PURESTORAGE®

Tech Talk

“Why Pure”

PURESTORAGE®

Better Science

Justin Emerson

Technology Evangelist - Pure Storage

From transistors to bits

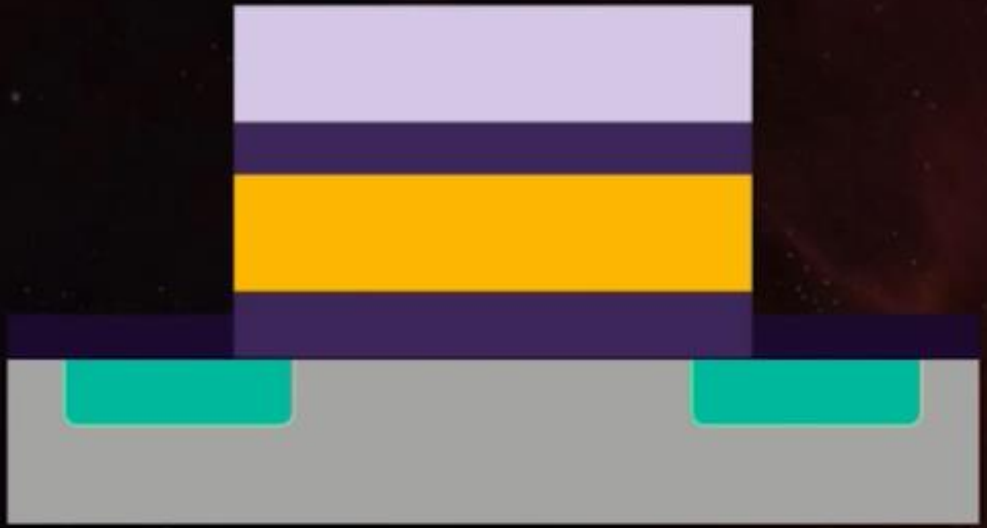

©2022 Pure Storage, Inc.

From transistors to bits

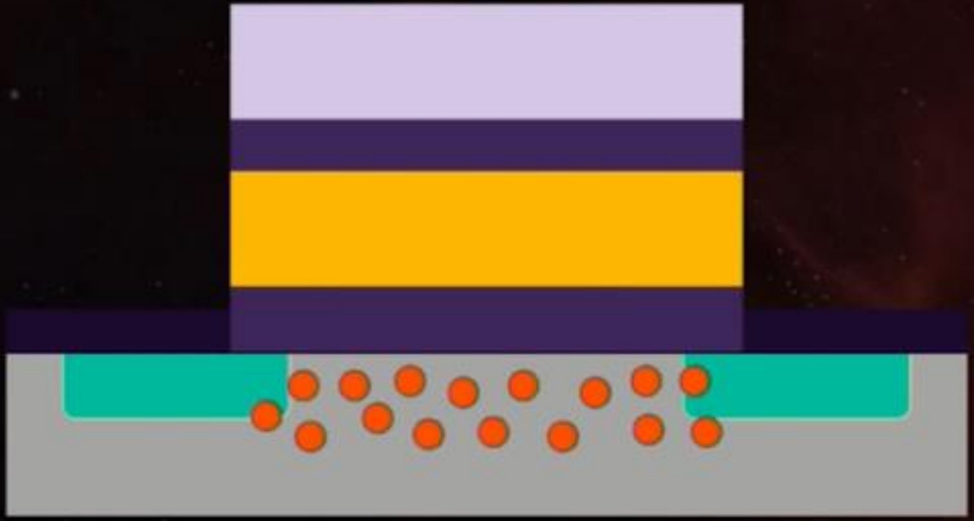

©2022 Pure Storage, Inc.

From transistors to bits

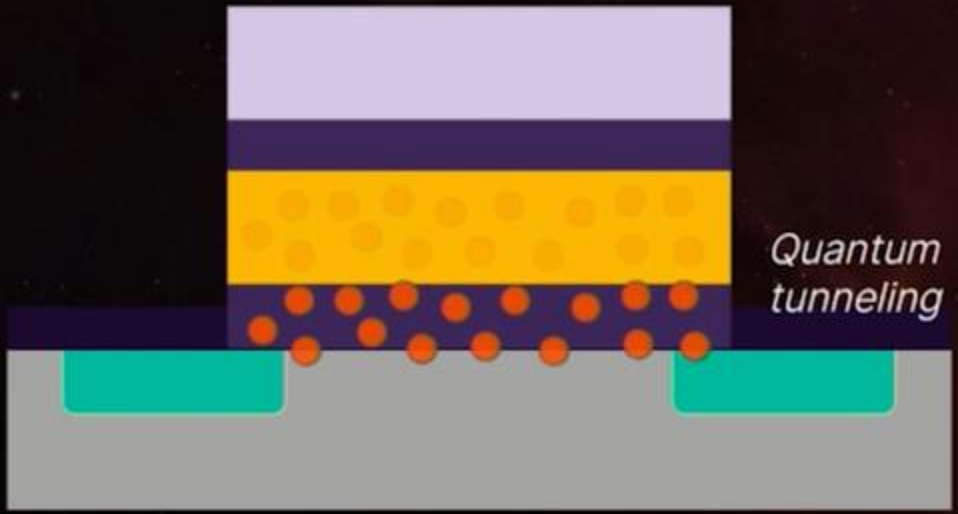

©2022 Pure Storage, Inc.

From transistors to bits

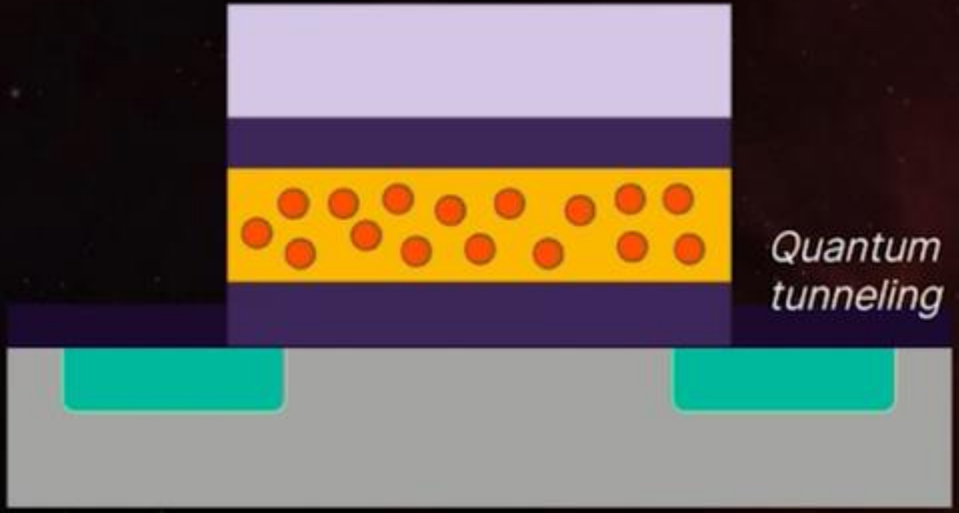

Charge trap flash cell

©2022 Pure Storage, Inc.

From transistors to bits

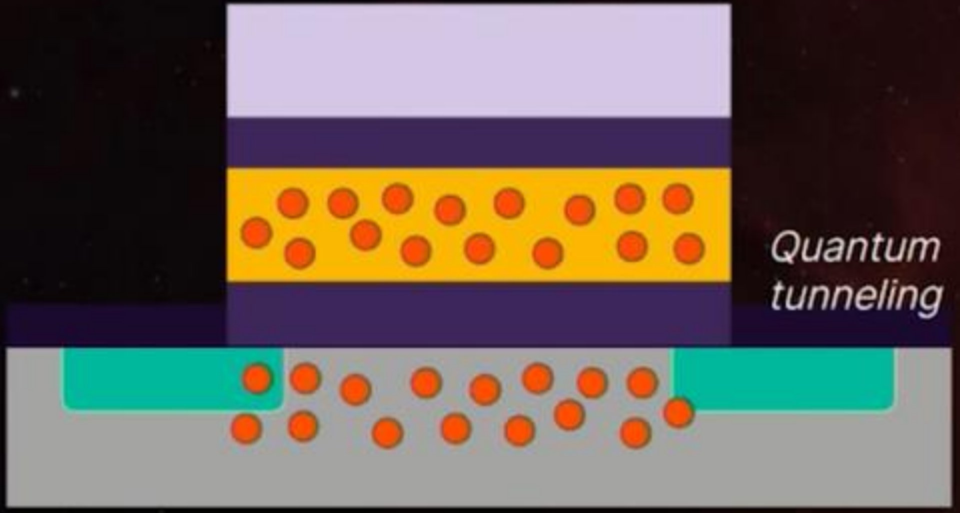

Charge trap flash cell

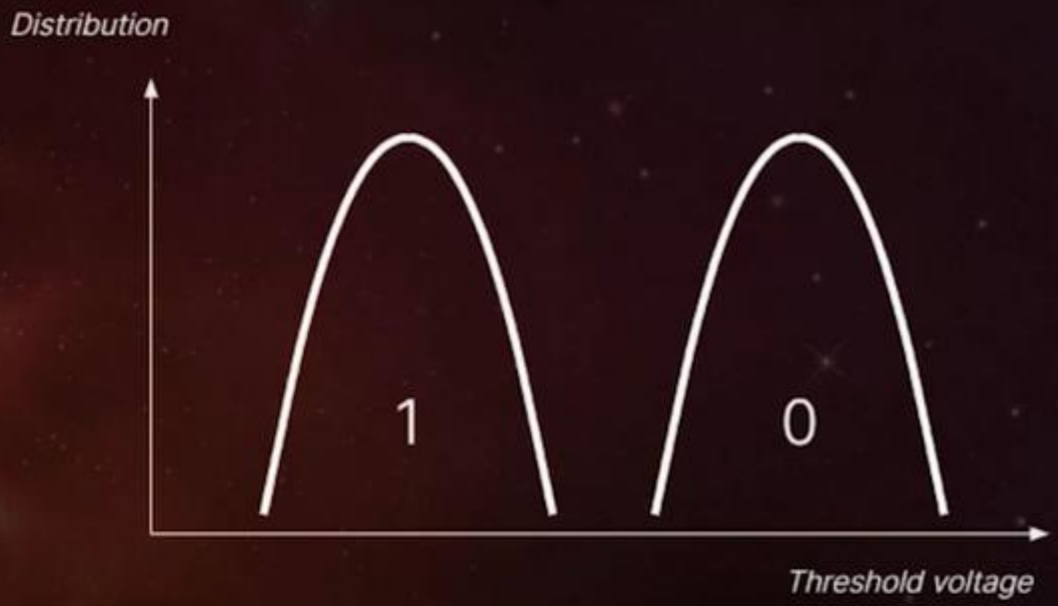

©2022 Pure Storage, Inc.

