



SOFTWARE-ENABLED FLASH™

Software-Enabled Flash™ Has Arrived

Eric Ries, SVP, KIOXIA America

Earle F. Philhower, III, Sr. Manager, KIOXIA America

On behalf of the Software-Enabled Flash Project

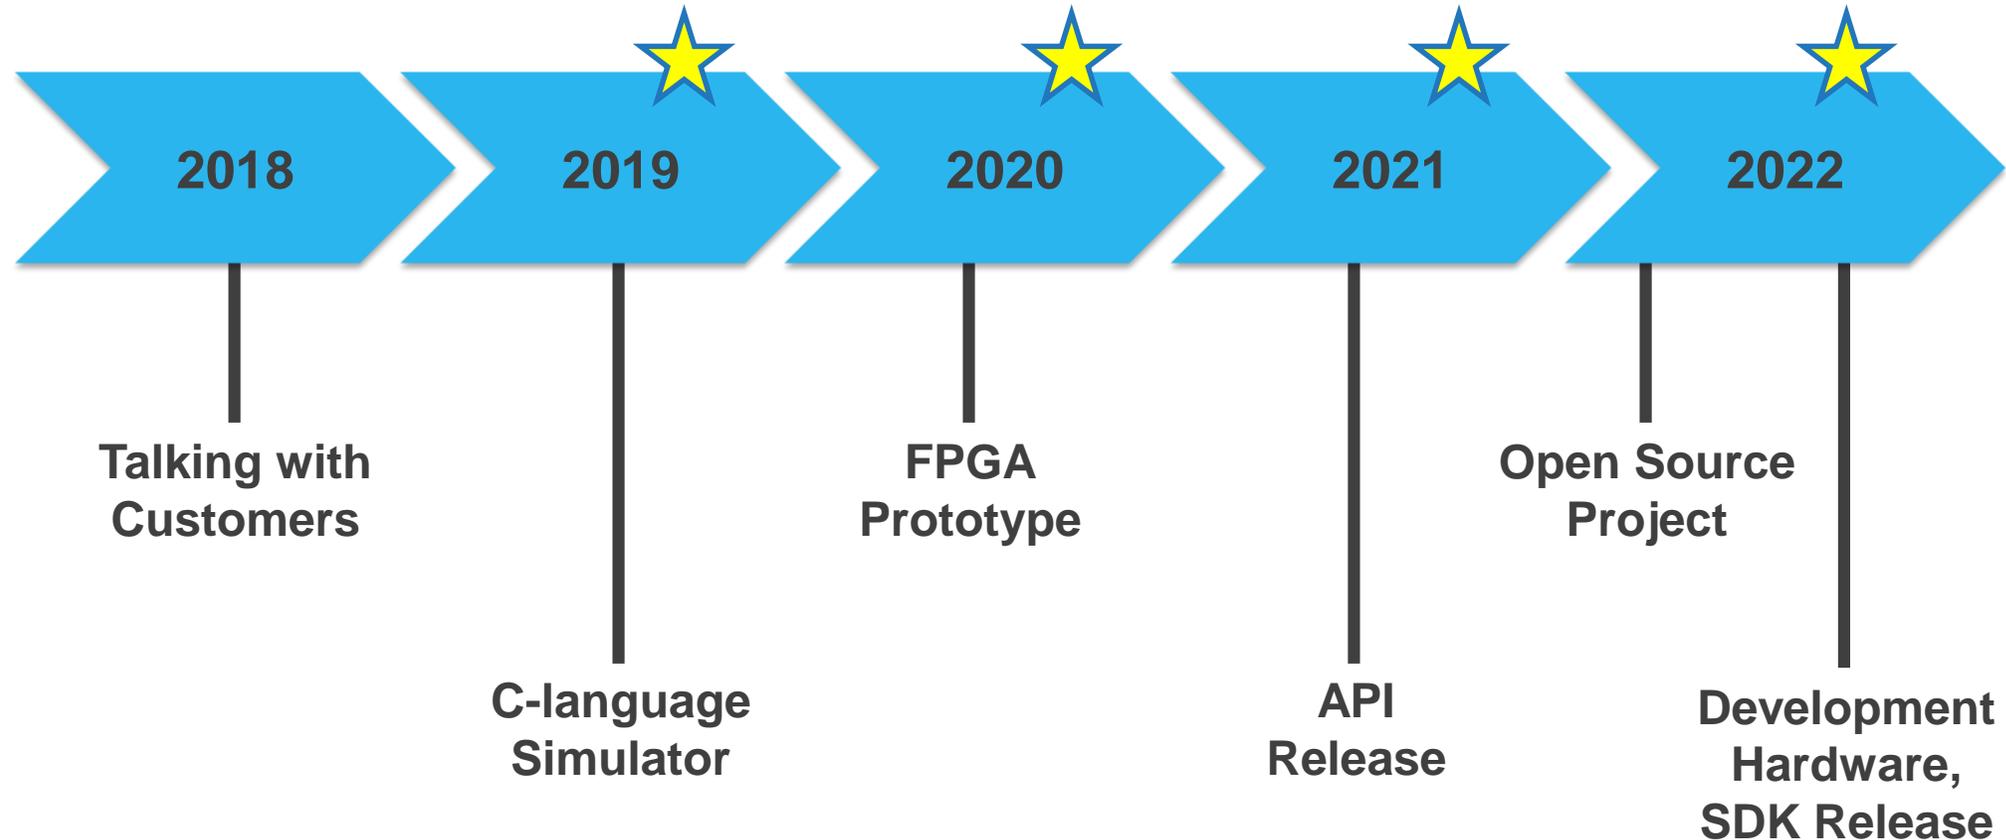
Open Source Software-Enabled Flash™ Project

