



Flash Memory Summit

Better Science with DirectFlash


Justin Emerson

Principal Technology Evangelist - Pure Storage




Flash is at an Inflection Point


How we wish things worked:




Except it's more like:



But also changes over time:



And degrades with usage:



THE REALITIES OF NAND

Every generation gets:

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

How do systems solve these challenges?



The Problems with QLC

HIGHER DENSITY ALSO MEANS LOWER PERFORMANCE AND ENDURANCE

Media Performance (lower write latency = higher performance)			Media endurance (# P/E cycles)	
~0.5 ms	High performance can be utilized in specific areas, but not for bulk data	SLC	>100k	High endurance necessary for extreme applications (e.g. aerospace)
~1.2 ms	Latency is great, but no longer cost competitive	MLC	10k	Easy to manage endurance, but no longer cost competitive
~2.4 ms	Latency is manageable for performant systems	TLC	3k	Low endurance tolerable with careful system design
10-20 ms	Write latency is very high, requires optimization	QLC	<1k	Endurance too low to be manageable for most Enterprise storage systems

Pure's solution to reliability, performance, and efficiency: DirectFlash

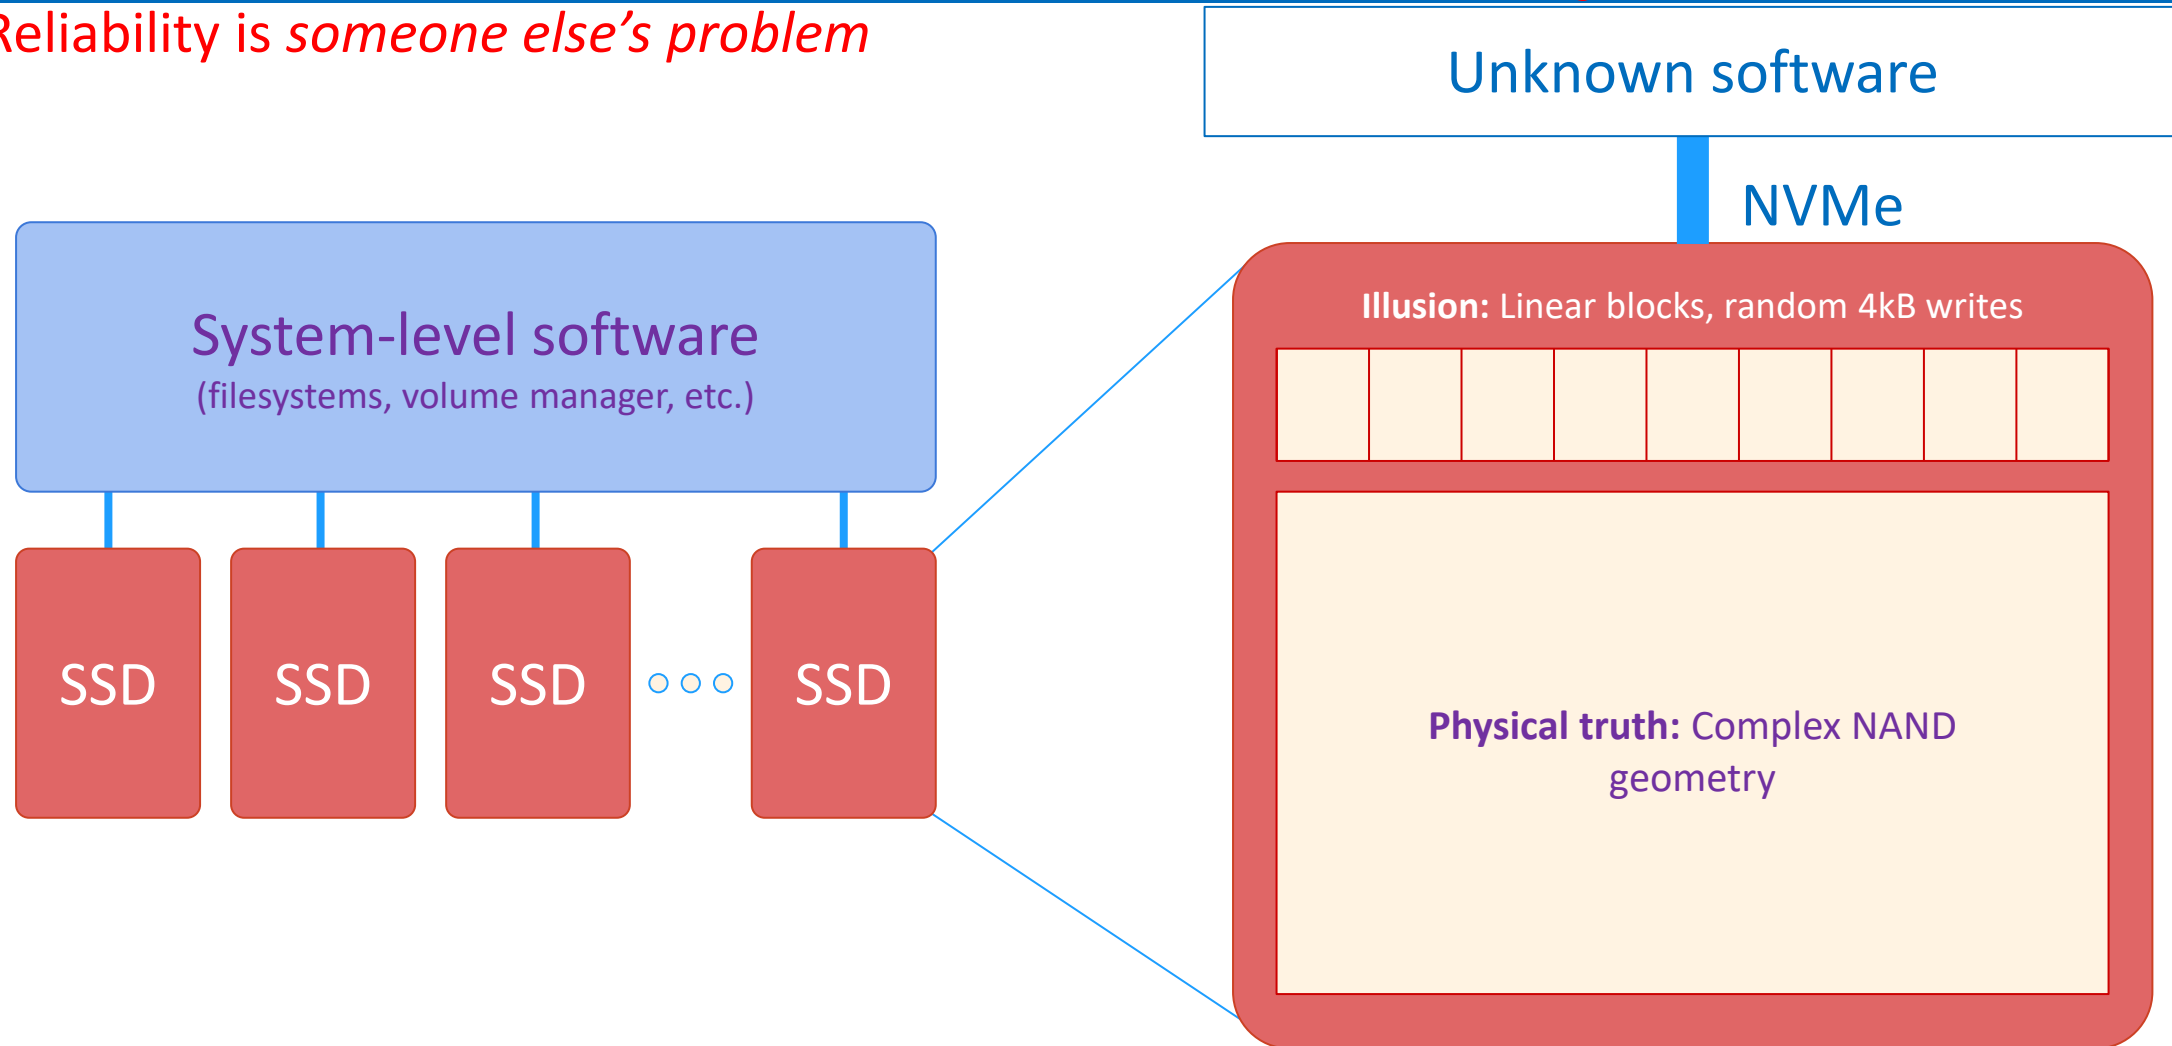
QLC Problems to Solve in All Flash Systems:

- 1 **Low endurance**
=> minimize write amplification
- 2 **Long program/erase times**
=> control tail latencies
- 3 **Media overheads**
=> create efficient end-to-end mappings



Should we treat SSDs as a commodity?

