KIOXIA

G2M RESEARCH

The Next Flash Revolution at Scale: Open Source Software + Software-Enabled Flash

G2M Research Webinar

Wednesday Nov 17, 2021

G2M Research Introduction and Ground Rules

Mike Heumann
Managing Partner, G2M Research

Webinar Agenda

- Ground Rules and Webinar Topic Introduction (G2M Research)
- Requirements of flash at scale
- Live survey
- Software-Enabled Flash™ technology introduction
- Meeting the needs of flash at scale
- Audience Q&A
- Wrap-Up

Open-Source Storage Software at Scale

- Open-source storage software packages are utilized by hyperscalers and cloud service providers (CSPs) in their solutions:
  - Ceph
  - Hadoop
  - OpenStack

  * CloudStack

  * Lustre

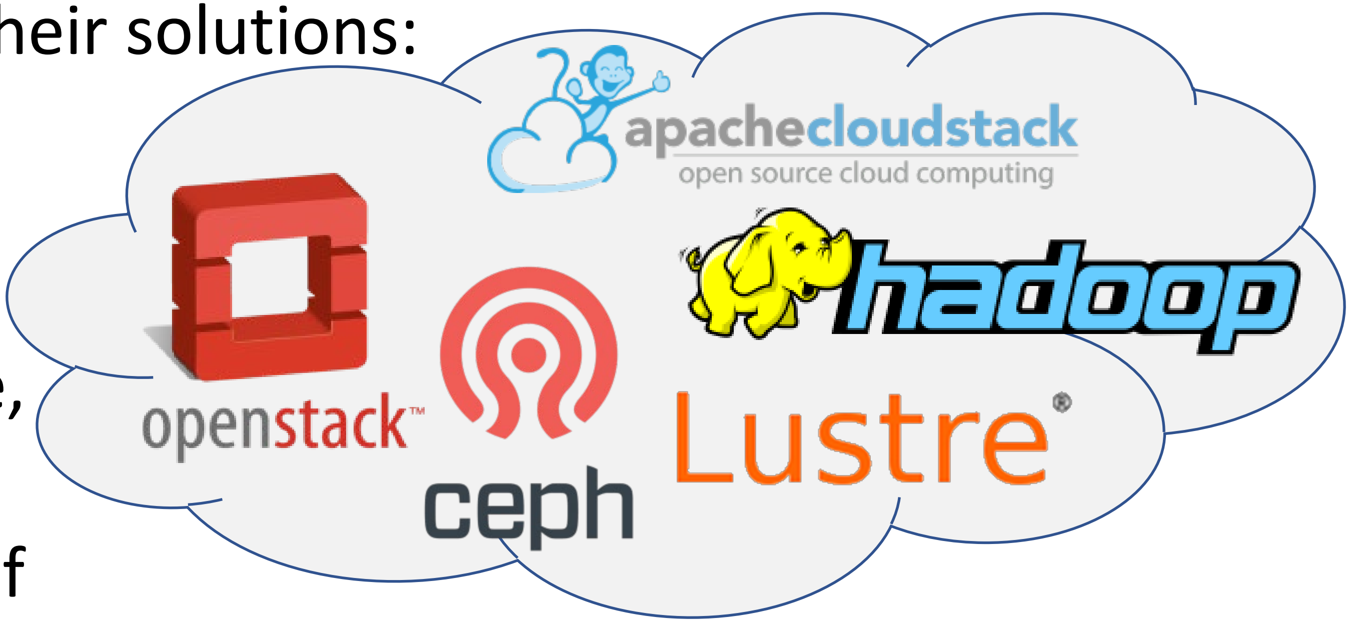

- These packages are highly tunable, customizable, and extensible
- The applications that run on top of these storage packages often have very different performance needs
- Extending this "tunability" to the flash media itself could significantly increase the performance of these applications

** Names and logos are the property of the respective companies*

Kioxia/G2M Research Webinar on Software-Enabled Flash™ Memory

Sean Stead
Senior Staff
www.kioxia.com

KIOXIA

Mike Heumann
Principal Analyst
www.g2minc.com

G2M RESEARCH

The Next Flash Revolution at Scale: Open Source Software + Software-Enabled Flash™ Technology

Memory & Storage Strategy Division

KIOXIA America, Inc.

© 2021 KIOXIA America, Inc. All Rights Reserved.

Change is needed in cloud flash storage

© 2021 KIOXIA America, Inc. All Rights Reserved.

Flash Storage

A legacy drive issue . . .

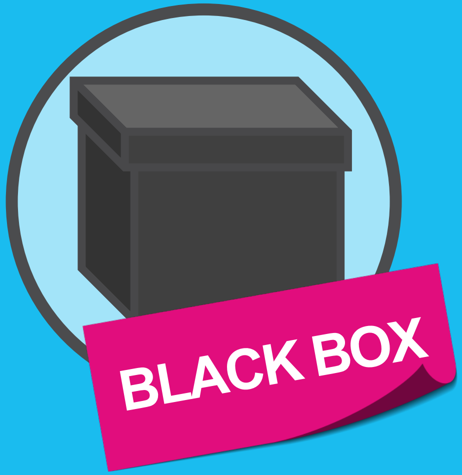

BLACK BOX

INFLEXIBLE

LIMITED

© 2021 KIOXIA America, Inc. All Rights Reserved.

Legacy Protocols:

- SAS
- SATA
- Block
- ZNS

Protocols fixed by device type can never adapt to changes…

© 2021 KIOXIA America, Inc. All Rights Reserved.

