

Data Scrubbing and Rebuild Using Offload

Aug 7, 2025

Devesh Rai

KIOXIA America, Inc.



the Future of Memory and Storage

Host Orchestrated Compute Offload Building Blocks

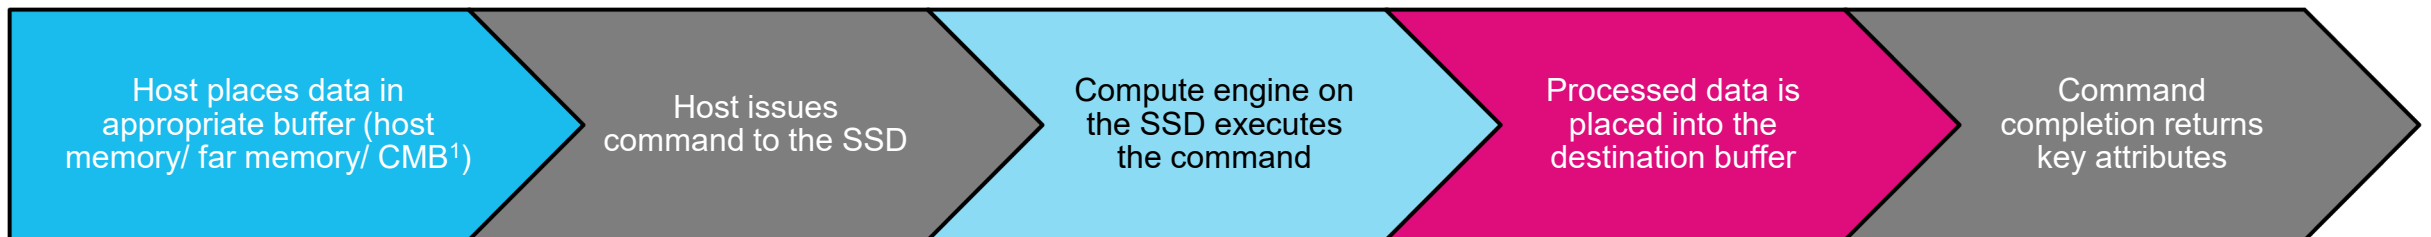
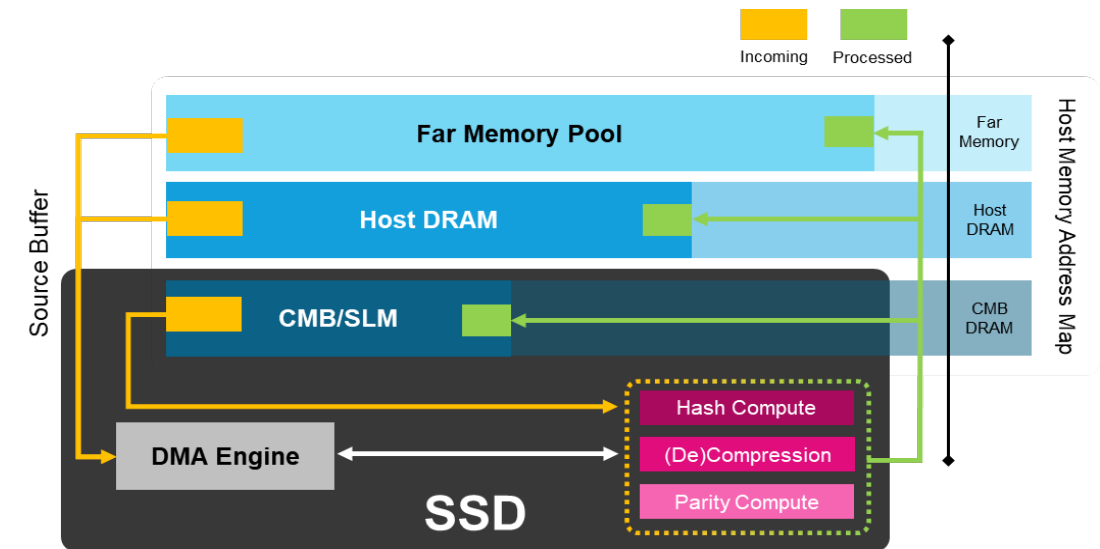
Power-efficient
compute
engines