Reliability is *someone else's* problem





Should we treat SSDs as a commodity?

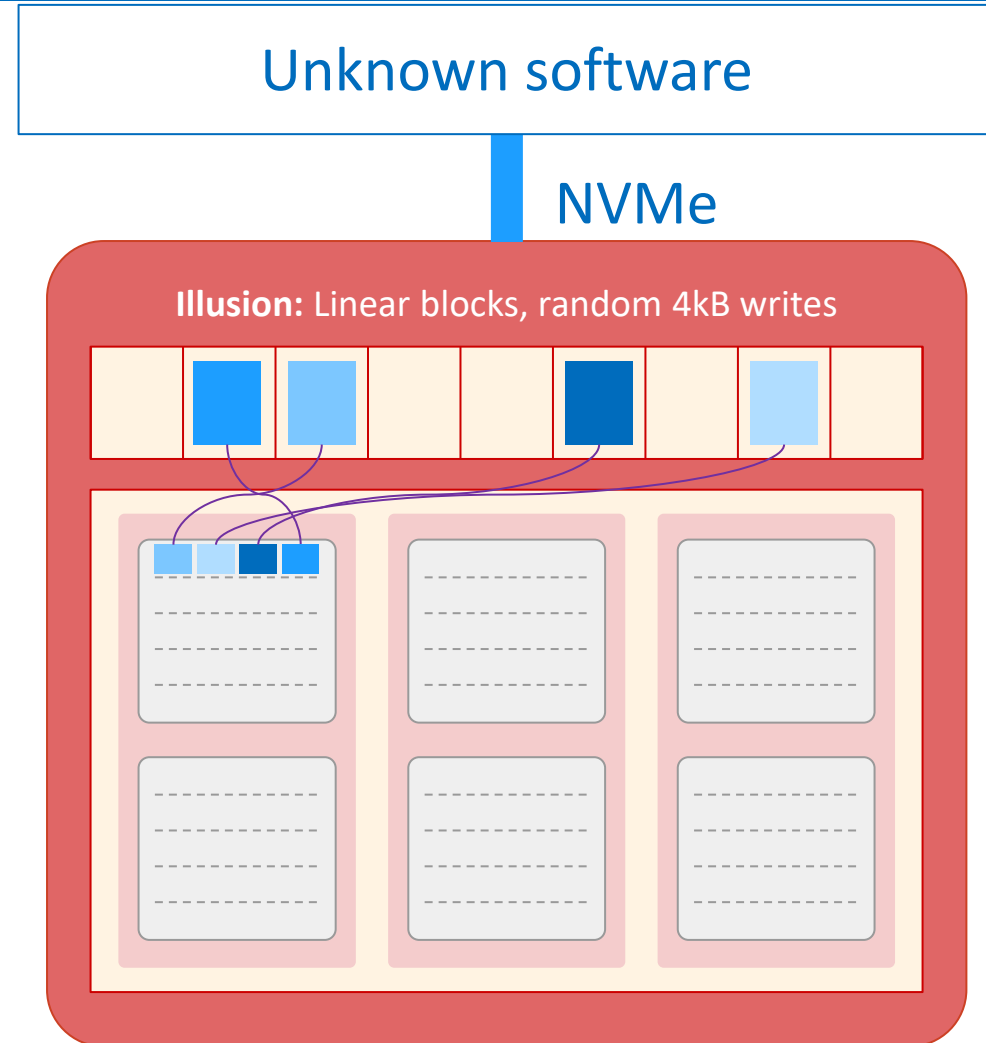
Reliability is *someone else's problem*

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.





Should we treat SSDs as a commodity?

