TOSHIBA

A Toshiba Perspective on HDD Capacity, Innovation and Technology Evolution

Toshiba Electronics Europe GmbH

August 2026

Scope of Disclosure	Public
Head of Information Owner Section	**Igor Misko - Senior Product Marketing Engineer**

Toshiba HDD Business History: Products and Technology

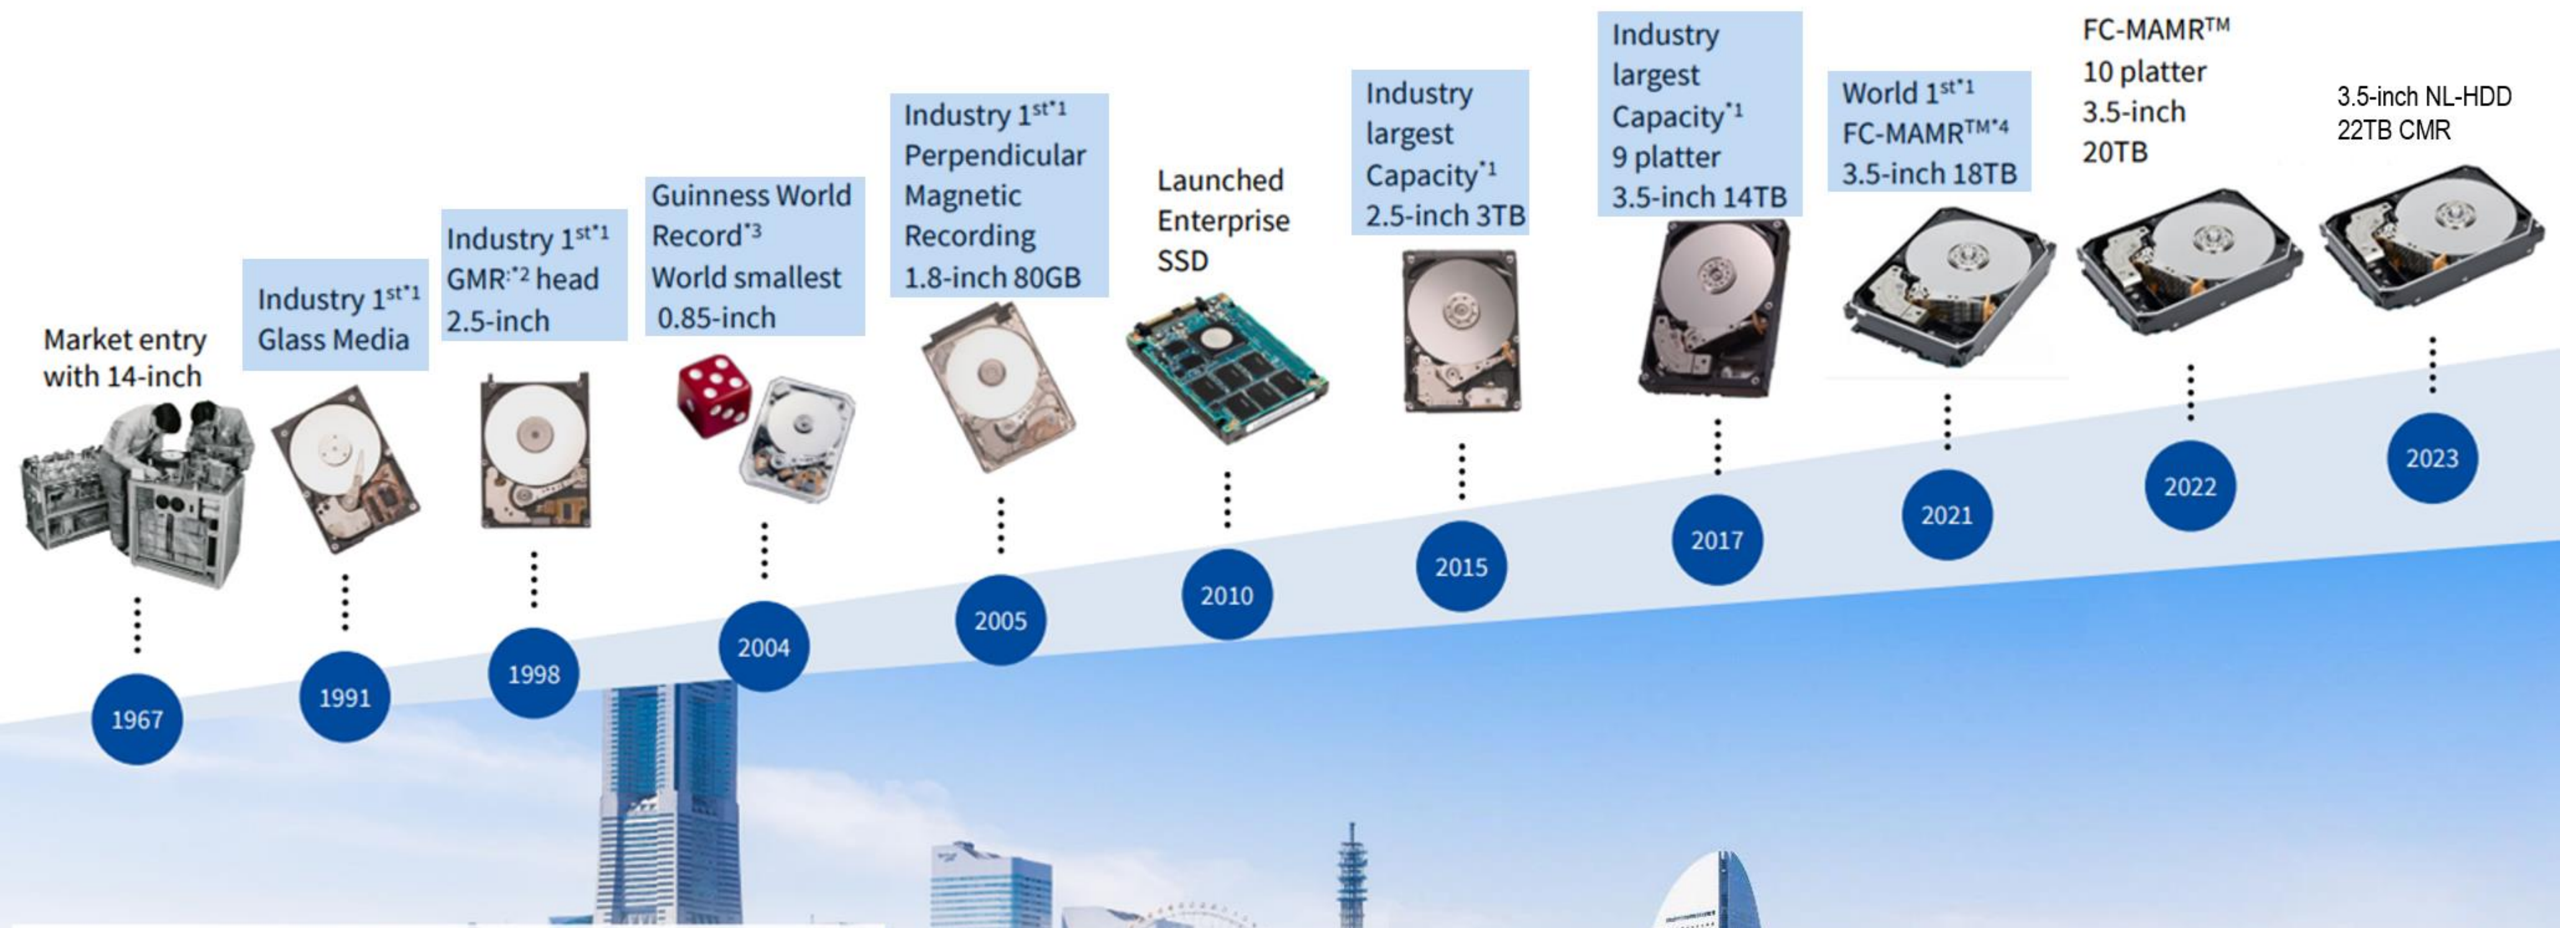

*1: Source: Toshiba, at the times of press releases
*2: GMR: Giant Magneto Resistive
*3: Guinness World Records is registered trademark of Guinness World Records Ltd.
*4: FC-MAMR™: Flux Control - Microwave Assisted Magnetic Recording
*5: MAS-MAMR: Microwave Assisted Switching - Microwave Assisted Magnetic Recording

Supporting industry needs with leading-edge HDDs since 1967

HDD capacity increased ~1,000× every two decades

Three milestones reveal a repeated capacity leap — 20 MB → 20 GB → 20 TB.

1984

20 MB

EARLY HDD MILESTONE

×1,000 →

2002

20 GB

ENTERPRISE SCALE

×1,000 →

2022

20 TB

MODERN CAPACITY

1,000,000× IN 38 YEARS

20 MB → 20 TB

IF THE PATTERN CONTINUED

2042 ≈ 20 PB

Illustrative trajectory — not a product forecast

Each leap multiplied usable capacity by ~1,000 — one millionfold from 1984 to 2022.