Cloud industry issues with fixed protocols:

1. Functions limited by device protocol
2. Difficult to adapt to changing workloads
3. A complex supply chain

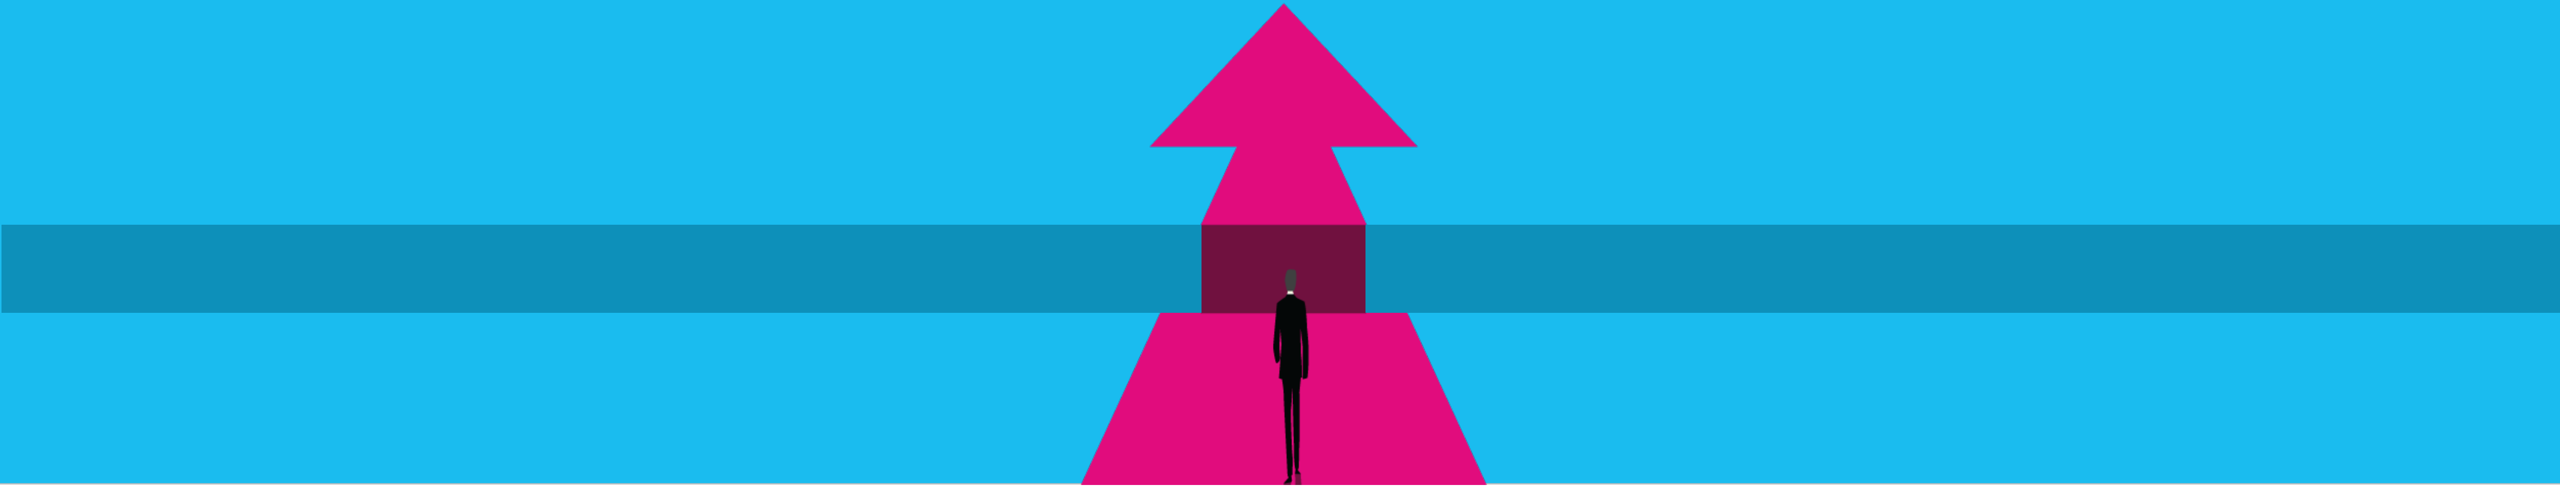

KIOXIA

© 2021 KIOXIA America, Inc. All Rights Reserved.

