Western Digital.

Tech Brief

Continuous Innovation for Highest Capacities and Lower TCO

October 2024

Hard Disk Drives (HDDs) remain the foundation of the data center and will continue to be the dominant storage media for the foreseeable future. Cloud-based workloads continue to drive demand for capacity-optimized HDDs that support the growing data needs of cloud and hyperscale storage, massive scale-out high density data centers, content delivery networks, commerce and XaaS. Cloud service provider business models rely on continued HDD capacity increases to efficiently meet data center storage density requirements as data growth increases. The burgeoning generative AI industry will need tremendous amounts of data storage to train AI models, and even more storage capacity to retain the outputs of these generative models.

For over 50 years, Western Digital has enabled the data revolution by innovating a variety of technologies to grow HDD capacity on a regular cadence, meeting this increasing demand. IDC estimates that the amount of data created by endpoints, the edge and at the core will reach 393.9 zettabytes (ZB) annually by 2028[1], and worldwide installed storage capacity will exceed 19 ZB by that year[2]. To meet the requirements of this Zettabyte Age, new HDD technologies will be required to expand the amount of data in the data center while lowering the system-level Total Cost of Ownership (TCO).

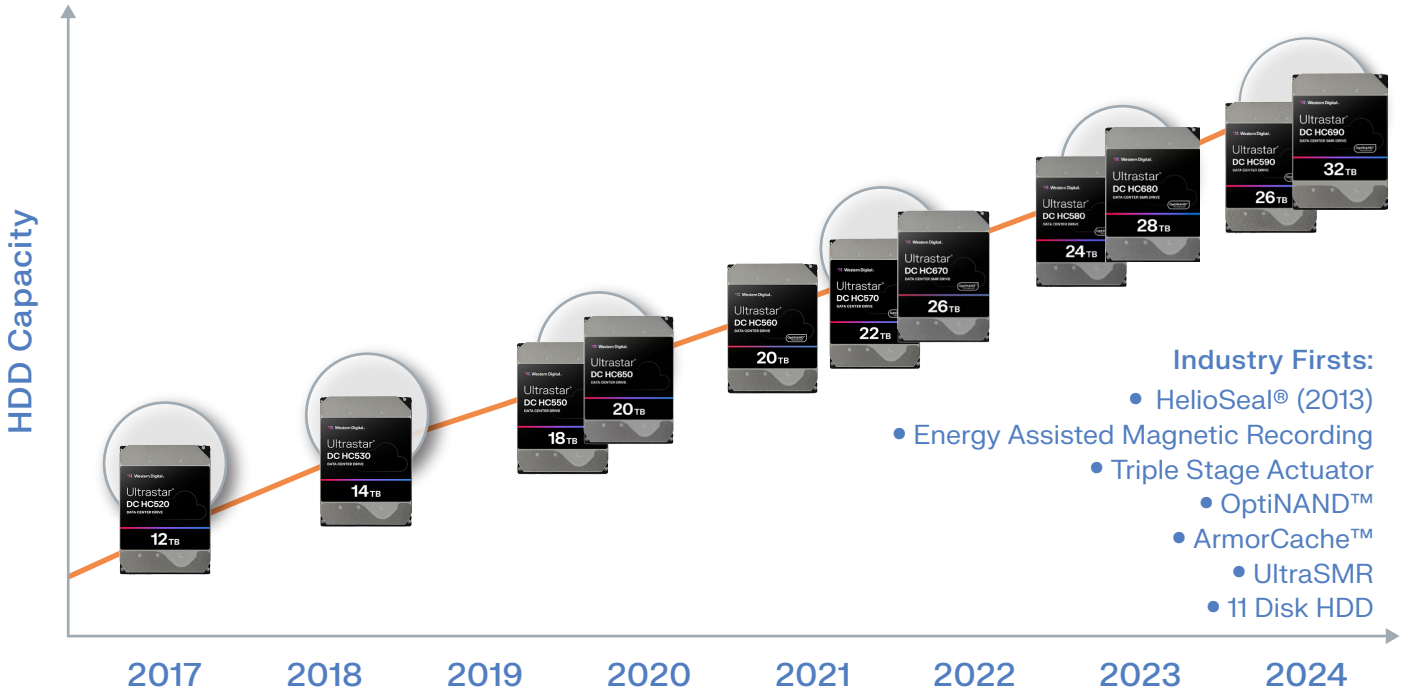

A data center HDD is a complex piece of hardware and software with nearly 300 components and more than a million lines of software code. It is a technically sophisticated system with multiple interdependent technologies working together to deliver high capacity data storage, performance and reliability. These innovative technologies do not exist simply for their own sake; every innovation is in service of delivering what customers need to turn their critical data into valuable information.