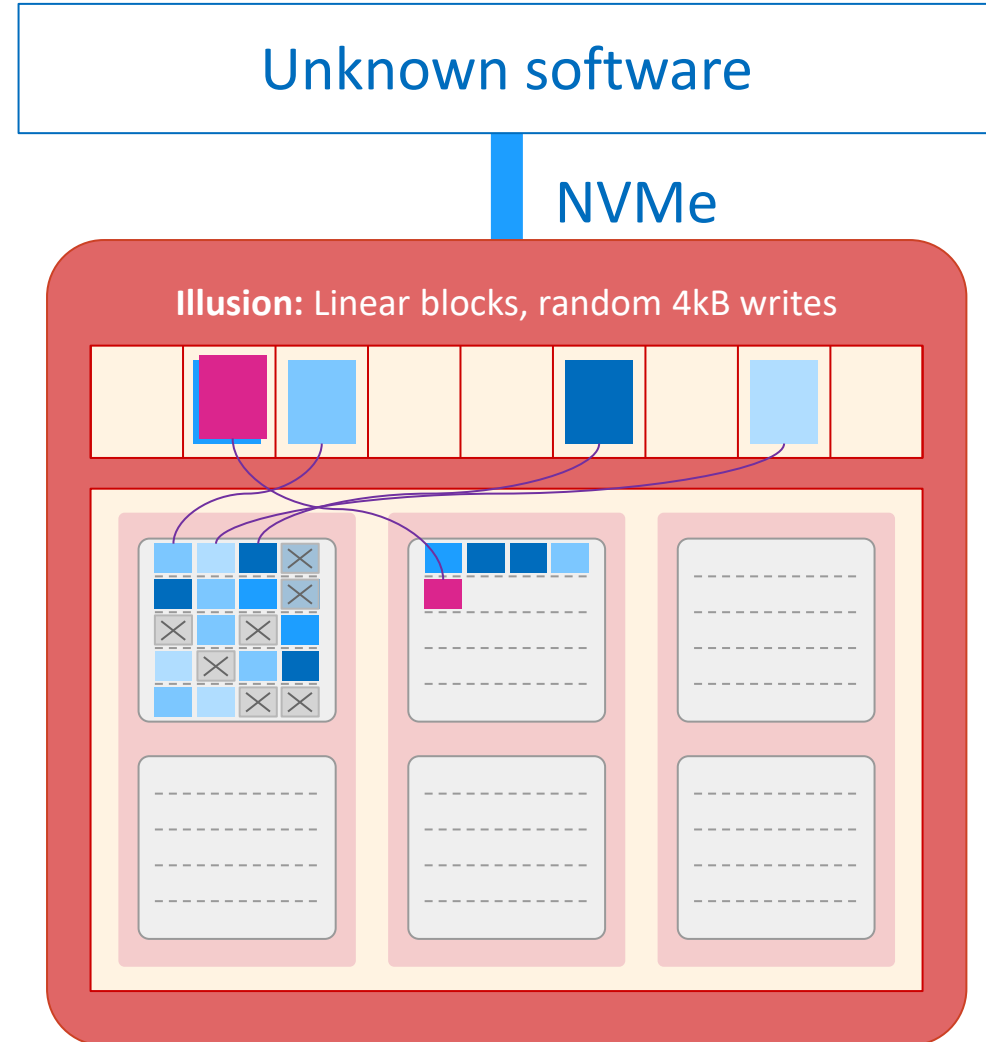
Reliability is *someone else's problem*

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.





Should we treat SSDs as a commodity?

Reliability is *someone else's problem*

1

Low endurance

=> minimize write amplification

2

Long program/erase times

=> control tail latencies

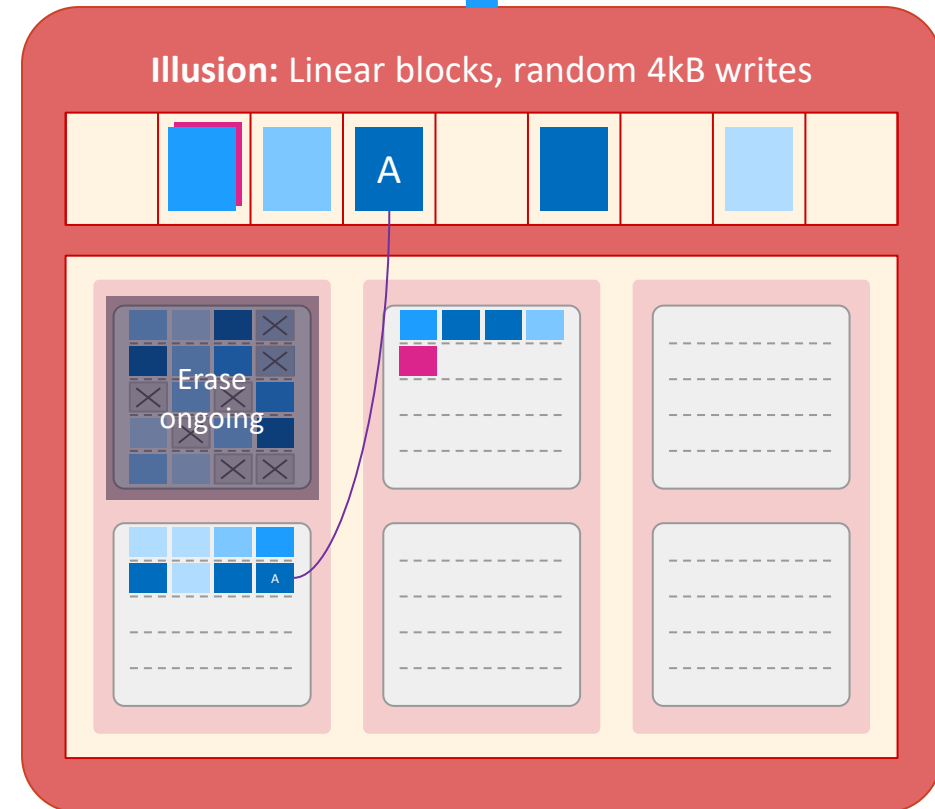
Applications have **no visibility** into the placement of blocks on physical media, and reads may be stuck behind slow (10s milliseconds)