TOSHIBA

Roadmap of Larger Capacity HDDs for Data Centers

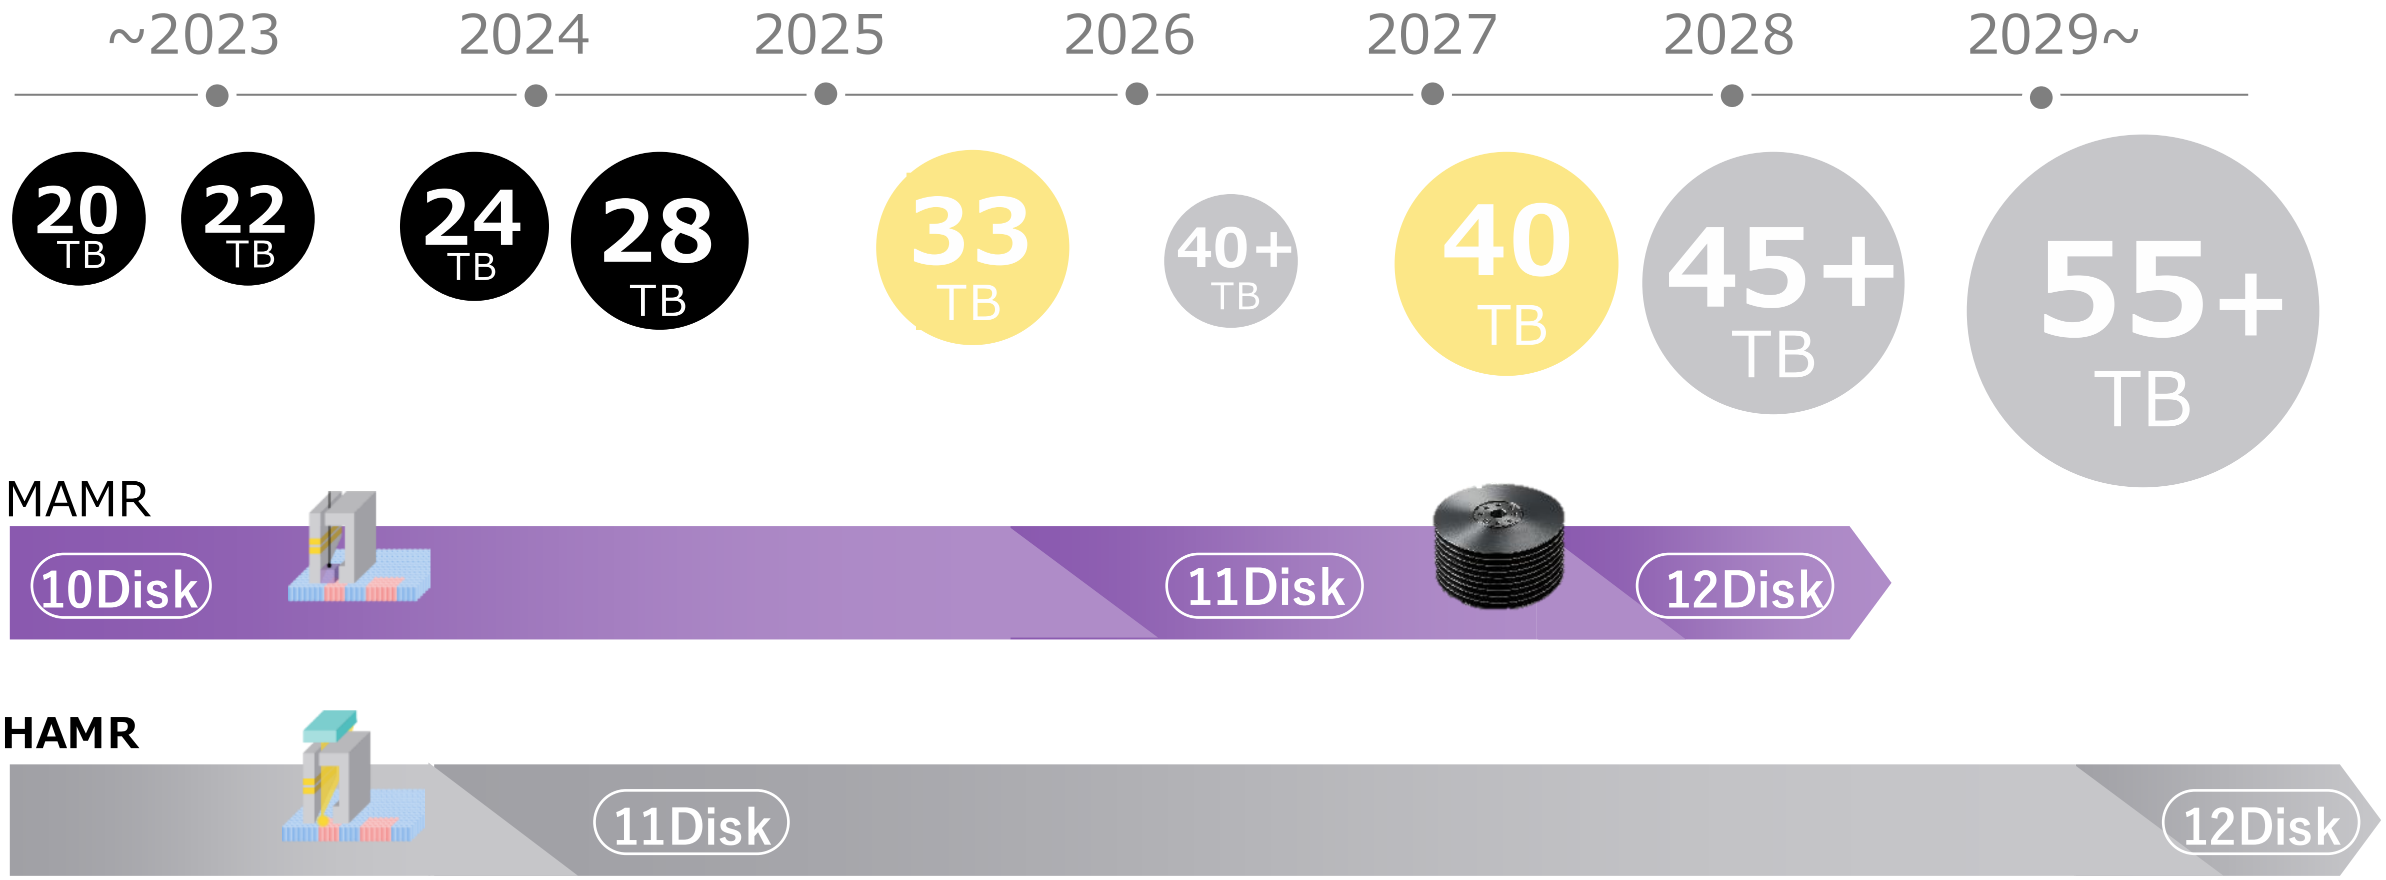

Continue to launch advanced, new generation-technologies to promote larger storage capacity

Increasing Capacity Requires Multiple Technology Levers

More capacity in the existing 3.5-inch form factor helps reduce the infrastructure footprint per PB.