- **Founded April 2022**
- **Software-Enabled Flash specifications and software released to OSS**

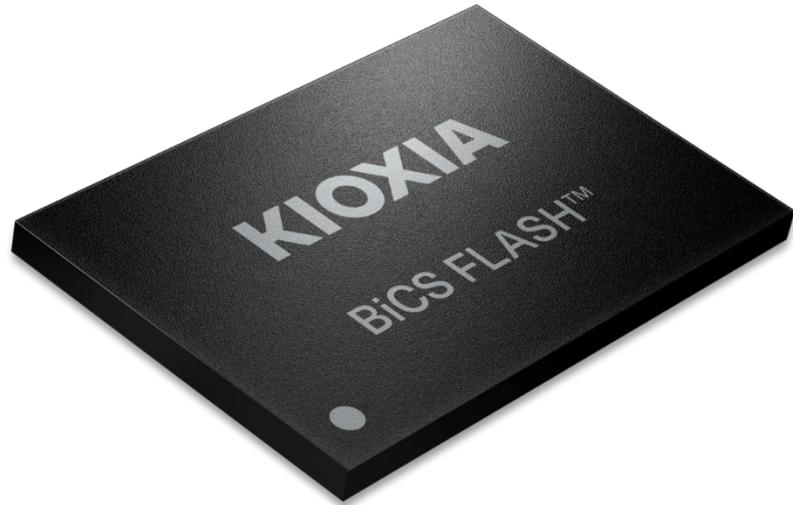


<https://softwareenabledflash.org>

How did we get here?



A Different Way of Thinking About Flash



- Drop legacy storage paradigm
- Deliver full parallelism of flash
- Explicit controls over isolation, queueing modes
- Application defined latency outcomes