“Software-Defined *you name it*” rules the rack

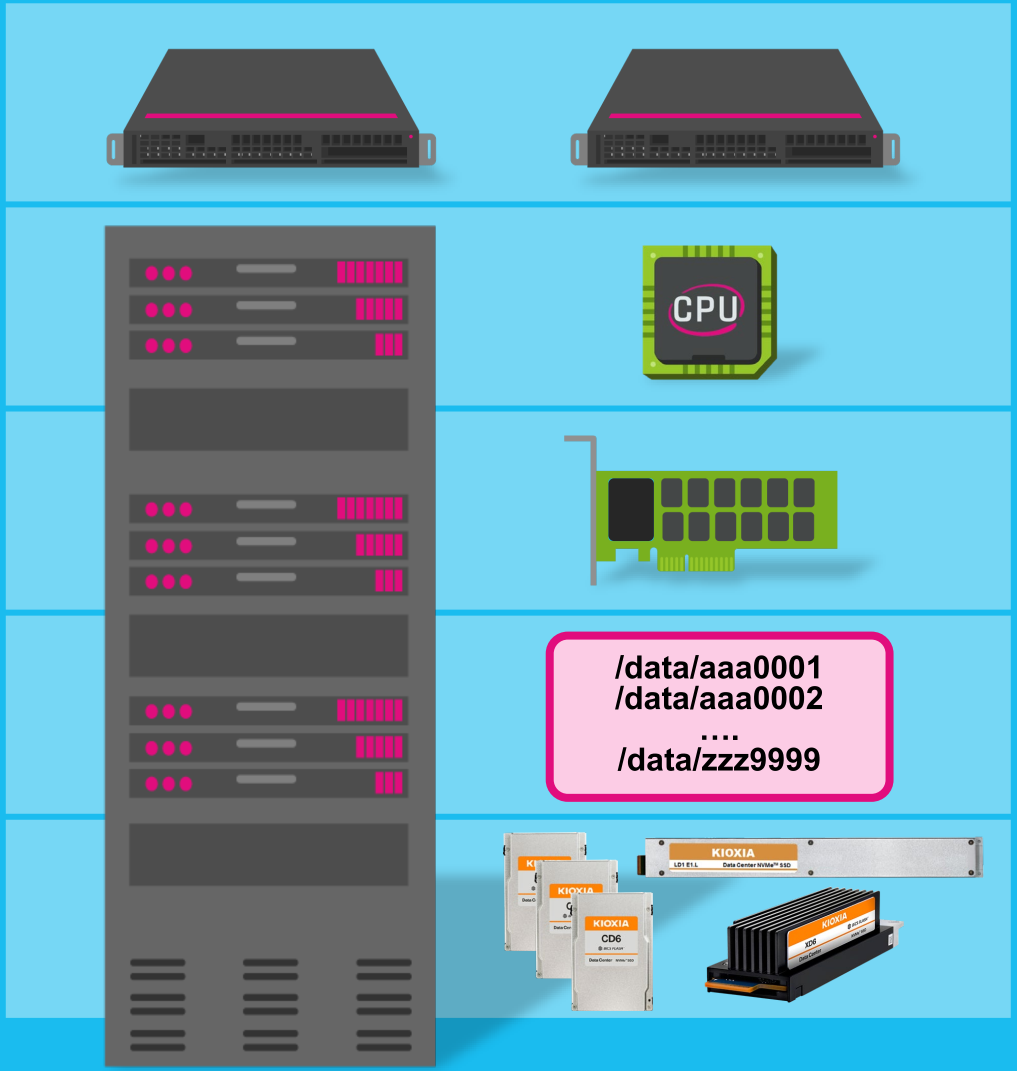

Software-Defined Networking

Software-Defined Computation

Software-Defined Hardware

Software-Defined Storage

It’s now time to apply the same logic to flash…

© 2021 KIOXIA Corporation. All Rights Reserved.

What if Flash was Software Defined?

© 2021 KIOXIA America, Inc. All Rights Reserved.

Introducing;

Software-Enabled Flash™ Technology

A host-managed, software defined approach to flash hardware

Fundamentally redefines the relationship between host and solid-state storage

- Brings control of media to the host
- Host applications have complete control over storage functionality and behavior
- Solves legacy overhead problems and enables new features
- Maximizes flash flexibility, performance and parallelism…

in other words, VALUE!

© 2021 KIOXIA America, Inc. All Rights Reserved.

INNOVATE

Software-Enabled Flash™ technology

Designed for flash storage innovation

- Sheds the legacy HDD paradigm
- Defines protocol through a software API
- Adapt and customize to your storage requirements
- Easily change protocols with a simple driver change
- Develop customized implementations from the host side

A powerful way forward for flash.

© 2021 KIOXIA America, Inc. All Rights Reserved.

BENEFITS

Advantages of a software defined-flash solution

Deploy the full power and capabilities of flash

- Single adaptable flash device hardware SKU (TCO)
- Fast time to market with new flash geometries
- Easily adapted to new application requirements
- Reference source code and large code reuse
- SEF handles the low level flash for you

© 2021 KIOXIA America, Inc. All Rights Reserved.

Audience Survey #1

Which open-source storage platforms are utilized in your organization? (**choose all that apply**):

Ceph:	43%
CloudStack:	0%
Hadoop:	29%
Lustre:	14%
OpenStack:	71%
Other:	29%
None of the above:	14%

Software-Enabled Flash™ Technology is software-defined flash

VALUE

Time to market

Accelerates transition to newest flash generations

Resource allocation

Focus development where it matters most

Maximize your flash

Adapt and change flash as workloads change

Redefines the relationship between host and solid-state storage…

© 2021 KIOXIA America, Inc. All Rights Reserved.

A Software Enabled API

1 A software enabling technology built around an open source flash-native API

2 It delivers a rich interface of functions and tools to simplify storage innovation

3 The API abstracts flash details enabling future generations of flash to work without code changes

4 Allows any flash vendor to build and optimize their flash to the API

5 KIOXIA will be providing sample source code, documentation, and libraries

Software-Enabled Flash Technology

is not software but it is

software enabled.

© 2021 KIOXIA America, Inc. All Rights Reserved.

Software and Hardware Working Together

DESIGN

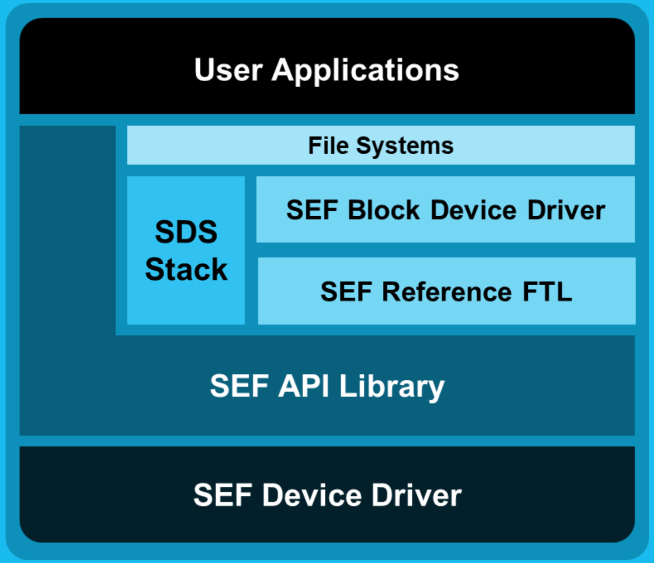

User Applications

File Systems

SDS Stack

SEF Block Device Driver

SEF Reference FTL

SEF API Library

SEF Device Driver

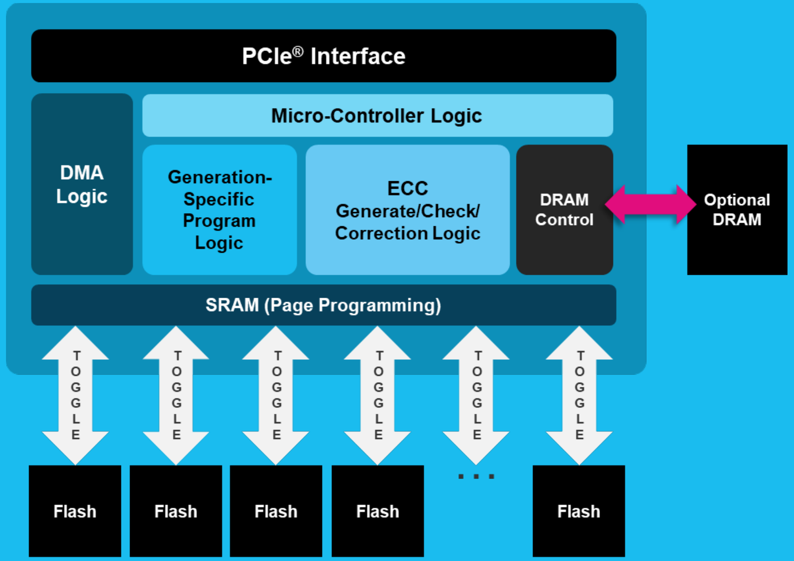

PCIe® Interface

DMA Logic

Micro-Controller Logic

Generation-Specific Program Logic

ECC Generate/Check/ Correction Logic

DRAM Control

Optional DRAM

SRAM (Page Programming)

TOGGLE

Flash

Flash

Flash

Flash

. . .

