

THNSNJ480PCS3 480 Gbyte Recording Capacity 2.5-inch Equipped with SATA Interface for Enterprise Use

●SAKAI Yoshimichi ●OHYAMA Ryuichiro ●MATSUDA Yoshiharu

The increasing volumes of data accompanying the rapid expansion of cloud services and big data in recent years have given rise to the need for data center systems with higher speed and larger capacity in order to efficiently provide network services. Solid-state drives (SSDs) have come into widespread use as storage devices for servers and storage systems in data center systems due to their superior memory access performance compared with hard disk drives (HDDs).

In response to these circumstances, Toshiba has developed the THNSNJ480PCS3, a 2.5-inch enterprise SSD with a capacity of 480 Gbytes equipped with a Serial Advanced Technology Attachment (SATA) to realize easier handling, as an addition to its lineup of high-end enterprise SSDs incorporating Serial Attached SCSI (Small Computer System Interface) (SAS) technology.

> 1. Introduction

In recent years, the volumes of data handled by cloud computing, big data analytics, the Internet of Things (IoT) and other Internet-based solutions have grown explosively. Servers and storage systems are becoming more important because of the need to handle ever greater amounts of data efficiently. Indeed, servers and storage systems will drive the future evolution of the information society. So far, enterprise servers and storage systems have relied on HDDs that have advantages over SSDs in terms of cost per byte and capacity. However, users who prioritize the total cost of ownership (TCO) demand SSDs that excel not only in access performance but also in power efficiency (data rate per watt)⁽¹⁾.

While some enterprise users who have been using SSDs

favor a highly reliable SAS interface⁽²⁾, many prefer SSDs with a SATA interface that combines cost effectiveness and high reliability. In particular, start-up enterprises and fast-growing data centers tend to choose a SATA interface, which is widely used in PCs, emphasizing the balance between power efficiency and cost.

To address these market needs, Toshiba has developed 2.5-inch enterprise SSDs with a SATA interface that incorporates an original controller featuring a powerful error correction code (ECC). The new SATA interface SSDs are available in capacities of 480 Gbytes, 240 Gbytes and 120 Gbytes⁽¹⁾.

This article provides an overview of the THNSNJ480PCS3 480-Gbyte SATA interface SSD (Figure 1) as well as the underlying technologies used to achieve high performance and reliability.




Figure 1 THNSNJ480PCS3 480 Gbyte recording capacity 2.5-inch SSD for enterprise use —

The enterprise SSD provides higher performance than client SSDs. Despite a greater component count, the 2.5-inch SSD has a thickness of only 7 mm owing to high-density board assembly.

> 2. Overview of the THNSNJ480PCS3

Table 1 shows the main specifications of the THNSNJ480PCS3.

The THNSNJ480PCS3 is Toshiba's first-generation enterprise SSD with a SATA interface. This enterprise SSD consists of multi-level cell (MLC) NAND flash memories fabricated using a 19-nm process and a controller with powerful ECC that Toshiba originally developed for client SSD applications⁽³⁾. Its host interface is compliant with SATA Generation-3.

Although the THNSNJ480PCS3 has a total NAND flash capacity of 512 Gbytes, the published user area specification is 480 Gbytes because of a pre-allocated area for overprovisioning. On an SSD, garbage collection frees up storage space by collecting stale (invalid) pages scattered over blocks into new physical data blocks. Since the overprovisioning space is an extra free area on a drive, it helps improve the efficiency of garbage collection, which in turn improves sustained random write performance (when

a random-write workload is in full saturation). While the access performance in the initialized state is important for client SSDs, the sustained performance of continuous write operations is a prime concern for enterprise SSDs. Therefore, the THNSNJ480PCS3 provides overprovisioning space. Furthermore, the THNSNJ480PCS3 incorporates power backup technology in order to supply backup power temporarily in the event of a sudden drop in supply voltage during a data write, thereby maintaining the integrity of data received from a host computer.