DRAM
bandwidth
saving

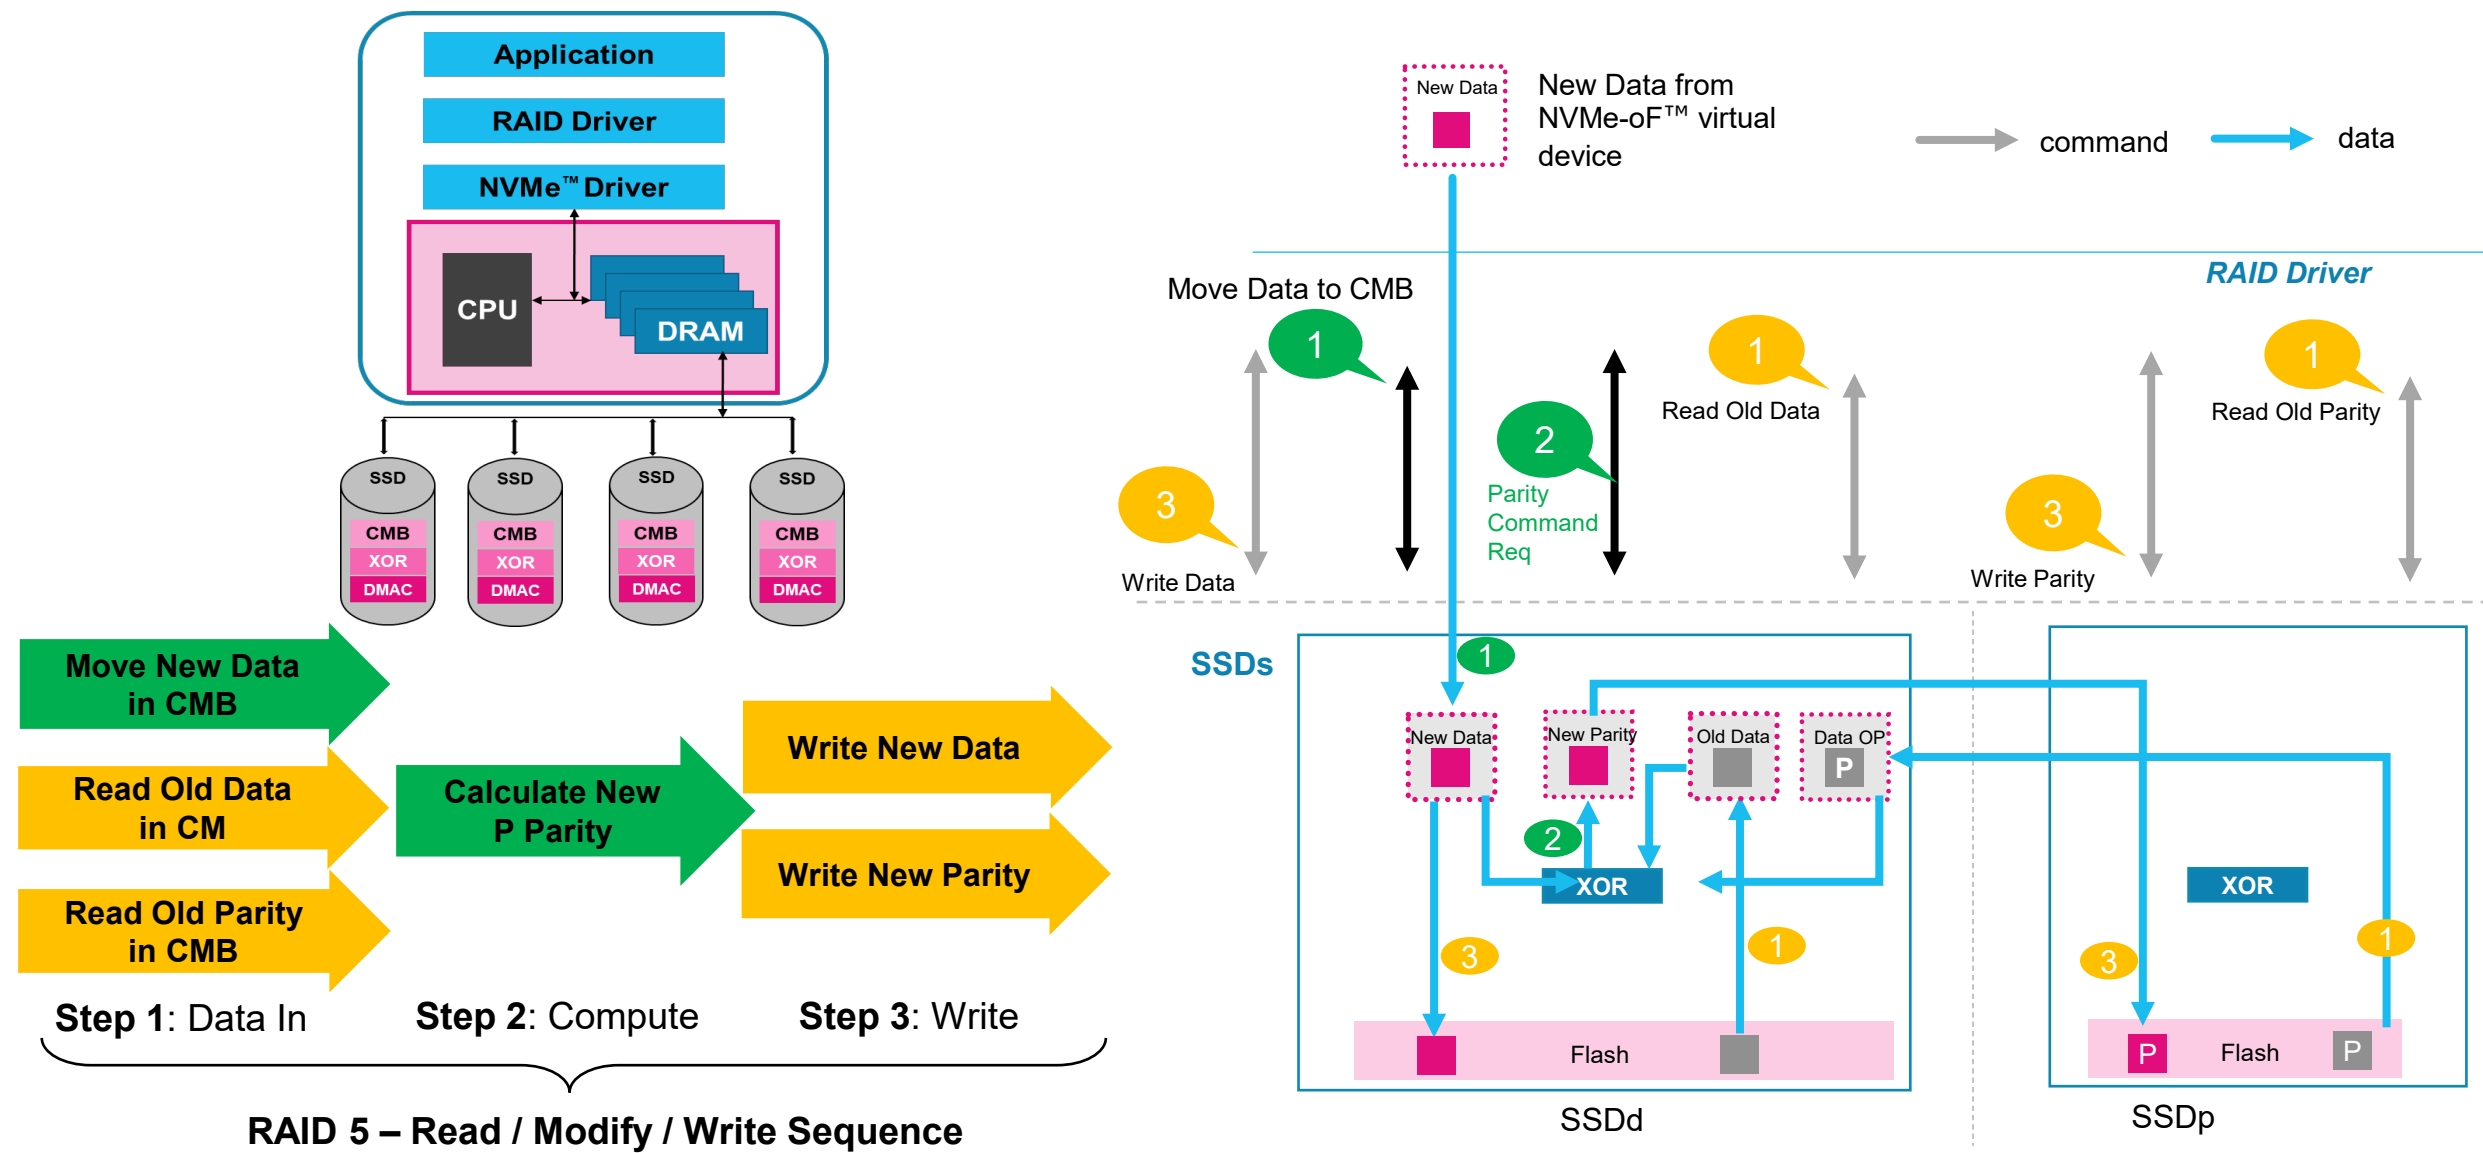
Host
orchestrated
standard based

Applications of Offload Engines

- ❑ **Hash/CRC³:** Dedupe, Object/File signature/scrubbing, buffer integrity
- ❑ **(De)Compression:** Compression with levels, decompress and filter
- ❑ **Parity Compute:** Erasure code (EC), compare, **Data scrubbing, RAID Rebuild**



Command and Data Flow Example for RAID 5 Write



