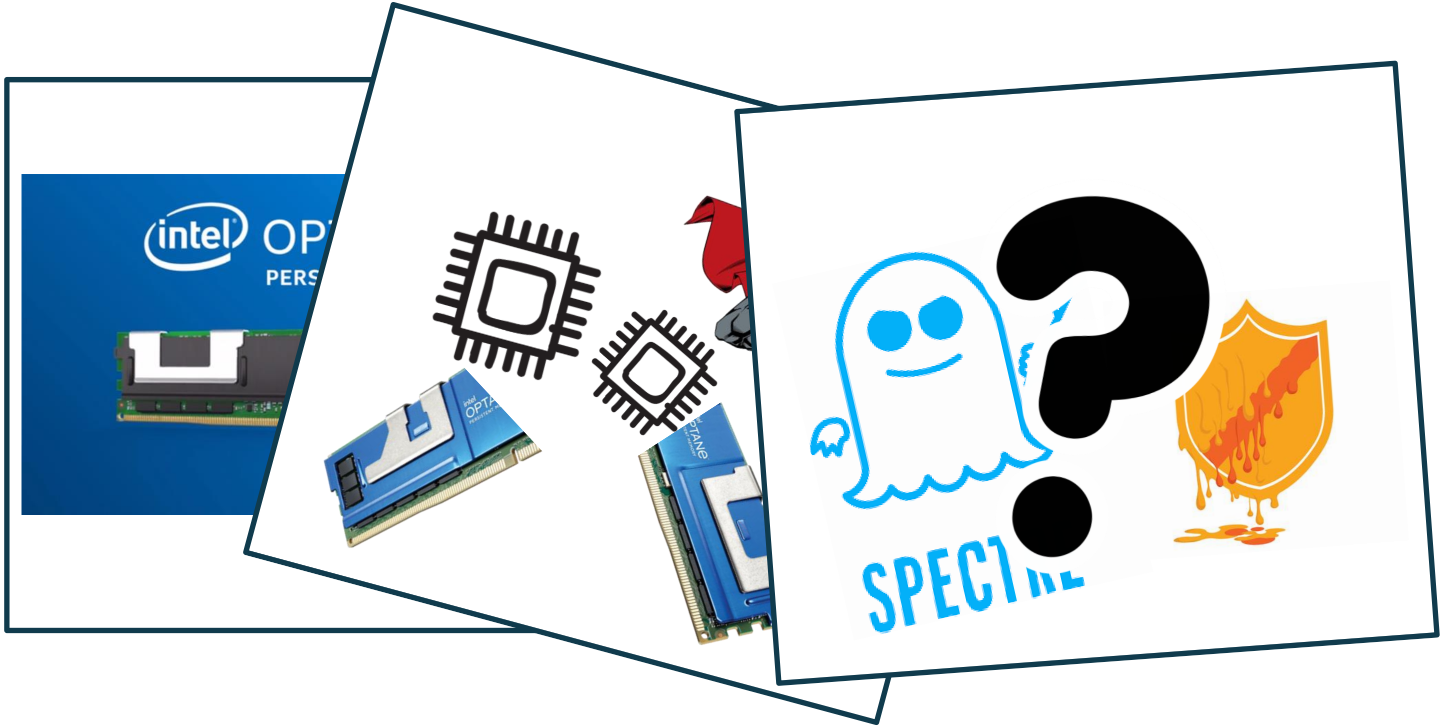

Summary

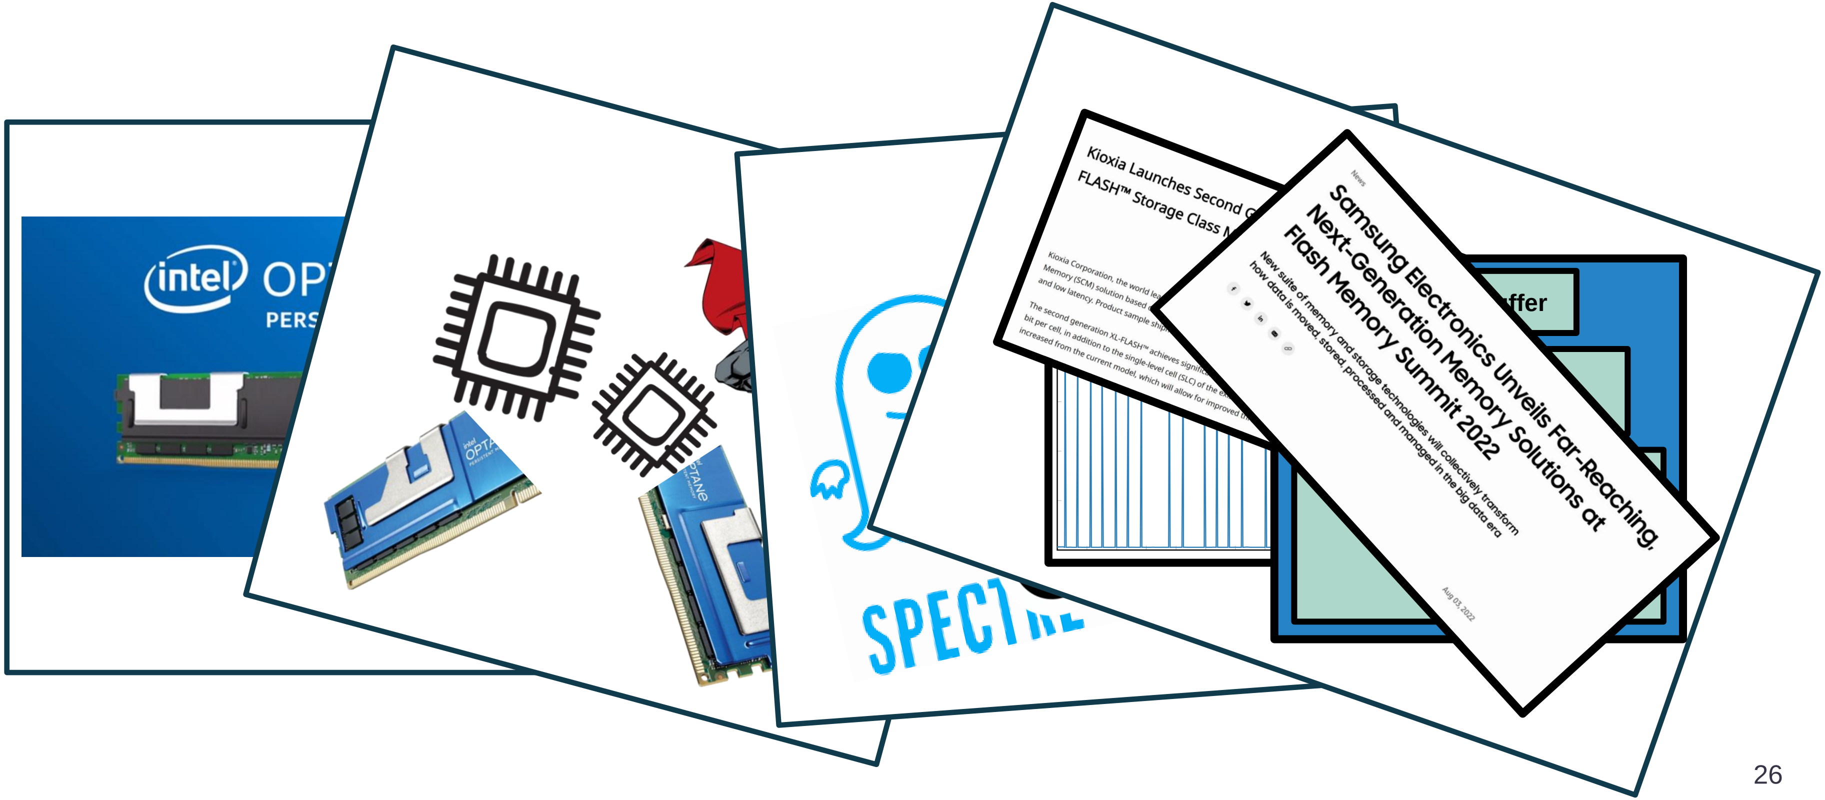

Kioxia Launches Second G

FLASH™ Storage Class M

Kioxia Corporation, the world lea

Memory (SCM) solution based o

and low latency. Product sample shipm

The second generation XL-FLASH™ achieves significa

bit per cell, in addition to the single-level cell (SLC) of the ex

increased from the current model, which will allow for improved th

News

Samsung Electronics Unveils Far-Reaching, Next-Generation Memory Solutions at Flash Memory Summit 2022

New suite of memory and storage technologies will collectively transform how data is moved, stored, processed and managed in the big data era

Aug 03, 2022

Side-Channel Attacks on Optane Persistent Memory

Sihang Liu, **Suraaj Kanniwadi**, Martin Schwarzl, Andreas Kogler,
Daniel Gruss, Samira Khan

Usenix Security Symposium 2023