The Ultrastar® DC HC590 26TB[3] CMR HDD and Ultrastar DC HC690 32TB SMR HDD represent the next step in an innovation story stretching back decades. More recently, the growth of cloud storage/compute, social media, and generative AI have changed the game, driving the need for maximum capacity HDDs. Starting with the world's first helium-filled HDD, and extending the promise of helium-filled storage capacities with the world's first 11-disk HDD, and a host of technologies in between, Western Digital is continuously innovating to meet the ever-expanding needs of a data-driven world.

1. Source: IDC Global Datasphere Forecast, 2024-2028, May 2024, US52076424
2. Source: IDC Worldwide Global StorageSphere Forecast, 2024-2028, June 2024, US52312824
3. One terabyte (TB) is equal to one trillion bytes. Actual user capacity may be less due to operating environment.

HelioSeal® Technology

HelioSeal technology is the foundation for Ultrastar high capacity HDDs. Western Digital is the industry leader in helium, having shipped more than 125 million helium HDDs to date. The Ultrastar DC HC590 and DC HC690 are Western Digital's 10th generation drive with HelioSeal. HelioSeal provides a less dense atmosphere inside the HDD case, virtually eliminating aerodynamic turbulence. This allows read/write mechanisms to track more precisely and reliably over storage media, enabling higher performance, lower power and higher capacity with an increased number of media disks within the enclosure.

HelioSeal carries another benefit. Being a sealed enclosure, the HDD is less vulnerable to environmental contamination issues than air-filled drives, and less affected by the altitude limitations of air-filled drives. This property to altitude helped make history in 2017, when the Event Horizon Telescope collected the data leading to the world's first ever image of a black hole[4]. The massive amounts of data generated by high-altitude observatories require local storage, but air-filled HDDs were less reliable to operate at high altitudes, due to aerodynamic issues created by the lower air pressure. HelioSeal HDDs provided the storage necessary to collect and analyze the data that produced this historic image.

125M

Since its introduction in 2013, more than 125 million HelioSeal drives have been deployed through ten generations of proven reliability.

Triple Stage Actuator

Western Digital was first to integrate the triple stage actuator (TSA) in 18TB CMR and 20TB SMR HDDs. A TSA makes use of three pivot points: The Voice Coil Motor (VCM) Actuator, the milliactuator and the microactuator. The milliactuator and microactuator each has a pair of piezoelectric elements (or "piezo elements") attached to the different components of the suspension. When a voltage is applied, one piezo element expands while the opposite piezo element contracts in each actuator, causing the milli- and microactuators to change the angle of the head relative to the track. The TSA can be visualized as the difference between having an arm with only a shoulder (the VCM), to having an arm with a shoulder (VCM), elbow (milliactuator), and wrist (microactuator), three independent areas of movement.

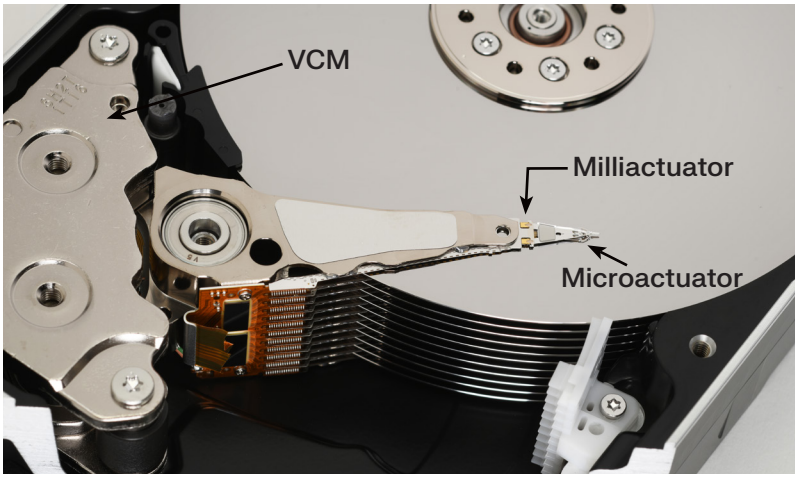

The TSA with three pivot points can increase granularity during track following and seeking. During track following the actuator stays on a particular data track to read/write data. The TSA decreases head position error to help mitigate the effect of vibrations from the internal spinning disks, seek-induced vibrations, as well as external system vibrations. This allows for more precise positioning of the head on the track. In addition, the microactuator can adjust the head angle for more optimal positioning over the tracks across the entirety of the actuator stroke.

