Trend

Trends in Storage Product Technologies and Toshiba's Approach

● YANAGI Shigenori ● OHSHIMA Shigeo ● ITO Jun

Toshiba has remained at the forefront of the storage industry for more than 40 years and offers a wide range of storage technologies, including hard disk drives (HDDs), NAND flash memories, and solid-state drives (SSDs), as well as storage solutions applying these technologies. In the present era, often referred to as the "big data era," digital technologies are showing unprecedented expansion and storage products have become one of the most vital components of society.

We will continue to support the broad spectrum of requirements for storage products and solutions by developing and implementing various technologies to improve device storage capacity, processing performance, reliability, and availability.

Expanding client SSD market driven by lightweight and small SSDs

Characteristics of SSDs

An SSD is a storage device that uses NAND flash memories to store data. Generally, SSDs have an electronic interface compatible with HDDs.

Table 1 briefly describes the advantages of SSDs over HDDs. The main characteristics of SSDs are as follows:

(1) Since SSDs have no mechanical (moving) components, they are more resistant to shock and vibration than HDDs.

(2) Since SSDs access data electronically without any mechanical activity, they have less latency than HDDs that cause mechanical overhead such as head seek times.

(3) SSDs mainly consist of NAND flash memories and other semiconductor devices and thus provide greater flexibility in form factor design than HDDs. This makes it easier to reduce the size and weight of SSDs.

At present, SSDs are broadly divided into two categories according to target applications: client SSDs for notebook PCs and other consumer electronics, and enterprise SSDs for use in servers and other high-end equipment.

At first, client SSDs were used as replacements for HDDs, primarily in notebook PCs. With the increasing capacity of NAND flash memories due to the evolution of process technology, SSDs are now finding their way into various applications.

Because of their small size and weight, SSDs are commonly used in mobile devices, including notebook PCs, tablets and camcorders, and because of their high shock and vibration resistance, SSDs are often selected for use in in-flight entertainment systems. In recent years, data centers have also been driving the demand for SSDs because of their high data throughput.

Table 1. Main advantages of SSDs compared with HDDs

Characteristic	Description
High read/write performance	SSDs access data electronically at high speed without any mechanical activity.
Shock and vibration resistance	Since SSDs have no moving components, they are not affected by physical shocks.
Low power consumption	Since SSDs have no motors, they consume less power and thus generate less heat per data throughput.
Small and lightweight	With no mechanical restrictions, SSDs provide greater flexibility in form factors; it is easier to reduce the size and weight of SSDs.
Silent operation	Since SSDs have no moving components, they operate silently.

Evolution of form factors

As described above, SSDs provide great flexibility in form factor design, which makes it possible to reduce the size and weight of SSDs. **Figure 1** shows the evolution of client SSDs in terms of form factors.

Since early SSDs were designed as drop-in replacements for traditional HDDs, most came in the same 2.5-inch (or 3.5-inch) form factor. These SSDs could have been made smaller, but manufacturers chose to make SSDs physically compatible with HDDs.

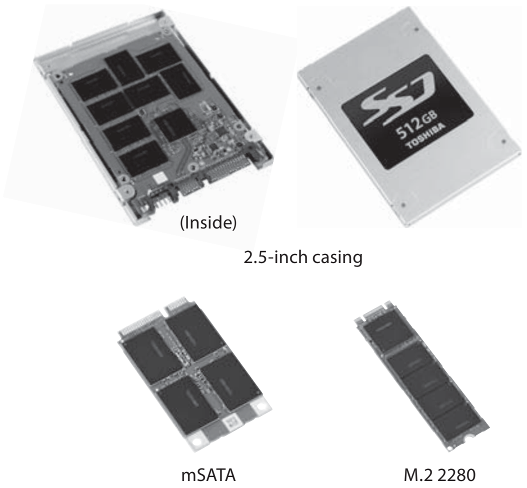

Form factor	Width (mm)	Length (mm)	Height (mm)
2.5-inch casing	69.85	100	7
mSATA	30	50.95	3.95
M.2 2280	22	80	2.3 (single-sided) 3.5 (double-sided)

Figure 1. Trends in reduction of form factor of client SSDs.
In addition to SSDs in 2.5- and 3.5-inch form factors compatible with HDDs, SSDs with card form factors such as mSATA and M.2 are now available.

Later, SSDs with a card form factor called mini Serial Advanced Technology Attachment (mSATA) appeared. mSATA SSDs consist of NAND flash memories mounted on a small printed circuit board (PCB) and look like DRAM modules. mSATA SSDs have contributed to reducing the size and thickness of notebook PCs and other SSD applications. In recent years, a long thin card form factor called M.2 has been standardized and used in various small devices. Smaller SSDs provide an additional benefit in that mobile devices can allow more space for a larger battery to increase the run-time.

