

Delivering Improved Performance and Power Efficiency with Next-Generation BG4 Series Client NVMe™ SSDs

Advanced BG4 Series SSD Capabilities in Comparison to the BG3 Series

Yasuyuki Nagahara, Grant Van Patten, Scott Harlin

Introduction

Toshiba Memory introduced its fourth generation of single package ball grid array (BGA) solid-state drives (SSDs), called the BG4 Series, at the Consumer Electronics Show 2019. This new series of client NVMe™ SSDs utilize the company's latest 96-layer BiCS FLASH™ 3D technology and is designed to fuel the future of mobile computing and IoT devices. The BG4 Series delivers better performance and larger maximum capacities than previous NVMe SSDs in this product category¹, and represents one of the smallest removable drives currently available. BG4 Series SSDs include flash memory and a new controller in one encompassing 16mm x 20mm package, enabling larger amounts of flash memory to be added to smaller and thinner devices while extending battery life to improve the mobile experience.

Similar to previous generations, the BG4 Series leverages the Host Memory Buffer (HMB) feature to maintain high-performance without the use of integrated DRAM (Figure 1). The feature uses a portion of host memory to manage flash memory within an SSD, and delivers similar performance as SSDs with DRAM when the feature is turned on. This cost-effective DRAM-less design has resulted in one of the world's thinnest SSDs that delivers a heightened mobile user experience at a fraction of the power requirement when compared to other NVMe-based SSDs.

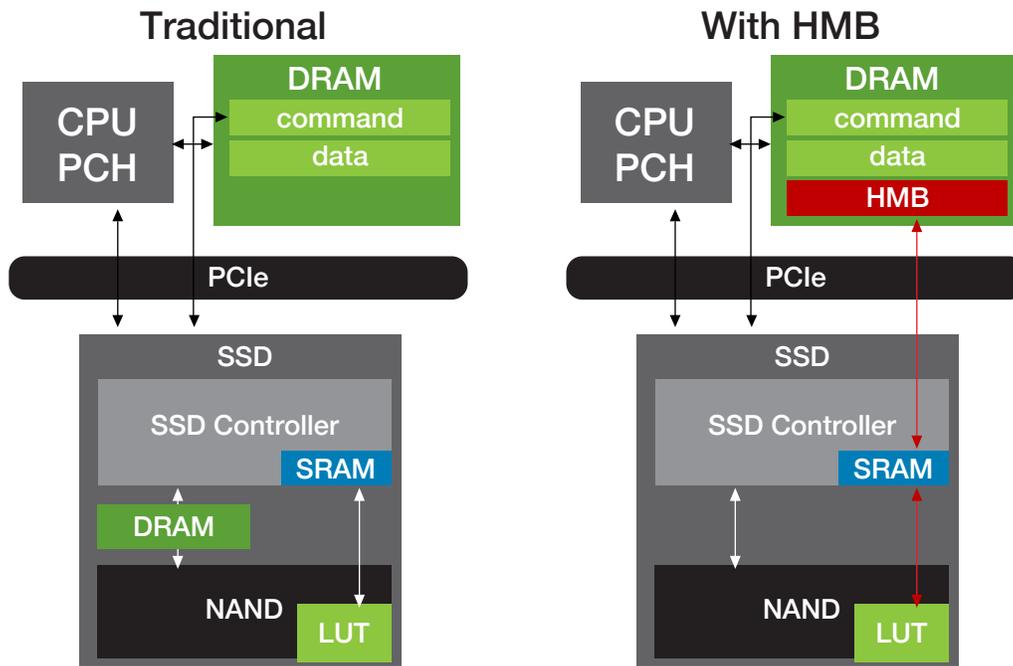


Figure 1 depicts a traditional client SSD, as well as one with HMB included

When compared to its predecessor (the BG3 Series), the BG4 Series delivers significantly faster random read performance and twice the capacity using less power. This technical brief presents improvements to the BG4 Series relating to performance and power efficiency when compared to the BG3 Series.