Vendor Neutral

Managed under the Linux Foundation



Open Source Governance Processes

Open Technical Steering Committee (TSC) meetings

Specification built for multiple implementations

Flash technology independent
Controller technology independent

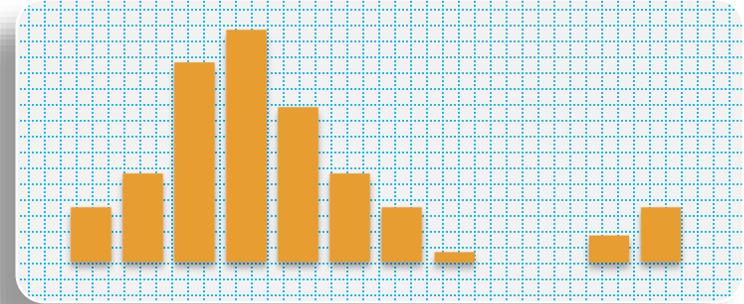
Open Source API and SDK

Available today on GitHub

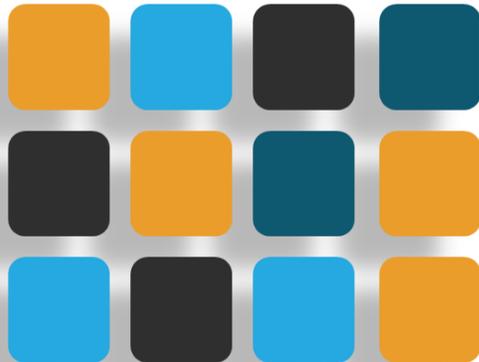
Developer Control



Isolation



Latency Outcomes



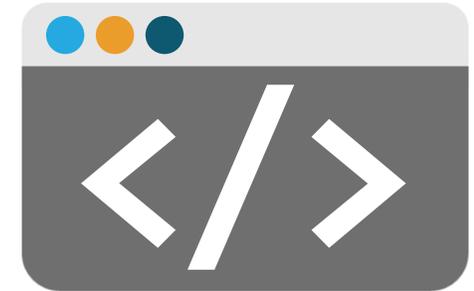
Data Placement



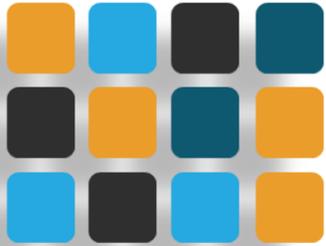
Copy Offload

Maximize Value of Flash Memory

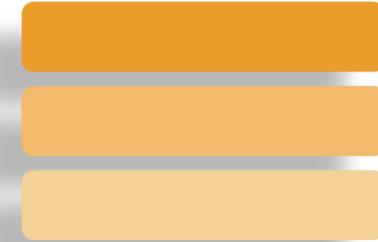
- **Application-defined data layout for maximum performance**
- **Workload optimized resource allocation**
- **Explicit requests for flash behavior**