Fail in Place Mechanism at the SSD Level

KIOXIA Built-in SSD Level Failure Mitigation

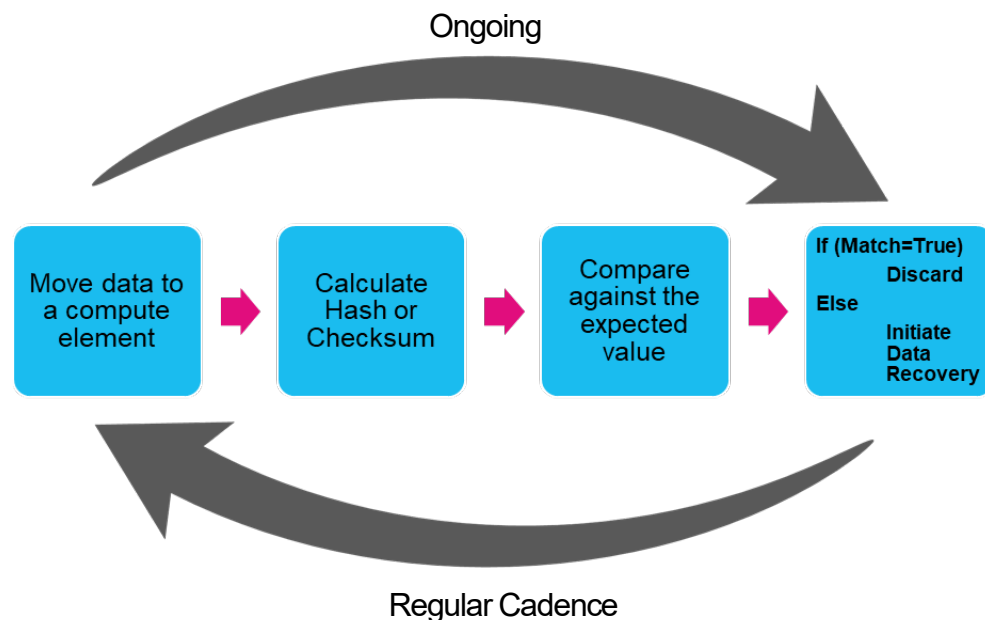
Failure Mitigation Features	Description
Wear Leveling	Internally developed “Always on” algorithm to allow even media wearing
Die Fail Tolerance and Recovery	Enterprise and data center SSDs historically supports die fail tolerance and recovery
Customizable End-of-Life Behavior	When and what actions to take upon reaching certain thresholds
Overprovisioning (OP) and Available Spare	Available Spare, critical warning in SMART ¹ / health information Read Only mode if spare capacity falls below threshold
Dual-Port Functionality in Enterprise SSDs	Ensures high availability in case of path failure

Data verification mechanism still required to ensure data integrity at File/Object/Stripe level.