operations to conflicting die.

Unknown software

NVMe

Illusion: Linear blocks, random 4kB writes



Read of A stalled behind erase in same die



Should we treat SSDs as a commodity?

Reliability is *someone else's problem*

1

Low endurance

=> minimize write amplification

2

Long program/erase times

=> control tail latencies

3

Media overheads

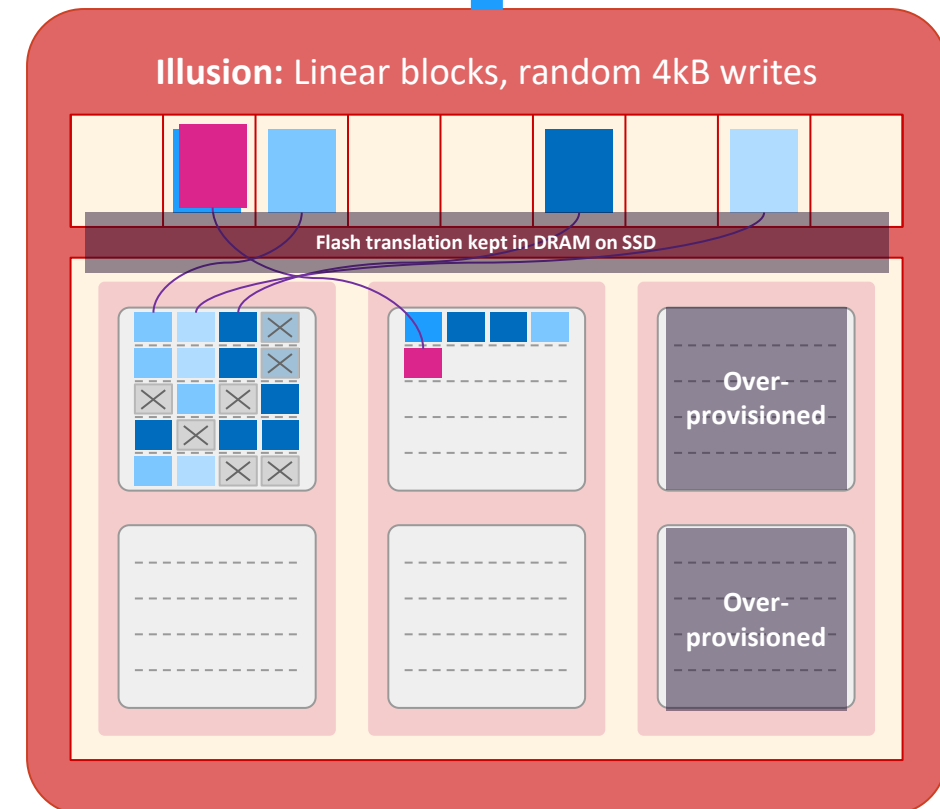
=> create efficient end-to-end mappings

Supporting random access is expensive.

A 10PB (raw) system has 10TB of SSD

Unknown software

NVMe





Should we treat SSDs as a commodity?