Product Overview

The BG Series, developed by Toshiba Memory, represents their first product with HMB capabilities, and delivers very fast SSD performance² in support of this feature. The following is a short overview outlining the product capabilities for the BG3 Series and BG4 Series of client NVMe SSDs:

BG3 Series

The BG3 Series was introduced to the industry in August 2017 as Toshiba Memory's third generation of single package BGA SSDs. The series is driven by 64-layer BiCS FLASH 3D technology developed by Toshiba Memory and designed to be smaller, lighter, faster and more power-efficient than previous generations, mostly to enable the evolution of mobile computing and IoT devices. The BG3 Series features HMB within the NVMe revision 1.2.1 specification to deliver high-performance without using integrated DRAM.

Available in 128GB³, 256GB and 512GB capacities, the BG3 Series supports PCI Express® (PCIe®) Gen3 x2 lanes. All three capacity models are available in a surface-mount M.2 1620 single package (16mm x 20mm) or a removable M.2 2230 module (22mm x 30mm) to enable mobile device design flexibility. As real estate in mobile devices can be limited and challenging, the BG3 Series offers very thin SSDs supporting 1.3mm profiles for 128GB and 256GB capacities, and 1.5mm profiles for 512GB capacities.



Based on internal benchmark testing conducted by Toshiba Memory in a lab environment using the CrystalDiskMark⁴ tool (explained further in Section 4, Test Criteria), the BG3 Series with 512GB capacities and two PCIe Gen3 lanes delivered:

- Sequential read performance up to 1,635 megabytes per second (MB/s)
- Sequential write performance up to 482 MB/s
- Random read performance up to 174,000 input/output operations per second (IOPS)
- Random write performance up to 31,000 IOPS
- Performance-power ratio up to 0.97mW divided by MB/s

BG4 Series

The BG4 Series is Toshiba Memory's fourth generation of single package BGA SSDs and announced to the industry in January 2019. BG4 utilizes 96-layer BiCS FLASH 3D technology and also leverages the HMB feature, but within the NVMe Revision 1.3b specification. This product doubles the lane count with support for PCIe Gen3 x4 lanes, and delivers more performance in the same power envelope when compared to other products².

Capacities supported within the BG4 Series include 128GB, 256GB, 512GB and 1024GB (approximately 1TB), with each capacity model available in a BGA M.2 1620 surface-mount design or an M.2 2230 removable module. Similar to the BG3 Series, the BG4 Series represents one of the thinnest SSDs available delivering 1.3mm profiles for models up to 512GB capacities, and 1.5mm profiles for 1024GB capacities.



Based on internal benchmark testing conducted by Toshiba Memory in a lab environment using the CrystalDiskMark tool, the BG4 Series with 1024GB capacities and four PCIe Gen3 lanes delivered:

- Sequential read performance up to 2,336 MB/s
- Sequential write performance up to 1,815 MB/s
- Random read performance up to 349,000 IOPS
- Random write performance up to 56,000 IOPS
- Performance-Power ratio up to 0.78mW divided by MB/s

In a quick side-by-side comparison, the BG4 Series provides significant benefits over the BG3 Series⁵:

Metric	BG3 Series	BG4 Series	BG4 Advantage
PCIe Gen3 Lanes	X2	X4	Double the lanes – less traffic congestion
Maximum Capacity	512GB	1024GB	Twice the SSD density to store data
BiCS FLASH™ Support	64-layers	96-layers	33% improved NAND voltage change ⁶
512GB BGA Z-height	1.5mm	1.3mm	Smaller form factor for mobile devices
Sequential Read	up to 1,635 MB/s	up to 2,336 MB/s	~43% faster
Sequential Write	up to 482 MB/s	up to 1,815 MB/s	~277% faster
Random Read	up to 174,000 IOPS	up to 349,000 IOPS	~100% faster
Random Write	up to 31,000 IOPS	up to 56,000 IOPS	~80% faster
Performance-Power Ratio	up to 0.97mW/MB/s	up to 0.78mW/MB/s	~20% decreased power use

Description of Benchmarks

Benchmark tests were conducted by Toshiba Memory in a lab environment that compared storage performance and the power efficiency between BG Series generations. For this requirement, the BG4 Series and BG3 Series were tested for comparison covering bandwidth (sequential read and write performance), IOPS (random read and write performance), and power efficiency (via performance-power ratios). The performance tests utilized the industry-standard CrystalDiskMark tool that is widely accepted and used to test the performance of solid-state storage by reading and writing through the file system. The power test was developed and implemented using an in-house developed system and software.