During seek the actuator arm is moving from one track to another. As the arm approaches the target track, the milliactuator/microactuator are turned on in order to align to the target track faster. Finer positioning control enables the ability to arrive on-track more quickly, while reducing ringing and vibration.

These improvements in precision and performance enable higher tracks per inch (TPI). Higher TPI leads to higher areal density, which allows for more data on each disk and higher capacities.

4. https://blog.westerndigital.com/helium-filled-hdd-black-hole-image/

Energy-Assisted Magnetic Recording

Western Digital's 18TB CMR and 20TB SMR HDDs were the industry's first to use EAMR technology. Energy-assisted Perpendicular Magnetic Recording (ePMR) allows capacities to scale beyond legacy Perpendicular Magnetic Recording (PMR). ePMR was productized as part of Western Digital's research and characterization of EAMR technologies such as Microwave Assisted Magnetic Recording (MAMR) and Heat Assisted Magnetic Recording (HAMR) technologies.

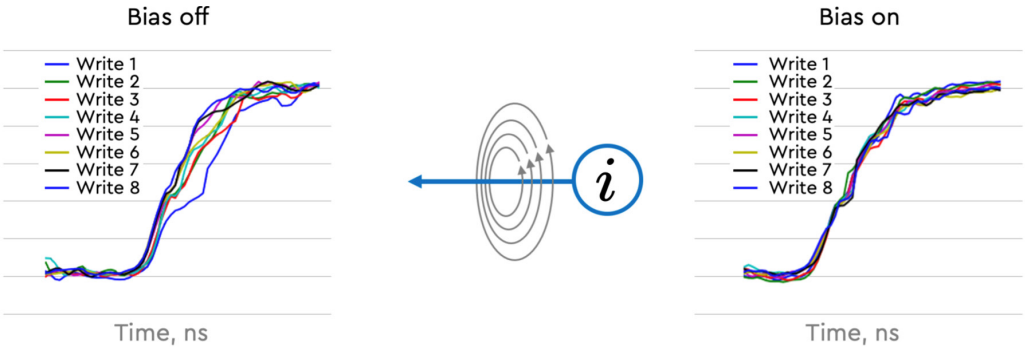

During write operations, recording heads do not reach saturation in a consistent way, and thus provide an inconsistent magnetic field on the media. This produces distortion in write currents, which is characterized as 'jitter'. ePMR applies an electrical current to the main pole of the write head throughout the write operation. This current generates an additional magnetic field which creates a preferred path for the magnetization flip of media bits. This, in turn, produces a more consistent write signal, significantly reducing jitter. When jitter is reduced it is possible to minimize the space between bits written, thus increasing BPI and areal density.

OptiNAND

Western Digital first launched OptiNAND technology with the Ultrastar HC550 20TB CMR HDDs for flash-enhanced drives by integrating NAND flash with traditional spinning disk media. The resulting HDD architecture incorporates innovative changes to the firmware algorithm and system-on-a-chip (SoC). These unique advances provide a time-to-capacity advantage over relying entirely on recording subsystem advances, allowing new capacities to reach market sooner and offering TCO improvements to our customers.