Platform	Recording	Capacity	Disks	TB/Platter
MG11 CMR	FC-MAMR	24TB	10	2.40
MA11 SMR	FC-MAMR + SMR	28TB	10	2.80
M12 CMR	FC-MAMR	28TB	11	2.55*
M12 SMR	FC-MAMR + SMR	34TB	11	3.09

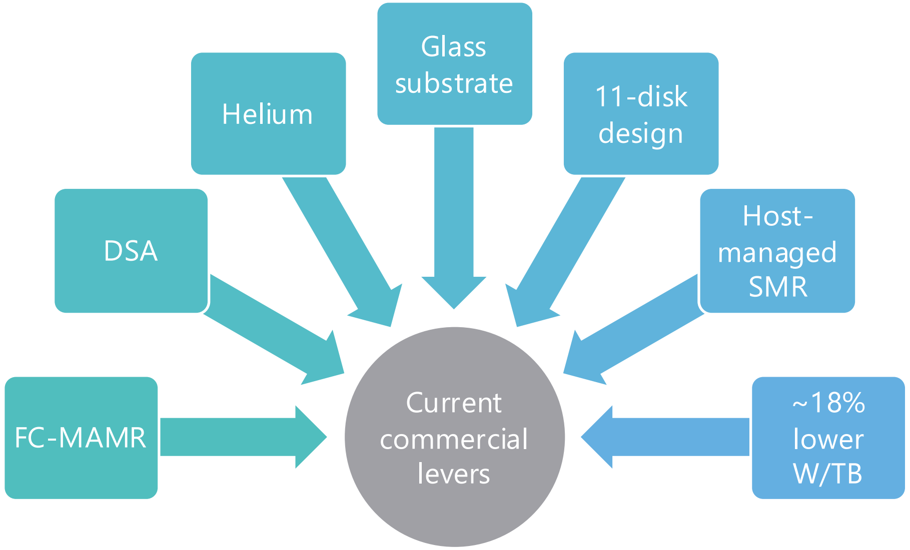

More capacity is not delivered by one breakthrough...

It is delivered by recording technology, media, mechanics and power-efficiency improvements working together.

MAMR (Microwave Assisted Magnetic Recording) Technology

•FC-MAMR in mass production. MAS-MAMR development for next gen.

Flux control type (FC-MAMR™)

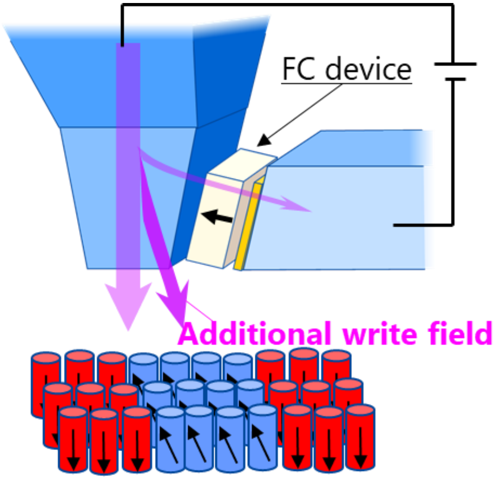

News release

TOSHIBA Toshiba Electronic Devices & Storage Corporation

Toshiba Announces 18TB MG09 Series Hard Disk Drives

3rd-generation 9-disk Helium-sealed design and innovations in energy-assisted recording help customers achieve new levels of storage density and power efficiency

February 18, 2021
Toshiba Electronic Devices & Storage Corporation

Tokyo, Japan —February 18, 2021 —Toshiba Electronic Devices & Storage Corporation (Toshiba) announces the 18TB[1] MG09 Series HDD, Toshiba's first HDD models with energy-assisted magnetic recording. The MG09 Series features Toshiba's third-generation, 9-disk Helium-sealed design and Toshiba's innovative Flux Control – Microwave Assisted Magnetic Recording (FC-MAMR) technology, to advance Conventional Magnetic Recording (CMR) density to 2TB per disk, achieving a total capacity of 18TB.

Sample shipments of 18TB MG09 Series HDD to customers are expected to start sequentially at the end of March 2021.

https://toshiba.semicon-storage.com/ap-en/company/news/news-topics/2021/02/storage-20210218-1.html

World's first commercialization of energy-assisted magnetic recording with spintronics

Resonant switching type (MAS-MAMR)

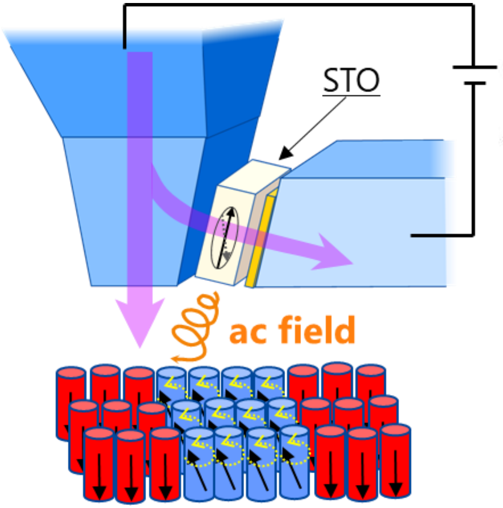

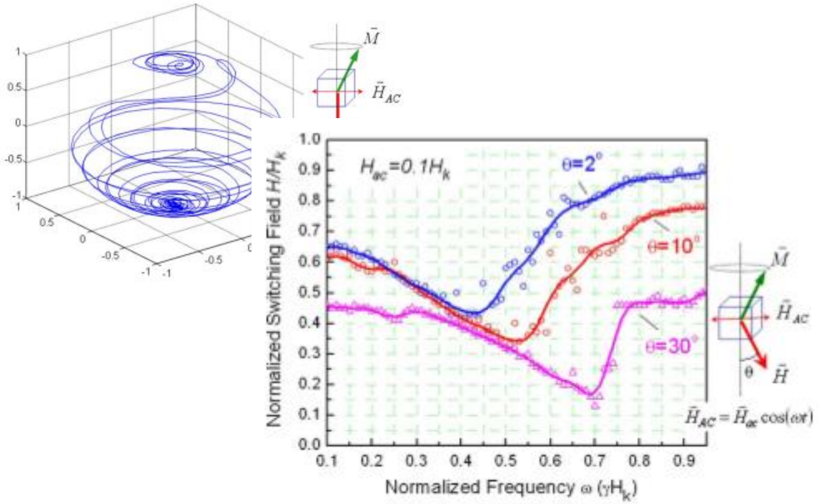

J. G. Zhu et. al., *IEEE Trans. Magn.*, 44, 1, 2008.

	In-gap device	Operating principle	Mag. state in the device	Medium optimization
FC-MAMR	Flux control device：FC device	Spin transfer torque	Reversal (Change flux flow)	**Not necessary**
MAS-MAMR	Spin torque oscillator：STO	Spin transfer torque	Oscillation (Generate ac field)	Necessary (FMR phenomenon)