From one bit to many

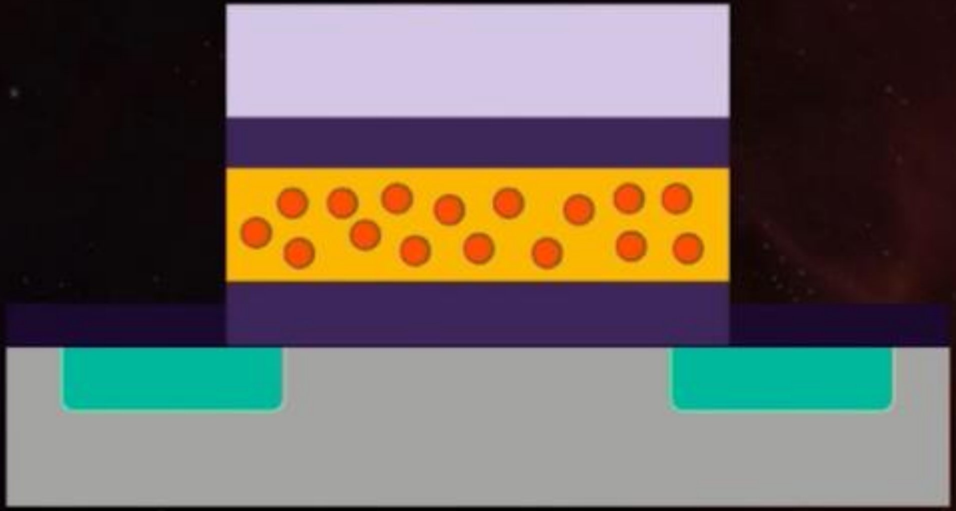

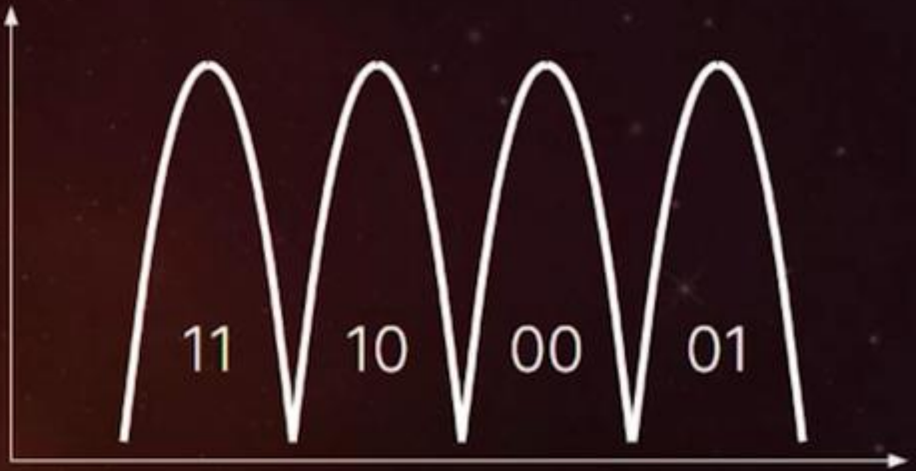

©2022 Pure Storage, Inc.

From one bit to many

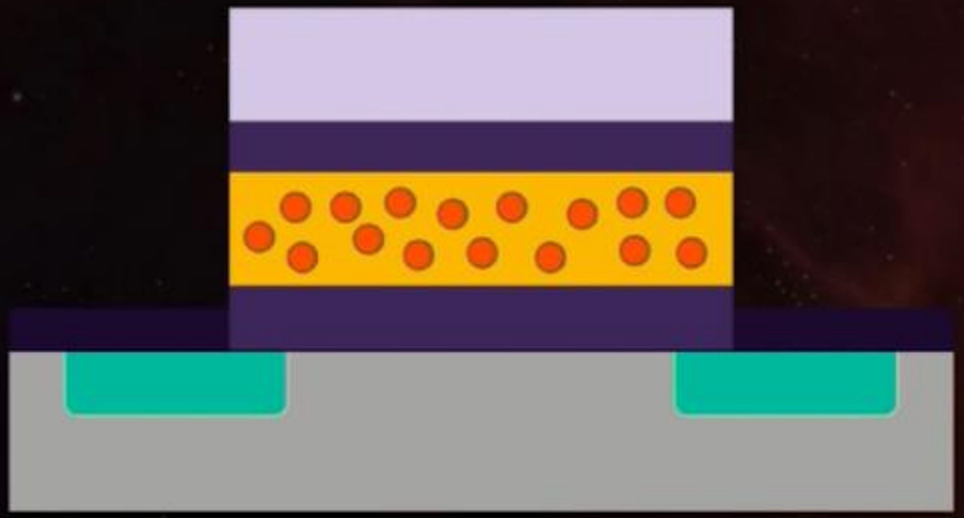

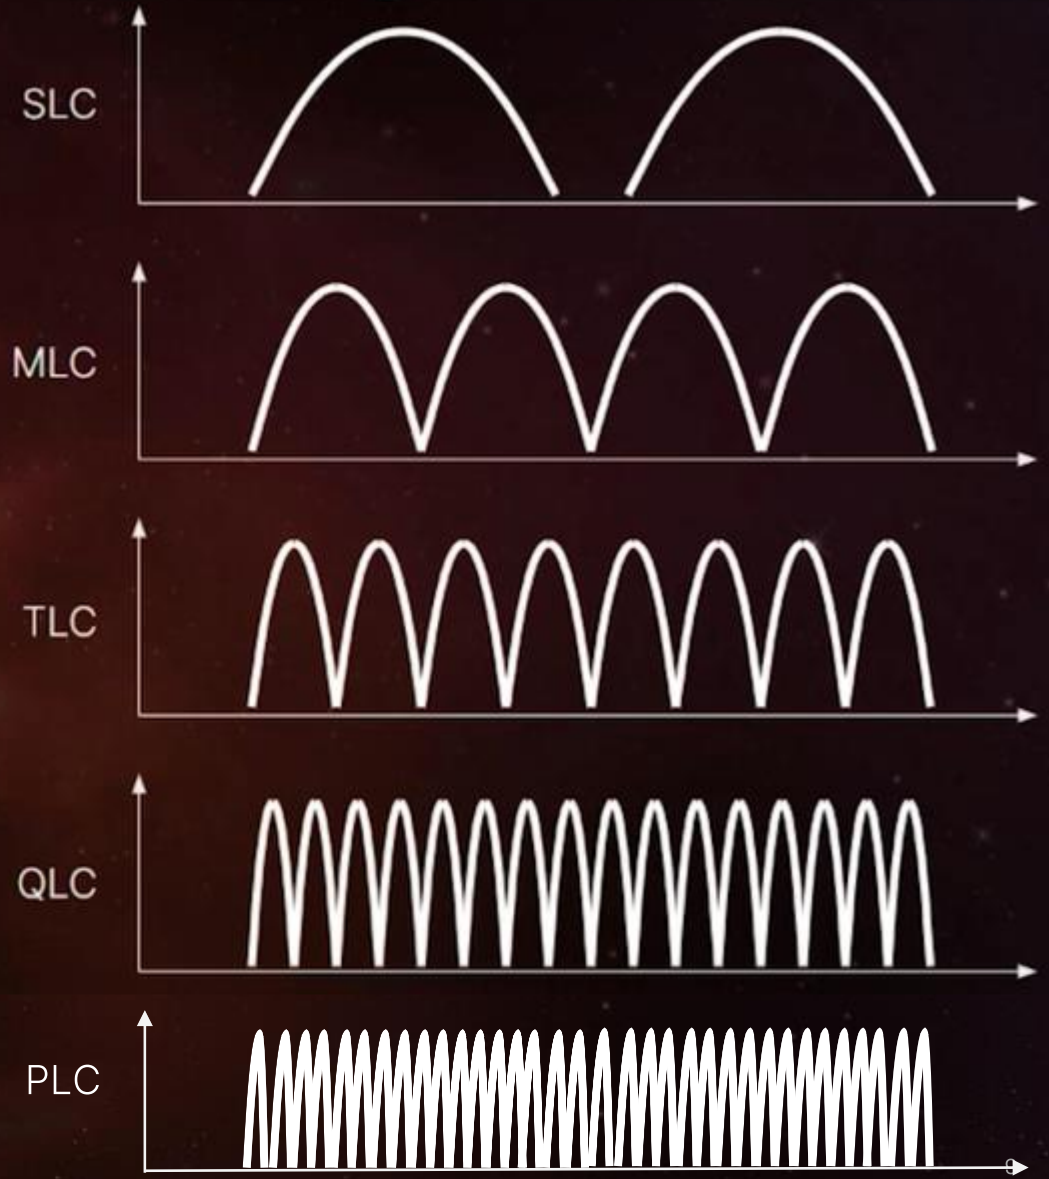

©2022 Pure Storage, Inc.

