

Western Digital[®]

Optimizing SSDs for Multiple Tenancy Use

Liam Parker

*Senior Technologist, SSD Systems
Engineering*

August 7, 2018



Flash Memory Summit



Forward-Looking Statements

Safe Harbor | Disclaimers

This presentation contains certain forward-looking statements that involve risks and uncertainties, including, but not limited to, statements regarding solid-state storage technology and product development efforts, business strategy, growth opportunities, market positioning, data growth and its drivers, demand for digital storage and market trends. Forward-looking statements should not be read as a guarantee of future performance or results, and will not necessarily be accurate indications of the times at, or by, which such performance or results will be achieved, if at all. Forward-looking statements are subject to risks and uncertainties that could cause actual performance or results to differ materially from those expressed in or suggested by the forward-looking statements.

Key risks and uncertainties include volatility in global economic conditions, actions by competitors, business conditions, growth in our markets, pricing trends and fluctuations in average selling prices, and other risks and uncertainties listed in our filings with the Securities and Exchange Commission (the "SEC") and available on the SEC's website at www.sec.gov, including our most recently filed periodic report, to which your attention is directed. We do not undertake any obligation to publicly update or revise any forward-looking statement, whether as a result of new information, future developments or otherwise, except as required by law.



Overview

- Is physical data separation everything?
 - NVM Sets : A write to one set does not impact a read from another.
- Can we improve performance & endurance of multi-tenancy SSDs without physical data separation?
- Show how we can improve random write speeds for multi-tenancy variable active workloads by using a feature available in some Western Digital SSDs.

Multiple Tenancy Requirements

- SSDs used in cloud computing...
 - Multiple customers (tenants) per SSD.
- New protocol to support multiple tenants:
 - NVM sets, Namespaces.
- Customers want...
 - Consistency of service.
 - High Quality of Service (QoS).
- Cloud providers want...
 - High endurance & performance.
 - Consistency.
 - Lower costs (i.e. lower OP / TCO).
- Reality is...
 - Tenants can change behavior over time.
 - Makes consistency and lower cost harder to achieve.



- Using NVMe, tenants' data can be separated according to namespace/stream.


Logical & Physical Separation

- Logical / block separation : Drive is separated into distinct partitions each used by a different tenant.
 - Pros:
 - Simple
 - Cons:
 - Potential for tenant’s data mixing within die.
 - Lower QoS due to contention between tenants.
- Physical die separation :
 - Pros:
 - No die contention between tenants. Better QoS.
 - Cons:
 - Lower maximum per tenant bandwidth.
 - Die failure has higher cost.
 - Cannot use unutilized resources for other tenants.

Example Random write workload on 7% OP drive:

Tenant	Write Rate	Logical Range
A	Ra	La
B	Rb < Ra	Lb > La
C	Rc < Rb	Lc > Lb

Logical / Block Separation	Physical Die Separation
OP shared across drive. Tenants written to different NAND blocks. Tenants share die.	OP different (tuned) for each tenant. Tenants data written to different die.



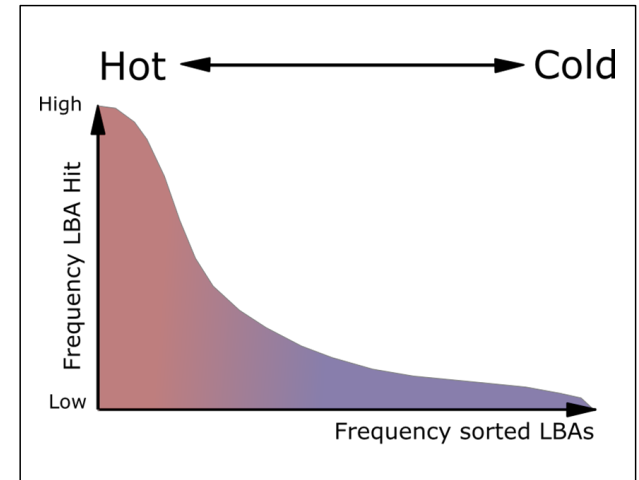
Fix It Automatically in the SSD...