What Makes an HDD Enterprise-Class?

Rotational Vibration (RV) Sensors

An enterprise JBOD with 60, 90 or 108 drives becomes a vibration generator. A single drive vibrating affects neighboring drives. Toshiba MG drives integrate RV sensors that continuously monitor rotational vibration and adjust head positioning accordingly. Multiple sensors detect shock and vibration, helping maintain performance and reliability in dense storage systems.

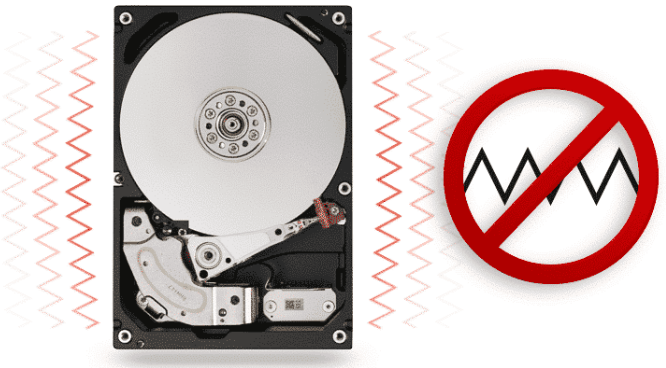

Dual-Stage Actuator (DSA)

Why DSA Exists?

- Tracks became narrower
- Bits became smaller
- Areal density increased massively

At 24 TB and beyond, the read/write head must follow tracks with nanometer-level precision.

A conventional voice-coil actuator alone is no longer accurate enough.

Persistent Write Cache (Power-Loss Protection)

When an HDD acknowledges a write, some data may still be located in cache memory. If power suddenly disappears, conventional drives risk losing those pending writes. Toshiba's Persistent Write Cache (PWC) technology stores cached data in non-volatile flash memory and commits it to media once power returns.

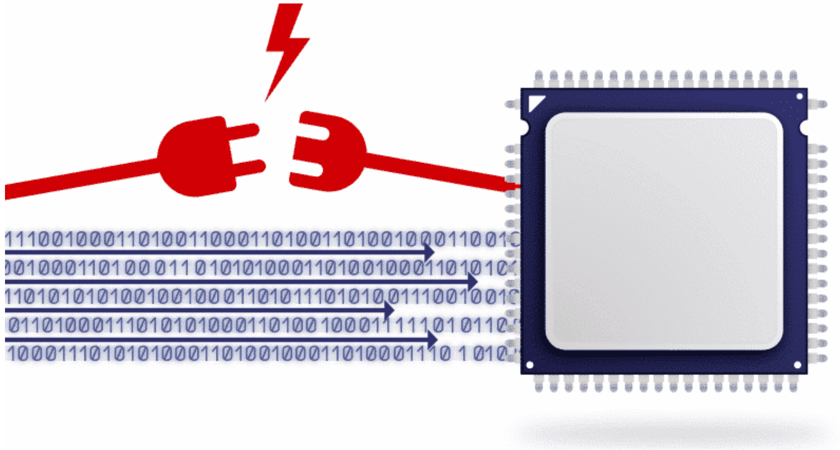

Precision, Reliability & Data Protection Technologies

A 2016 HDD fleet remains in operation a decade later

2,700 MG04 enterprise drives continue to deliver 16 PB of raw storage.