There is no free lunch

How we wish things worked:

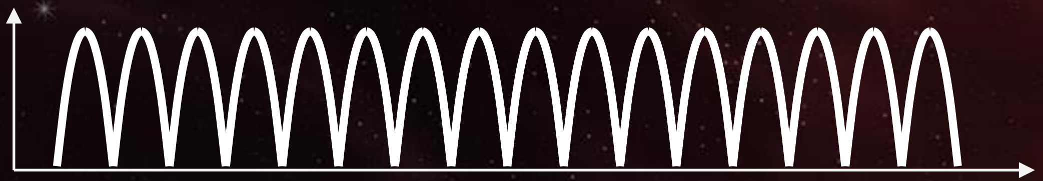

Except it’s more like:

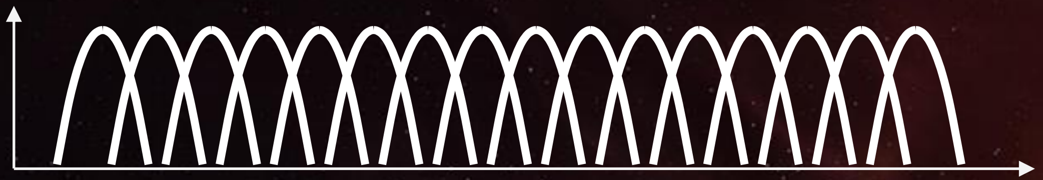

But also changes over time:

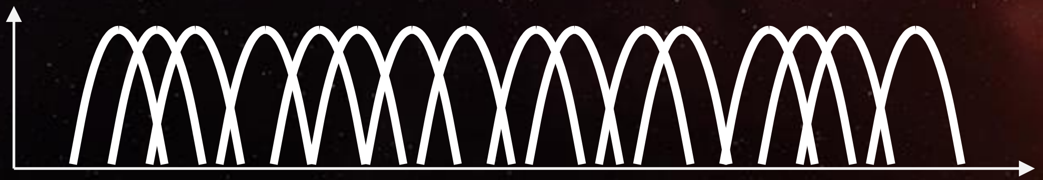

And degrades with usage:

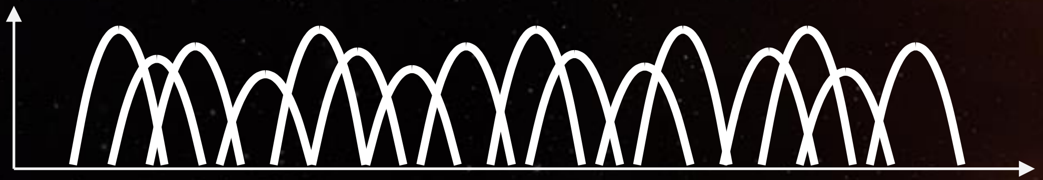

The realities of NAND

- Lower endurance
- More difficult to read
- Slower to program, erase
- Worse retention
- More caveats and quirks

Every generation gets worse

How can future **systems** solve these challenges?

Should we treat SSDs as a commodity?

Reliability is *someone else's problem*

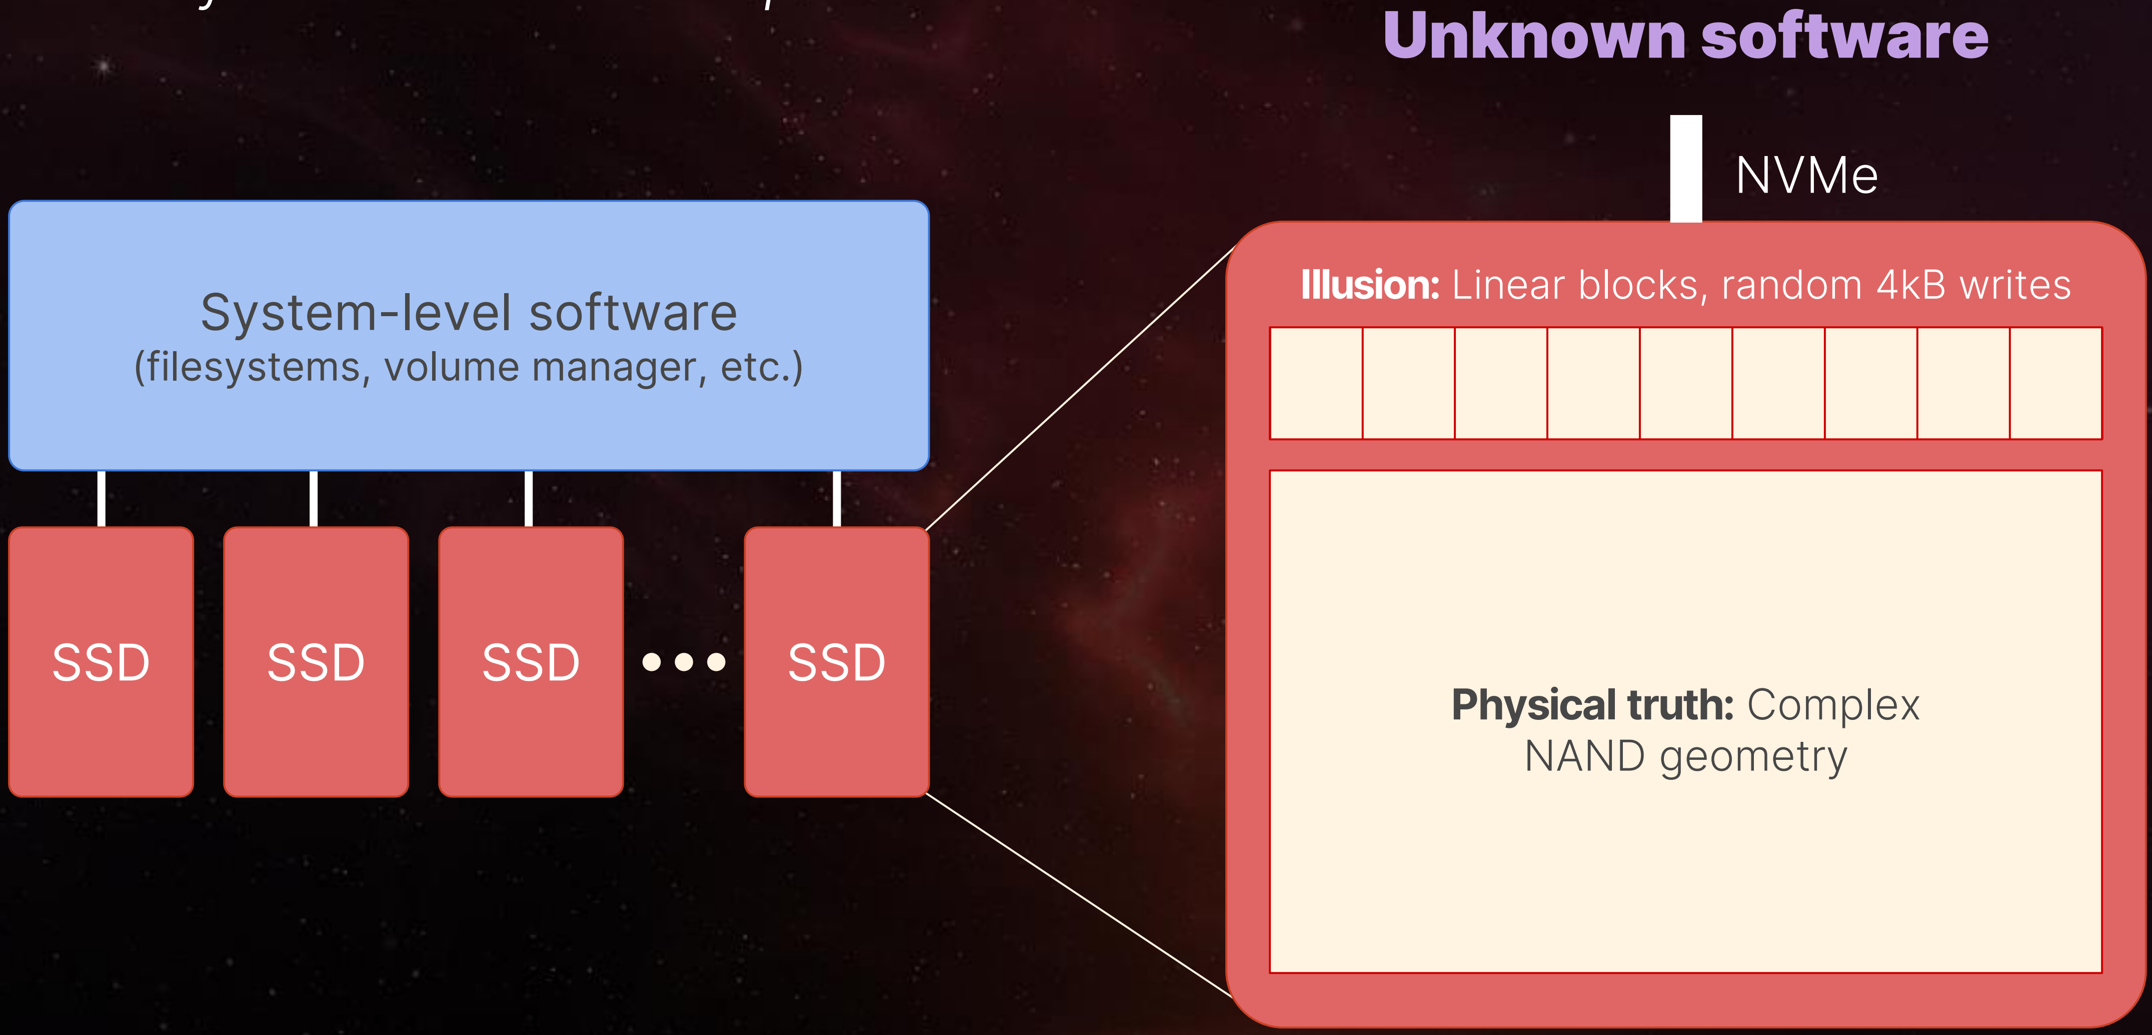
Unknown software
NVMe
System-level software
(filesystems, volume manager, etc.)
SSD
SSD
SSD
SSD
Illusion: Linear blocks, random 4kB writes
Physical truth: Complex NAND geometry

©2022 Pure Storage, Inc.