Multiprotocol, in software



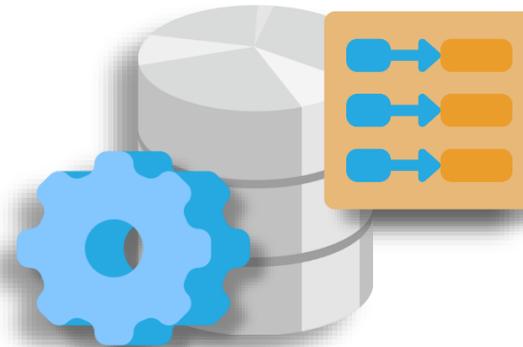
Block



Zoned



Smart FTL



+ others

Software-Defined Flash vs. Self-Built Drives

Self-Built Drive

Software-Enabled Flash™



Latency, Performance and Isolation

Optimized Garbage Collection

Application Optimized Protocols

Board Layout, Certifications

Flash Controller and Firmware

Flash Timings and Protocols

Physical Flash Layout

Latency, Performance and Isolation

Optimized Garbage Collection

Application Optimized Protocols



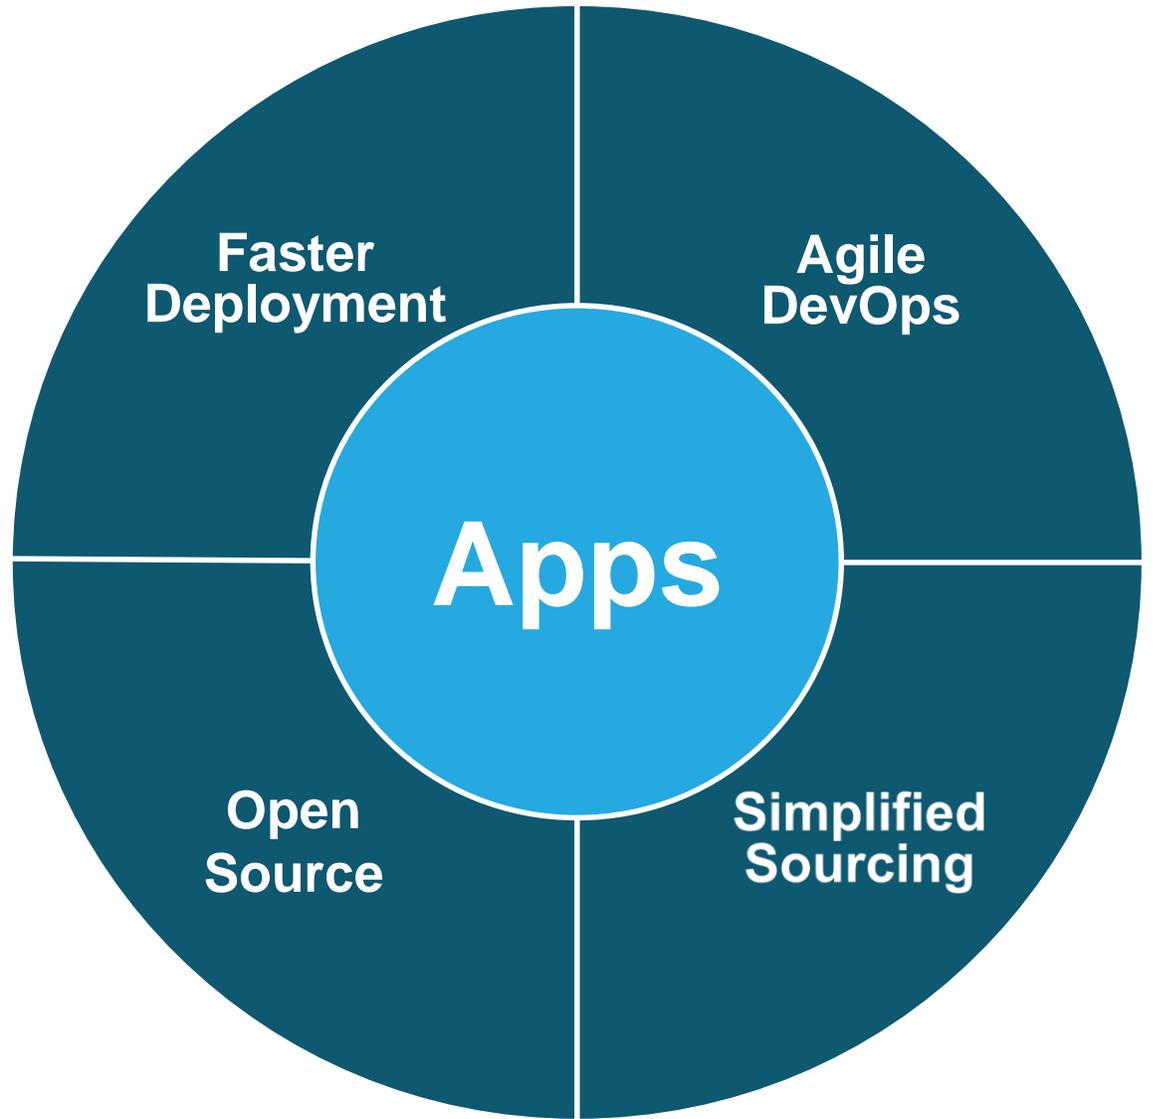
SOFTWARE-ENABLED FLASH™

Faster Flash Deployment and Re-Deployment

- **Explicit** behavior requests, not implicit
- **Simplified** migration between techs
- **Avoid** rewriting I/O stacks
- **Abstract** flash management tasks