Client SSDs are expected to become even smaller and lighter. **Figure 2** shows two SSDs developed by Toshiba: a single-package SSD, as well as the M.2 2230 SSD, a PCB-based module with a single-package SSD. These SSDs were showcased at the Consumer Electronics Show (CES) in Las Vegas, U.S.A, in January 2015. Housed in a ball grid array (BGA) measuring 16 mm × 20 mm × 1.65 mm, the single-package SSD provides a maximum capacity of 256 Gbytes. It is expected to be selected for use in various applications, including mobile devices (see pages 21–24).

Evolution of host interfaces

SSDs have no mechanical (moving) components and thus cause no mechanical overhead like HDD seek times. Therefore, SSDs deliver high data rates. In order to obtain the best performance from the fast internal

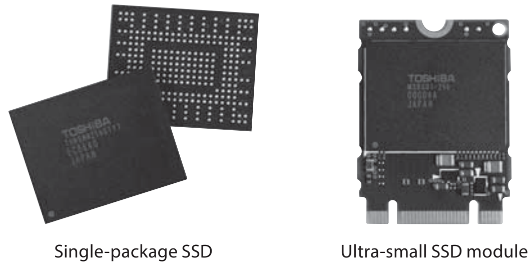

Form factor	Width (mm)	Length (mm)	Height (mm)	
			128 Gbytes	256 Gbytes
Single-package SSD	16	20	1.40	1.65
Ultra-small SSD module	22	30	2.20	2.45

Figure 2. Single-package SSD and ultra-small SSD module.
These SSDs were showcased at 2015 International CES. Client SSDs will continue to become smaller and lighter and will be used in mobile and many other applications.

data rates, manufacturers have developed faster host interfaces to achieve higher I/O throughput.

Figure 3 shows the recent evolution of host interfaces for client SSDs. Since early SSDs were primarily designed as HDD replacements, SSDs generally have a host interface compatible with HDDs. Around 2010, the most commonly used interface was SATA Gen. 2 with a specified transfer rate of 3 Gbits/s. The next revision of SATA, Gen. 3, runs at double the transfer rate (6 Gbits/s) and is most widely used in currently available SSDs. Recently, SSDs with an even faster interface called PCI Express(†) (PCIe(†)) have appeared and found initial use in high-end applications requiring higher performance. PCIe(†) allows the transfer rate to be increased not only by increasing the operating frequency but also by increasing the number of data lanes. PCIe(†) is expected to supersede SATA in the years ahead.

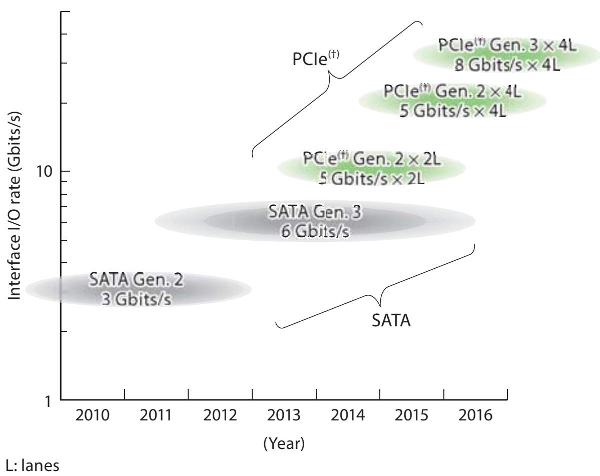

Figure 3. Trends in interface of client SSDs.
At present, SATA is the most commonly used interface for client SSDs. It may be replaced by higher-speed PCIe(†) in the future.

Technical trends in flash storage for enterprise use

Opportunities for flash-based storage in the era of the data deluge

One forecast says that the volume of data generated worldwide will grow to 44 zettabytes (zetta: 10^{21}) by 2020[(1)] (**Figure 4**). However, it will be impossible to store all that data with the estimated total storage capacity available in the world in 2020, considering the projected capital investment in plants that manufacture magnetic disks for HDDs and NAND flash memories for SSDs. Valuable information obtained by analyzing huge amounts of data is an essential ingredient of big-data businesses. In order to achieve successful big data analytics, it is necessary to store large volumes of data temporarily and analyze them in real time. Thus, the amount of computing required is also likely to explode. This implies that the metrics employed to measure storage performance will undergo a fundamental change, moving from a capacity-oriented storage model to a model that prioritizes the input/output (I/O) throughput and processing capacity. Under these circumstances, storage devices with one of the following characteristics are needed:

(1) Storage devices with high-speed I/O and memory suitable for storing a huge volume of data temporarily to process them
(2) Storage devices with high-speed computing performance specifically designed to analyze a huge volume of data
(3) Storage devices with flexibility in scaling up storage according to the changing and increasing requirements, as well as high-speed I/O and memory

Because of high-density, high-capacity and random-access performance advantages, NAND flash memories are suited to satisfy these requirements.