Should we treat SSDs as a commodity?

Reliability is *someone else's problem*

1

Low endurance
=> minimize write amplification

Random overwrites create fragmentation, and individual SSDs **lack context** to separate block-level lifetimes, resulting in high write amplification.

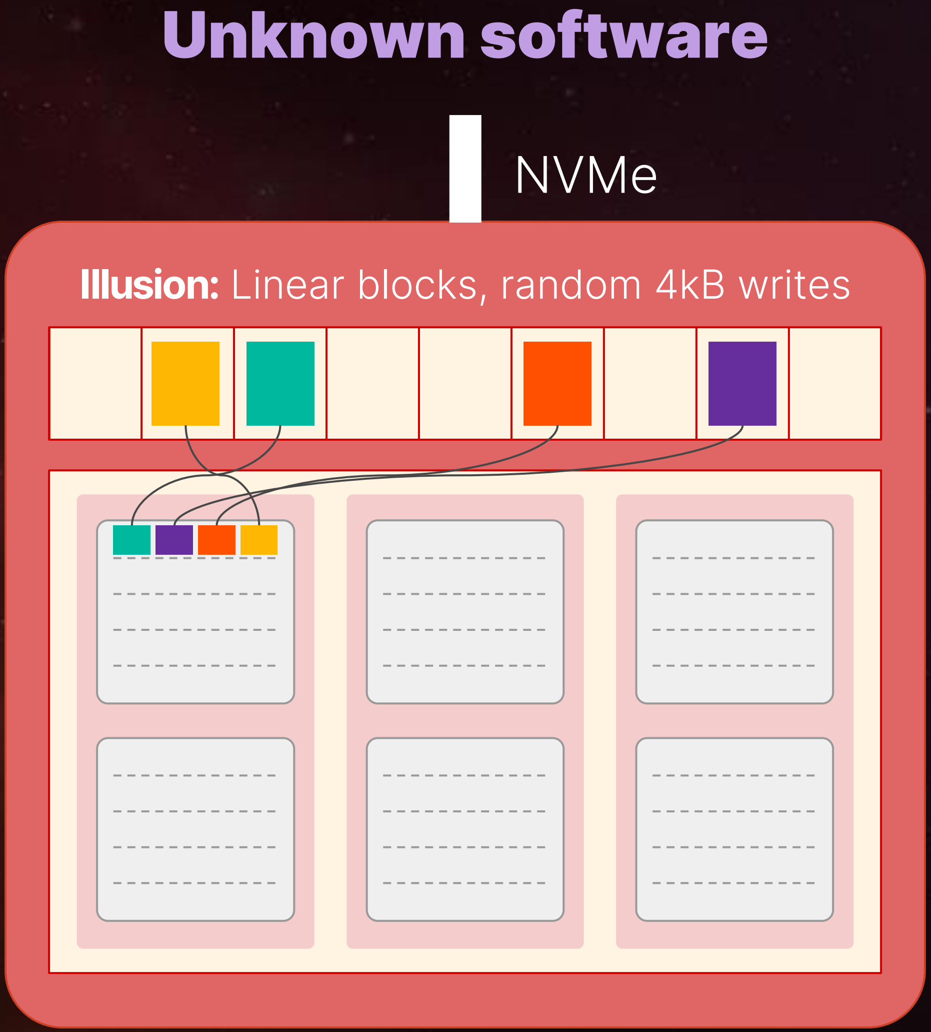

©2022 Pure Storage, Inc.

Should we treat SSDs as a commodity?

Reliability is *someone else's problem*

1

Low endurance
=> minimize write amplification

Random overwrites create fragmentation, and individual SSDs **lack context** to separate block-level lifetimes, resulting in high write amplification.

Unknown software

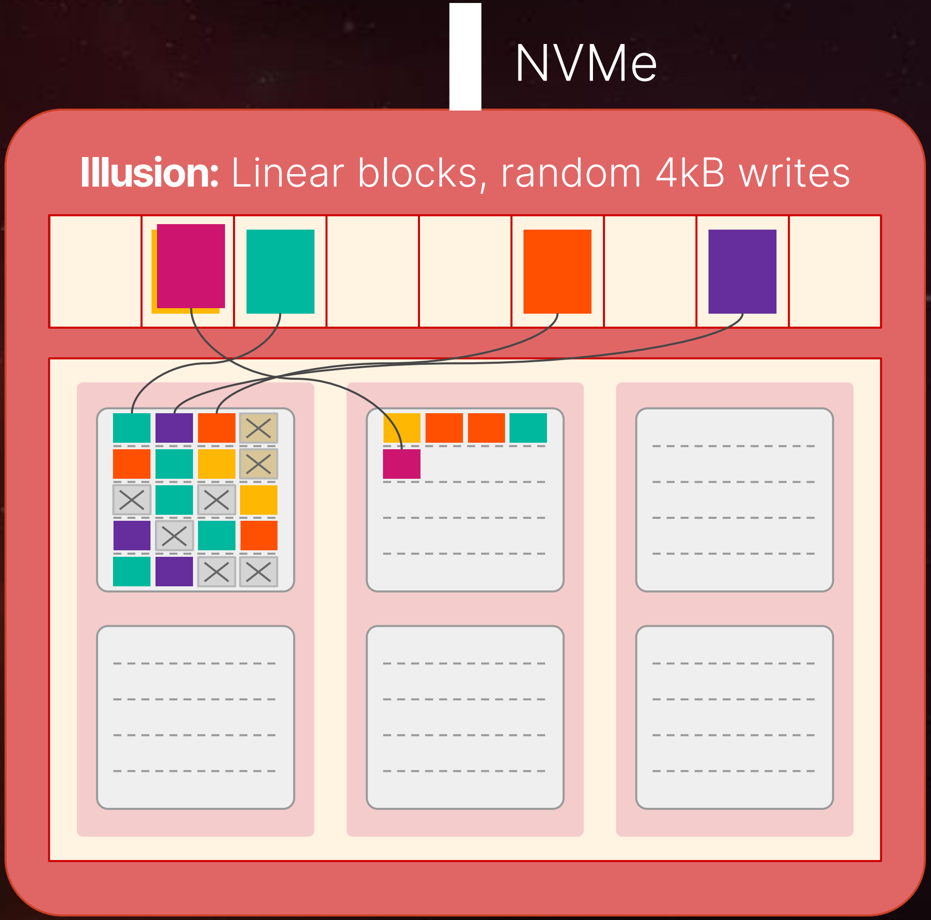

Should we treat SSDs as a commodity?

Performance is *someone else's problem*

1

Low endurance
=> minimize write amplification

2

Long program/erase times
=> control tail latencies

Applications have **no visibility** into the placement of blocks on physical media, and reads may be stuck behind slow (10s milliseconds) operations to conflicting die.

Unknown software

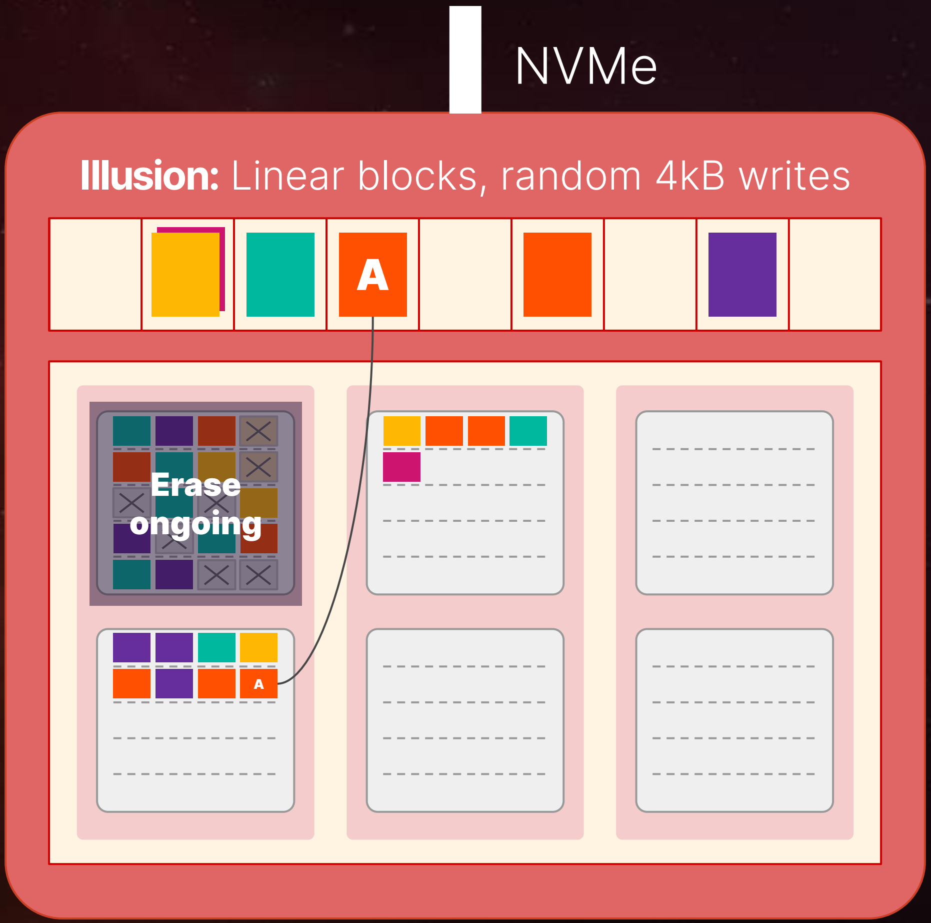

Read of A stalled behind erase in same die

©2022 Pure Storage, Inc.