Table 1 Main specifications of THNSNJ480PCS3

Characteristics		Specification
Storage capacity		480 Gbytes
NAND process		19 nm
Host interface		SATA Generation-3 (6 Gbits/s)
Access performance	Sequential reads (Mibytes/s)	500
	Sequential writes (Mibytes/s)	400
	Random reads (kIOPS)	75
	Random writes (kIOPS)	12
Power efficiency (kIOPS/W)		18.8
Data reliability (bit error rate)		1×10^{-17}

MiB: mebibyte = 2^{20} bytes
IOPS: Input Output per Second

> 3. Performance Improvement

Generally, if a host computer continuously performs random write operations to an SSD, the input/output operations per second (IOPS) value gradually decreases and levels off to a sustained level of performance (Figure 2). The sustained random-write performance is the minimum IOPS, which is an important measure of performance for enterprise SSDs. The THNSNJ480PCS3 contains a controller that was originally designed for 19-nm NAND flash memories for client SSD applications. The new enterprise SSD uses the methodologies described below in order to improve the sustained IOPS performance for 4-Ki random block writes (kibi = 2^{10} ; 4 Ki = 4 Kbytes). Conventional SSDs maintain a lookup table (LUT) that




Figure 2 Performance trend in continuous random write operation —

Random write operations cause the IOPS value to gradually decrease and level off to a sustained level of performance.

maps logical block addresses (LBAs) from a host computer to physical addresses on a NAND flash array, and keep the mapping table managed by the LUT on a NAND flash memory. Client SSDs that do not have a dynamic random access memory (DRAM) use a static random access memory (SRAM) embedded in the SSD controller as a cache for the mapping table.

Unlike conventional SSDs, the THNSNJ480PCS3 has a discrete large DRAM chip that serves as a cache for the mapping table. Using DRAM as a mapping table cache helps improve the cache hit rate over client SSDs, shortening overall access time. Additionally, the higher cache hit rate helps reduce the number of accesses to the DRAM that holds a mapping table, making it possible to allocate memory bandwidth to accesses to the NAND flash array.

Consequently, the THNSNJ480PCS3 delivers a sustained 4-Kbyte write performance of 12,000 IOPS, compared to the 5,000 IOPS of client SSDs.

> 4. High-Density Board Assembly

Although a casing with a thickness of 9.5 mm was previously commonly used for 2.5-inch HDDs and SSDs, demand for a 7-mm casing has been increasing in recent years. The THNSNJ480PCS3 contains eight NAND flash memories, and a supercapacitor to maintain data integrity in the event of a sudden drop in supply voltage. The THNSNJ480PCS3, which also contains a backup power supply circuit, has a higher bill of materials than client SSDs. Soldering electronic components on both sides of a printed circuit board (PCB) helps increase the flexibility for component placement, but at the expense of heat shock resistance (i.e., reliability) and assembly productivity. To avoid these downsides, Toshiba aimed to achieve single-sided PCB assembly. Conventional through-hole PCBs could not be used for single-sided assembly because of their low space utilization factor for components and signal traces. Therefore, Toshiba utilized a built-up PCB that provides higher flexibility for component placement and trace routing. An eight-layer PCB was chosen to achieve high routing efficiency (Figure 3).




Figure 3 Cross-sectional outline of built-up multilayer printed circuit board (PCB) —

The eight-layer built-up PCB makes it possible to form vias between layers, significantly increasing the flexibility for component placement and trace routing on the top layer.

Because the use of a built-up PCB made it possible to minimize the number of vias on the top surface, it also improved the efficiency of component placement. Additionally, the built-up PCB also helped reduce the number of through-holes across the board, significantly improving the routing flexibility on each layer. Furthermore, the use of physically small resistors and capacitors (0.6 mm × 0.3 mm) helped reduce overall component footprint. As a result of the foregoing, single-sided PCB assembly has been achieved.

> 5. Thermal and Shockproof Solutions

5.1 Thermal Design

One of the challenges that must be solved to achieve high-density board assembly is prevention of the degradation of the SSD reliability due to the heat generated by the controller and other on-board components.

High-density board assembly means a SSD controller that dissipates considerable heat must be placed close to a NAND flash array and a supercapacitor that has narrow operating temperature ranges. Therefore, it is necessary to remove the heat dissipated by the controller and reduce the amount of heat that propagates to the NAND flash array and the supercapacitor. As a solution, the THNSNJ480PCS3 was designed to conduct the controller heat to the SSD casing in order to remove heat through its entire surface using the fan air of a system in which the SSD will be installed. Toshiba performed thermal flow analysis to simulate temperature profiles (Figure 4) to create a casing design suitable for heat removal and to determine the specifications for and the amount of thermal conductive sheet required.