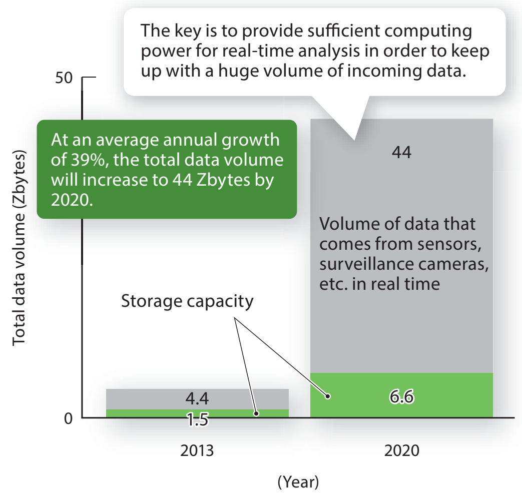

* Created based on "The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things, Sponsored by EMC (2014-04)"[(1)], from IDC's Digital Universe

Figure 4. Dramatically increasing gap, global data volume vs. storage capacity.
The volume of data generated worldwide is predicted to grow to 44 Zbytes by 2020. The storage industry is facing the challenge of developing new technologies to process and analyze such a huge amount of data in real time.

IT market segments and their requirements for SSDs

The information technology (IT) market is broadly divided into three segments, each of which has different requirements (**Figure 5**). The requirements for on-premises deployments include high reliability, availability and maintainability. On-premises systems use the

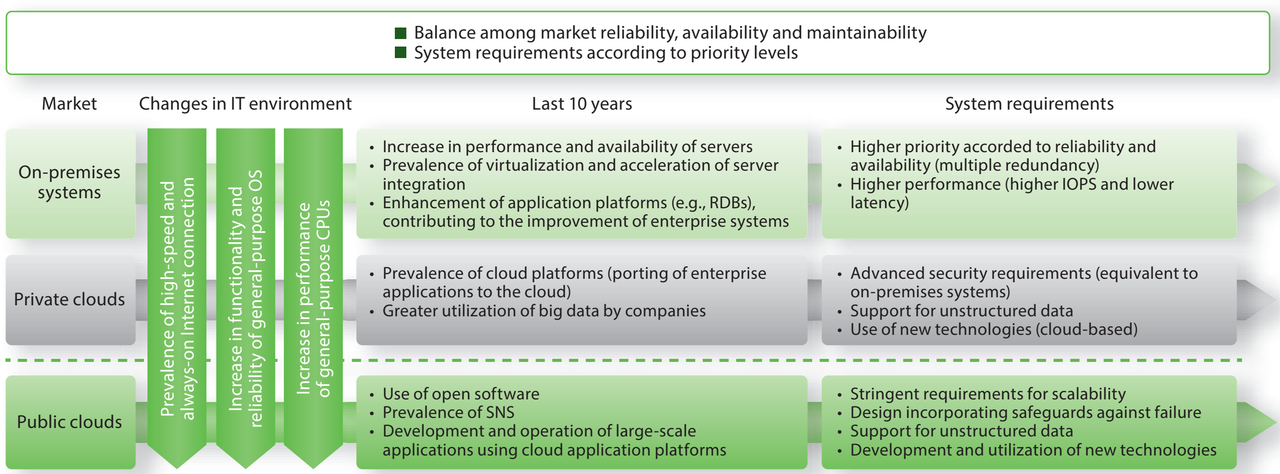

OS: operating system RDB: relational database IOPS: input/output operations per second SNS: social networking service

Figure 5. Changes in IT environment and in storage system requirements.
Availability, reliability, data throughput, maintainability and other requirements are changing in each IT market segment.

conventional centralized data processing system that is built to order according to the target services by using high-performance servers and large storage systems. Private clouds are large infrastructure operated solely for a single organization and consist of multiple general-purpose servers. Specific tasks are handled on private clouds at low cost. Public clouds seek to achieve economies of scale as they are run as businesses that offer shared services at pay-as-you-go rates.

The requirements for servers and storage vary according to the nature of data handled by them, and a wide variety of products are available, ranging from high-performance and high-reliability models to those that are aimed at achieving adequate performance at the minimum total cost of ownership (TCO).

The following paragraphs discuss the performance required for servers and storage to fulfill the specific needs of the three market segments (**Figure 6**). High-performance, high-reliability and long-life products with high added value are offered for the on-premises segment of the IT market. The throughput of an on-premises system can be increased through vertical scaling (**see Column**), but vertical scaling requires a costly investment to retrofit a system. For adequate vertical scaling, it is necessary to estimate a throughput requirement accurately and to determine system specifications needed to meet that requirement. The 44 zettabyte data deluge mentioned above is just one forecast; an investment must be made, based on an accurate forecast.

Private and public clouds have an advantage in that they can be scaled horizontally in a flexible manner as the throughput requirement grows. This makes private and public clouds more advantageous than an on-premises approach in cases where it is difficult to forecast the throughput requirement accurately. For private clouds, servers are selected, considering a trade-off between cost and performance to meet the throughput requirement. For public clouds, low-cost commodity products tend to be favored.