Improve application results



Built for Storage Innovation



Open ecosystem for multi-vendors



Data placement control



Isolation and latency



Faster time-to-market than DIY

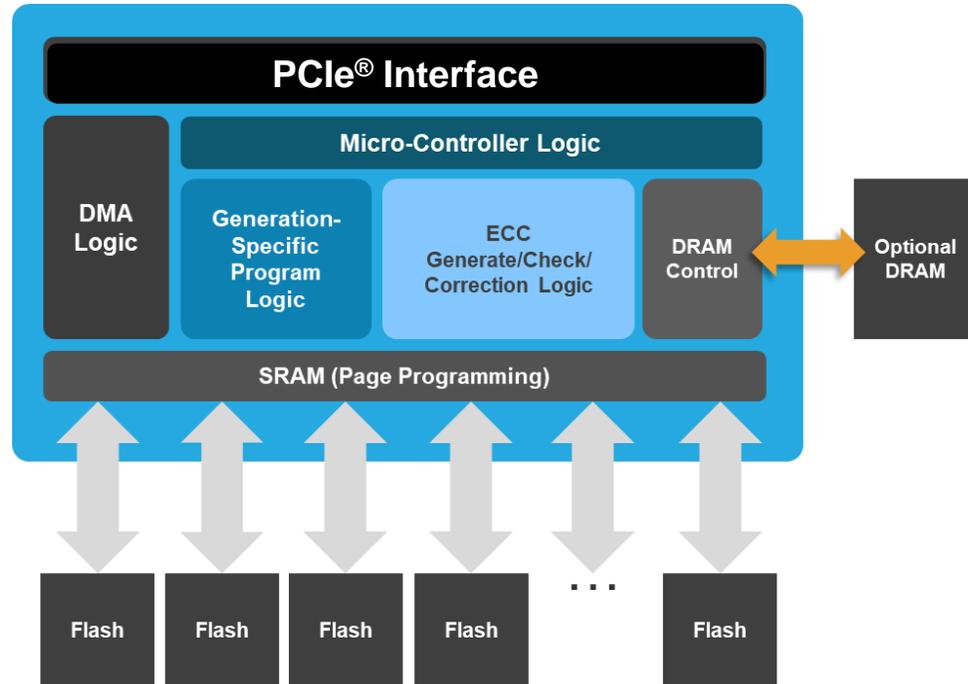


Improve application results

Hardware and Software

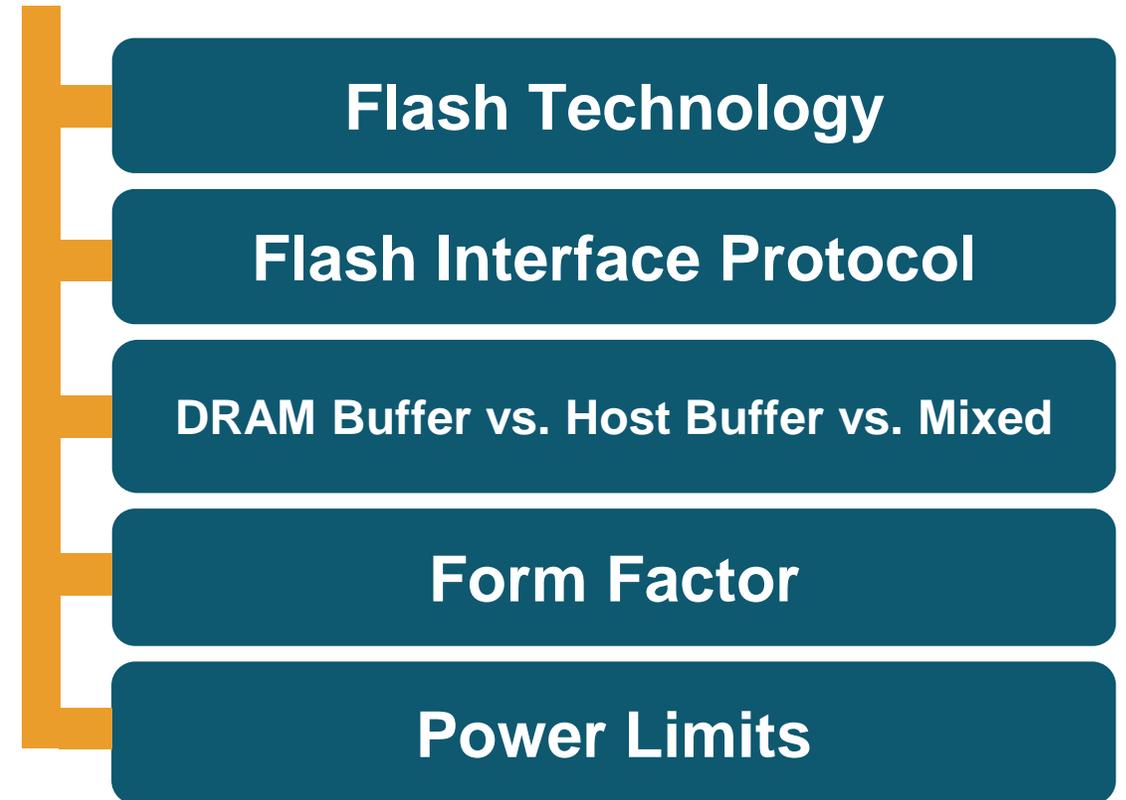
Working Together

Custom Engineered Hardware Platform



* Software-Enabled Flash compatible controllers available from Marvell™.

Vendor Configurable



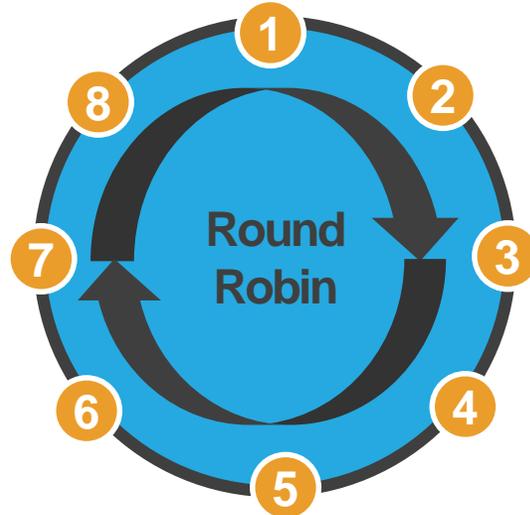
* PCIe is a registered trademark and/or service mark of the PCI-SIG

* The Marvell name and logo is a trademark and/or service mark of Marvell Technology, Inc.