Test Criteria

The performance results were based on the highest values in each series (BG3 and BG4), comparing the best performance of both drives. To measure just the performance, and not power-efficiency, the hardware and software equipment used for the benchmark tests included the following:

- **Test PC:** Intel® Core™ i7-7700 processor with dual CPUs, 3.60GHz frequency, four (4) processing cores, and 8GB of DRAM (4GB*2)
- **Operating System:** Microsoft® Windows® 10 version 1803 (Build 17134.112)
- **Target Use Case:** General PC Usage
- **Storage Devices:** Toshiba Memory BG4 Series SSD with 1024GB capacity, PCIe Gen3 x4 lanes and Toshiba Memory BG3 Series SSD with 512GB capacity, PCIe Gen3 x2 lanes
- **Benchmark Software:** CrystalDiskMark 6.0.0 (an open source benchmark tool commonly used to test the performance of solid-state storage)
- **Test Disclaimer:** Test results will vary based on test software and hardware. Normal published results have typically used benchmark software to isolate each performance metric. For the BG4 versus BG3 comparison, CrystalDiskMark 6.0.0 was used to consistently simulate PC usage, delivering ease of repeatability by anyone testing these platforms.

When measuring power efficiency, as outlined in Section 3, Toshiba Memory connected internal equipment to measure power and performance. An internal program was also developed to run read and write operations continuously in order to obtain multiple stable power data for statistical analysis. The hardware and software equipment used for this benchmark test included the following:

- **Test PC:** ASRock™ DeskMini 110, CPU Intel Pentium® processor G4560 @ 3.50GHz, 4GB of DRAM
- **Operating System:** Ubuntu® 14.04 LTS 64-bit

4a Set-up & Test Procedures

The test environment was set-up using the hardware and software equipment outlined above. Dual CPUs within the Intel platform host the Windows application on one and the CrystalDiskMark benchmark tool on the other to avoid contention from either application. Without the two CPU testbeds, the applications could interfere with memory and/or CPU cycles, compromising the test results.

The CrystalDiskMark benchmark tool was used to conduct and determine results from sequential read and write tests, and random read and write tests using settings that were selected from a drop-down menu as outlined in Section 5, Test Results. One BG4 Series SSD with 1024GB capacity and PCIe Gen3 x4 lane support was tested against one BG3 Series SSD with 512GB capacity and PCIe Gen3 x2 lane support. The tests were run with the Host Memory Buffer feature invoked and with an access range of 32GiB⁷.

The BG4 and BG3 SSDs were each tested five times and the highest value was recorded. These tests and analysis of sequential and random read-write performance, as well as power efficiency, help to showcase how BG4 Series SSDs deliver higher access to data and lower power consumption when compared to the previous BG3 generation.

Test Results

The test results encompass five benchmark tests that were conducted with the highest values presented in Chart 1 below. A description of each test is also included.

