

Data Storage Strategies

Wim De Wispelaere, VP Strategic Initiatives, Western Digital

Javier González, Principal Engineer, Samsung Electronics

- NAND, PCIe® and NVMe® have changed storage drastically
 - The storage bottleneck is moving from the device into higher storage layers; kernel, filesystem, application
- Better data placement coordination between applications and devices is required. This drives innovations at the interface level
 - NVMe: Streams, ZNS, KVSSD, (FDP)
- Storage architectures move to disaggregation and data processing in infrastructure or devices
 - Fabrics, xPU (DPU, IPU, TPU), Computational Storage
- Customer adoption requires standardization, software ecosystem and vendor alignment



D2PF overview

Industry Collaboration


D2PF: Data Placement Processing and Fabrics Foundation



Western Digital.

SAMSUNG

- Announced a Joint Collaboration focused on Data Placement, Processing and Fabrics on March 29th
- Initial focus will be Zoned Storage on 3 tracks




D2PF Mission Statement and Objective

To foster an **open** ecosystem around **advanced storage devices** so that their advantages can be utilized and **accessible for everyone.**

- SNIA:
 - alignment on use cases and device properties
- LF
 - Open-Source software components to facilitate adoption and integration
- Community
 - Open community for all interested
 - Join us at D2PF to broaden the community

Membership Makeup

Invitation Open to all!



Consumers	<ul style="list-style-type: none">• Users of fully baked solutions• Seeking Tutorials
Tinkerers	<ul style="list-style-type: none">• Power Users of integrated solutions• File and fix bugs/edge-cases• Benchmark creators/verifiers
Builders	<ul style="list-style-type: none">• Integrators with keystone projects• Commit features and improvements to keystone projects• Add Value to their projects
Founding Members	<ul style="list-style-type: none">• Core Committers to keystone projects• Real world at-scale use-cases



Example Solution: Data Placement

Example solution: Zoned Namespaces

Reduce WAF and OP

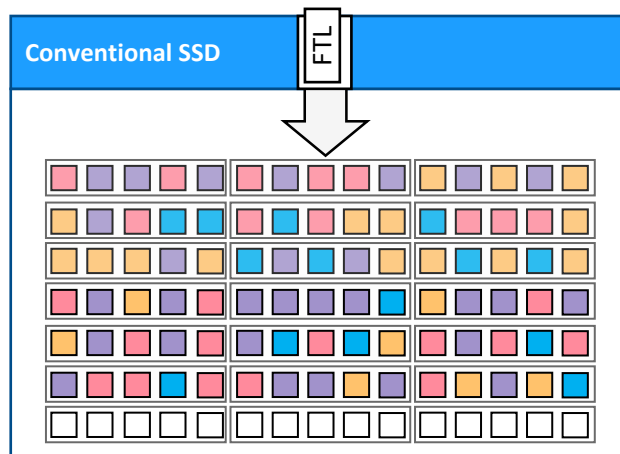
Remove device GC

Higher density


Higher endurance

Improved QoS

**ZNS Enables Faster
QLC Adoption**







- ✓ Support for random write workload
 - Device-Managed data placement - may cause noisy neighbor
 - Includes extra NAND (OP) for flash
 - Implicit support for multi-tenancy
- ✗ Higher capacity requires more OP NAND and DRAM
- ✗ Flash management and host accesses affect QoS



- ✓ Focus on sequential workloads
 - Host-managed data placement
 - No OP and lower DRAM requirement as host manages device
 - Explicit support for multi-tenancy
- ✓ Higher drive density
- ✓ Sustain high performance and improve QoS

ZNS: Higher Density and Endurance



*Up to 25% higher
usable capacity*

Lower write amplification allows
ZNS SSD to expose most of the
RAW NAND as Usable Capacity