Flash

- Data placement
- Workload / tenant isolation
- Latency outcome control
- Buffer management

- Flexible DRAM configurations
- Host CPU offload
- Flash generation abstraction
- Advanced die time scheduling

(PCIe is a registered trademark of PCI-SIG)

© 2021 KIOXIA America, Inc. All Rights Reserved.

API + SDK

Software API Benefits and Values

- Faster application time-to-market
- No need to build from scratch
- Base for new development
- New options for Virtualized/Containerized application classes
- Deploys the full power and capabilities of FLASH hardware
- The SDK will allow developers to evaluate Software-Enabled Flash without writing a single line of code

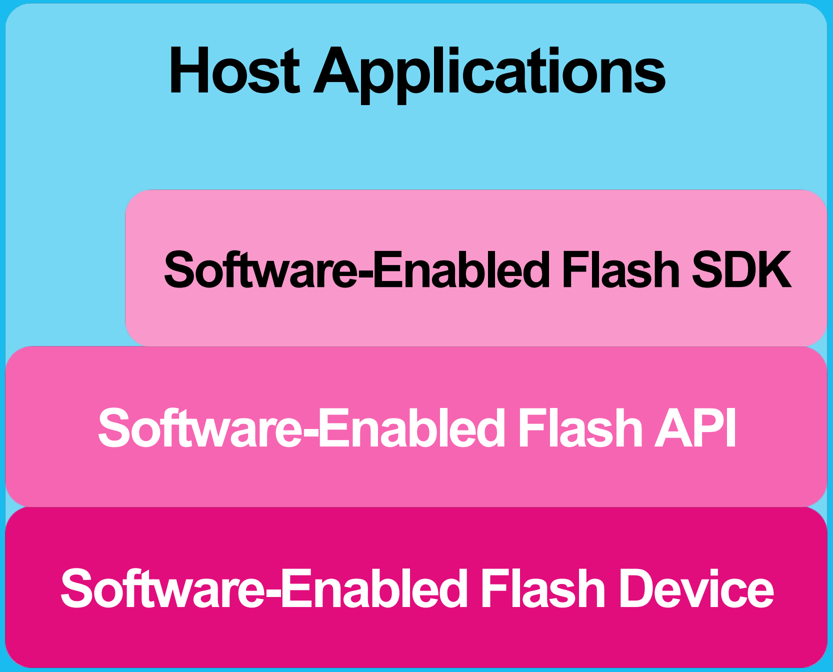

Host Applications

Software-Enabled Flash SDK

Software-Enabled Flash API

Software-Enabled Flash Device

© 2021 KIOXIA America, Inc. All Rights Reserved.

ABSTRACTION

Abstract the flash with a software API

- Low-level flash management
- Flash type, vendor & generation
- Hardware isolation of flash die
- Expose/manage parallelism
- Read/write
- Latency
- Die time management
- Offloading host resources
- Lifecycle management

Flash programming, ECC, bad blocks

Noisy neighbors

Workload or performance lag

Garbage collection / wear leveling

© 2021 KIOXIA America, Inc. All Rights Reserved.

FUNCTIONS

Software-Enabled Flash™ Capabilities

1. Multi-protocol
2. Data placement & workload isolation
3. Latency control & queuing

© 2021 KIOXIA Corporation. All Rights Reserved.

1 Multi-protocol

A **SIMPLIFIED** supply chain
using a **Software-Defined** Flash Device

Sources a single flash device

Inventory of SEF devices

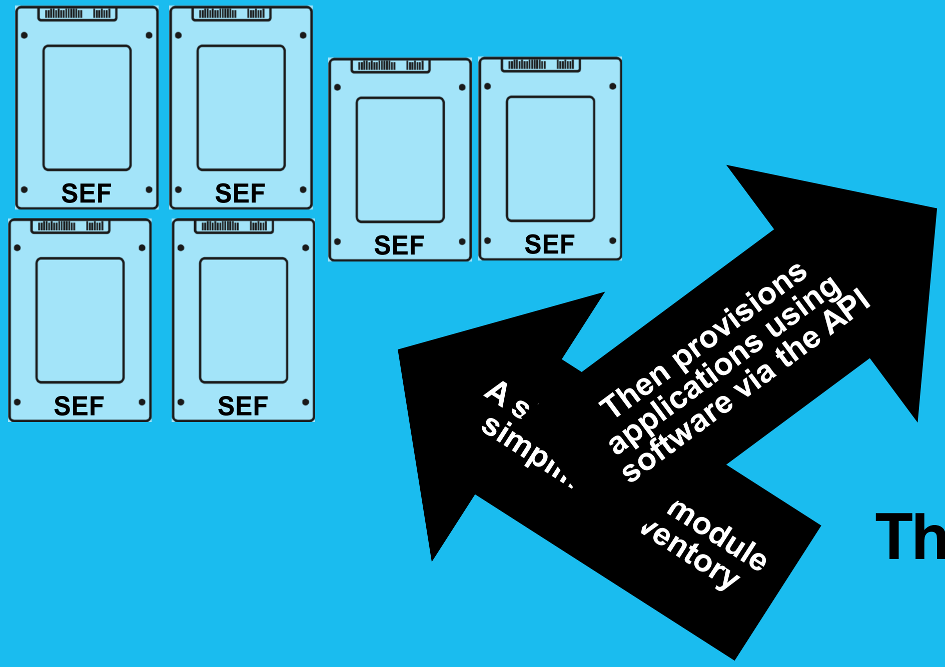

As **applications change**
simply change the software
to **redefine storage** functionality