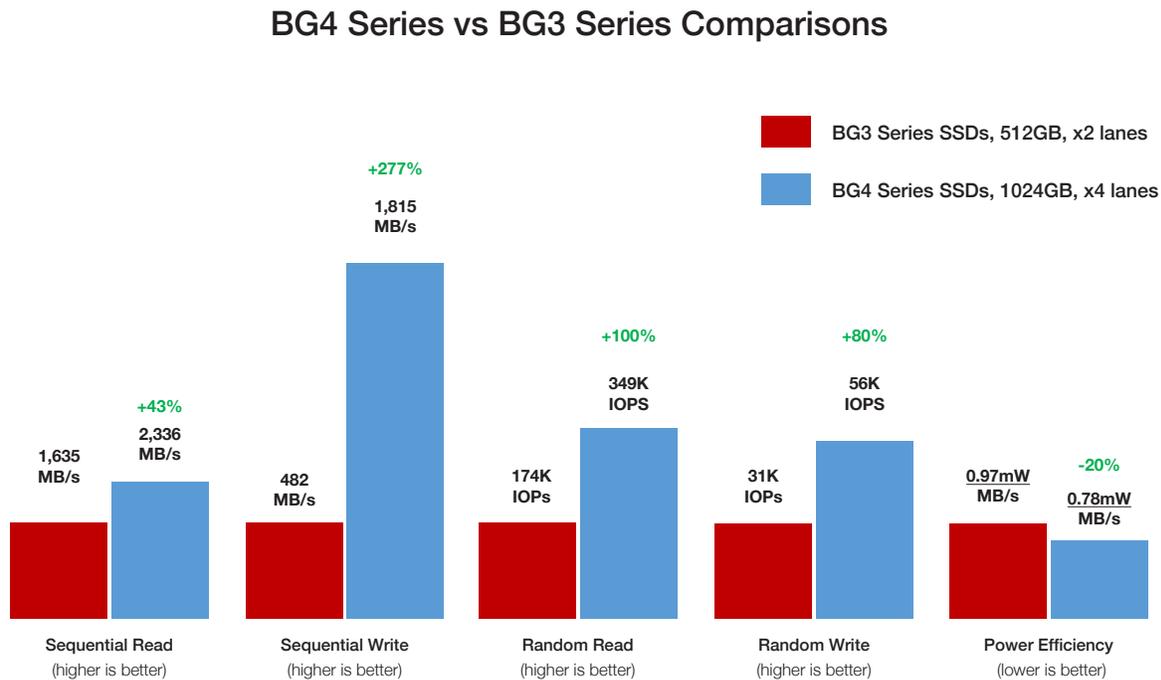


Chart 1: CrystalDiskMark 6.0.0 test results in a Microsoft Windows environment

5a Sequential Read

Sequential read is a data access pattern whereby large contiguous blocks of data, 128KiB⁸ in this test case, are read from adjacent logical block addresses of flash memory in sequential order. The term is used mostly for benchmarking, and speed is measured in MB/s. This kind of performance testing is read-intensive (100% read and 0% write) at the application level (Table 1).

SSD Tested	Highest Result (5 runs)	BG4 Advantage
BG3 Series	1,635 MB/s	
BG4 Series	2,336 MB/s	+43%

Table 1: Sequential read comparison between BG3 Series and BG4 Series SSDs (higher is better)

5b Sequential Write

Sequential write is a data access pattern whereby large contiguous blocks of data, 128KiB in this test case, are written to adjacent logical block addresses of flash memory in sequential order. The term is used mostly for benchmarking, and speed is also measured in MB/s. This type of performance testing is write-intensive (0% read and 100% write) at the application level (Table 2).

SSD Tested	Highest Result (5 runs)	BG4 Advantage
BG3 Series	482 MB/s	
BG4 Series	1,815 MB/s	+277%

Table 2: Sequential write comparison between BG3 Series and BG4 Series SSDs (higher is better)

5c Random Read

4K random read is a data access pattern whereby small blocks of data, 4KiB in this test case, are read from random logical block addresses of flash memory in the device being tested. The term is used mostly for benchmarking, and speed is measured in IOPS. This type of performance testing is read-intensive (100% read and 0% write) at the application level (Table 3).

SSD Tested	Highest Results (5 runs)	BG4 Advantage
BG3 Series	174,000 IOPS	
BG4 Series	349,000 IOPS	+100%

Table 3: Random read comparison between BG3 Series and BG4 Series SSDs (higher is better)

5d Random Write

4K random write is a data access pattern whereby small blocks of data, 4KiB in this test case, are written to random logical block addresses of flash memory in the device being tested. The term is used mostly for benchmarking, and speed is measured in IOPS. This type of performance testing is write-intensive (0% read and 100% write) at the application level (Table 4).