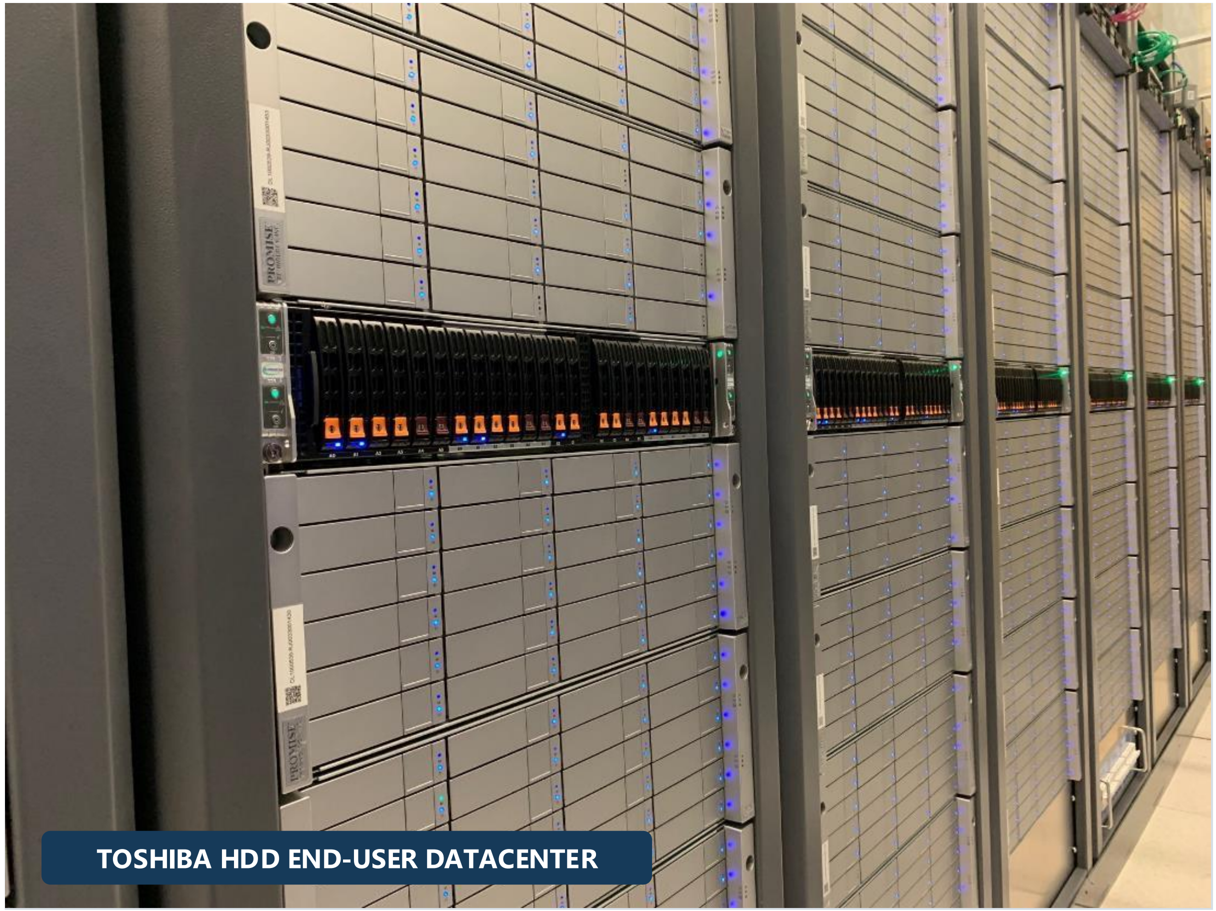

TOSHIBA HDD END-USER DATACENTER

10 YEARS

INSTALLED 2016 • STILL OPERATING

14	2,700	16 PB
RACKS	HDDs	RAW STORAGE

SYSTEM PROFILE • MG04 • 6 TB HDDs

FIVE-YEAR RELIABILITY • 2,700 DRIVES

ORIGINAL ESTIMATE	OBSERVED MG04	MG11 REFERENCE
60	~10	0.35%
FAILURES	FAILURES	SPECIFIED AFR
	≈0.07% AFR	≈47 failures

Observed AFR assumes full 5-year exposure. MG11 projection is statistical, not guaranteed.

Higher capacity does not have to come at the expense of reliability.

Why HDD Reliability Depends on Workload

The right workload and environment can significantly influence HDD service life

HDD reliability is a system outcome

RANDOM I/O

MORE SEEKS → MORE ACTUATOR ACTIVITY → MORE HEAT
→ MORE SERVO CORRECTIONS → MORE STRESS

High seek intensity increases mechanical and thermal demand.

SEQUENTIAL I/O

FEWER SEEKS → STABLE OPERATION → LOWER TEMPERATURES
→ LOWER STRESS → LOWER OBSERVED AFR

Sequential placement creates steadier operating conditions.

Workload is a reliability lever

Lower seek intensity can reduce the stresses associated with higher observed AFR.

WHERE SEQUENTIAL WORKLOADS APPEAR

Object storage, data lakes, backup repositories and HM-SMR deployments favor sequential placement—reducing mechanical stress while maximizing capacity efficiency and service life.

HDD reliability also depends on system design and operating conditions. That is exactly why we need real-world validation. And this brings us to the Toshiba HDD Innovation Lab in Düsseldorf.

Toshiba HDD Innovation Lab Düsseldorf

What customers can evaluate

- Enterprise storage architectures
- HDD qualification
- RAID / HBA validation
- ZFS and software-defined storage
- Performance
- Power consumption
- Thermal behavior
- AI and data-lake storage

Infrastructure

- Toshiba enterprise HDD portfolio
- Multiple controllers
- Multiple storage platforms
- High-speed networking
- Benchmark automation
- Real-world workload testing

Customer benefit

- Reduced deployment risk
- Faster qualification
- Better architecture decisions
- Improved TCO understanding
- Evidence-based storage selection

Storage decisions should be measured, not assumed.

Questions?

Thank you.

12-disk stacking technology

Toshiba has successfully verified the industry's first*1 12-disk stacking technology for high-capacity hard disk drives.

Plans to introduce new 40TB-class HDDs in 2027, utilizing 12-disk stacking technology and combining it with MAMR.

5 core technologies leaded to 12-Disk stacking HDD:

- Integrated Stacking System
- Lite-Rotor
- Tough Enclosure
- Ultra-thin
- Glass Technology

"Integrated Stacking system" has achieved 8% Scale down per generation in height dimension (B). Supported by Toshiba corporate laboratory, manufacturing assembly technology will evolve.