Although Hadoop[†] is a well-known platform for parallel distributed computing for big data analytics, it has drawbacks in terms of cluster size scalability for an on-premises approach and in terms of a provision for improving computing power for private clouds. Both on-premises systems and private clouds are based on platforms designed to offer a wide variety of applications and services. Using Hadoop[†] for them is becoming increasingly inadequate because Hadoop[†] is an extremely large software system, making it hard to predict the size of its deployment. Consequently, hori-

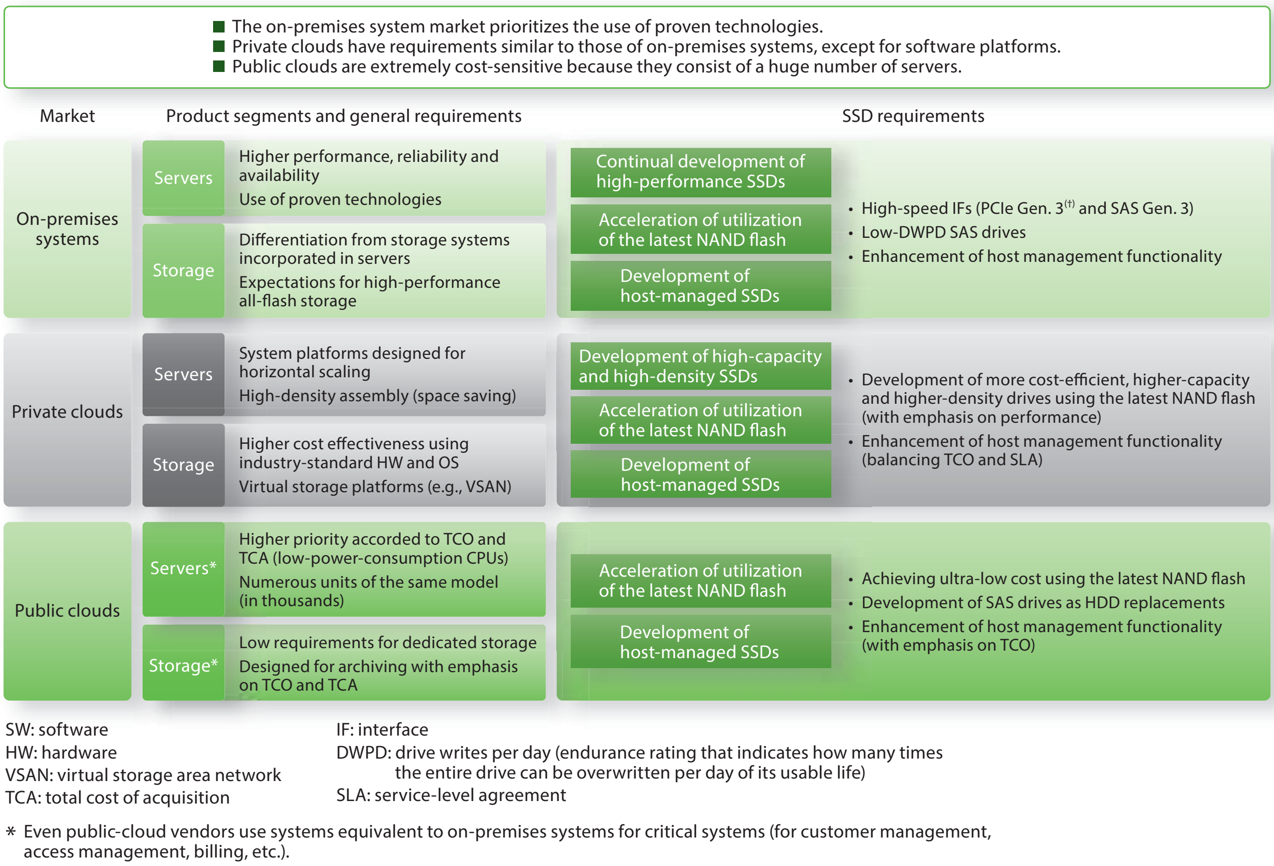

Figure 6. General requirements for each IT market segment and requirements for SSDs.
Manufacturers need to address different requirements for SSDs according to the target IT market segments and product segments.

Column Vertical scaling (scale-up) vs. horizontal scaling (scale-out)

There are two methods for increasing the performance and capacity of servers and storage devices: horizontal and vertical scaling.

To scale vertically (scale up) means to replace individual servers or storage devices with upper models or adding new components to them in order to increase performance or capacity. To scale horizontally (scale out) means to add more nodes to a network in parallel without changing the performance or capacity of individual servers or storage devices in order to increase the overall performance or capacity (**Figure A**).

There are trade-offs between these two models. Either one is favored over the other, depending on a system's application area and architecture (**Table A**). With the prevalence of the Internet and the evolution of software technology, horizontal scaling is becoming suitable for many applications, including cloud computing.