Should we treat SSDs as a commodity?

Efficiency is *someone else's problem*

1

Low endurance
=> minimize write amplification

2

Long program/erase times
=> control tail latencies

3

Media overheads
=> create efficient end-to-end mappings

Supporting random access is expensive. A 10PB (raw) system has **10TB of SSD DRAM** and **2PB of hidden flash**.

Unknown software

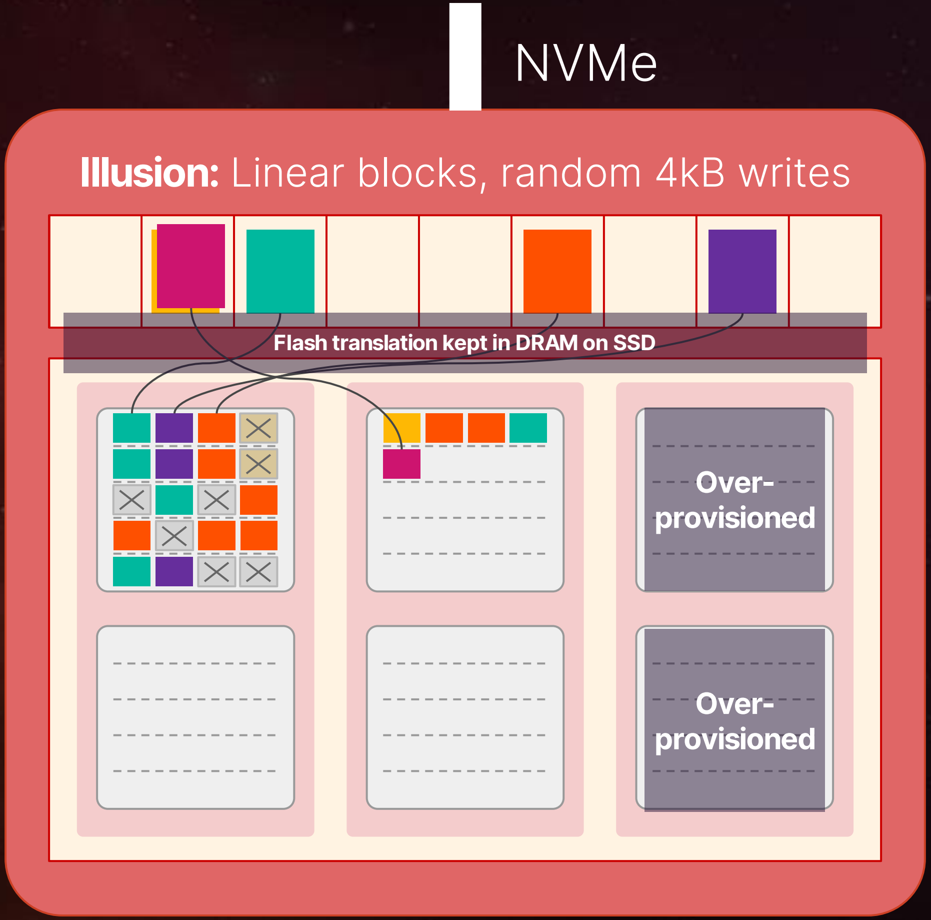
NVMe
Illusion: Linear blocks, random 4kB writes
Flash translation kept in DRAM on SSD
Over-provisioned
Over-provisioned

©2022 Pure Storage, Inc.

Should we treat SSDs as a commodity?

Efficiency is *someone else's problem*

1

Low endurance
=> minimize write amplification

2

Long program/erase times
=> control tail latencies

3

Media overheads
=> create efficient end-to-end mappings

Supporting random access is expensive. A 10PB (raw) system has **10TB of SSD DRAM** and **2PB of hidden flash**.

Unknown software

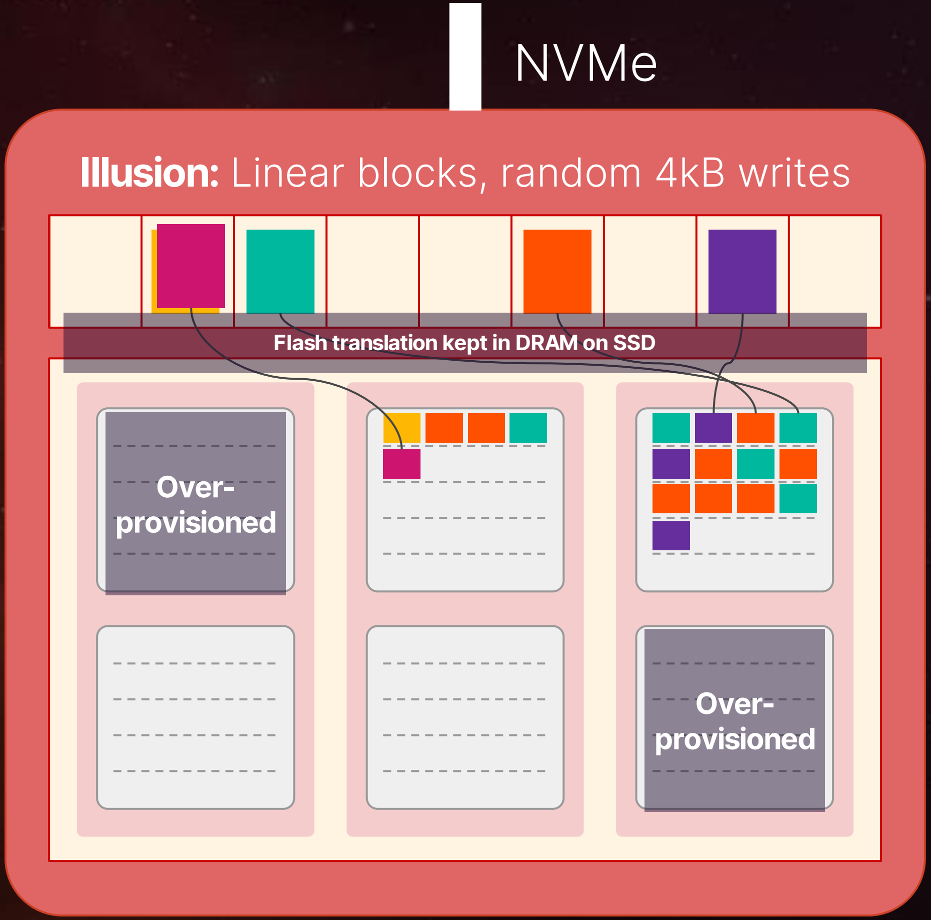
NVMe
Illusion: Linear blocks, random 4kB writes
Flash translation kept in DRAM on SSD
Over-provisioned
Over-provisioned

©2022 Pure Storage, Inc.

Modern SSDs are engineering marvels

But large-scale **systems** must go further

©2022 Pure Storage, Inc.

DirectFlash extends flash lifetime

1

Low endurance
=> minimize write amplification

DirectFlash enables Purity SW to colocate data and metadata with similar expected lifetimes, aligning to the underlying physical NAND geometry.