SSD Tested	Highest Results (5 runs)	BG4 Advantage
BG3 Series	31,000 IOPS	
BG4 Series	56,000 IOPS	+80%

Table 4: Random write comparison between BG3 Series and BG4 Series SSDs (higher is better)

5e Power Efficiency

To determine the power efficiency and the respective performance-power ratio between the BG4 Series and BG3 Series, the sequential read and write performance, as well as capacity, were tested with an internally-developed program that ran the read and write operations continuously to obtain stable power data. The power consumed was recorded every 10 milliseconds (ms) over a 500ms range, creating 50 tested values that were condensed down to 3 to 5 primary values in a linear approximation. The maximum power consumption value recorded is the final power result. The recorded values covering both sequential read and write performance over a 500ms range are reflected in Figures 2 and 3 respectively, and linearly displayed using less data points.

In the two power-efficiency charts below, the products are represented as follows:

- **BG3 Series SSDs with two PCIe Gen3 lanes**
- **BG4 Series SSDs with two PCIe Gen3 lanes⁹**
- **BG4 Series SSDs with four PCIe Gen3 lanes**

Sequential Read Power Efficiency

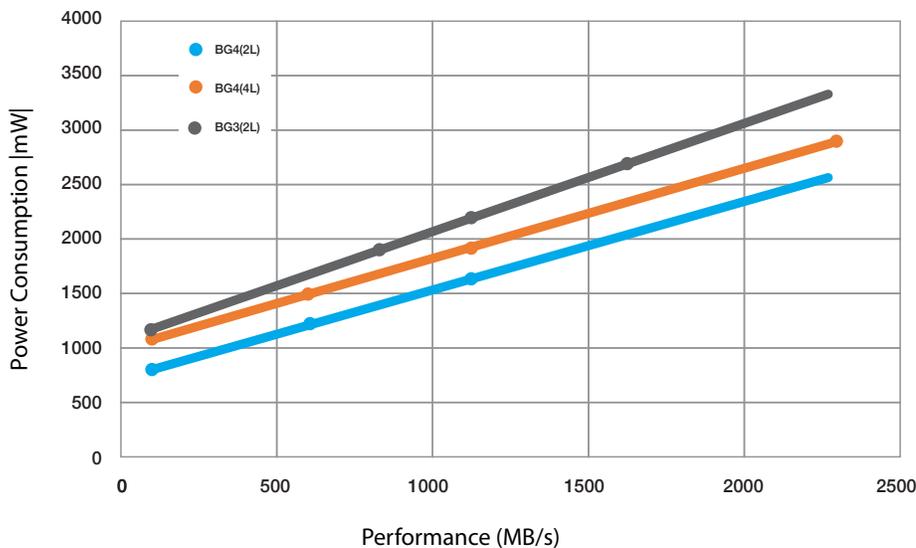


Figure 2: BG4 vs BG3 sequential read power efficiency results

RESULTS:

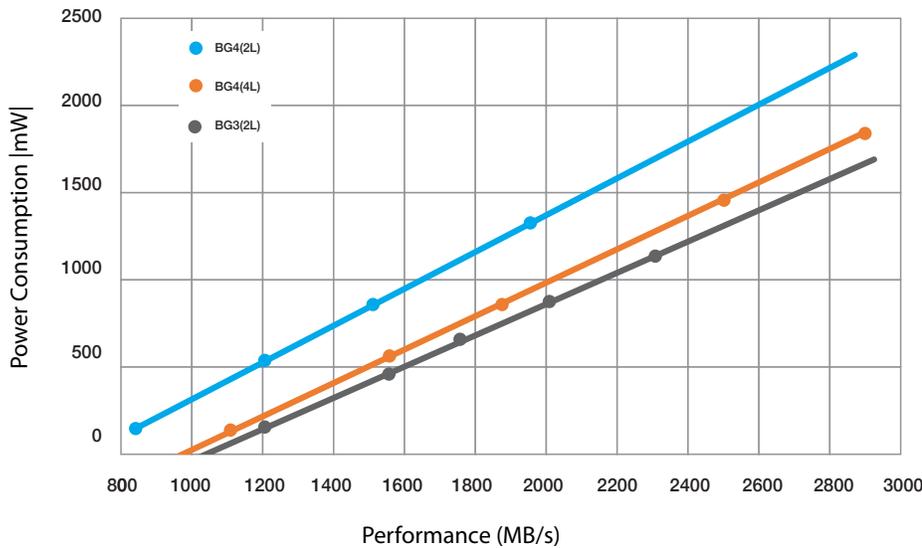
BG3 (2L) = 0.97mW / MB/s

BG4 (2L) = 0.78mW / MB/s

BG4 (4L) = 0.78mW / MB/s

Power consumption between BG3 (2L) and BG4 (2L) at the same performance = 635mW reduced