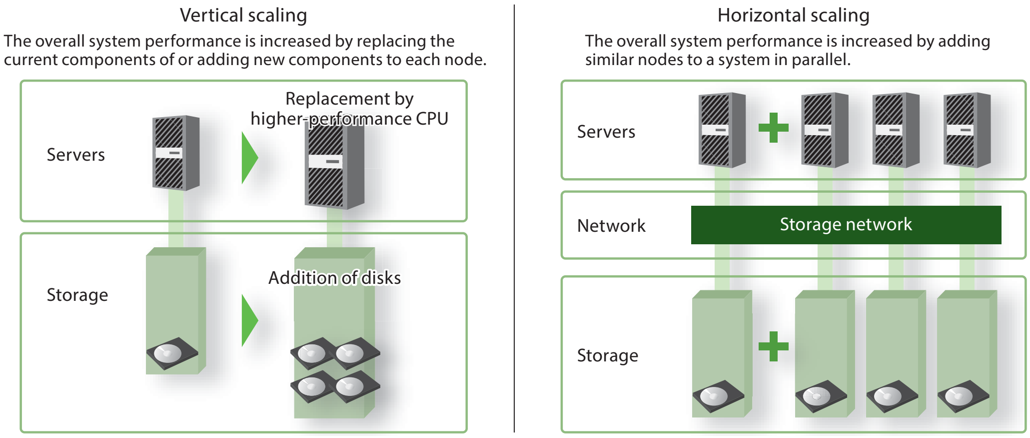

Figure A. Concepts of vertical and horizontal scaling

Table A. Advantages and disadvantages of vertical and horizontal scaling

Characteristic	Vertical scaling	Horizontal scaling
Pre-deployment planning	Planned ahead meticulously	Adjustable as needed
System expansion	Difficult	Easy
Possible application areas and architectures	General	Limited
Implementation cost	Needs initial funding for maximum system organization that will possibly be required in future	Allows iterative expansion of a system as business needs grow
Management cost	Low (fewer nodes to be managed)	High (more nodes to be managed)

zontally scalable servers and storage devices that are optimized for parallel distributed processing are often used for on-premises systems and private clouds (see pages 13–16).

Technical trends in the host interface for storage devices

In the 1990s, Small Computer System Interface (SCSI) was commonly used as a host interface for enterprise HDDs, whereas the Advanced Technology Attachment (ATA) interface was employed for HDDs for client PCs. Around 2006, a shift from a parallel interface to a serial interface began, and now HDDs have either a Serial Attached SCSI (SAS) or SATA host interface. SSDs, a major flash-based storage, came in two types, enterprise SAS SSDs and client SATA SSDs, which had the same operating specifications as for HDDs.

Although SAS and SATA were commonly used in early SSDs, later models have come to exhibit read and write performances (notably, random-access performances) two orders of magnitude higher than those of HDDs. As a result, the protocol overhead for SAS and SATA host-to-device communications became no longer negligible in comparison to the overall response times of SSD commands. Under these circumstances, PCIe[(†)], a physical interface that had been widely used for high-bandwidth graphics cards, and NVM Express[(†)] (NVMe[(†)]), a logical interface that allows for considerable enhancement in the parallel processing of SSD commands, have been employed as new host interfaces of SSDs. These host interfaces contribute to obtaining the optimal speed performance from SSDs.

Although the SAS, SATA and PCIe(†) protocols use a logical block size that is a multiple of 512 bytes, applications do not generally handle logical block addressing. There are layers that bridge between applications and a storage device, such as a filesystem that translates uniquely defined objects called filenames and logical block addresses, and data redundancy provided by the Redundant Array of Independent Disks (RAID). As semiconductor devices shrink in size and improve in performance, novel storage products have emerged that are designed to incorporate the functionality of all these layers into a storage device in order to let applications read and write directly from/to object storage. Object storage is considered suitable for long-term and cloud storage purposes because of its advantages, including the ability to hide hardware-specific details from users, excellent scalability, and an architecture that eliminates the needs for a filesystem and RAID hardware on the host.

In an object storage system, a filesystem runs on the CPU incorporated in a conventional HDD or SSD instead of the host. Object storage can be viewed as a small server as it consists of Ethernet(†) interfaces, CPUs, memories and storage. Object storage can open up various possibilities not only as a storage medium but also as a server (see pages 9–12).

SSD product segments

In the past, limited write endurance was a concern in the case of using NAND-based storage. In order to secure adequate write endurance, early SSDs used single-level-cell (SLC) NAND flash memories. However, after using SSDs for a few years, the market learned that NAND write endurance would not impose a significant limit on the lifetime of SSDs, and as a result, multi-level-cell (MLC) NAND flash memories came to be commonly used for SSDs. In recent years, SSDs have appeared that consist of triple-level-cell (TLC) NAND flash memories with lower write endurance.