Hardware Enabled Latency Control



Massively parallel, per-die I/O queues



Hardware-enforced I/O prioritization
Multiple, programmable scheduling modes



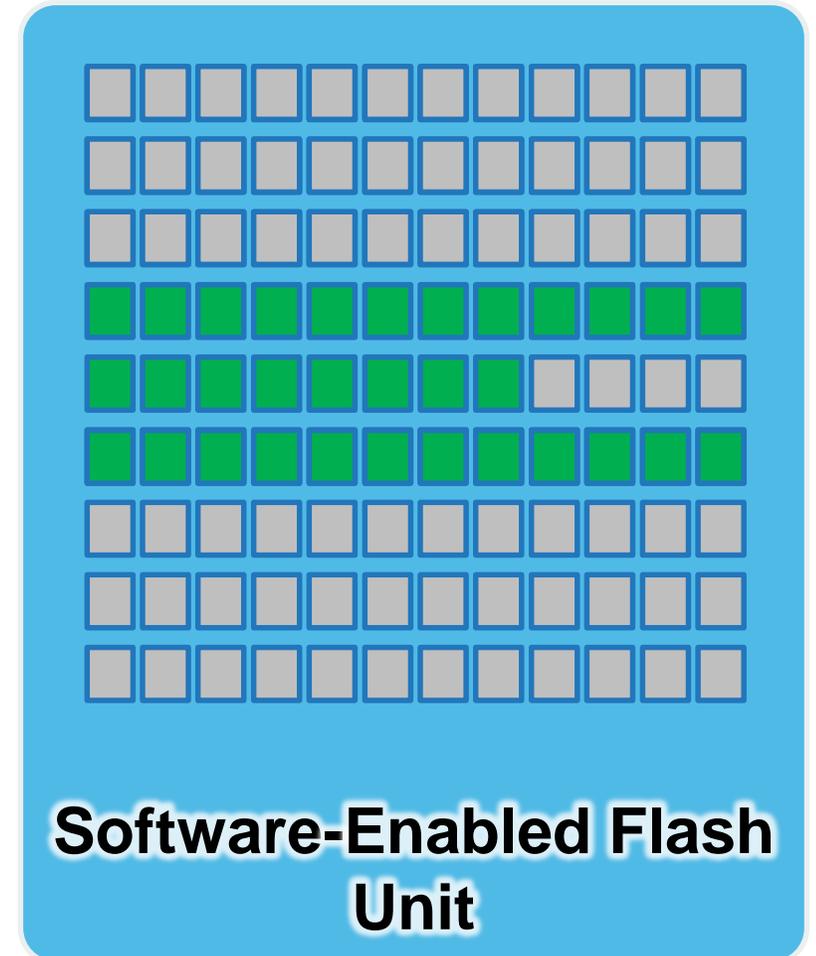
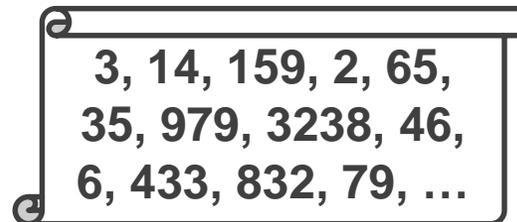
Die-Time Weighted Fair Queueing
Erase, Program, Read weights configurable

Copy Offload and Application-Driven GC

- Powerful intra-drive data copy primitive
- Single command does it all
 - Application specifies blocks to copy
 - Drive manages entire operation



- Applications
 - Database compaction
 - Log-Structured merge trees
 - Range copy for object storage
 - (and of course, efficient garbage collection)