Reliability 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

Unknown software

NVMe



Modern SSDs are engineering marvels

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

DirectFlash extends flash lifetime

1

Low endurance


=> minimize write amplification

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

Purity software

NVMe

Direct physical control and visibility into flash



DirectFlash improves Performance

2

Long program/erase times

=> control tail latencies

~3x lower write amplification

=> 3x fewer write operations

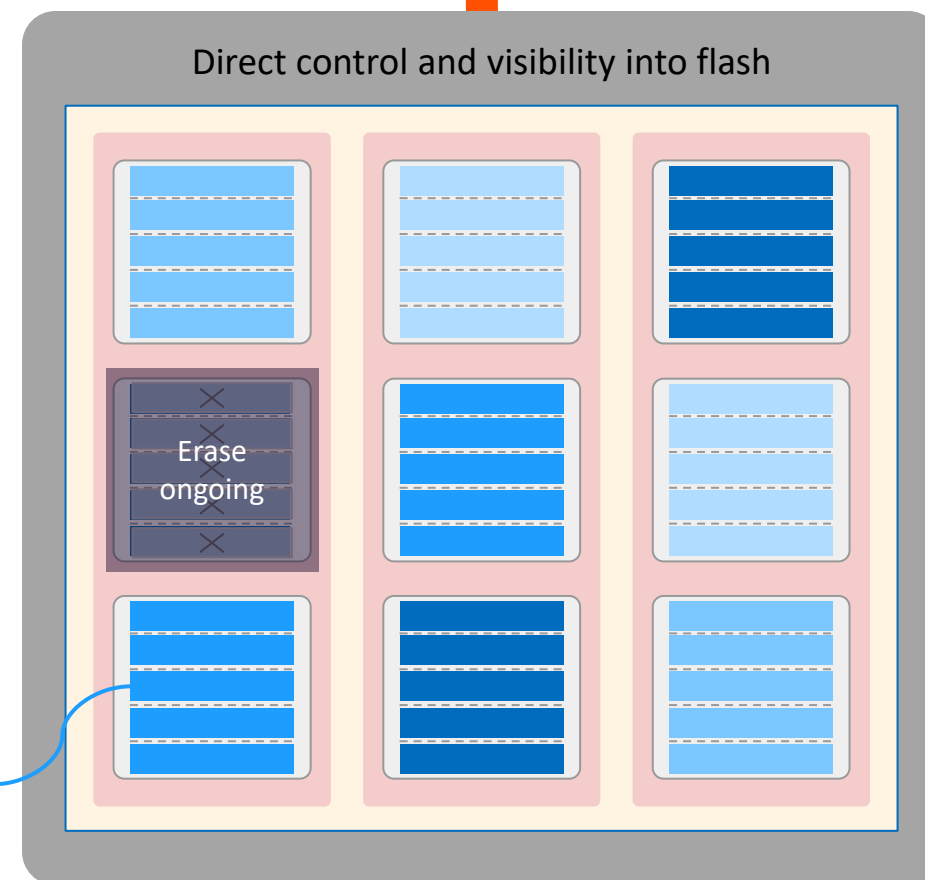
DirectFlash also provides granular controls over data placement and scheduling, enabling Purity to mitigate the impacts of long program/erase times

Concurrent reads served from system-level parity

Purity software

NVMe

Direct control and visibility into flash



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.

Purity software

NVMe


Direct control and visibility into flash



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. As a result, our modules use significantly less power than commodity alternatives.

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






DirectFlash Reliability: Proven at Scale





DirectFlash Reliability: Proven at Scale



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




Flash Memory Summit

DirectFlash MODULES


//TLC

//QLC


THE WORLD'S FIRST SOFTWARE-DEFINED FLASH MODULE




GLOBALY
SOFTWARE-DEFINED




RELIABILITY
3-4X Lower ARR



NO
HIDDEN
FLASH



DETERMINISTIC
LATENCY




ULTRA
DENSE





DirectFlash is Better Science

Shipping product at scale since 2017



DirectFlash powers Pure's all-flash products



FlashArray //C



FlashArray //X



FlashArray //XL



NEW FlashBlade //S

Your Data is the System - not the Hardware



Deploy Once, Upgrade In Place



No Planned Downtime



No Data Migration



Full Performance, Non-Stop



Modernize Hardware and Software



Multi-Generational Lifespan



Investments Protected

Come talk to us if you're...



Flash Memory Summit

A Hyperscaler looking for a different approach to cloud storage...

An Enterprise looking for fast, reliable storage with industry-leading efficiency...

A Developer looking for highly scalable, programmable storage for your applications...

Uncomplicate Data Storage, Forever



Flash Memory Summit