Furthermore, NAND-based SSDs provide great flexibility in design-for-performance and design-for-product lifetime. Manufacturers address the diverse needs of the IT market by offering SSDs with different performances and product lifetimes for each category of host interfaces (**Table 2** and **Table 3**). Such product offerings are unique to SSDs that have limited write endurance. However, this causes the number of product models to increase, making product management more complicated and less flexible.

Price reductions are important targets in the development of SSDs. With the aim of reducing SSD prices, the goals of our SSD development roadmap include: technologies to increase the density and utilization of NAND flash memories, features for increasing the design flexibility to improve performance and product lifetime (see pages 17–20), and reductions in the parts cost overhead caused by increased capacities.

Toshiba's enterprise SSDs and application technologies for NAND flash memories

Although SSDs first emerged as HDD replacements, changes in the IT market environment, notably the data deluge, have caused customers' expectations of Toshiba to increase. This section describes NAND-based application products designed to meet these expectations. Memory-based application products have tremendous possibilities, and developers need flexible ideas and the ability to turn them into reality.

Our NAND flash memory development draws on the integration and reliability technologies built up over decades of HDD development as well as the high-performance and high-reliability technologies acquired through experience with rapidly evolving SSDs. Additionally, new value will be added to NAND flash memories in order to contribute to the development of and innovation in the IT industry (**Figure 7**).

Table 2. Product segments estimation of enterprise SSDs (Serial Attached Small Computer System Interface (SAS) type and Serial Advanced Technology Attachment (SATA) type), 2015–2016

Interface	Category	DWPD	Form factor	Capacity	Power consumption (average)	Performance (random reads)	Expected applications
12 Gbits/s SAS	High Endurance	25	2.5-inch 15 mm-thick	200 Gbytes–1.6 Tbytes	9–14 W	200 kIOPS	• Hierarchical storage (all-flash and hybrid) • Mid-range to high-end servers
	Mid-Endurance	10		400 Gbytes–3.2 Tbytes			
	Value Endurance	3		400 Gbytes–4 Tbytes			
	Read-Intensive	1					
6 Gbits/s SATA	Mid-Endurance	5–10	2.5-inch 7 mm-thick	200 Gbytes–2 Tbytes	2–5 W	90 kIOPS	• General servers
	Value Endurance	3		240 Gbytes–3.2 Tbytes			• Hyper-scale systems • Low-end servers
	Read-Intensive	<1		480 Gbytes–4 Tbytes			

T: tera (10^{12})

* The above table is based on estimates by Toshiba.

Table 3. Product segments estimation of enterprise SSDs (PCI Express[(†)] type), 2015–2016

Category	DWPD	Form factor	Capacity	Power consumption (average)	Performance (random reads)	Expected applications
High Performance (add-in cards)	10	FHFL (×16) FHHL (×8/×16) HHHL (×8)	6.4 Tbytes 3.2 Tbytes	25–50 W	<3.2 M IOPS (×16) <1.6 M IOPS (×8)	• Server-side caches • Application accelerators
	5					
	1 - 3					
Mid-Range Performance	10	2.5-inch (×4, 15 mm-thick)	400 Gbytes–3.2 Tbytes	9–25 W	600–800 kIOPS	• Mid-range to high-end servers
	3		400 Gbytes–4 Tbytes			
	1					
Low-End Performance	5	2.5-inch (×4, 7 mm-thick)	400 Gbytes–2 Tbytes	6–12 W	150–300 kIOPS	• Hyper-scale systems • Low-end servers • Microservers
	3	M.2 (22110) 2.5-inch (×2 and ×4, 7 mm-thick)	400 Gbytes–4 Tbytes	6–9 W	150–200 kIOPS	
	<1					

FHFL: Full-height Full-length FHHL: Full-height Half-length HHHL: Half-height Half-length

*1: ×2 to ×16 denote the number of lanes.
*2: The above table is based on estimates by Toshiba.