Data Scrubbing Using Offload

Data verification using Data Scrubbing in Conventional Setup

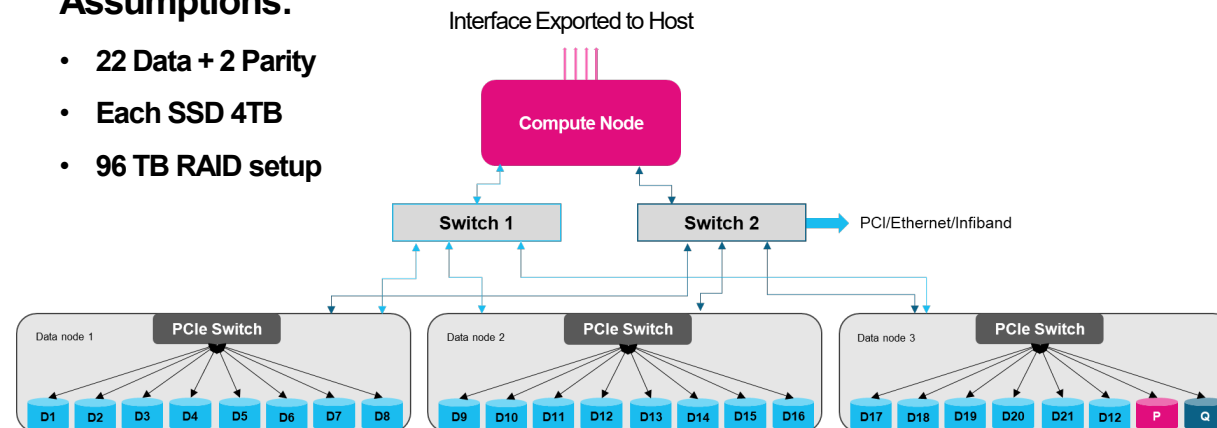
- **Data Scrubbing:** Early detection and correction of errors
- **Method employed :** Hash, checksum or RAID technology
- **Supported Level :** File, Object and RAID Stripe



All data movement during scrubbing operation is an overhead penalty paid to ensure data integrity

Assumptions:

- 22 Data + 2 Parity
- Each SSD 4TB
- 96 TB RAID setup



Compute node performing data and parities verification for one stripe using RAID

1. $P + D1 + D2 + D3 + D4 + \dots + D22 = 0$
2. $Q + g1.D1 + g2.D2 + g3.D3 + \dots + g22.D22 = 0$

For each Scrubbing Cycle:

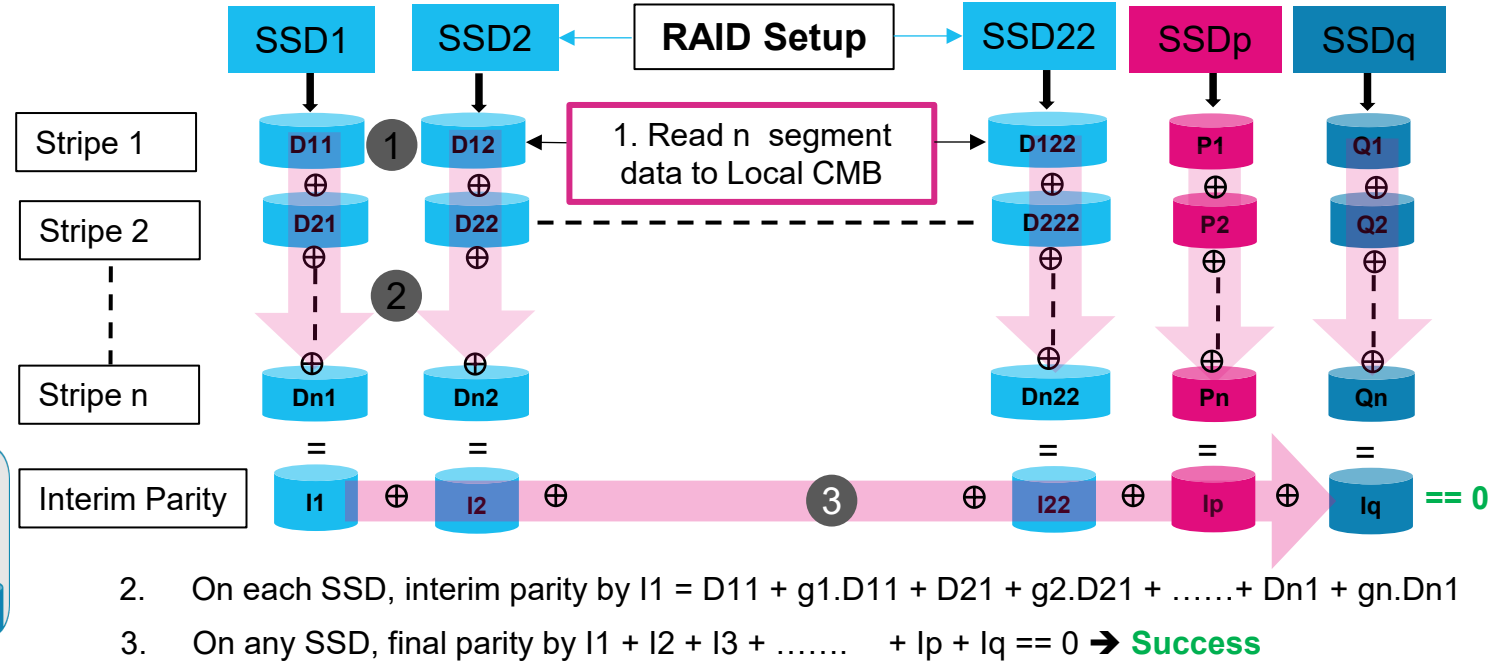
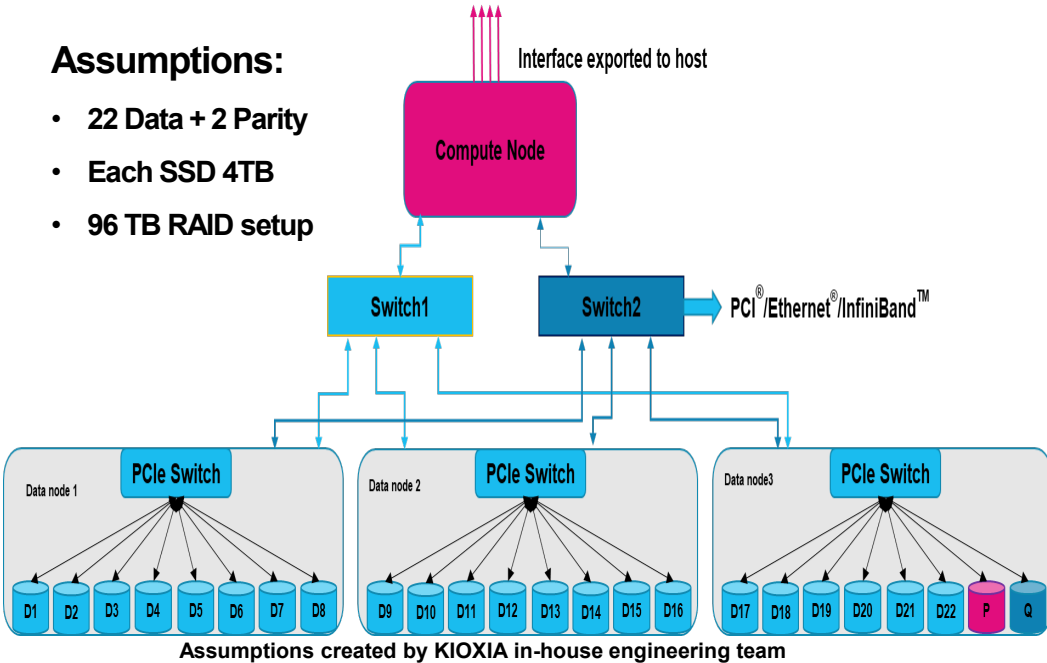
- 96TB data moves over PCIe[®], network and CPU
- 192TB through memory subsystem

Assumptions created by KIOXIA in-house engineering team.

Error Detection using Data Scrubbing with RAID/EC Offload

Assumptions:

- 22 Data + 2 Parity
- Each SSD 4TB
- 96 TB RAID setup



- Using 3 step process, system resource usages reduced by ~99%
- No data passes through CPU and DRAM on compute node
- For n stripes, only one stripe moves over network and PCIe®
- **A sustainable and scale out solution**
- $Dnm \Rightarrow$ Data segment D of nth stripe on SSDm

System Resources to Scrub 96 TB (subject to change)	Conventional Method	KIOXIA Data Scrubbing Offload	% of Savings
Data passes through CPU for parity computation	96 TB	0 TB	100%
Data passes through DRAM controller	192 TB	0 TB	100%
Data passes through network	96 TB	0.4 TB	99%
Data passes through PCIe® interface	96 TB	0.4 TB	99%