Sequential Write Power Efficiency



RESULTS:

BG3 (2L) = 1.16mW / MB/s

BG4 (2L) = 0.98mW / MB/s

BG4 (4L) = 1.09mW / MB/s

Power consumption between BG3 (2L) and BG4 (2L) at the same performance = 420mW reduced

Figure 3: BG4 vs BG3 sequential write power efficiency results

Given the test results depicted in Figures 2 and 3 above, power efficiency between the BG4 Series and BG3 Series were tested for sequential read and write performance and capacity against power consumption. The results are depicted in Tables 5 and 6 below as follows:

SSD Tested	Highest Result (5 runs)	BG4 Advantage
BG3 Series	0.97mW / MB/s	
BG4 Series	0.78mW / MB/s	-20%

Table 5: Sequential read performance-power ratio between BG3 and BG4 Series SSDs (lower is better)

SSD Tested	Highest Result (5 runs)	BG4 Advantage
BG3 Series	1.16mW / MB/s	
BG4 Series	0.98mW / MB/s	-16%

Table 6: Sequential write performance-power ratio between BG3 and BG4 Series SSDs (lower is better)

From the results of the benchmark tests, the BG4 Series delivered significant performance and power consumption advantages over the previous generation BG3 Series as follows:

- 43% higher sequential read performance
- 277% higher sequential write performance
- 100% higher random read performance
- 80% higher random write performance
- 20% improvement in the performance-power ratio

Additional Features

The BG4 Series provides unique features when compared to the prior BG3 Series generation. On the surface, the BG4 Series delivers twice the storage capacity, from 512GB to 1024GB, enabling more data to be stored. The series also delivers twice the PCIe lanes which data can travel. Below the surface, the BG4 Series includes such features as power loss signaling, HMB data integrity, and support for the NVMe 1.3b specification including namespace management commands, sanitize and telemetry.

Power Loss Signaling (PLS) is a solution to ensure that SSD integrity is maintained when a power loss event occurs. If a power loss is detected by the subsystem in the host, any acknowledged data that is in transit from the controller to non-volatile memory will be flushed using emergency power from the host. This capability is particularly effective for boot drives that need to ensure data integrity in the event of a system power loss.

The BG4 Series also improves the Host Memory Buffer technology with enhanced reliability that protects against host DRAM failures. It utilizes a hardware module that performs error protection for host DRAM. All HMB data that resides in host DRAM or SRAM support error correction coding which implements one bit error correction and two bits error detection (SECCDED)¹⁰.

In support of the NVMe 1.3b specification, the BG4 Series supports namespace management commands that enable data to be organized as objects as opposed to traditional block or file-based approaches. This type of object storage eliminates the need for a hierarchical structure as it simplifies data access by placing everything in a single namespace. An identifier assigned to each object makes it easy to index or retrieve data, or find a specific object, such as a photo or video. Specification support also enables the BG4 Series to delete the mapping table and erase all blocks that have been written to, using the Sanitize Command.

Additionally, the BG4 Series supports telemetry via the NVMe 1.3b specification by capturing, analyzing and reporting SSD and telemetric data to users. Through this capability, BG4 Series SSDs deliver valuable insights regarding storage and application activities, and provide suggested ways to optimize the system and performance including application workload monitoring, infrastructure troubleshooting, and historical or real-time analysis.

Summary

As demonstrated by the internal benchmark tests, the BG4 Series delivers higher performance and an improved performance-power ratio in one of the industry's densest SSD packages. It improves on all four-corner performance benchmark metrics over the BG3 Series by delivering much higher sequential and random read-write performance while improving the performance-power ratio by 20%.