Figure 4 Results of simulation of temperature profiles of PCB assemblies and outer cases obtained by thermal flow analysis —

The new SSD exploits system fan air to eliminate controller-generated heat from the SSD casing, thereby reducing the temperature rise of the neighboring NAND flash array and supercapacitor.

5.2 Shock and Vibration Resistance

Since there has been user demand for enterprise SSDs that have higher shock and vibration resistance than HDDs, we decided on the following specifications:

Shock resistance (Op.): 9,800 m/s² (1,000 G), pulse width = 0.5 ms

Shock resistance (Non-op.): 9,800 m/s² (1,000 G), pulse width = 0.5 ms

Vibration resistance (Op.): 21 m/s² (2.17 Grms), 100 to 800 Hz

Vibration resistance (Non-op.): 159 m/s² (16.3 Grms), 20 to 2,000 Hz

SSD failures due to the application of shock and vibration primarily occur when considerable bending of the PCB has caused a solder joint of the controller or a NAND flash memory to crack, resulting in an open circuit.

To prevent this problem, we have attached heat conductive sheets to both sides of the PCB in order to immobilize it inside the casing, thus reducing its bending. In order to meet the specification requirements, stress distributions due to the application of shock were simulated using nonlinear analysis (Figure 5) to find a casing shape that effectively suppresses the bending of the PCB.

This structure has an added benefit in that it further enhances heat removal because it helps dissipate heat from both sides of the PCB.




Figure 5 Results of simulation of stress distributions at time of application of shock obtained by nonlinear analysis —

The PCB is immobilized from both sides in the SSD casing in order to reduce its bending.

> 6. Data Integrity Technology

The controller in the THNSNJ480PCS3 incorporates Toshiba's original ECC technique called Quadruple Swing-By Code (QSBC). In addition to this robust ECC capability, the THNSNJ480PCS3 combines wear leveling (a function to equalize program/erase cycles across the NAND flash memory) and background patrol (a function to check data integrity and correct errors as they happen). Consequently, the THNSNJ480PCS3 provides 24-hour operation for five years and daily data rewriting of the entire drive capacity.

Furthermore, for enterprise applications, it is crucial to maintain the integrity of data received from a host computer in the event of a sudden drop in supply voltage during active operation. To ensure that all data from the host computer is written to a NAND

flash memory, the backup power supply circuit has a large-value electric double layer capacitor (Figure 6) When the supply voltage has dropped without any notification from the host computer, the SSD switches to the backup power supply to maintain data integrity.

As one of the measures to achieve long-term reliability of five years, the supercapacitor is placed farthest from the controller that generates the most heat among the on-board components.




Figure 6 Component mounting side of PCB —

The SSD incorporates a capacitor and a back power supply circuit to maintain data integrity in the event of a sudden drop in supply voltage.

> 7. Conclusion

The demand for enterprise SATA interface SSDs is expected to grow further to meet expanding data center requirements. Higher-capacity SSDs are also required to handle increasing amounts of data. Toshiba will continually develop new SSDs in a timely fashion using each new generation of high-capacity NAND flash memory and launch them in the expanding SSD market.

[1] Definition of capacity: Toshiba 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 = 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.

References

- (1) MK4001GRZB Solid-State Drive for Enterprise Use Achieving High Performance and High Reliability
TOSHIBA Review 66, 8, 2011, pp. 40-43
- (2) 1.6 Tbyte SSD for Enterprise Use Applying MLC NAND Flash Memory
TOSHIBA Review 68, 9, 2013, pp. 46-48
- (3) High-Performance Control Technologies for SSD Using 19 nm-Generation NAND Flash Memory for PC Market
TOSHIBA Review 67, 12, 2012, pp. 43-46



SAKAI Yoshimichi
Memory Div.



HOYAMA Ryuichiro
Memory Div.



MATSUDA Yoshiharu
Storage Products Div.