A very flexible, geometry agnostic and adoptable solution for early error deduction

Rebuild Using RAID Offload

Efficient Rebuild of 2 SSDs using RAID offload

1. Read good data segments of one RAID strip on compute node

2. Compute these two equations in order

a. $D1 = (g2.D2 + g3.D3 + + g22.D22 + Q)/g1$

b. $P = D1 + D2 + D3 + D4 + D22$

3. Write D1 and P to new SSDs on respective stripe

4. Repeat steps 1-3 for each RAID stripe generation

Without offload

1. Read Dm to Dn segments in the CMB of one SSD in each data node

2. Calculate node specific intermediate data D_{in} and parity P_{in} . For example, on node 1

1. $P_{i1} = D2 + + D8$

2. $D_{i1} = (g2.D2 + + g8.D8) / g1$

3. Read intermediate data D_{in} and parity P_{in} from each data node on compute node and compute D1 followed by P parity

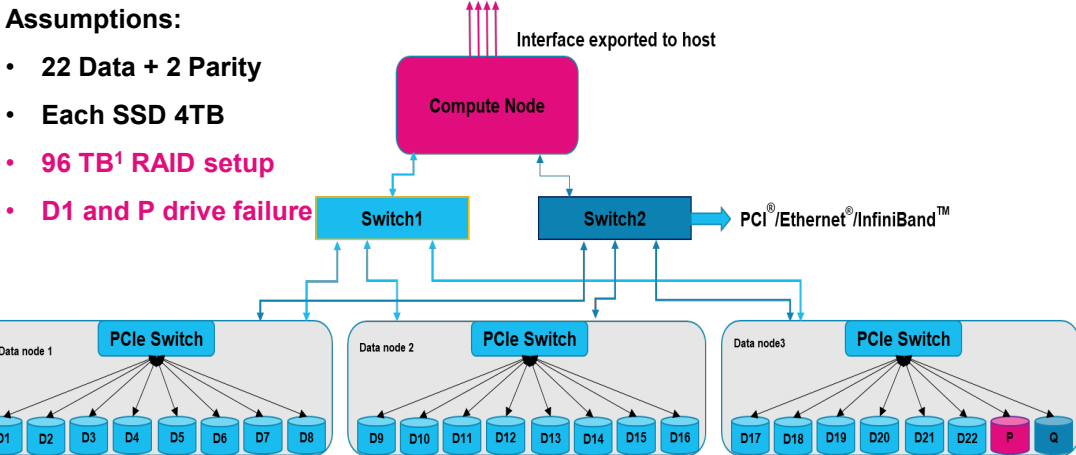
1. $D1 = D_{i1} + D_{i2} + D_{i3}$

2. $P = D1 + P_{i1} + P_{i2} + P_{i3}$

4. Write P and D1 to new SSDs on respective stripe

5. Repeat steps 1-4 for rebuilding each RAID stripe

With offload



System Resources to build two drives (subject to change ²)	Conventional Method	KIOXIA Data Rebuild with Offload	% of Savings
Data passes through CPU for parity computation	88 TB	24 TB	73%
Data passes through DRAM controller	192 TB	64 TB	66%
Data passes through network	96 TB	32TB	66%
Data passes through PCIe® interface	96 TB	108TB	-12%

Rebuild with offload can saturate write line speed

Summary

- **SSD level data loss mitigation to**
 - Avoid rebuild process specially with high-capacity SSDs
 - Recovery of data even in case of die failure
- **Efficient data scrubbing method to detect failure early while saving significant system resource cost**
- **Efficient rebuild process saving significant system resource cost and time**
- **Data scrubbing and rebuild is based on KIOXIA RAID Offload Technology, and ‘Best of Show’ award winner at the Future of Memory and Storage (FMS) 2024 conference**

KIOXIA

Definition of capacity: KIOXIA defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of $1\text{GB} = 2^{30} = 1,073,741,824$ bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, such as Microsoft Operating System and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

All company names, product names and service names may be trademarks of their respective companies.

Images are for illustration purposes only.

© 2025 KIOXIA America, Inc. All rights reserved. Information, including product pricing and specifications, content of services, and contact information is current and believed to be accurate on the date of the announcement, but is subject to change without prior notice. Technical and application information contained here is subject to the most recent applicable KIOXIA product specifications.