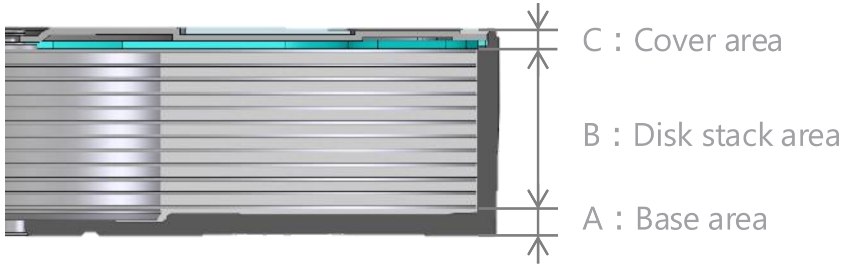

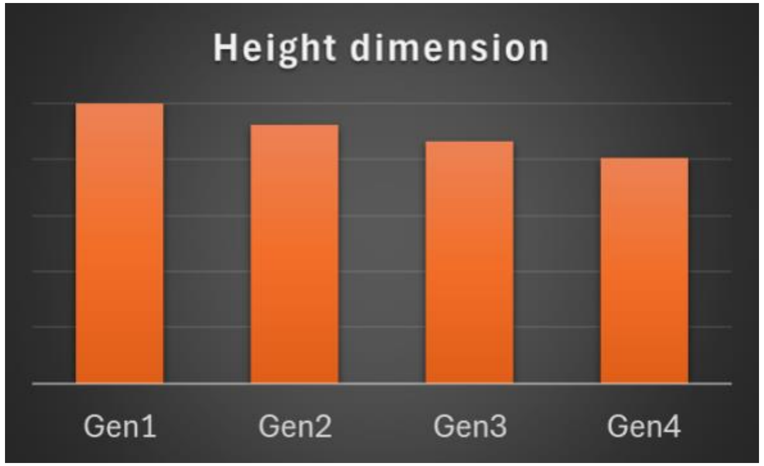

"Lite-Rotor" technology reduces the mass by approximately 12% compared to conventional products.
This contributes to improved shock performance.

Mechanical innovation continues to create room for future capacity growth.

Advanced Power Modes

Table 4.14 State definition – APM / EPC

Power Mode		Description	
Use APM	Use EPC	Heads Location	Motor
Active		Flying (Load)	Spinning
Performance Idle			
Active Idle	Idle_a		
Low Power Idle	Idle_b	On ramp (Unload)	
N/A	Idle_c / Standby_y		(at Low Speed)
Standby	Standby_z		Spin Down
Sleep (by Command)			

- Standby_Z Condition Recovery Time
  It reports on the return time from Standby_Z State to Active State. The unit is ms.
  This field always reports 61A8h (= 25 000ms = 25s).
- Standby_Y Condition Recovery Time
  It reports on the return time from Standby_Y State to Active State. The unit is ms.
  This field always reports 1388h (= 5000ms = 5s).
- IDLE_A Condition Recovery Time
  It reports on the return time from IDLE_A State to Active State. The unit is ms.
  This field always reports 064h (= 100ms).
- IDLE_B Condition Recovery Time
  It reports on the return time from IDLE_B State to Active State. The unit is ms.
  This field always reports 04B0h (= 1200ms = 1.2s).
- IDLE_C Condition Recovery Time
  It reports on the return time from IDLE_C State to Active State. The unit is ms.
  – This field always reports 1388h (= 5000ms = 5s).

Table 2.5 Power consumption (MG11ACA24T/22T/20TE)

Item		Current (typical)		Watts (typical)
		+5V DC	+12V DC	
Spin Up	Peak (*1)	0.46 A	1.47 A	-
	Maximum DC (*2)	0.22 A	1.35 A	-
Random Read 4KB Q1	Peak	0.96 A	1.87 A	-
	Average	0.31 A	0.56 A	8.23 W
Random Read 4KB Q16	Peak	0.97 A	1.86 A	-
	Average	0.33 A	0.61 A	8.81 W
Random Write 4KB Q1	Peak	0.84 A	1.86 A	-
	Average	0.34 A	0.38 A	6.22 W
Random Write 4KB Q16	Peak	0.84 A	1.83 A	-
	Average	0.34 A	0.38 A	6.21 W
Sequential Read	Peak	0.82 A	0.62 A	-
	Average	0.65 A	0.35 A	7.35 W
Sequential Write	Peak	0.87 A	0.51 A	-
	Average	0.71 A	0.35 A	7.65 W
Random Write/Read 4KB Q1	Peak	0.96 A	1.85 A	-
	Average	0.32 A	0.55 A	8.11 W
Random Write/Read 4KB Q16	Peak	0.98 A	1.84 A	-
	Average	0.32 A	0.63 A	9.15 W
Idle_A	Average	0.16 A	0.30 A	4.35 W
Idle_B	Average	0.15 A	0.20 A	3.03 W
Idle_C	Average	0.15 A	0.16 A	2.65 W
Standby	Average	0.11 A	0.003 A	0.58 W

Power savings depend on matching idle duration with acceptable recovery latency.