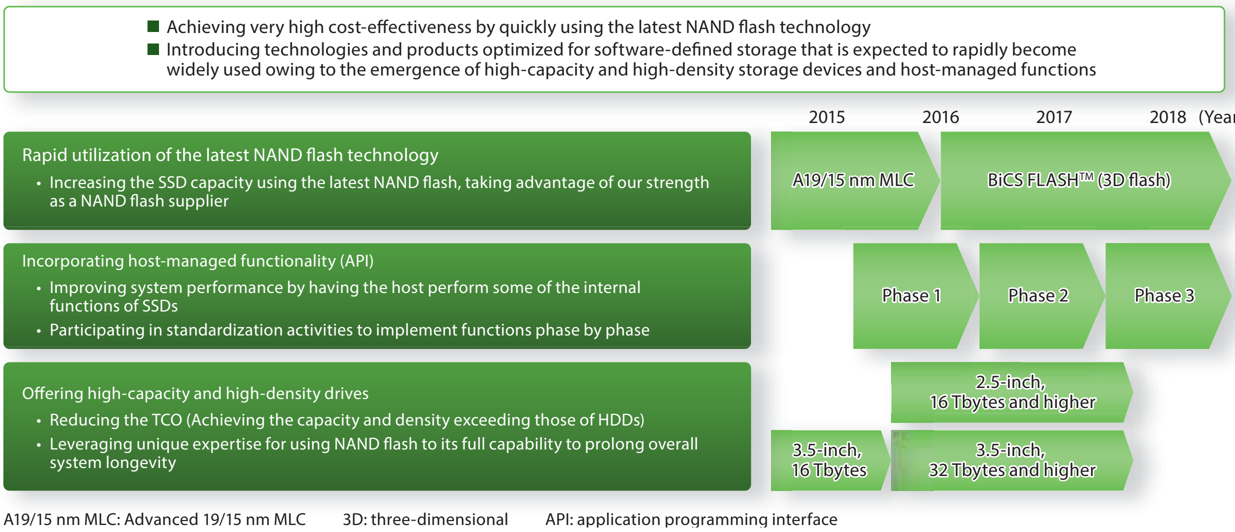

Figure 7. Toshiba's development strategy for enterprise SSDs, 2015-2018.
Toshiba aims to meet the needs for software-defined storage by improving the performance of SSDs through timely use of the latest NAND flash technology and by implementing and standardizing new functions.

Technical trends in HDDs

Changing environment surrounding HDDs

HDDs have long been the leading external storage device not only for notebook and desktop PCs but also for servers because of their good balance between performance and the cost per GB, but the position of HDDs in the storage market is now undergoing a major change.

Since HDDs spin magnetic disks and position the read/write head across the spinning disks using spindle motors, they cannot compare favorably with NAND flash and other semiconductor memories in terms of read and write speeds. The primary advantage of HDDs lies in the cost per GB due to their high areal density (see pages 25-28). Although the price of NAND flash memories has been declining year by year, they still cost over 10 times more than 3.5-inch HDDs that are available at three cents per GB. In order for HDDs to continue to be used as a storage device, it is necessary to keep their cost per GB at least 5 to 10 times less than that of NAND flash memories. **Figure 8** shows the trend in areal density from the past to the future.

To maintain the cost-per-GB advantage, the areal density of HDDs must be increased at an annual rate of over 15%.

Technologies for increasing the areal density

We are planning to adopt a shingled magnetic recording (SMR) technology to increase the areal density (see pages 29-32). SMR is expected to provide an approximately 20% increase in areal density. Other new record-

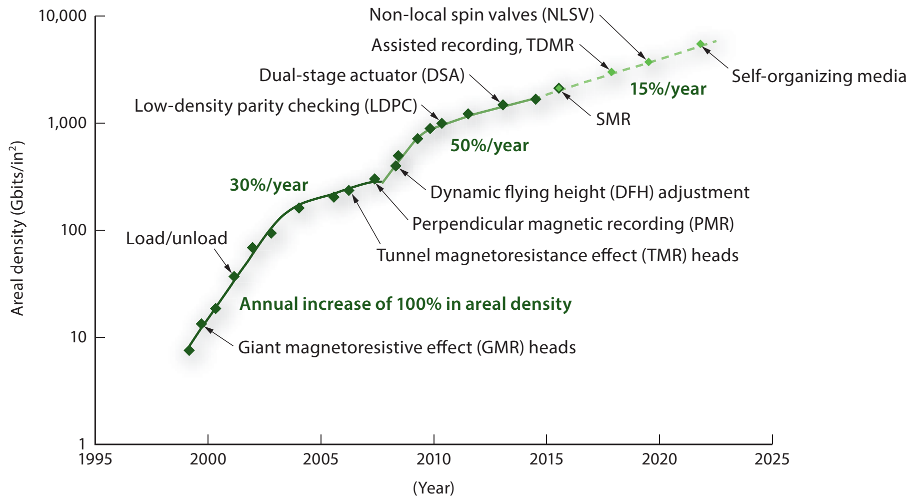

Figure 8. Trends in areal density of HDDs and innovation technologies.
The areal density of HDDs must be increased at 15% annually in order to maintain the cost-per-GB advantage over NAND flash memories.

ing technologies that are being developed include two-dimensional magnetic recording (TDMR), 4 Kbytes-per-sector formatting and track-level error correction code (ECC). By combining these technologies with the advancement of head and disk technologies, we will aim to maintain the annual rate of increase in areal density at 15% or so until 2025.

We are planning to launch a 2.5-inch HDD that provides a capacity of 1 Tbyte (tera: 10^{12}) with a single disk in 2016. This is double the capacity of today's mainstream 2.5-inch HDDs that provide a capacity of 500 GB per disk. Furthermore, we are aiming to realize a 3.5-inch nearline HDD with a capacity of 16 Tbytes (consisting of eight disks) at a cost per GB of one cent by 2019 (**Figure 9**).