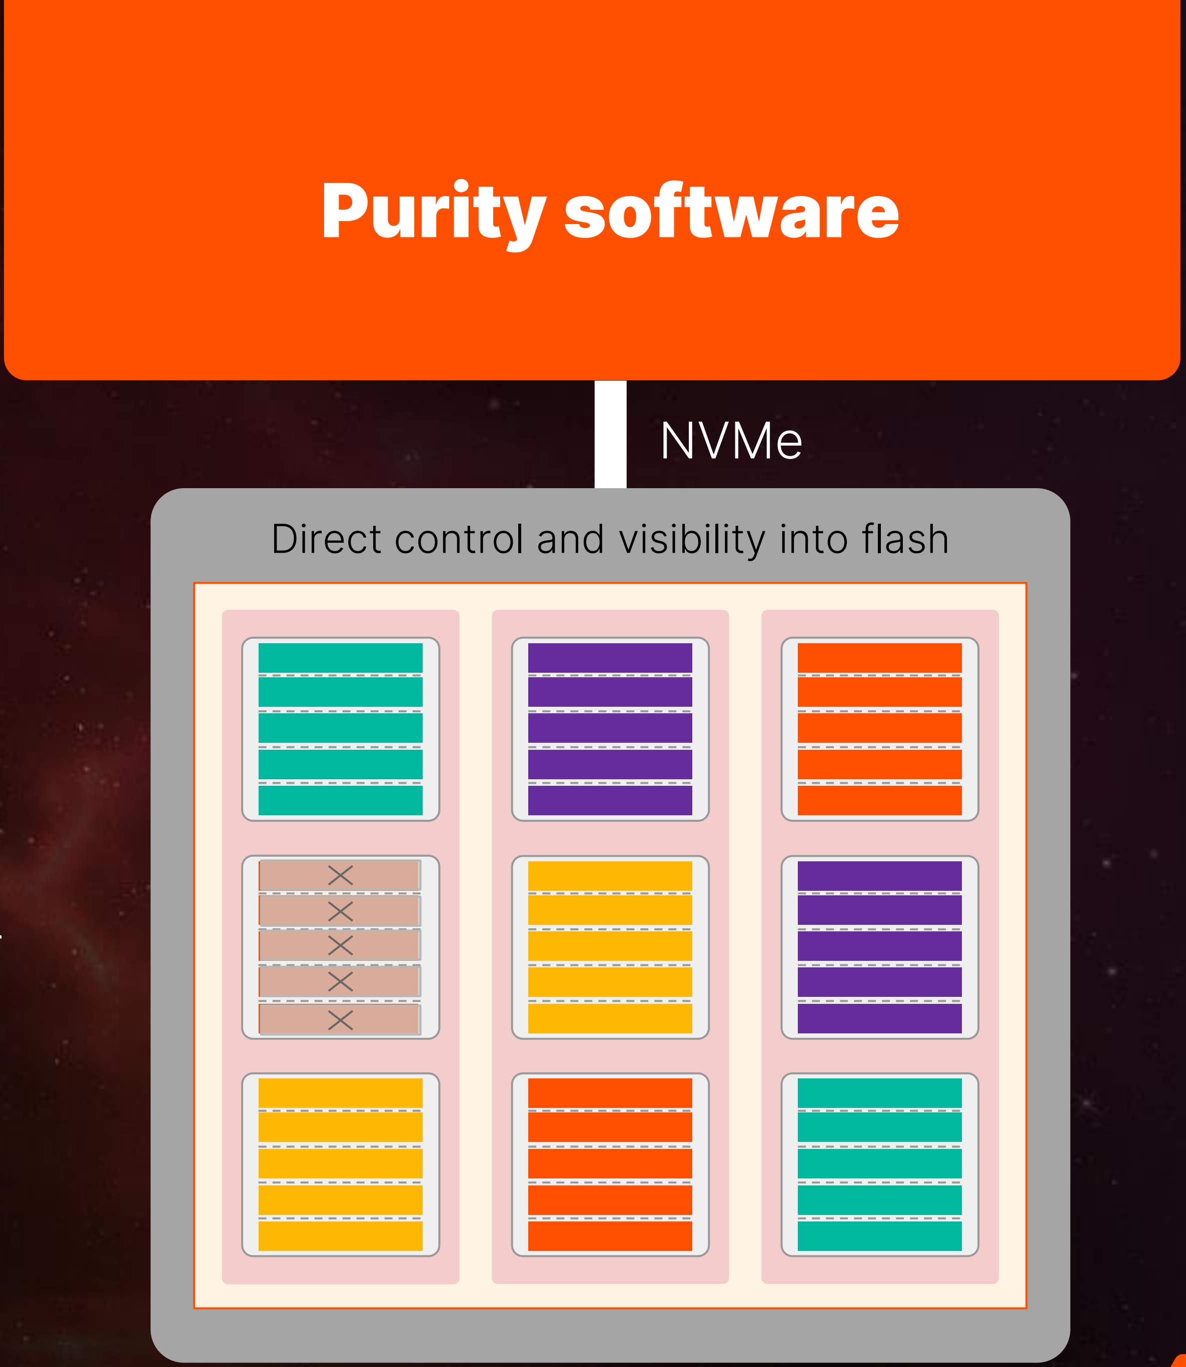

©2022 Pure Storage, Inc.

DirectFlash Reliability: Proven at Scale

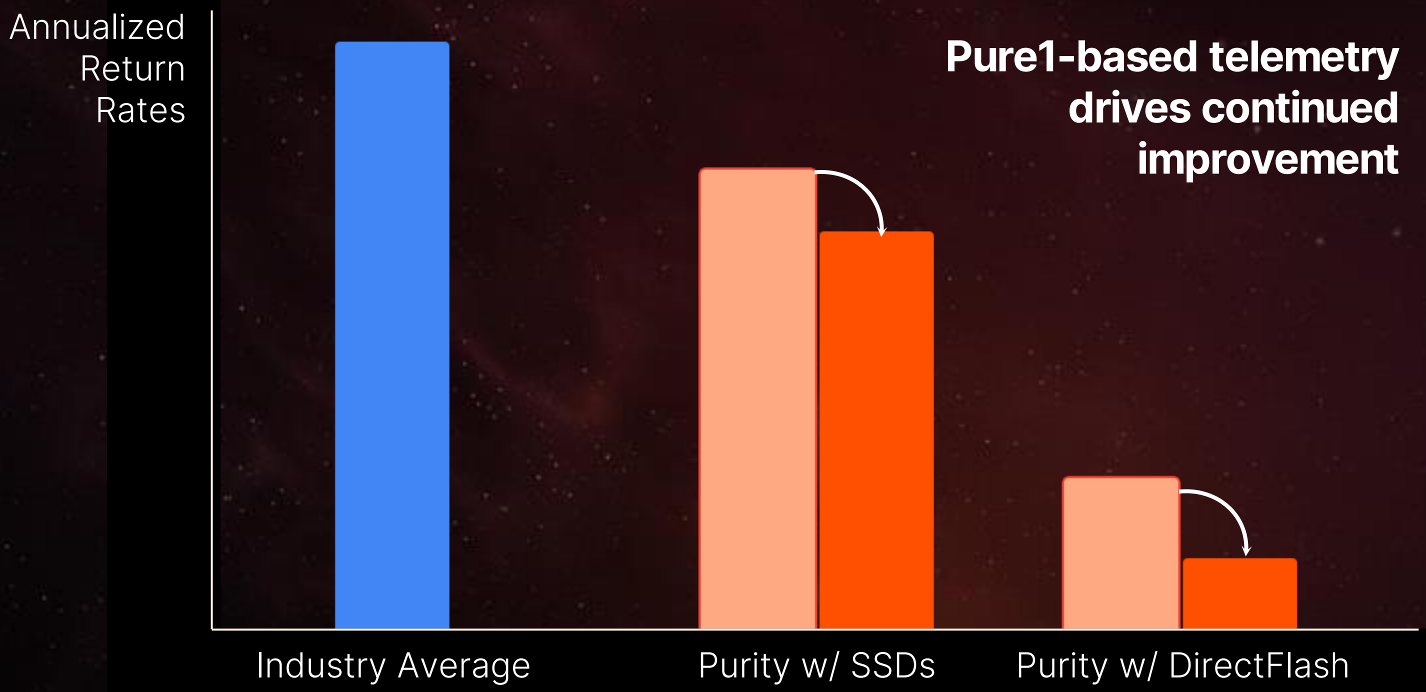

Pure1-based telemetry drives continued improvement

DirectFlash improves reliability ~3x over flash-optimized software on SSDs

©2022 Pure Storage, Inc.

DirectFlash improves performance

2

Long program/erase times
=> control tail latencies

~3x lower write amplification
=> 3x fewer write operations

DirectFlash also provides granular controls over data placement and scheduling, enabling Purity to mitigate the impacts of long program/erase times.