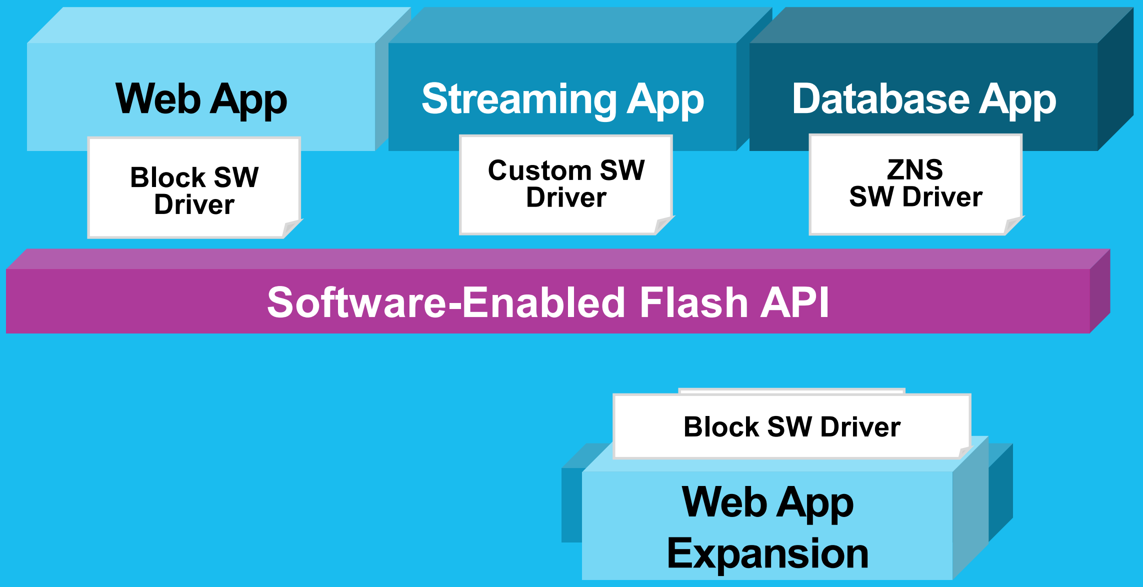

This new approach improves **TCO**
by sourcing a single storage module
and adapting its capability through software

© 2021 KIOXIA America, Inc. All Rights Reserved.

2 Data Placement and Workload Isolation Defined by the Host

Noisy Neighbor Problems ?

Isolate tenants from each other

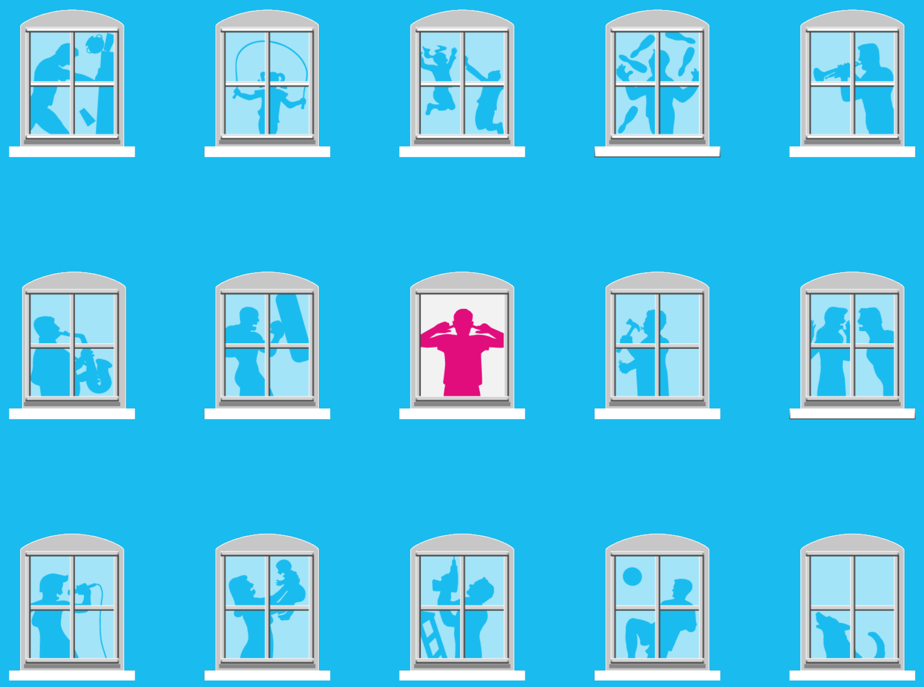

Noisy Neighbor Control

Define Data Placement

Define and manage HW and SW QoS domains

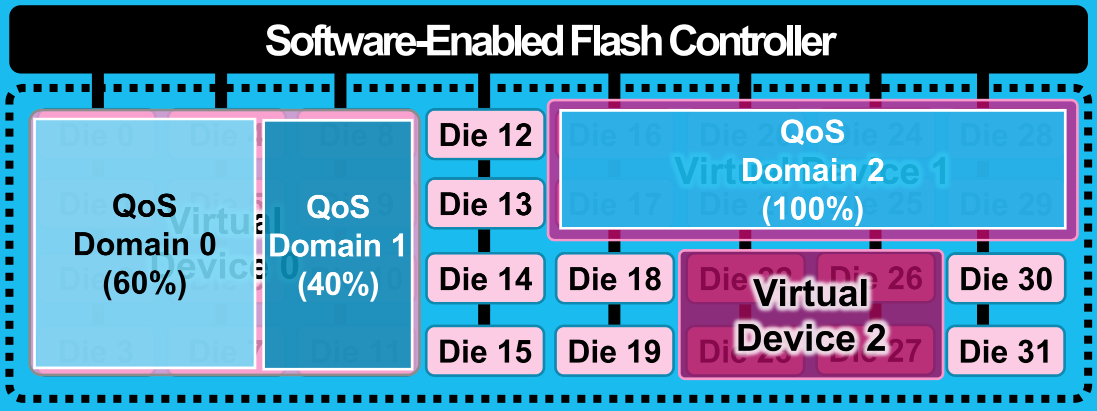

Flash Die Isolation
(Virtual Device)

Hardware Isolation

Flash Block Isolation
(QoS Domain)

Software Isolation

- Unique security encryption by QoS domain
- Dynamically reconfigurable overprovisioning for each namespace
- Manage I/O resources of flash die on an application-defined basis

© 2021 KIOXIA America, Inc. All Rights Reserved.