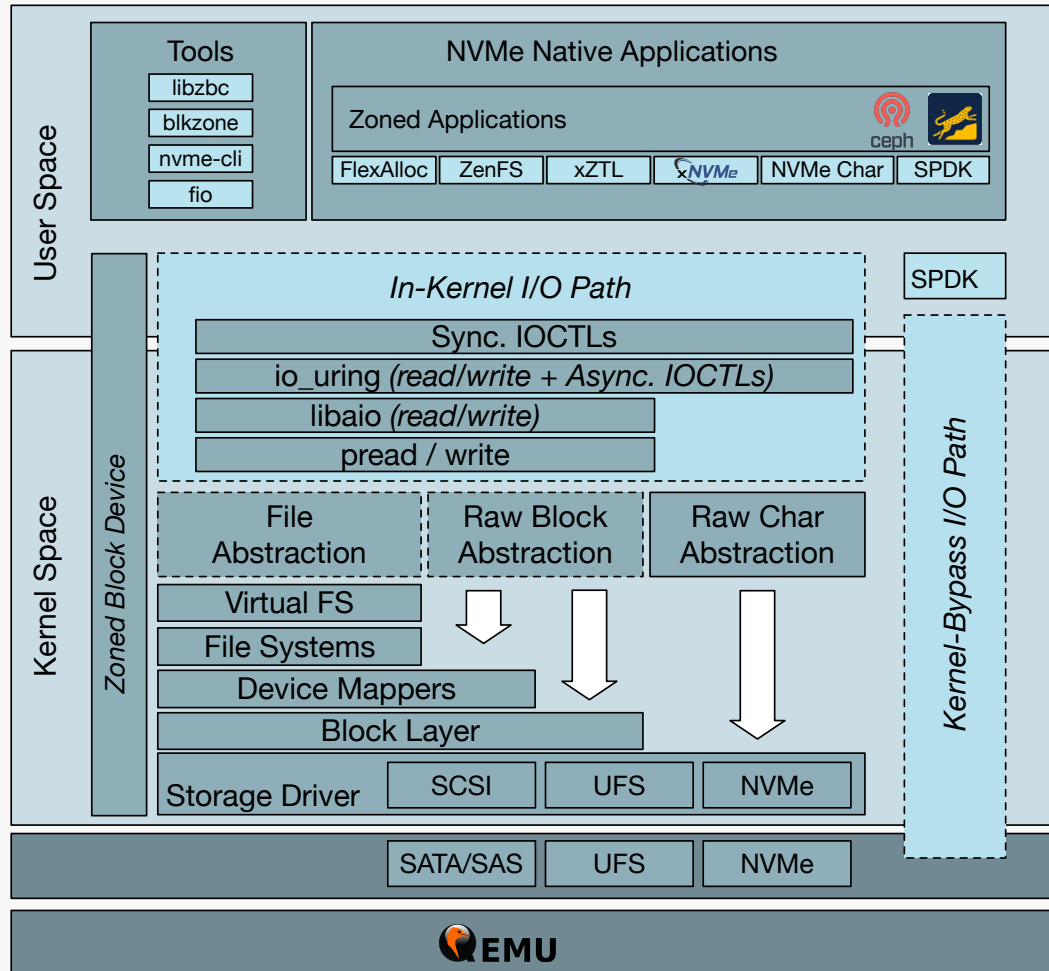
*Up to 4.37x higher
endurance*

Lower write amplification and
sequential writing leads to
higher endurance of a ZNS SSD

Example	0.8 DWPD	2 DWPD	3.5 DWPD
	7,680 GB	6,400 GB	8,192 GB
	3,840 GB	3,200 GB	4,096 GB
	1,920 GB	1,600 GB	2,048 GB
	960 GB	800 GB	1,024 GB

ZNS: Standards & Ecosystem


Linux[®] Zone Support



Standard Advancements

NVMe

- Core spec available
- Small TPs in progress



SNIA Zoned TWG

- Working on spec to align industry around zoned models
- Target specification this year