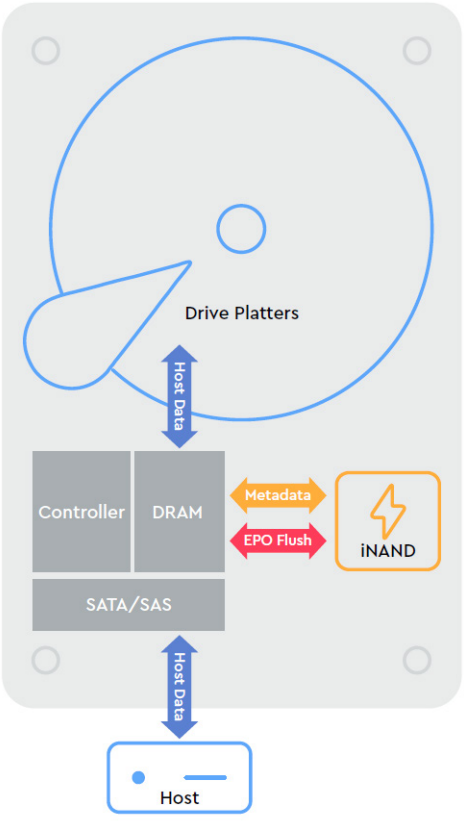

OptiNAND increases capacity first by moving some metadata from the media to the NAND flash device. Repeatable runout (RRO) is the portion of the position error signal that is repeatable for every spindle revolution. RRO metadata is generated in the factory during manufacturing. In prior generation HDDs, the RRO metadata would be stored on a disk, whereas OptiNAND stores this data in flash, freeing up disk space for the customer while enabling faster data access. Additionally, write operations are recorded to reduce adjacent track interference (ATI). In prior generation HDDs, write operations were recorded at the track level, while refreshes were done for entire tracks. OptiNAND records write operations in flash at the sector level. This metadata is used to refresh sectors instead of whole tracks. Eliminating excess refreshes allows tracks to be placed closer together without performance loss, thereby increasing capacity.

The addition of a flash device to the HDD also provides feature benefits. OptiNAND is not a hybrid HDD architecture, where the flash devices is ordinarily in the data path. However, data in DRAM that may need to be quickly stored in the case of an emergency power off (EPO) event can quickly be written to the flash device. This capability enables features like ArmorCache and UltraSMR.

ArmorCache

Throughout the history of hard disk drives, customers have been forced into a difficult decision about how to handle cached data. With Write Cache enabled (WCE) mode, the drive will achieve optimal write performance by accepting commands, signaling command completion to the system after caching data in DRAM, and writing them to media as soon as possible. However, if power is unexpectedly lost between the command being reported as complete and the data being written, the data will be lost. In Write Cache disabled (WCD) mode, the drive does not signal command completion for any writes that it cannot complete. This ensures that data is never lost, but it also means that the drive cannot cache any commands that it cannot guarantee will be written, which significantly limits write performance.

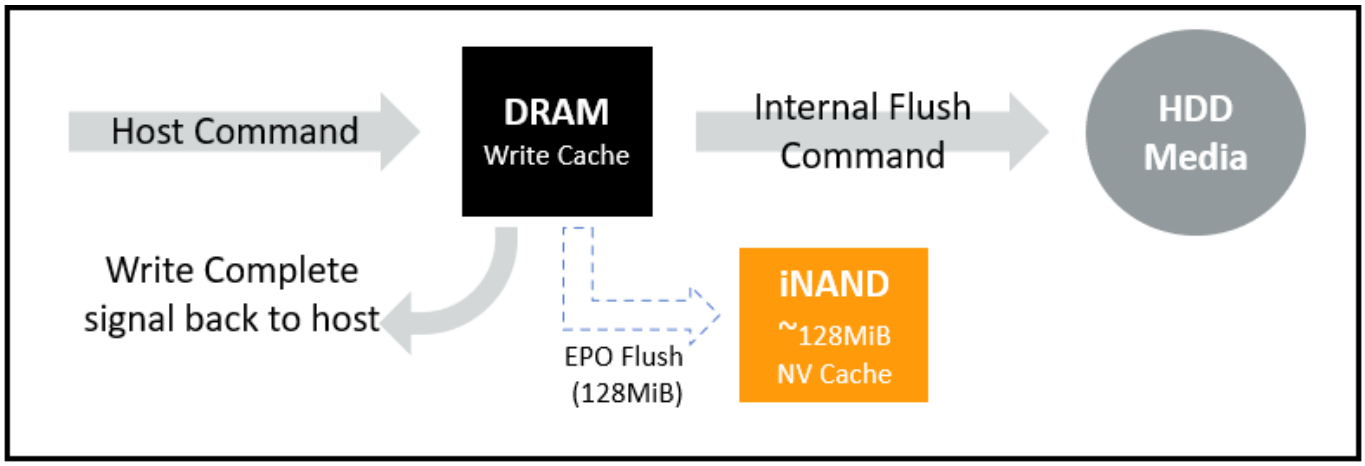

With the ArmorCache feature, this choice is no longer necessary. When operating in WCE mode, drives can ensure that the DRAM cache will be safely written upon loss of power, and no data is lost. When operating in WCD mode, a larger cache size guarantees more data will be safely written upon power loss, and the drive operates with WCE-equivalent performance. Whichever the system chooses, ArmorCache technology combines the performance of WCE mode with the data protection of WCD mode, offering the user the best of both worlds.