Software-Enable Flash™ technology deploys different protocols spanning 100's to 1,000s of devices to the individual flash die level

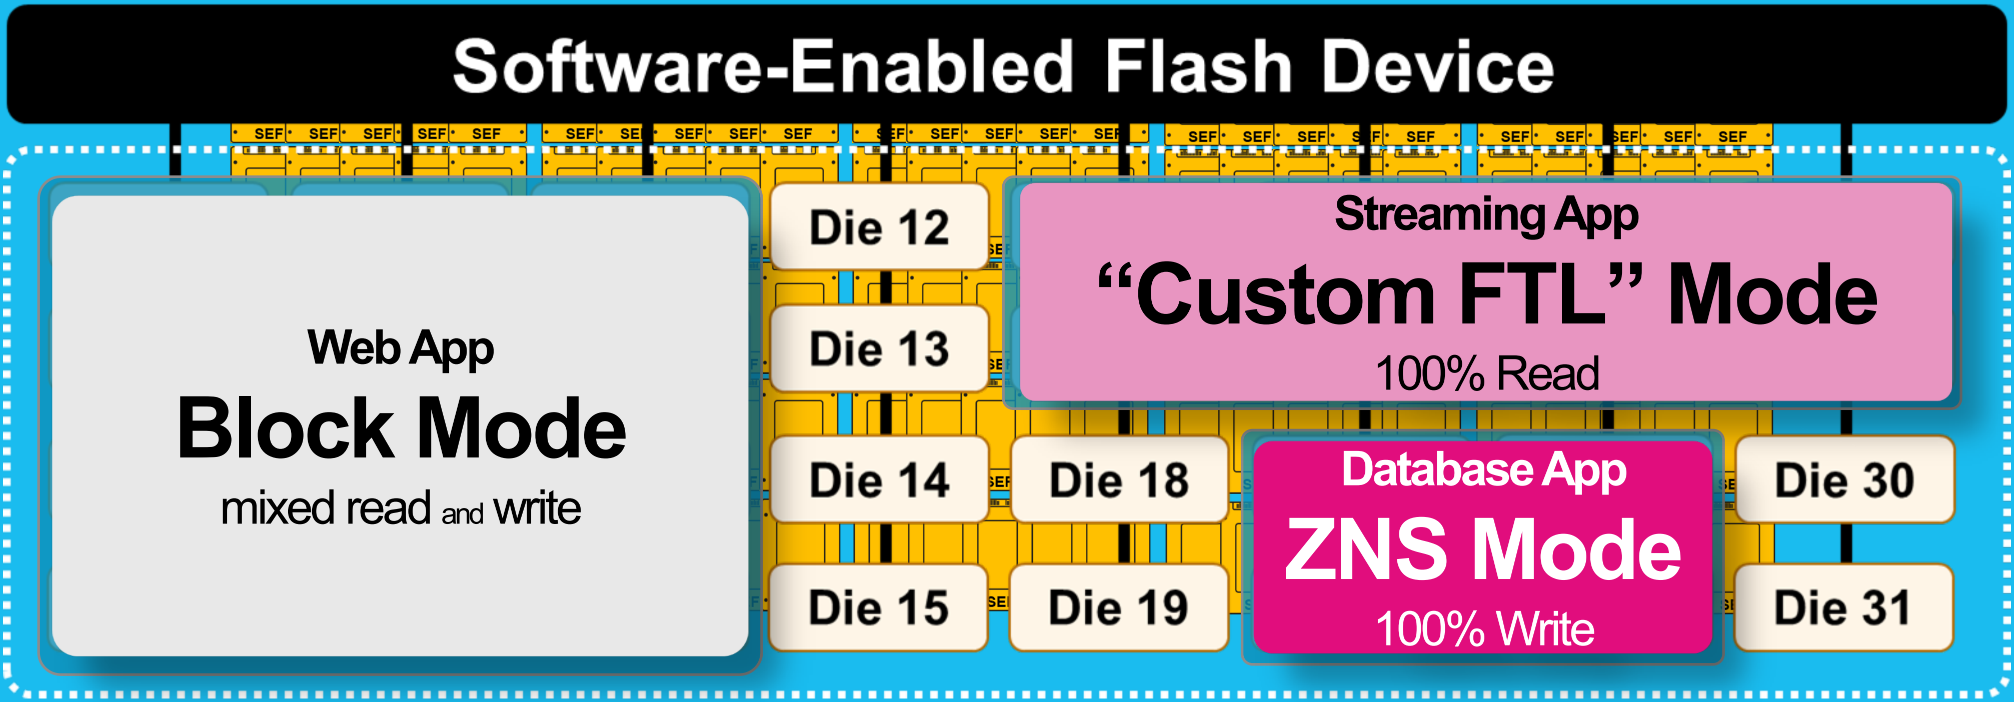

Workloads are physically isolated from each other

© 2021 KIOXIA America, Inc. All Rights Reserved.

Data Placement, Workload Isolation and Multi-Protocol

DEMO

© 2021 KIOXIA America, Inc. All Rights Reserved.

Audience Survey #2

What is the average size of your open-source storage software based storage pools? (pick one answer):

• Greater than 5 PB:	0%
• Between 1PB and 5PB:	0%
• Between 250TB and 1PB:	0%
• Between 50TB and 250TB:	0%
• Less than 50TB:	20%
• Don’t know:	80%

3 Latency Managed with Software Controls

Which is better ?

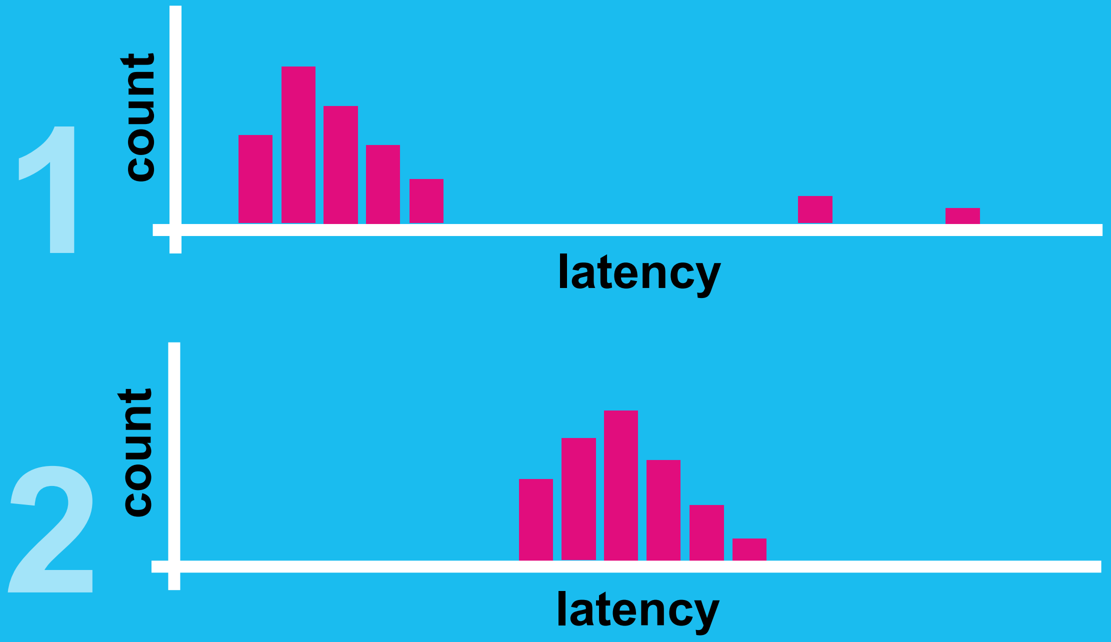

- Latency is driven by application requirement
- Faster IO (lower latency) with some outliers
- Consistent latency (with a little lower IO rates)
- Software-defined means you can tune the SSD to meet need