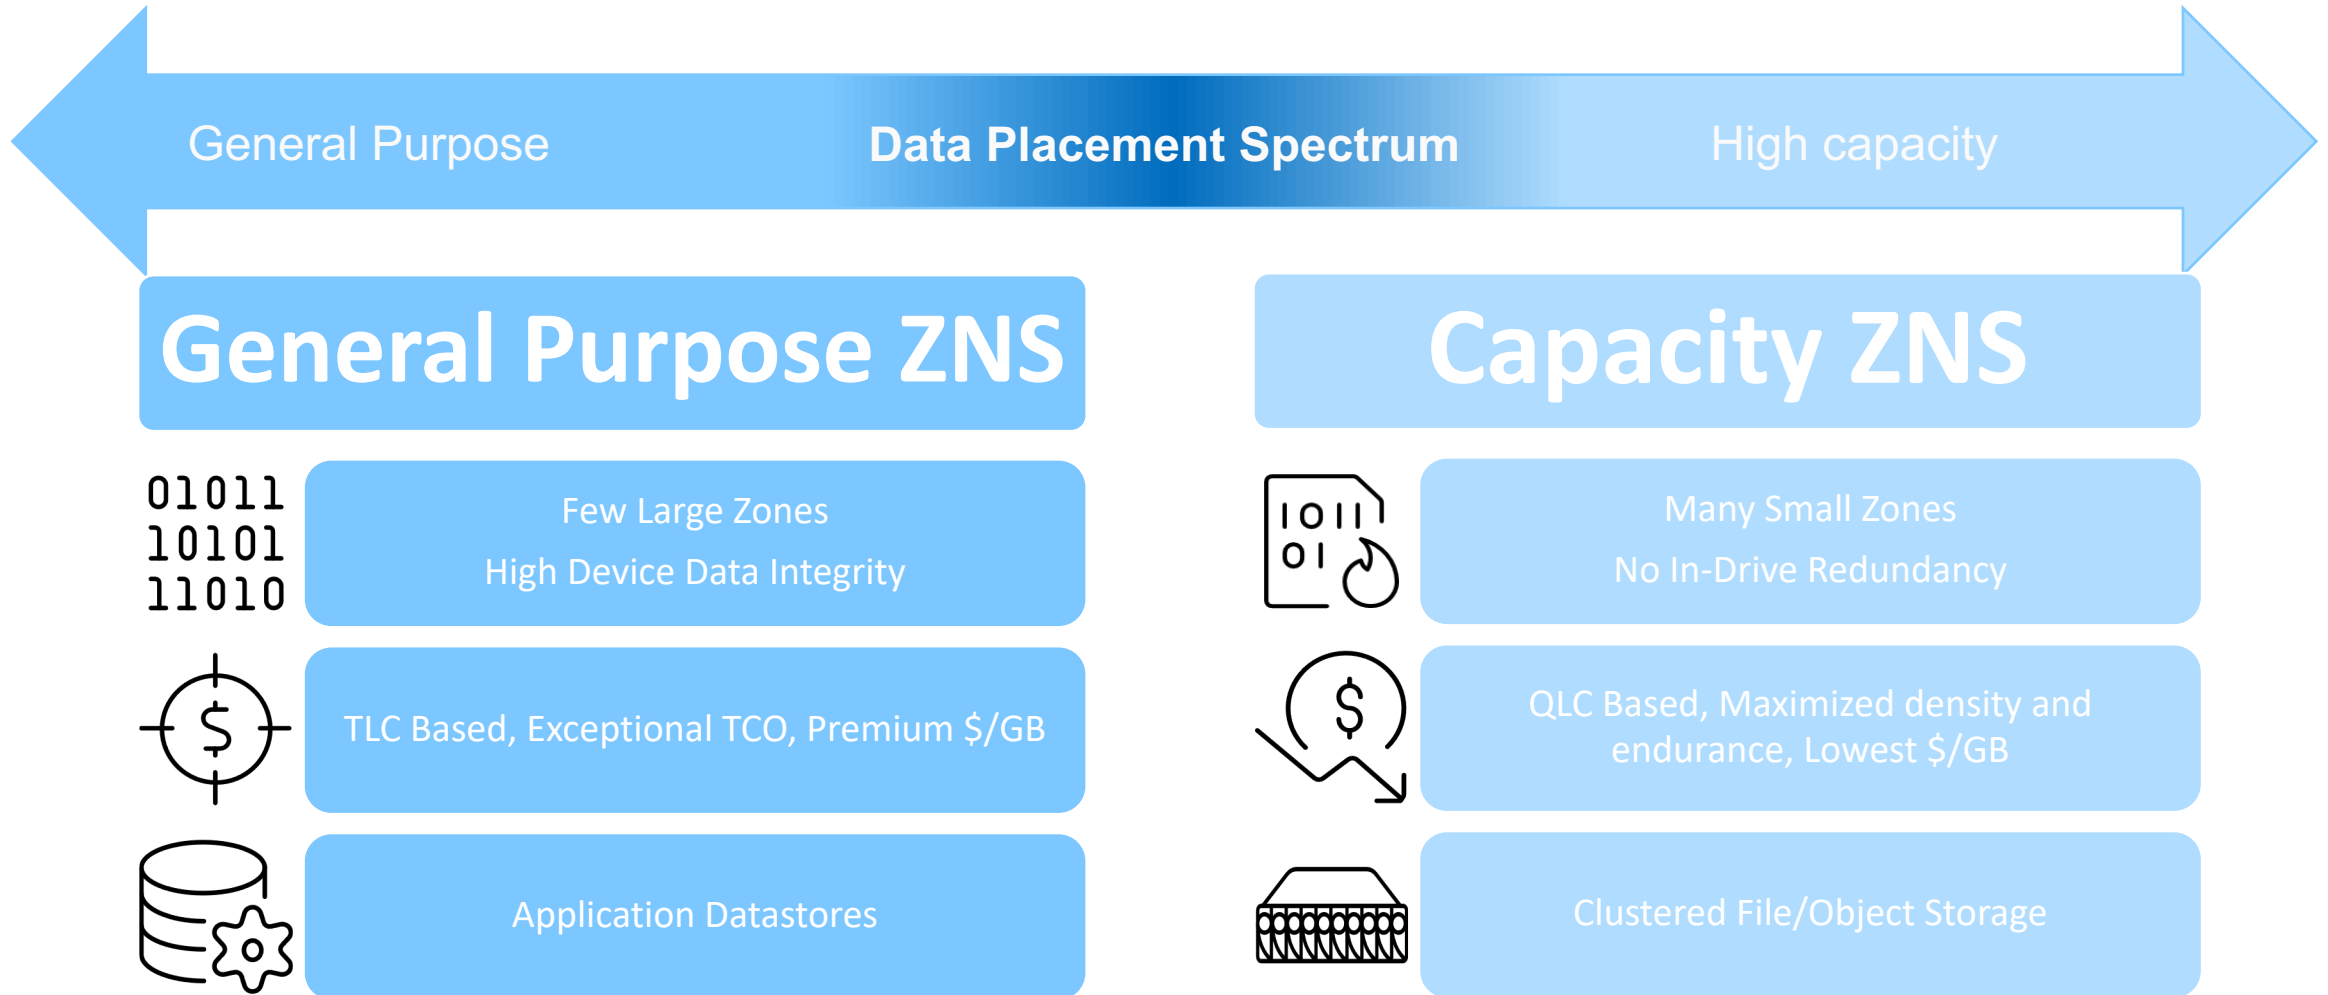
D2PF

- Linux Foundation Initiative to aid adoption of new interfaces



ZNS: Use-case specification alignment

Example: SNIA Zoned Storage TWG





D2PF Future Scope

Building blocks of the future

What is possible when we use a combination of advanced data technologies?

Data Placement
enabling technologies

Examples: ZNS, Ongoing
TPs

Higher density

Higher endurance

Improved QoS

Fabrics enabling
technologies

Examples: NVMe/TCP

Flexible Storage/CPU ratio

Smart NIC processing

Cloud native architecture

Computational Storage
enabling technologies

Examples: CS in SNIA,
Ongoint TPs

Native device throughput

High parallelism / scalability

CPU offload

Technology Showcasing and Enablement

Examples: SMRC



Conclusion

Conclusion



Flash Memory Summit

- New generation NAND storage devices require data placement coordination, fabrics and infrastructure processing to unfold their full potential
- Use case adoption requires standards, common device properties and open software stack support
- D2PF is being established to facilitate adoption on all levels.
- Get in touch with us:
 - wim.de.wispelaere@wdc.com
 - javier.gonz@samsung.com



Thank You

Western Digital and the Western Digital logo are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries.

Samsung and the Samsung logo are trademarks or registered trademarks of Samsung Electronics Co., Ltd.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

The NVMe word mark is a trademark of NVM Express, Inc.

PCIe is a registered trademark and/or service mark of PCI-SIG in the U.S. and/or other countries.

All other marks are the property of their respective owners.