Simple (but application-controlled)
Garbage Collection

Hardware Enabled Tenant Isolation

Virtual Device

Die-level isolation

Complete physical separation
of data

User-configurable at deployment

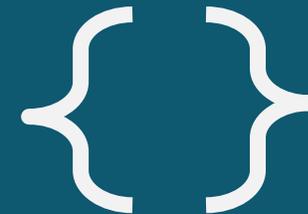


Quality of Service Domain

Workload-level isolation

Complete logical separation of
data per superblock

Isolated garbage collection,
overprovisioning



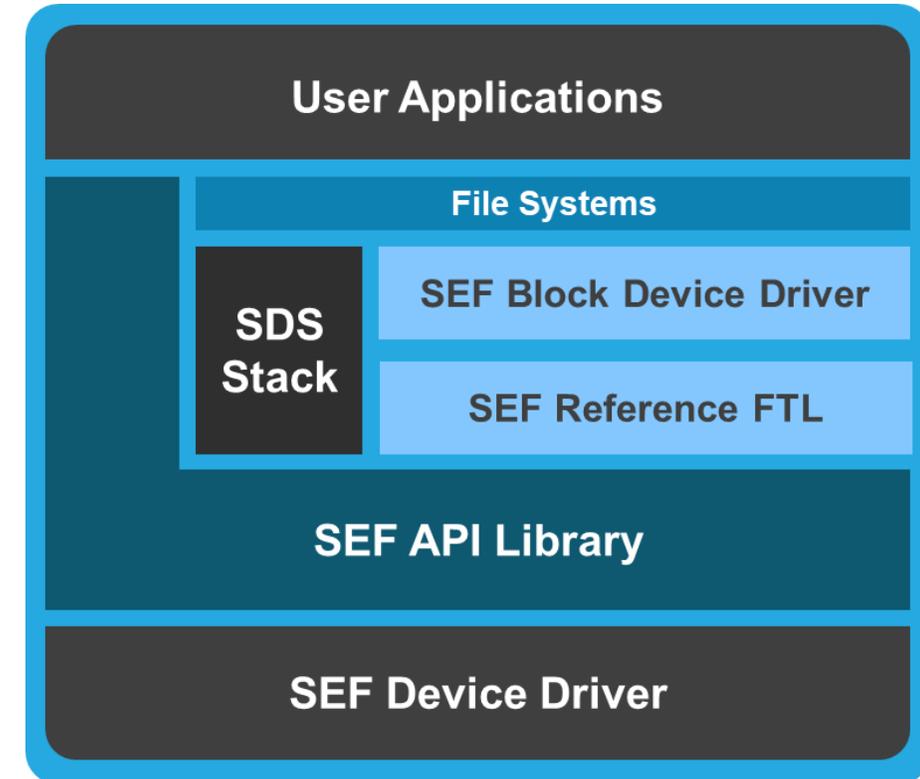
Customizable Software Stack

Open source software stack

- Choose the level of abstraction
- Leverage the work of others

Built for customization

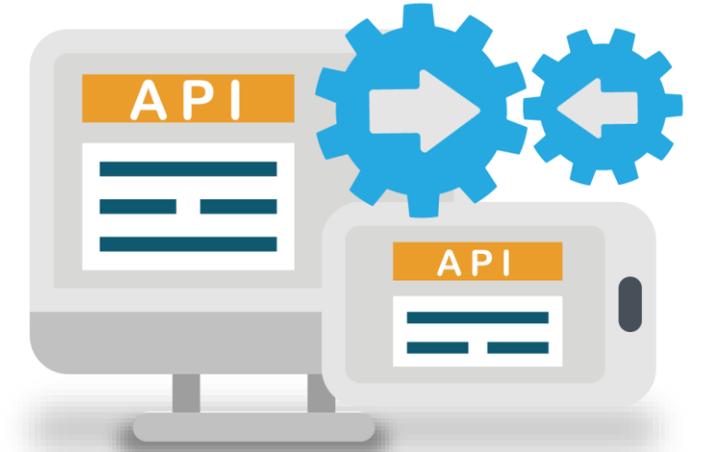
- Quick to get prototypes
- Full source to customize to unique needs



Source Code Included

API Header and Library

- Provides thin, user-level access to device
- Low-level interface for multiple vendors



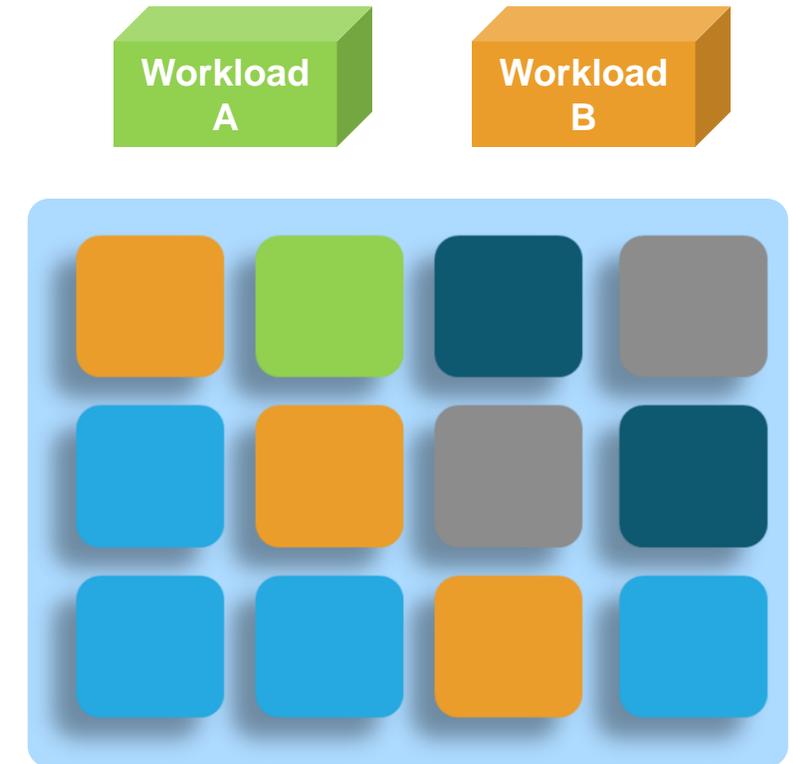
Software Development Kit

- User-level FTL
- Virtualization device drivers
- FIO testing tool
- CLI management interface



Easy “Bake Your Own FTL”