- Tenant's with different write patterns create a non-uniform workload.
- But all tenants have the same OP for logical separation on a standard SSD.
- Physical separation have pre-determined OP per tenant that cannot change.

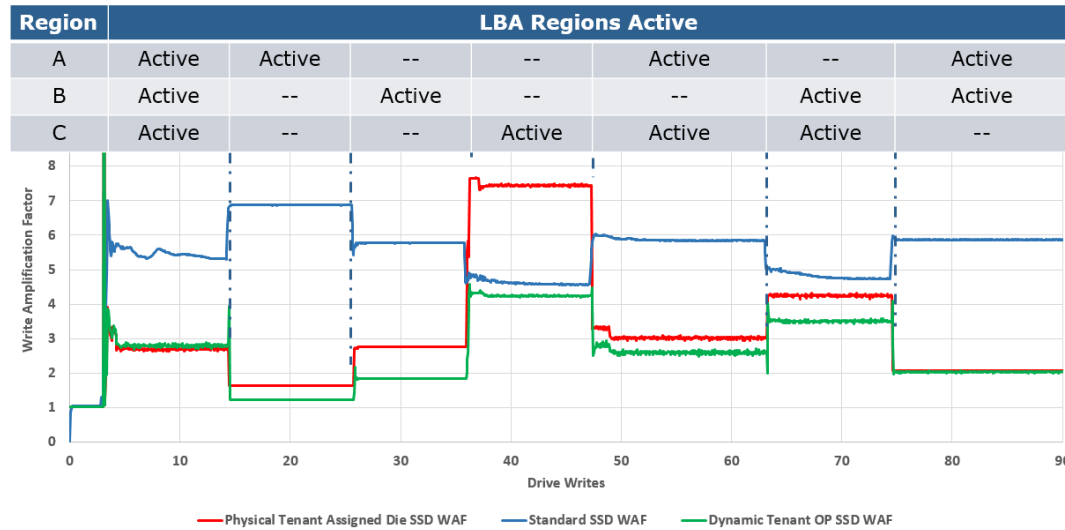
- Is there another way to approach the problem?
- What if the SSD could always give the OP to where it was needed?

Non-Uniform Workload in an SSD

- Write Amplification (WA) is a function of Over-Provisioning (OP).
- Incorrectly apportioned OP:
 - Higher Write Amplification (worse than uniform random).
 - Lower Performance.
 - Poorer Quality of Service.
 - Lower Endurance.
- Track “temperature” of all logical pages.
- Western Digital Guardian Technology Platform dynamically calculates optimal OP and assigns that OP across the drive.
- The write amplification is reduced and performance increased.



What can be achieved?



	Endurance Benefit (PEs)						
Physical	2x	4.2x	2x	0.6x	1.94x	1.19x	2.84x
Dynamic OP	1.95x	5.6x	3x	1.08x	2.25x	1.35x	2.89x

Wrap Up : Guardian Technology™ Platform

- Off the shelf drive with Dynamic OP allocation better than physical separation as tenants alter behaviour over time.
- Wear levelling always maintained across all physical blocks.
- Endurance of SSD improved.
- TCO improved.
- Power reduced due to lower GC.
- Performance increased due to lower GC.
- QoS improved compared to standard SSD.
- Guardian Technology™ benefits any non-uniform workload, not just multi-tenancy use.
- Note: Guardian Technology™ Dynamic OP allocation is not yet available in all Western Digital SSDs.

The Western Digital logo is displayed in a large, bold, white sans-serif font. The background of the slide is a dark, abstract composition of numerous thin, overlapping lines in shades of orange, red, and purple, creating a sense of motion and depth. The logo is centered horizontally and positioned in the upper half of the slide.

Western Digital®

Western Digital, the Western Digital logo, and Guardian Technology are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. All other marks are the property of their respective owners.