Value addition strategies for HDDs

Demand for solid-state hybrid drives (SSHDs) is growing, especially among notebook PC users. An SSHD combines a 2.5-inch HDD with 8 Gibytes (gibi: 2^{30}) of NAND flash memory that acts as a cache for the data stored on the HDD. SSHDs are cost-effective storage that delivers an improvement in read and write speeds that is difficult to achieve with the conventional HDDs. We have expertise in both NAND flash memories and HDDs. By leveraging our expertise, we will also develop HDD products that combine the advantages of both of them.

It has been a while since the term "object storage" became popular in the field of servers. In this field, a new type of HDD with an Ethernet(†) interface called an

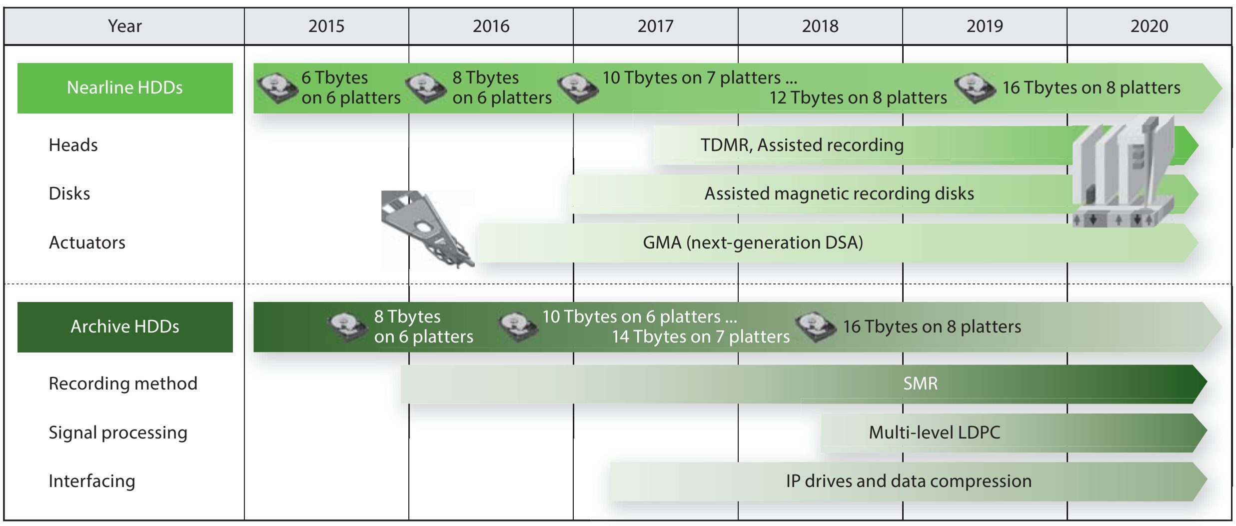

GMA: gimbal microactuator

Figure 9. Toshiba's development plan for nearline HDDs.
The cost per GB will be reduced by using SMR and TDMR.

Internet Protocol HDD (IP HDD) is being developed. IP HDDs convert a user file into a logical block address and a logical block length to relocate it instead of the host. This makes it easy to perform data compression inside IP HDDs and helps reduce the cost per GB.

HDDs can continue to evolve and come in many different variants. It is considered that today's HDDs have not achieved the maximum possible magnetic recording performance yet in terms of areal density because of the need to maintain a good balance between performance and capacity. It is important for us to deliver unique magnetic storage devices by taking advantage of our R&D resources.

Future prospects

We will continually improve HDDs, NAND flash memories, SSDs and other storage products in terms of their capacity, throughput, reliability, availability and other factors in order to provide storage solutions that satisfy a wide range of market requirements.

References

(1) Turner, V. et al. "IDC's Digital Universe, The Digital Universe Opportunities: Rich Data and Increasing Value of the Internet of Things, Sponsored by EMC (2014-04)". IDC White paper #IDC_1672. Accessed May 25, 2015. http://idcdocserv.com/1678.

- *PCI Express and PCIe are registered trademarks of PCI-SIG.*
- *Hadoop is a trademark or a registered trademark of Apache Software Foundation in the U.S.A. and other countries.*
- *NVM Express and NVMe are trademarks of NVM Express, Inc.*
- *BiCS FLASH™ is a trademark of Toshiba corporation.*
- *Ethernet is a trademark of Fuji Xerox, Co., Ltd.*

YANAGI Shigenori

Technology Executive. Semiconductor & Storage Products Company. He is involved in the development of technologies for SSD and storage solution products.

OHSHIMA Shigeo

Technology Executive. Semiconductor & Storage Products Company. He is involved in the development of technologies and strategies for flash storage systems.

ITO Jun

Technology Executive. Semiconductor & Storage Products Company. He is involved in the development of HDD technologies.

Notes:
"-inch" means the form factor of HDDs or SSDs. It does not indicate drive's physical size.
This publication may include Toshiba's original estimation, and/or data based on Toshiba's evaluation.