KIOXIA

Technical Brief

Understanding TBW[1] versus P/E Cycles[2] in Managed Flash Memory[3]

How to Use TBW Properly to Determine Lifetime Estimation

Managed flash devices driven by the embedded Multi-Media Card[4] (e-MMC) or the Universal Flash Storage[5] (UFS) specifications were first primarily used in smartphone applications and have since expanded to many other applications, including personal computers and servers. Traditionally, PCs and servers have been using HDDs and SSDs for data storage, both of which include endurance metrics based on Terabytes[6] Written (TBW) and/or Drive Writes per Day[7] (DWPD), whereas e-MMC and UFS devices utilize Program/Erase (P/E) cycles.

While TBW and DWPD are useful parameters for comparing SSD reliability in PC and server applications, they may be inaccurate for determining lifetime estimations of e-MMC and UFS devices if the actual customer workload is different from either a JEDEC® defined workload[8] or a vendor-specific workload.

This technical brief addresses potential misunderstandings relating to lifetime estimations of e-MMC and UFS devices, and provides a recommended approach suited to many applications that use e-MMC or UFS memory.

Introduction to TBW and WAF

One option for determining the total endurance of a flash storage device is to multiply its density by its maximum P/E cycle count - commonly referred to as TBW. The estimated TBW result helps to predict the amount of data (in terabytes) that can be written to a drive before it fails. It is based on the predictability of flash memory endurance which decreases linearly as the device is written to over its lifetime.

When calculating TBW, an additional value called the Write Amplification Factor (WAF) needs to be considered, as the higher the WAF is, the lower the TBW will be. WAF is an estimation formula used to help define the ratio of physical writes made to flash memory divided by the logical writes delivered from the host, and is dependent upon the workload or size of data written from the host. In general, small block oriented writes to random logical block address locations will have a higher WAF than larger block writes to a sequential block address.

Terabytes Written

$$\textbf{TBW} = \frac{\textit{Flash Memory P/E Cycles x Flash Memory Density}}{\textit{WAF}}$$

If customers are interested in knowing the TBW of an e-MMC or UFS device, they would need to calculate the WAF for that specific application. If the TBW is published in a product specification, the supplier will generally use either a JEDEC defined workload or a vendor-specific workload. The actual customer workload, however, is the most accurate and recommended for determining the TBW of an e-MMC or UFS device.

How to use TBW: Example of Two Applications with Different Workloads

When TBW is calculated without the WAF considered, it could pose a high risk relating to the accuracy of device lifetime estimation. For example, consider the two use cases below. A commercial surveillance camera engineer would like to use a 128 gigabyte[6] (GB) UFS device from Vendor A for 4K video recording (Use Case 1). An enterprise network security system engineer would like to use the same 128GB UFS device from Vendor A for data logging (Use Case 2). The advertised endurance specifications for the 128GB UFS device from Vendor A include the following:

128GB Device from Vendor A:

For Use Case 1, the surveillance camera must record 180TB over the life of the product, and write sequentially using 16KB chunk sizes. Based on the advertised TBW that uses 16KB chunk sizes, the UFS device from Vendor A *can be used* since TBW = 192TB and will support the camera's requirement of TBW = 180TB as depicted below:

For Use Case 2, the enterprise network security system needs to support 150TB of data written by the host in its lifetime, but unlike Use Case 1, the host is randomly writing data using 4KB chunk sizes. Without taking the chunk size into account, the engineer may assume that 128GB UFS device from Vendor A can be used since the advertised TBW = 192TB (when the chunk size is 16KB) and the TBW is greater than the 150TB required. *This would be an incorrect assumption*.

Since the actual customer workload is mainly random access with smaller 4KB chunk sizes, the engineer must request a calculated TBW value from Vendor A that uses a workload with 4KB chunk sizes. In this scenario, the vendor measured the TBW with a 4KB chunk size and reported a reduction in TBW to 96TB. The reduction is due to the UFS not being able to handle small chunk sizes as well as the larger chunk sizes, and has a larger WAF as well. When those TBW values are compared and based on a correct 4KB workload, it is clear that the engineer *will not be able to use* the 128GB UFS device from Vendor A in the network security system design as shown below:

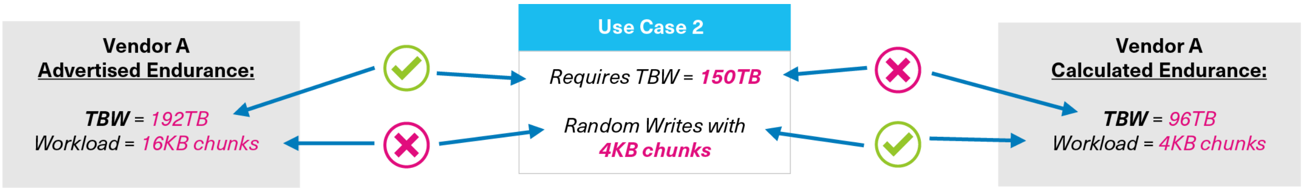

KIOXIA Provides Lifetime Estimation Analysis

KIOXIA can assist customers in this area by performing a lifetime estimation analysis of an e-MMC or UFS device and will make recommendations as necessary tailored to the customer's application.

Summary

KIOXIA offers a free value-added service to its customers by performing a lifetime estimation analysis of its deployed managed flash memory. Customers can take advantage of these services by using this KIOXIA link: https://customer-us.kioxia.com/inquiry/product.

NOTES:

[1] TBW or Terabytes Written measures how many cumulative writes that a drive can expect to complete over its lifespan.

[2] Program/Erase Cycles: The writing of data to one or more pages in an erase block, and the erasure of that block, in either order, as presented in the JEDEC SSD Specifications Explained, slide 8: https://www.jedec.org/sites/default/files/Alvin_Cox%20%5BCompatibility%20Mode%5D_0.pdf.

[3] Managed NAND flash memory integrates a controller with NAND flash memory into one package enabling internal management of memory.

[4] Embedded Multi-Media (e-MMC) is a specification developed by JEDEC. The current release is v5.1, published in February 2015.

[5] The Universal Flash Storage (UFS) v3.1 specification is the current release by JEDEC and published in January 2020.

[6] Definition of capacity - KIOXIA Corporation defines a kilobyte (KB) as 1,000 bytes, 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.

[7] Drive Write(s) per Day (DWPD): 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.

[8] The JEDEC 219 Solid-State Drive (SSD) Endurance Workload is defined in the JEDEC standard and is intended to mirror commonly used applications in either enterprise or client application environments, and may be applied, in some cases, with e-MMC and UFS storage.