Manage and Deliver Latency Outcomes

Application control over queuing

Three effective and useful advanced queuing modes

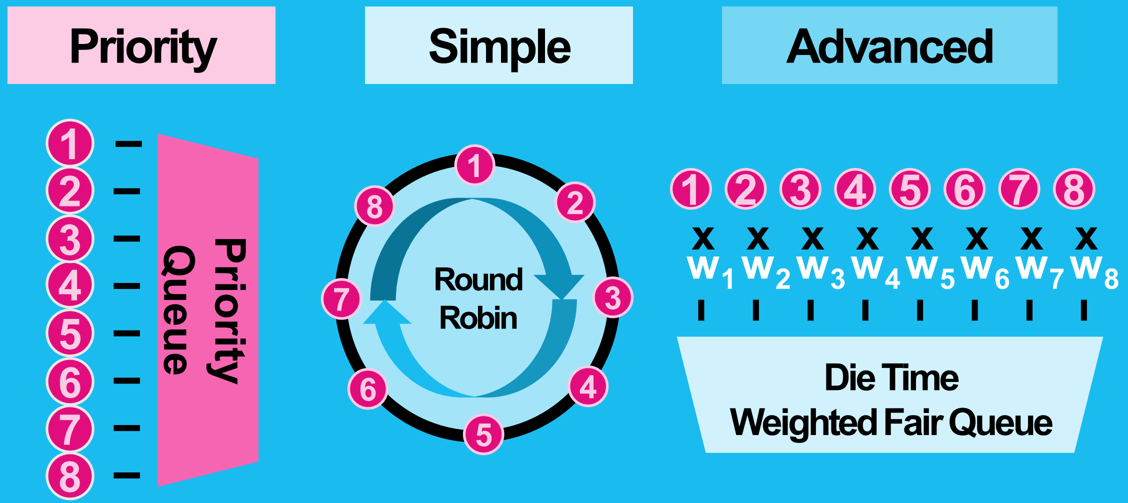

- Each virtual device has 8 FIFO command queues
- Host can specify a specific queue for each flash access command (read, program, erase) per QoS domain
- Scheduler handles suspend/resume for program & erase cmds

© 2021 KIOXIA America, Inc. All Rights Reserved.

Dynamic Latency Outcome Control

DEMO

Quality-of-Service Domains **isolate** and **manage** individual workloads on a shared system

Adjust **workload priorities** and **latency outcomes** as needed to optimize total application responsiveness

© 2021 KIOXIA America, Inc. All Rights Reserved.

MUCH MORE

There are many more “powerful and useful” Software-Enabled Flash™ capabilities to explore…

1. Nameless write
2. Unified write buffers
3. Asynchronous event handlers
4. Lock-less high performance FTL
5. Direct access read
6. Flash health and lifetime extension
7. Abstraction of flash generations
   - Deploy MLC, TLC and QLC media – with code re-use

© 2021 KIOXIA America, Inc. All Rights Reserved.

An Open Standard
for Vendors and Users

OPEN

Working to establish an OSS project for Software-Enabled Flash™ Technology to enable developers, vendors and technology co-travelers to collaborate freely

© 2021 KIOXIA America, Inc. All Rights Reserved.

OPEN SOURCE

SEF Open Source: Objective & Activity

Objective

- Appeals to large data center customers
- Encourages broad adoption
- Enable co-travelers & second sources
- Community membership driven

Activity

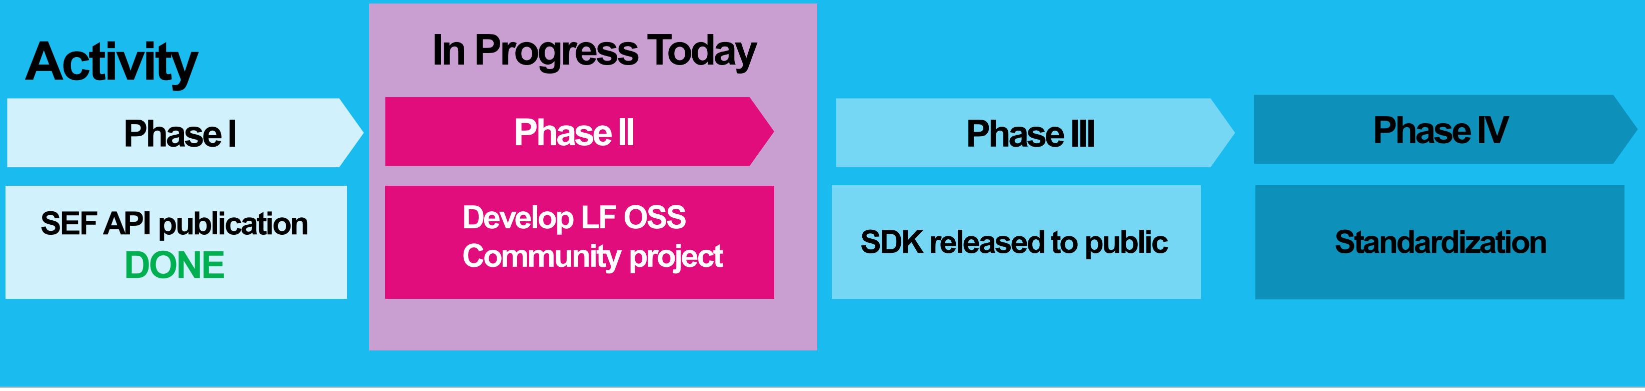

In Progress Today

Phase I

SEF API publication
DONE

Phase II

Develop LF OSS
Community project

Phase III

SDK released to public

Phase IV

Standardization

© 2021 KIOXIA America, Inc. All Rights Reserved.