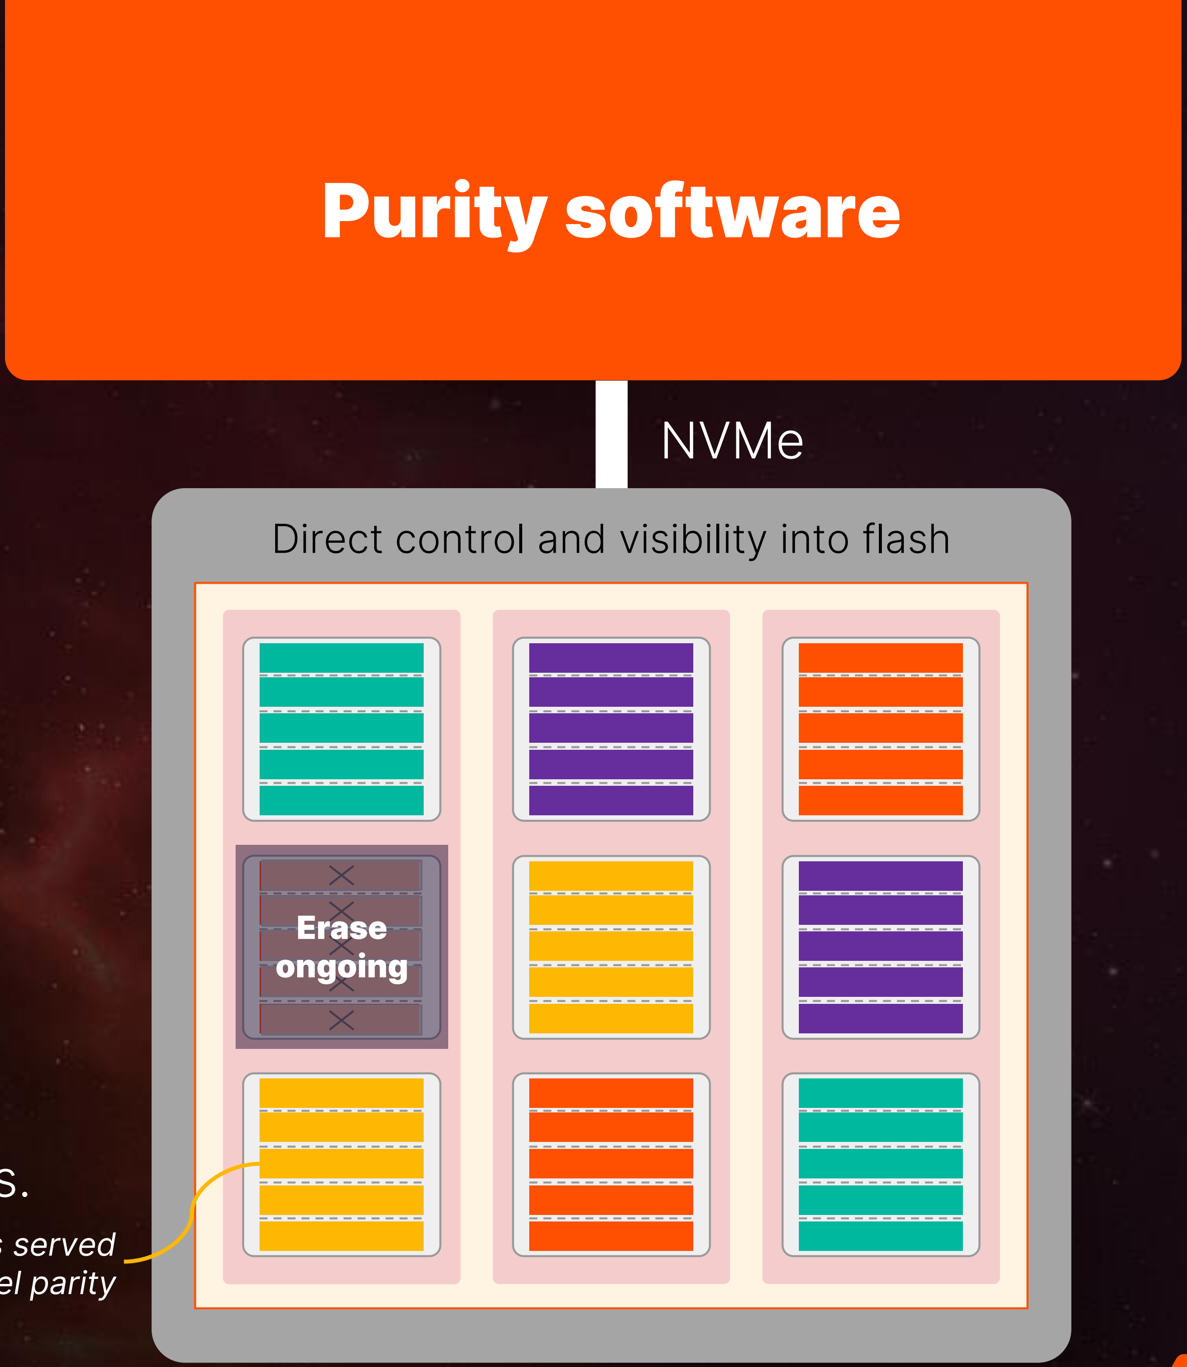

Concurrent reads served from system-level parity

©2022 Pure Storage, Inc.

DirectFlash improves efficiency

3

Media overheads
=> create efficient end-to-end mappings

DirectFlash eliminates unnecessary and inefficient drive-level mappings, enabling DRAM to be sized proportional to performance and removing flash over-provisioning from the DFM.

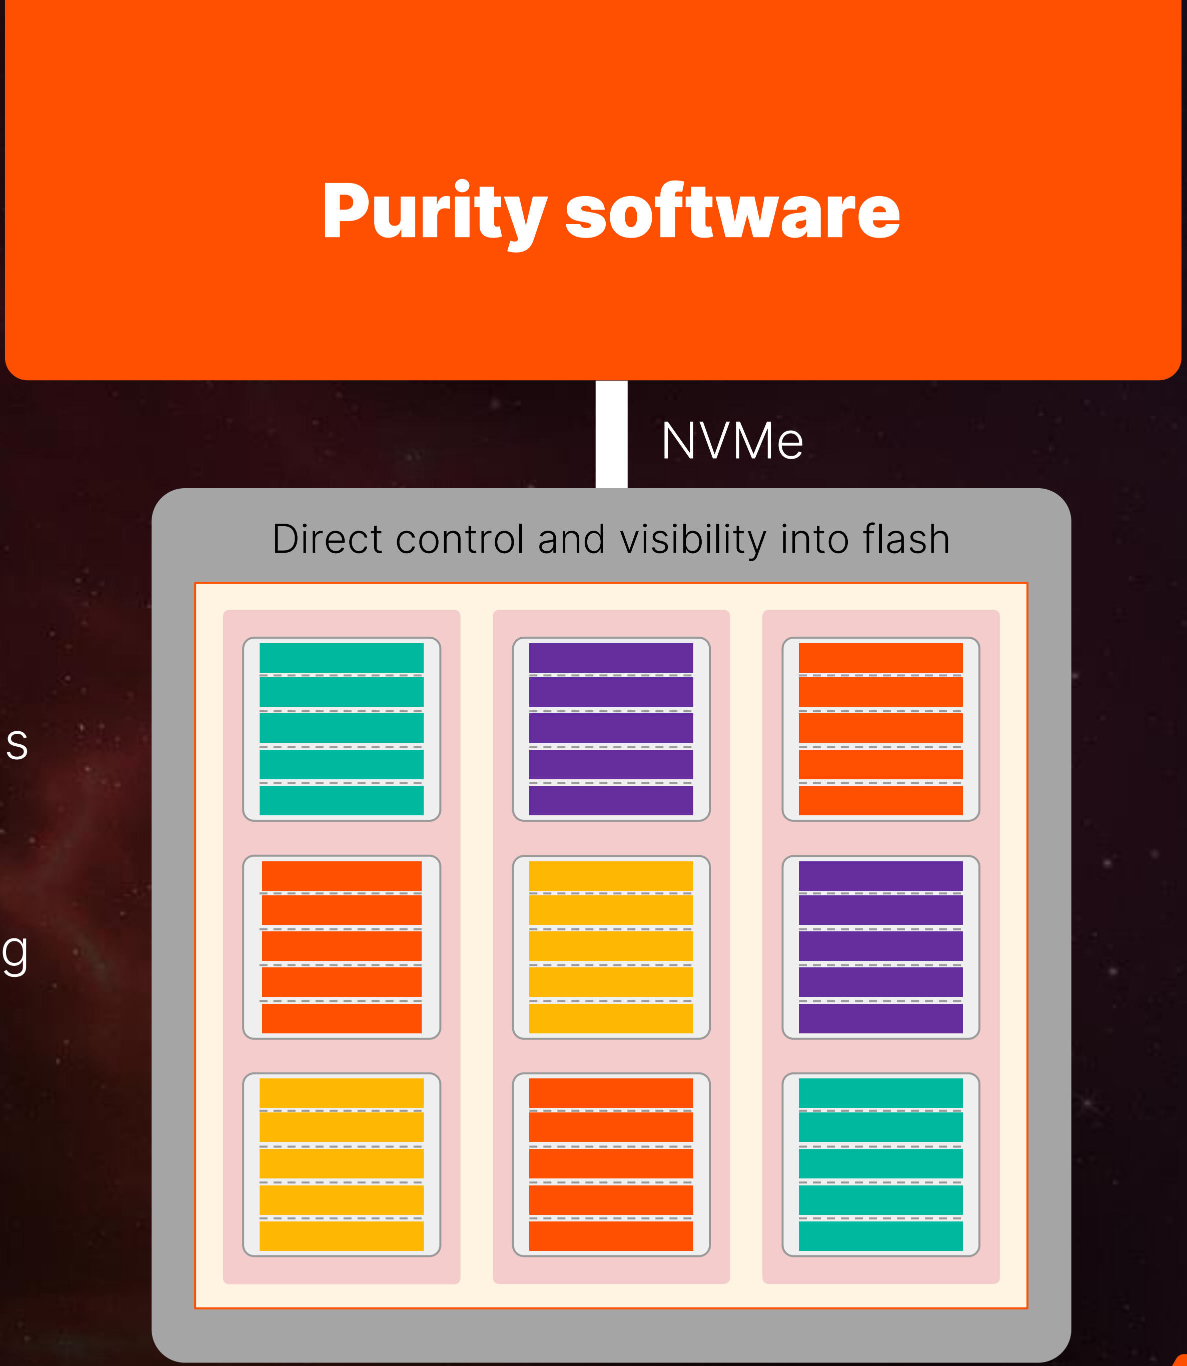

©2022 Pure Storage, Inc.

DirectFlash Efficiency Unlocks Scale

DirectFlash Modules (DFMs) require **40x less module-level DRAM**, while unlocking **20% more usable flash** as compared to SSD-based designs.

We are scaling DFM capacity with NAND fabrication technology, enabling new design options for all-flash storage systems.