Mechanical Platform for Multi-Platters

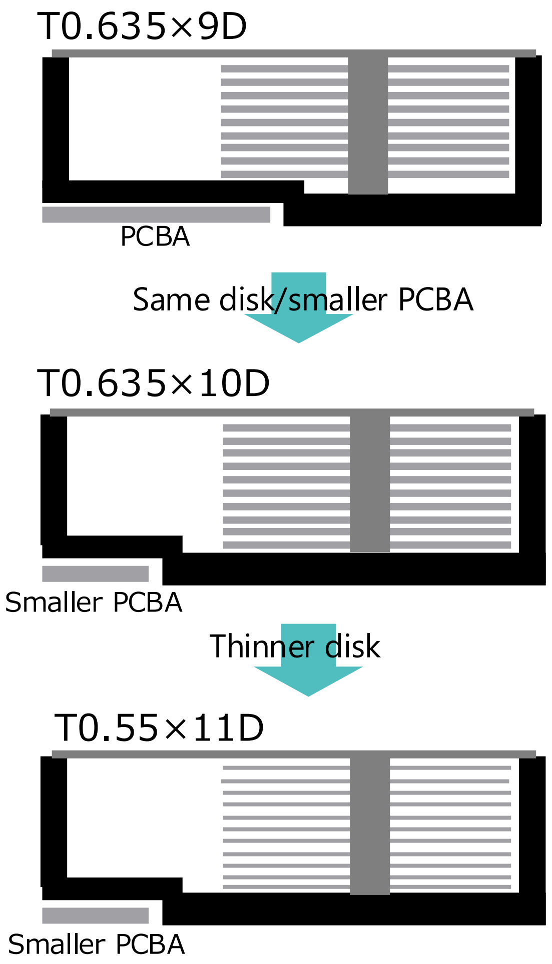

	～18TB	20TB～	30TB～
Mechanics	He Gen1	He Gen2	He Gen3
Number of Disk	**9**	**10**	**11**
Disk thickness	0.635mm	←	0.55mm
Op-Shock (Z-direction)	70G/2ms	50G/2ms	30G/2ms
Non Op-Shock	250G/2ms	200G/2ms	←

Every generation pushes mechanical limits further to increase capacity without changing the HDD footprint.

TOSHIBA

Source: Akikhiko Takeo at IDEMA Japan Symposium, 16th May 2024

“Ultra-thin” Technology

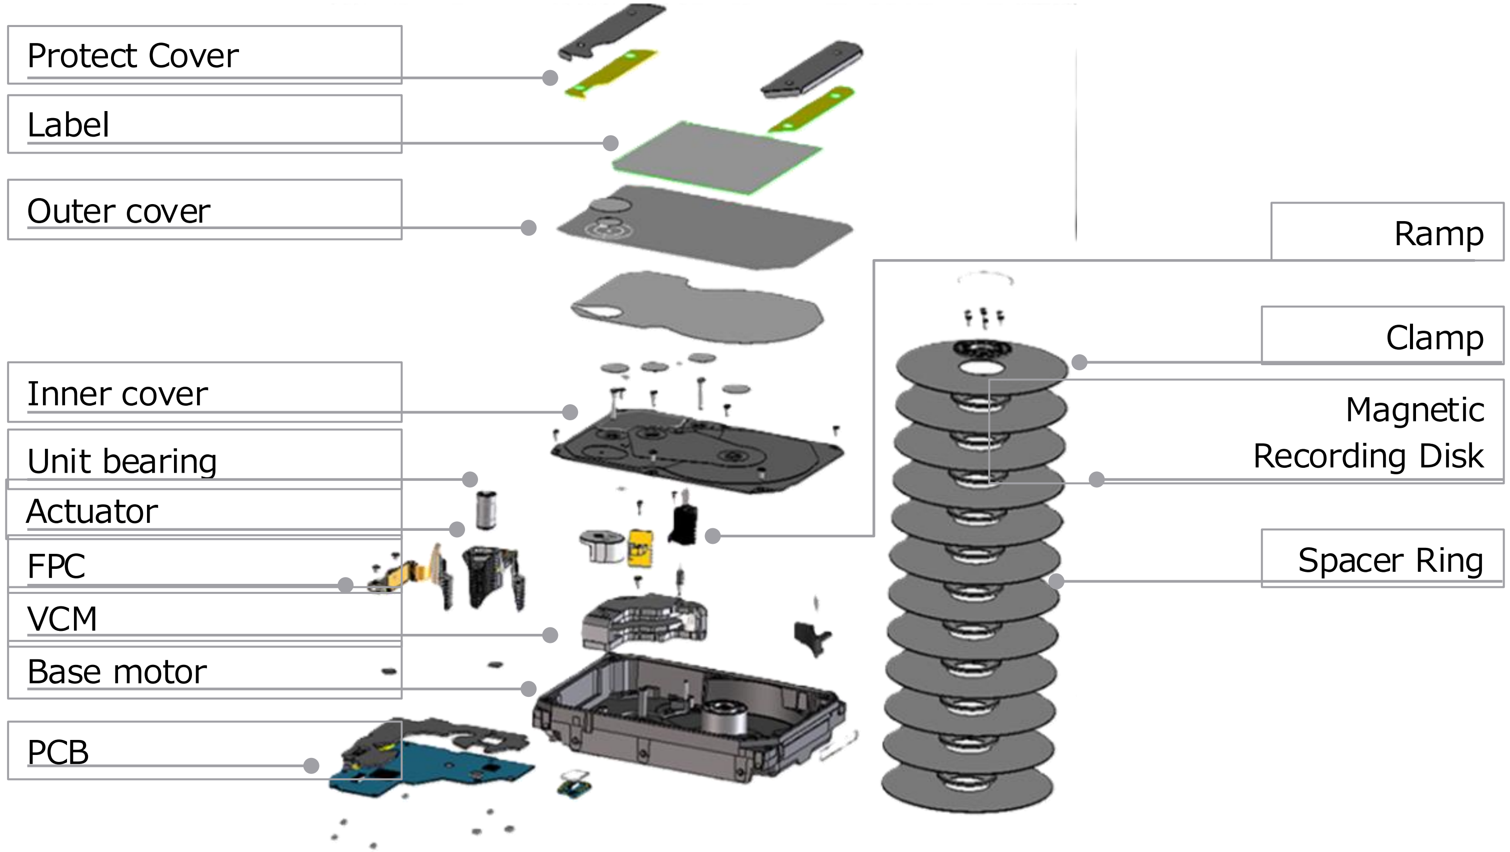

The design of all parts has been reviewed to achieve “ultra-thin”.