OSS MODULES

What will be released through the project?

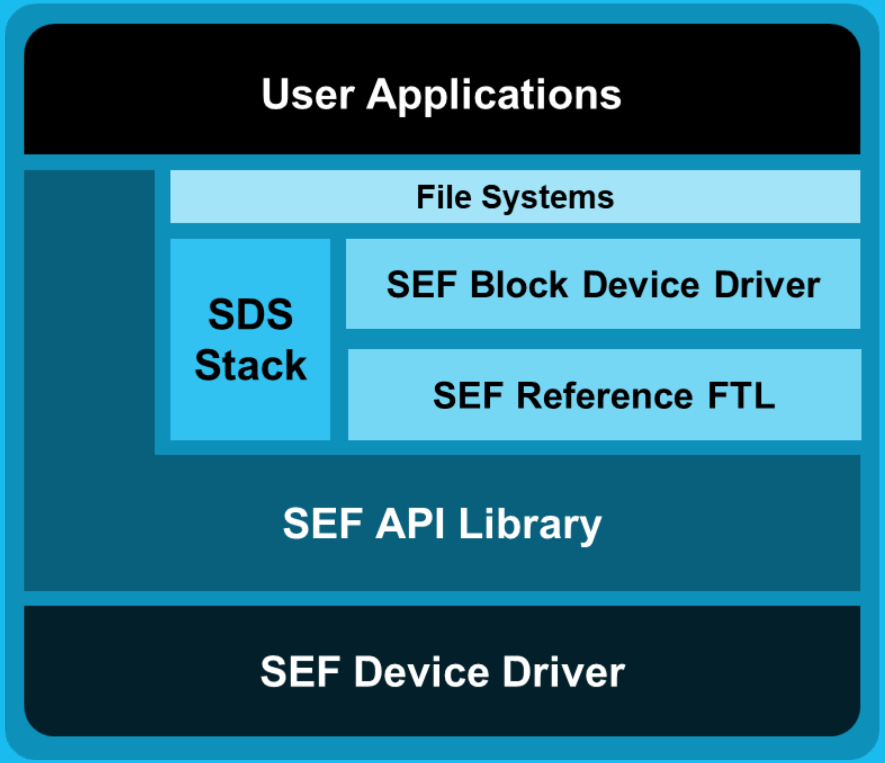

Documentation

- SEF API specification
- SDK document

Software Development Kit

- SEF Library
- SEF reference FTL
- SEF CLI
- SEF FIO
- SEF QEMU Block driver + ZNS
- NVMe™ Driver / IO_URING enhancements

NVMe is a trademark of NVM Express, Inc.

© 2021 KIOXIA America, Inc. All Rights Reserved.

Software-Enabled Flash™ technology

The Software-Defined Flash API

- No need to design new hardware from scratch
- SEF handles the low level flash so you don't have to

Full control over:

- Latency outcomes
- Die & block level data placement control
- Mode switching with multi-protocol support

- Reference source code examples accelerate adoption
- Resulting in TCO and TTM advantages

The **Open Source** SDK enables developers to **evaluate & test** Software-Enabled Flash technologies **without** writing a single line of code

© 2021 KIOXIA America, Inc. All Rights Reserved.

SUMMARY

Shed the legacy HDD paradigm – Software-Defined Delivers:

- Time to market with new flash; improves TCO
- Full control over latency outcomes
- Die & block level data placement control
- Mode switching of personalities and protocols

- A base for rapid development & customization
- Co-travelers supplying compatible controllers
- 2nd source vendors interested
- Potential to improve your TTM and TCO

Define data placement, latency control and device protocol with a Software-Defined API

Open Source project enables developers, second source vendors and industry co-travelers

Accelerating change through software…

© 2021 KIOXIA America, Inc. All Rights Reserved.

Software-Enabled Flash™ technology

a **force multiplier** in data center economics

© 2021 KIOXIA America, Inc. All Rights Reserved.

Want to know more?

INFORM

KIOXIA KIOXIA America Inc.

Corporate Group Information (KIOXIA Holdings Corporation)

R&D / Yokkaichi Plant / Company Information (KIOXIA Corporation)

United States (English)

Business Products Personal Products News About Us Contact Us

KIOXIA's Software-Enabled Flash Technology Now Available on GitHub

Flash Storage Development API Specification Delivers Unprecedented Digital Storage Host Control, Management

Read the Full Press Release >

New Literature

KIOXIA Software-Enabled Flash™ Technology

Tech Brief

How Can Hyperscale Environments Best Enable QLC Flash Memory?

New Video

What if Flash was Software Defined?

LF Live Webinar: What If Flash Was Software Defined?

More Literature and Videos

KIOXIA's Software-Enabled Flash™ Technology

Hyperscalers are looking to evolve and redefine the use of flash in their data centers. What is needed to accelerate this trend are tools that leverage the native speed and flexibility of flash, coupled with the ease of software-defined storage.

Website

www.SoftwareEnabledFlash.com

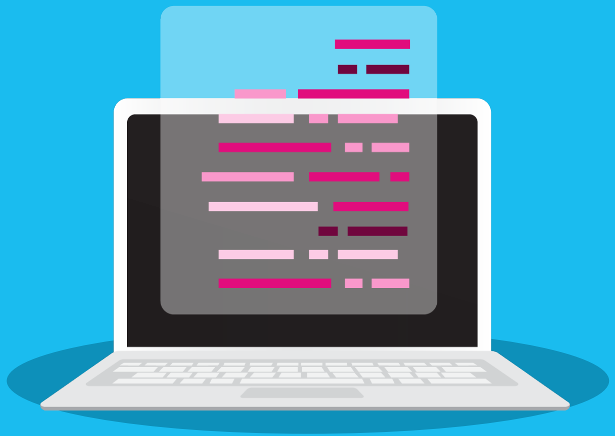

API

https://github.com/kioxiaamerica

Contact

sean.stead@kioxia.com

© 2021 KIOXIA America, Inc. All Rights Reserved.

Audience Q&A

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 1GB = 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.

© 2021 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.

G2M COMMUNICATIONS

Effective Marketing & Communications with Quantifiable Results