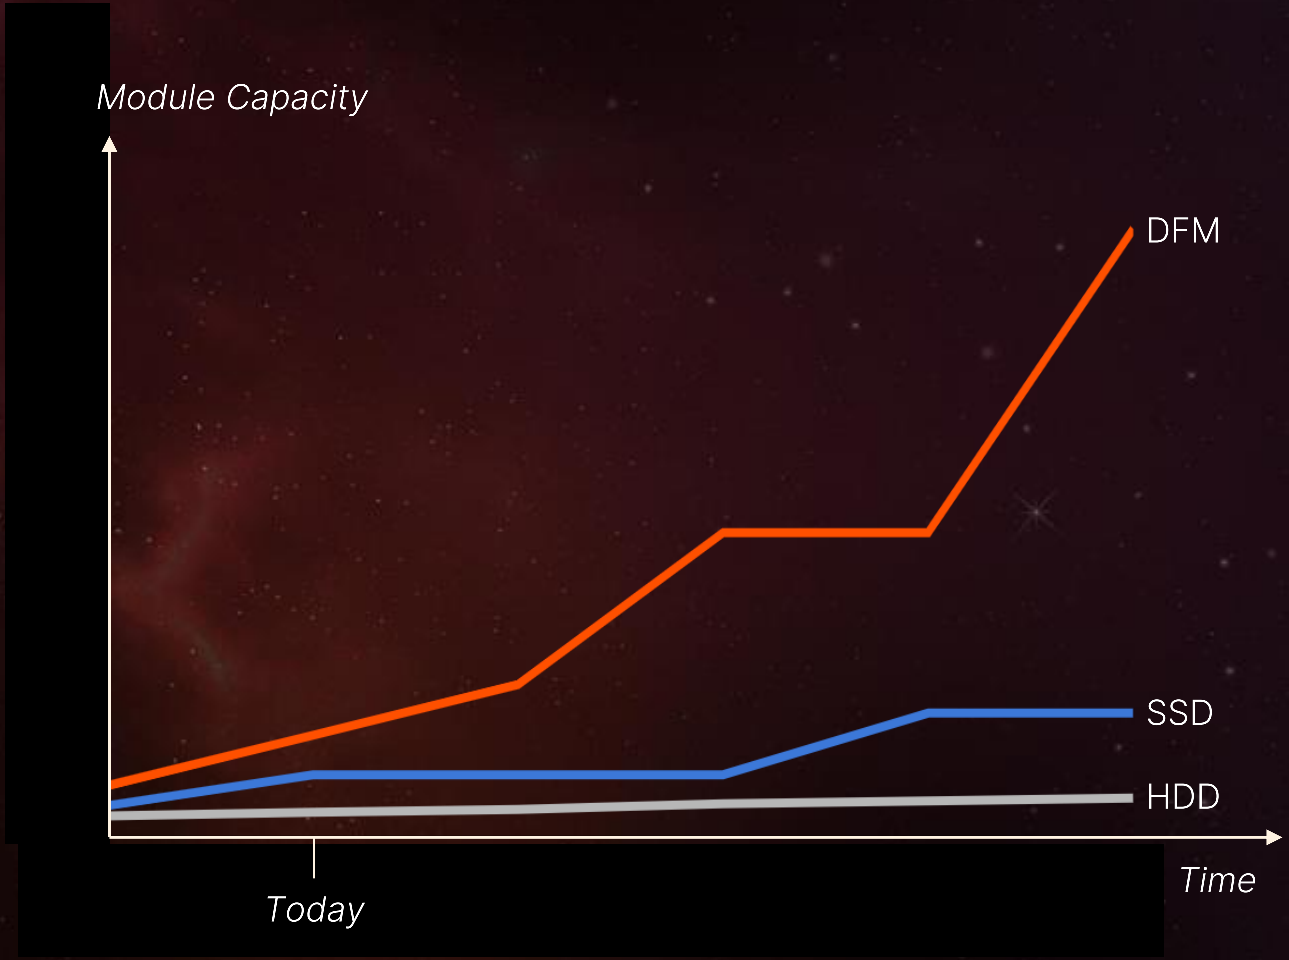

©2022 Pure Storage, Inc.

DirectFlash
MODULE

THE WORLD'S FIRST

SOFTWARE-DEFINED FLASH MODULE

GLOBALLY SOFTWARE-DEFINED

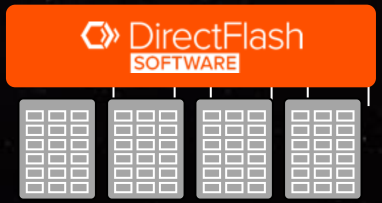

RELIABILITY

3-4X Lower ARR

NO HIDDEN FLASH NO WASTED DRAM

100% VISIBLE TO ARRAY

DETERMINISTIC LATENCY

ULTRA DENSE

4.4T TLC | 9T TLC | 18T TLC | 48T QLC | Larger DFMs Coming

©2022 Pure Storage, Inc.

DirectFlash is Better Science

AND WE HAVE A 10-YEAR HEAD START

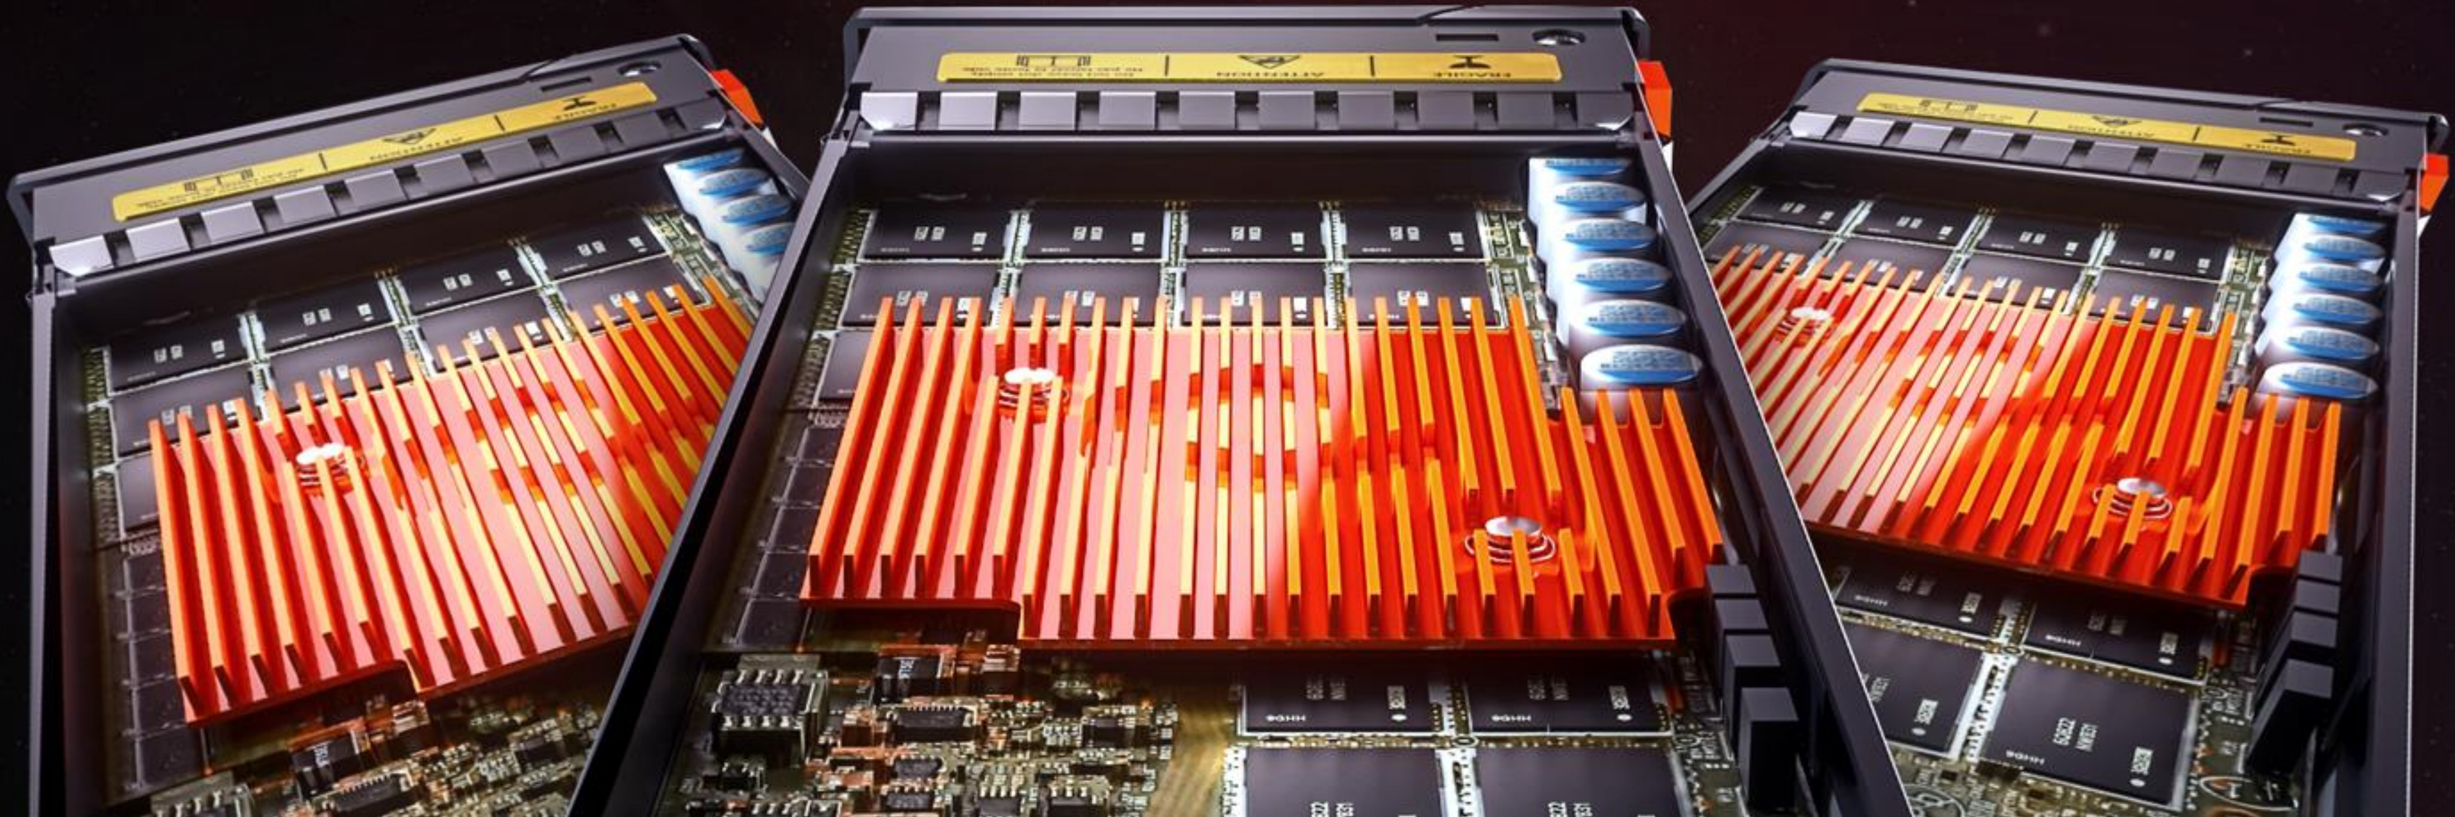

PURESTORAGE

Thank You!

홈페이지 : **www.purestorage.com/kr**

페이스북 : **www.facebook.com/purestoragekorea**

블로그 : **www.blog.naver.com/purestorage_korea**

유튜브 : **www.youtube.com/@PureStoragekr**

Uncomplicate Data Storage, Forever