



## Data Loss Mitigation with KIOXIA Enterprise and Data Center NVMe™ SSDs

Solid state drives (SSDs) use flash memory cells to store data and have high reliability when compared to hard drives or other forms of memory. However, SSDs can fail over time - just like any electronic device - and when they fail in a data center, a large amount of data can be lost. Depending on IT policies and failure preparedness, the loss of data can be easily recoverable from a RAID configuration, off-site backup or other options. Without mitigation processes in place, data losses can be disastrous especially if intellectual property (IP), customer information and/or financial data is involved.

At the device level, the NVMe specification that enables a host system to communicate with an SSD across a PCIe® bus includes several features for IT administrators to monitor the health of SSDs in their data centers. These features include the ability for NVMe SSDs to issue warnings as each approaches their respective end of life. KIOXIA, one of the leaders in PCIe NVMe SSDs, builds on the safety mechanisms within the NVMe specification with additional features that can extend NVMe SSD life, and enable data access after the drive reaches end of life.

This technical brief will cover these key capabilities within the current NVMe specification and a high-level explanation of the KIOXIA built-in NVMe SSD failure mitigation features.

### Monitoring SSD Health

In a data center, losing a drive or multiple drives without backup, redundancy or data protection can be devastating as customer records, critical business data, IP and other important files could be permanently gone within the blink of an eye, and its reach can affect other data centers and many users globally (Figure 1). SSDs in a fault tolerant configuration (such as RAID) provide essential redundancy to protect data stored on multiple drives.



Figure 1: a failed SSD can negatively affect a data center and its global reach

Before an actual drive failure occurs, the latest NVMe specification includes capabilities to keep IT administrators informed regarding the health of these NVMe drives by issuing advanced warnings that enable personnel to take preventative measures. The specification uses Self-Monitoring Analysis and Reporting Technology (SMART) health logs for these purposes, as well as system notifications called 'asynchronous events' where the SSD notifies the host of status changes, errors and health information. These events are messaged to the host until it takes the necessary steps

to clear the event. Notifications include events related to drive health when it reaches a critical state, as well as SSD health status and error codes provided to the host over the drive's life. Key codes and attributes include, but not limited to:

Attribute Name	Log Address	Description of SMART Health Log Event
<b>Critical Warning</b>	Byte 00 (Bit 0)	The available spare capacity has fallen below threshold (the spare area may be used up so the drive can no longer be written to)
	Byte 00 (Bit 01)	The SSD has reached a temperature lesser/greater than the defined minimum/maximum threshold
	Byte 00 (Bit 02)	SSD reliability has been compromised due to significant media errors
	Byte 00 (Bit 03)	The SSD has been placed in 'Read-only' mode
	Byte 00 (Bit 04)	The volatile memory device is failing
<b>Available Spare</b>	Byte 03	Shows the normalized percentage (0 to 100%) of the remaining spare capacity
<b>Available Spare Threshold</b>	Byte 04	An asynchronous event may occur when the available spare capacity falls below the value specified in this field (can be user set)
<b>Percentage Used</b>	Byte 05	Shows the vendor-specific estimate of the SSD's life based on actual usage and the manufacturer's prediction
<b>Media and Data Integrity</b>	Bytes 175:160	Shows the number of occurrences where the SSD detected an uncorrectable data integrity error

These health monitoring logs in the NVMe specification provide exceptional insight into the health of an SSD, enabling IT administrators to monitor a drive's health directly and take preventative measures (such as replacing an SSD) prior to a drive failure.

## Built-in SSD Failure Mitigation

KIOXIA SSDs feature a variety of built-in failure mitigation capabilities such as wear leveling, customizable end-of-life behavior, die failure recovery and dual-port functionality. These capabilities are standard in KIOXIA enterprise and data center PCIe NVMe SSDs, and include:

Failure Mitigation Feature	Description
<b>Wear Leveling</b>	The KIOXIA developed algorithm is 'Always On' and ensures that the SSD's media is wearing evenly so that premature drive failure can be prevented from overuse in a portion of the drive such as writing constantly to the same logical block address (LBA) range.
<b>Customizable End-of-Life Behavior</b>	KIOXIA enables IT administrators the ability to set up behaviors as the drive reaches specified SSD Life Left thresholds. This includes behaviors such as reducing the performance of the drive, sending asynchronous events to the host, and turning on the 'Read-only' mode. This is a significant benefit to IT administrators who can customize how, when and what actions to take upon reaching thresholds.
<b>Die Failure Recovery</b>	KIOXIA enterprise and data center PCIe NVMe SSDs implement data redundancy at the die level to support die failure recovery. This enables an SSD to sustain a flash memory die failure, recover from it and retain user data integrity. This implementation enables data recovery without the loss of drive functionality or uptime, and is ideal for mission critical applications.
<b>Dual-Port Functionality</b>	KIOXIA enterprise SSDs, such as the CM7 Series and FL6 Series, include dual-port functionality that enables a second path of access to the same drive in case a path failure occurs. This functionality is ideal for storage applications requiring redundancy, high-availability or protection against single-path failure.