UltraSMR

Shingled Magnetic Recording (SMR) technology enables the highest capacity HDDs, for the world's largest scale deployments where storage TCO is critical. Physically, this is done by writing data sequentially, then overlapping (or "shingling") data with an adjacent track of data. By repeating this process, more data tracks can be placed on each magnetic surface than with conventional magnetic recording (CMR). Once one track has been written, the recording head is advanced by only part of its width, so the next track will partially overwrite the previous one, leaving only a narrow band for reading. Much like shingles on a roof that cannot be repaired without lifting the overlapping shingle(s), this overlap requires special handling from the host to write to the drive along unique rules that don't allow for overwrite in place.

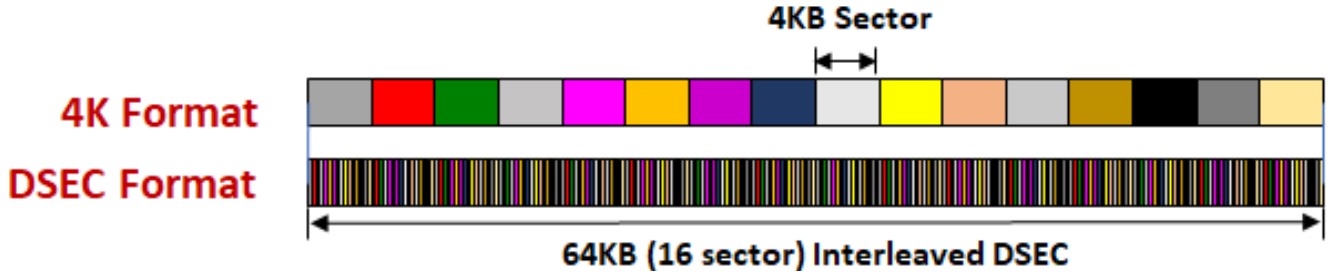

Starting with the 26TB Ultrastar DC HC670, UltraSMR expands upon SMR by utilizing novel data encoding techniques like Distributed Sector Format (DSEC) and Soft-Decoded Track ECC (sTECC) to increase the drive's error correction margin. These data encoding techniques reduce the likelihood of uncorrectable sector errors while simultaneously enhancing the ability of the drive to correct errors that would previously have been uncorrectable. This additional margin allows for tighter track pitch, increasing areal density. This leads to a greater capacity boost using UltraSMR relative to conventional SMR over CMR, increasing the TCO advantage justifying the adoption of SMR. The introduction of the 32TB Ultrastar DC HC690 HDD using UltraSMR offers a 6TB increase in capacity, or 23% larger, compared to the equivalent CMR HDD.

11-Disk Platform and Component Design

Innovations come in many varieties. But one of the most logical—and one of the key motivations for inventing HelioSeal—is that if you want to store more data, you put more disks into the HDD. Western Digital takes this idea to its next step with the Ultrastar DC HC590 and HC690 HDDs, the world's first 11-disk 3.5" form factor HDDs.

Adding an additional platter in the already crowded 3.5" HDD form factor is not trivial. To accomplish this, the spacing between disks was made narrower, which triggered multiple additional design changes. The TSA had to be redesigned to a low-profile version to fit within the narrower spacing, the base casting needed to be changed to provide more room for the additional disk, and the motor design changed to allow more room. Finally, the printed circuit board (PCB) needed to be redesigned and made to accommodate the new base casting design. In essence, the entire Z-height component stack needed to be altered to accommodate this new disk.

This was all done while maintaining excellent operational shock and vibration tolerance, and despite the added disk, power consumption is comparable to the previous generation 10-disk HDD.

Summary

HDDs are some of the most cutting edge, technologically advanced products made today. They incorporate truly interdisciplinary innovations in physics and engineering. Since the creation of the first HelioSeal HDD over a decade ago, Western Digital has been meeting the needs of our customers by innovating across this spectrum. From mechanical advances like the TSA and the groundbreaking 11-disk HDD, to recording technologies like ePMR and SMR, to novel advances pushing the envelope such as OptiNAND, ArmorCache, and UltraSMR that improve the efficiency of the mechanical and recording technology advances made. Each of these innovations helps provide high capacity, high performance, and reliable storage devices that result in our customers improving the TCO in their data centers.

Western Digital.

5601 Great Oaks Parkway
San Jose, CA 95119, USA
www.westerndigital.com