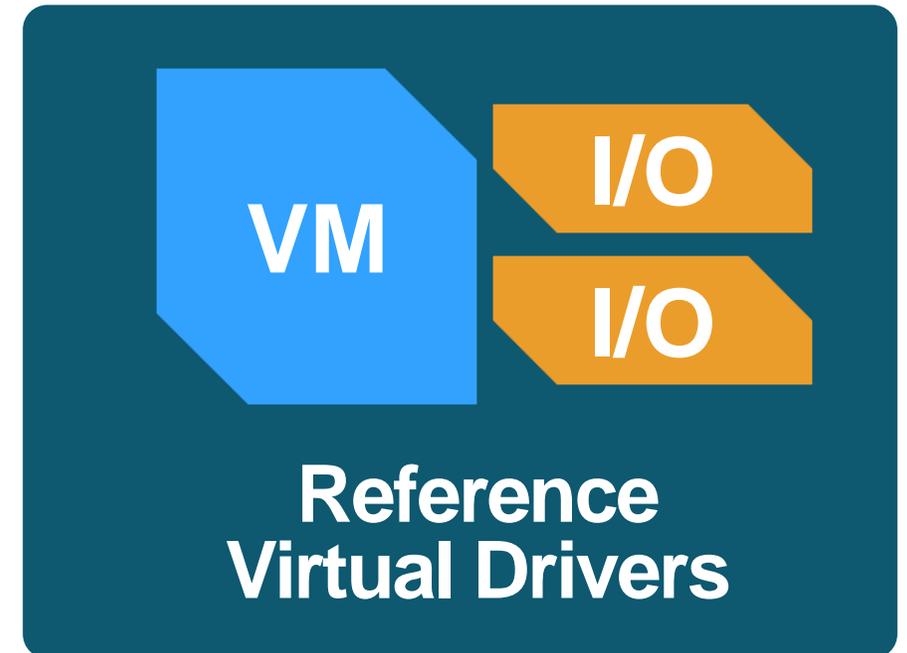
- Complete user-level FTL source code
- Runs on host, in each application
- Built for customization
 - Complete queueing/isolation capability
 - Application-optimized protocols
 - Fine-grained data placement
- Full documentation online, in code



Reference
Flash Translation Layer

Rapid Virtualization Deployment

- **QEMU VirtIO device drivers included**
 - ✓ Full source code
 - ✓ Traditional Block
 - ✓ Zoned Namespace (ZNS)
 - ✓ All queueing, isolation, offload



Proof of the Pudding is in the Eating...

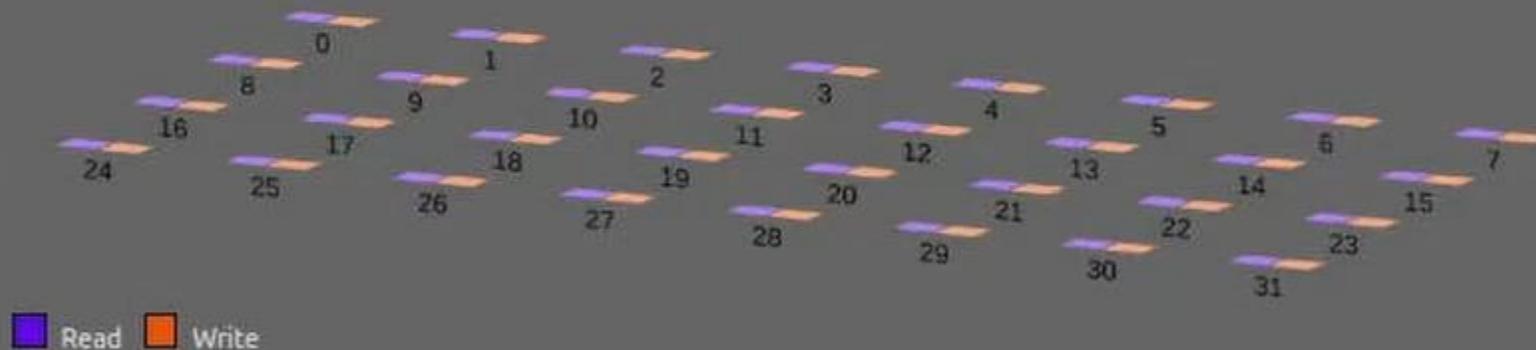
Software-Enabled Flash Demo

Start Demo

- ✓ Enable Workload 1
- ✓ Enable Workload 2
- ✓ Enable Workload 3

Software-Enabled Flash
Virtual Devices can isolate
workloads to individual flash
dies for complete
performance isolation

Usage



Redefining Flash for Storage Innovation



Focus on applications, not low-level flash operations



Better, tuned use of flash per application



**Isolation + Latency Control +
Placement + Offload + SDK**

= Better Application Results

Call to Action



Stop by the table, talk to our team



Birds of a Feather



**Join the project, get the code,
make something awesome.**



**Membership in the
Open Source project
is open**

Software-
Enabled
Flash
Needs
You!

SEF



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.

© 2022 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.