*NOTE: KIOXIA data center PCIe NVMe SSDs support single-port only.*

# KIOXIA Enterprise and Data Center PCIe NVMe SSDs

The CM7 Series is KIOXIA's next generation enterprise PCIe NVMe SSD product line. The series features significantly improved performance from PCIe 4.0 to PCIe 5.0, 30.72 terabyte<sup>1</sup> (TB) maximum capacity, dual-port for high availability, 1 Drive Write(s) Per Day<sup>2</sup> (DWPD) models for read intensive applications (CM7-R Series) and 3 DWPD models for mixed use applications (CM7-V Series), up to a 25-watt power envelope and a host of security options.

The CD8 Series is KIOXIA's data center PCIe NVMe SSD product line, which are suitable for scale-out and cloud applications. The series features significantly improved performance from PCIe 3.0 to PCIe 4.0, 15.36 TB maximum capacity, 1 DWPD models for read intensive applications (CD8-R Series) and 3 DWPD models for mixed use applications (CD8-V Series), up to a 20-watt power envelope and a host of security options.

The XD7P Series is KIOXIA's 2<sup>nd</sup> generation E1.S data center PCIe NVMe SSD product line designed to the Enterprise and Datacenter Standard Form Factor (EDSFF) E1.S specification and optimized for the OCP Data Center NVMe SSD specification, version 2.0. The series features 1 DWPD endurance and up to 7.68 TB capacity.

The FL6 Series is KIOXIA's 1<sup>st</sup> generation enterprise Storage Class Memory (SCM) SSD product line that delivers excellent system performance with low latency for applications where response time is critical. This dual-port, PCIe 4.0, NVMe 1.4 compliant SSD series features 60 DWPD endurance, up to 3.2 TB capacity, 29 μs read latency and 8 μs write latency.



**CM7 Series SSDs**

*PCIe 5.0 and NVMe 2.0 Specification Compliant*

**High-Performance**  
 SeqRead = up to 14,000 MB/s  
 RanRead = up to 2.7M IOPS  
 SeqWrite = up to 7,000 MB/s  
 RanWrite = up to 600K IOPS

**Endurance and Capacities**  
 1 and 3 DWPD options  
 1,600 GB - 30,720 GB capacities

**CD8 Series SSDs**

*PCIe 4.0 and NVMe 1.4 Specification Compliant*

**High-Performance**  
 SeqRead = up to 7,200 MB/s  
 RanRead = up to 1.25M IOPS  
 SeqWrite = up to 6,000 MB/s  
 RanWrite = up to 380K IOPS

**Endurance and Capacities**  
 1 and 3 DWPD options  
 800 GB - 15,360 GB capacities

**XD7P Series SSDs**

*PCIe 4.0 and NVMe 2.0 Specification Compliant*

**High-Performance**  
 SeqRead = up to 7,200 MB/s  
 RanRead = up to 1.54M IOPS  
 SeqWrite = up to 4,800 MB/s  
 RanWrite = up to 200K IOPS

**Endurance and Capacities**  
 1 DWPD  
 1,920 GB - 7,680 GB capacities

**FL6 Series SSDs**

*PCIe 4.0 and NVMe 1.4 Specification Compliant*

**High-Performance**  
 SeqRead = up to 6,200 MB/s  
 RanRead = up to 1.5M IOPS  
 SeqWrite = up to 6,200 MB/s  
 RanWrite = up to 400K IOPS

**Endurance and Capacities**  
 60 DWPD  
 800 GB - 3,200 GB capacities

## Summary

Through advanced drive health technologies provided by the NVMe specification, coupled with robust KIOXIA enterprise and data center SSDs, the potential of a surprise drive failure can be lowered significantly. When combined, these technologies can strengthen a data center infrastructure and deliver failure mitigation tools to help prevent against data loss and unnecessary downtime.

For more information on KIOXIA PCIe NVMe SSDs:

Enterprise SSDs: <https://business.kioxia.com/en-us/ssd/enterprise-ssd.html>

Data Center SSDs: <https://business.kioxia.com/en-us/ssd/data-center-ssd.html>

### NOTES:

<sup>1</sup> Definition of capacity - KIOXIA Corporation 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 1Gbit =  $2^{30}$  bits = 1,073,741,824 bits, 1GB =  $2^{30}$  bytes = 1,073,741,824 bytes and 1TB =  $2^{40}$  bytes = 1,099,511,627,776 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, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

<sup>2</sup> Drive Write(s) per Day: One full drive write per day means the drive can be written and re-written to full capacity once a day, every day, for the specified lifetime. Actual results may vary due to system configuration, usage, and other factors

### TRADEMARKS:

NVMe is a registered or unregistered trademark of NVM Express, Inc. in the United States and other countries. PCIe is a registered trademark of PCI-SIG. All other company names, product names and service names may be trademarks or registered trademarks of their respective companies.

### DISCLAIMERS:

© 2023 KIOXIA America, Inc. All rights reserved. Information in this tech brief, including product specifications, tested content, and assessments are current and believed to be accurate as of the date that the document was published, but is subject to change without prior notice. Technical and application information contained here is subject to the most recent applicable KIOXIA product specifications.