BG4 Series improvements to HMB technology enhance reliability and protect the SSDs against host DRAM failures through the addition of error correction. Support for the NVMe 1.3b namespace specification enables BG4 Series SSDs to organize data as objects for easier indexing and retrieval. Telemetry support has been added to the BG4 Series enabling storage and application valuable insights as well as suggested ways to optimize the system and performance. BG4 Series SSDs can also permanently erase data through sanitize support.

Delivering higher performance in a smaller physical footprint and thermal envelope enables smaller or more remote mobile device designs, as well as additional confidence that the system will be protected against data corruption or system crashes. By continuously pushing the envelope of what is possible with flash memory and SSD technology, Toshiba Memory continues to arm OEM customers with the technology they need to gain a competitive advantage. Targeted for boot applications, thin-light PC and embedded markets, BG4 Series SSDs are elevating the end-user experience with upgraded storage capacity and significantly improved performance, while continuing to make power efficiency a priority.

For more information, visit:

<https://business.toshiba-memory.com/en-us/product/storage-products/client-ssd.html>.

Disclaimer: Toshiba Memory America, Inc. (TMA) may make changes to specifications and product descriptions at any time. The information presented in this white paper is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Any performance tests and ratings are measured using systems that reflect the approximate performance of Toshiba Memory Corporation (TMC) products as measured by those tests. Any differences in software or hardware configuration may affect actual performance. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to any changes in product, component and/or hardware revisions, new model and/or product releases, software changes, firmware changes, or the like. TMA assumes no obligation to update or otherwise correct or revise this information.

TMA makes no representations or warranties with respect to the contents herein and assumes no responsibility for any inaccuracies, errors or omissions that may appear in this information.

TMA specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. In no event will TMA be liable to any person for any direct, indirect, special or other consequential damages arising from the use of any information contained herein, even if TMA is expressly advised of the possibility of such damages.

Notes

¹ Based on published performance and capacity specifications from competitive products and publicly available on vendor websites.

² Based on published performance specifications from competitive products and publicly available on vendor websites.

³ Definition of capacity: Toshiba Memory 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 1GB = 2³⁰ bytes = 1,073,741,824 bytes, 1TB = 2⁴⁰ 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.

⁴ CrystalDiskMark is an open source disk drive benchmark tool for Microsoft® Windows®. Based on Microsoft's MIT-licensed Diskspd tool, this graphical benchmark is commonly used for testing the performance of solid-state storage. It works by reading and writing through the file system in a volume-dependent way.

⁵ The data cited in the comparison chart are measurement results as tested by Toshiba Memory and do not guarantee similar competitive performance results as published in publicly-available vendor content.

⁶ BG3 Series NAND voltage change = 1.8V. BG4 Series NAND voltage change = 1.2V. Percentage change is 33%.

⁷ A gibibyte (GiB) is a unit of measure to determine capacity used in computing. The prefix gibi originated with the binary system for measuring data capacity and based on powers of two. One gibibyte equals 2³⁰ or 1,073,741,824 bytes = 1024 mebibytes.

⁸ A kibibyte (KiB) is a unit of measure to determine capacity used in computing. The prefix kibi originated with the binary system for measuring data capacity and based on powers of two. One kibibyte equals 2¹⁰, or 1024 bytes.

⁹ The two-lane power data cited in the comparison chart are power efficiency measurement results as tested by Toshiba Memory and do not guarantee similar competitive performance results as published in publicly-available vendor content.

¹⁰ SECCDED is terminology that stands for Single-bit Error Correction, Double-bits Error Detection.

ASRock is a trademark of ASRock, Inc. Intel, Pentium and Intel Core are registered trademarks or trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Microsoft, Windows and Windows Server are registered trademarks of Microsoft Corporation in the United States and/or other countries. NVMe is a trademark of NVM Express, Inc. PCI Express and PCIe are registered trademarks of PCI-SIG. Ubuntu is a registered trademark of Canonical Ltd. All other trademarks or registered trademarks are the property of their respective owners.