# UltraSMR Technology

## Introduction

Total cost of ownership (TCO) for storage is a critical metric for the demands of the modern data center. Every hard disk drive (HDD) in the data center requires power. Every watt of power those HDDs consume and transform into heat, requires cooling. Each HDD takes up physical volume and consumes a physical slot in a chassis—the "slot tax"—and must be supported by host bus adapter (HBA) port bandwidth. HDD TCO calculations are generally based upon metrics such as $/TB, W/TB, TB/in³, TB/slot, etc. Therefore, increasing capacity per drive typically reduces TCO.

Zoned Storage, which in HDDs refers to shingled magnetic recording (SMR) and, more specifically, host-managed shingled magnetic recording (HM-SMR), helps customers to better improve the TCO of the storage in their data centers. By overlapping the physical tracks on the media during write operations, an HM-SMR HDD results in higher capacity compared to a conventional magnetic recording (CMR) HDD of the same generation. While this requires software changes to properly exploit the capacity advantage of HM-SMR, the common TCO metrics are improved by Zoned Storage HDDs.

With UltraSMR, Western Digital has increased the capacity advantage of Zoned Storage HDDs relative to CMR HDDs. Previous generations of SMR HDDs on the market provided a capacity advantage of 11% or 2TB/HDD relative to CMR, while UltraSMR has extended this to 18% or 4TB/HDD in the same HDD generation. From a schedule perspective, SMR provided roughly one HDD generation capacity advantage relative to CMR. UltraSMR extends the capacity advantage to two HDD generations before the same capacity is available in CMR. Large customers, from a TCO perspective, can now justify the software investments needed to deploy Zoned Storage HDDs.

Western Digital's UltraSMR has accomplished this through a combination of recording subsystem technologies, advanced large block encoding and increased parity formats in the recording, the addition of OptiNAND™ as an enabling technology, and proprietary firmware and algorithms to make these all work together. The result is a capacity advantage that is driving TCO decisions for our customers towards Zoned Storage HDDs.

## Two-Dimensional Magnetic Recording

In nature, doubling of sensory organs is critical. Having two eyes, rather than one, allows us to see with depth perception. Is that lion in the distance coming towards me or moving away? Should I be scared? Having two ears, rather than one, allows us to identify location of sounds in our environment. I hear a stick snap and I need to determine whether the bear wanting to eat me is to my left or to my right. The ability to integrate data about our environment from two sources gives us much more information about what surrounds us in our environment.

An HDD read head is a sensor which detects the magnetic field of the track passing below. With a wide, well-defined track, a single read head can accurately provide data to the HDD controller to determine the bits originally written to the media. But as those tracks narrow and/or become shingled, and as the signal-to-noise ratio (SNR) lowers, a single read head may no longer be capable of recovering the written data. For an HDD, two heads are better than one!

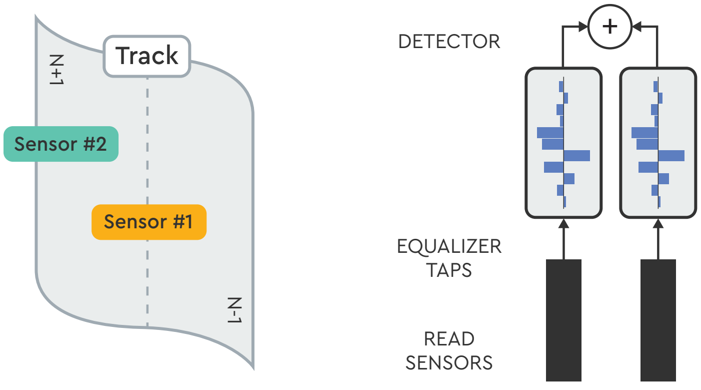


Figure 1: TDMR Head Skew and Noise Cancellation

Two-dimensional magnetic recording (TDMR) is an established technology that adds a second read head to the HDD slider as shown in Figure 1, offset from the first. This offset and the resulting differences in the signals provided by the two read heads allows the controller and firmware to combine the signals and filter out off-track noise, improving signal-to-noise ratio, and better cancel inter-track interference (ITI). Accordingly, tracks can be spaced closer together in CMR drives, and shingled tracks can overlap to a larger extent in SMR drives. This increases HDD tracks per inch (TPI), driving higher density and capacities.

## Distributed Sector Format

HDDs use extremely powerful low density parity check (LDPC) error correction algorithms to accurately read back written data from the media. As areal density rises, and TPI increases, the margin to read back sectors from the media reduces. Physical defects to the media and/or areas of weak write can overwhelm these algorithms, resulting in uncorrectable sector reads. This becomes exacerbated as track pitch decreases due to SMR.

To counteract this phenomenon, UltraSMR introduces distributed sectors (DSEC). As shown in Figure 2, distributed sectors spread data in logical sectors over a wider physical area, interleaving data from multiple sectors into one area. Distributed sector is ideal for SMR, as SMR forces sequential writes and does not allow update-in-place after a track is written. Thus, aggregating multiple sectors and distributing them into a larger block write will not incur performance penalties due to future read-modify-write operations, as update-in-place is not possible.

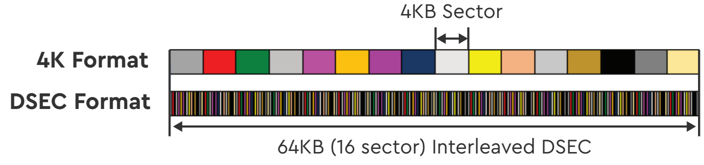


Figure 2: Traditional Sectors Distributed Over Larger Area

Physical defects or weak writes are generally localized. With standard sectors, that local disturbance can damage a large portion of one single sector, making it impossible to reconstruct, even if the sectors nearby are all readable. With distributed sectors, that local disturbance may damage a very small portion of many sectors, with each individual sector being easily readable with today's LDPC engines, as shown in Figure 3.

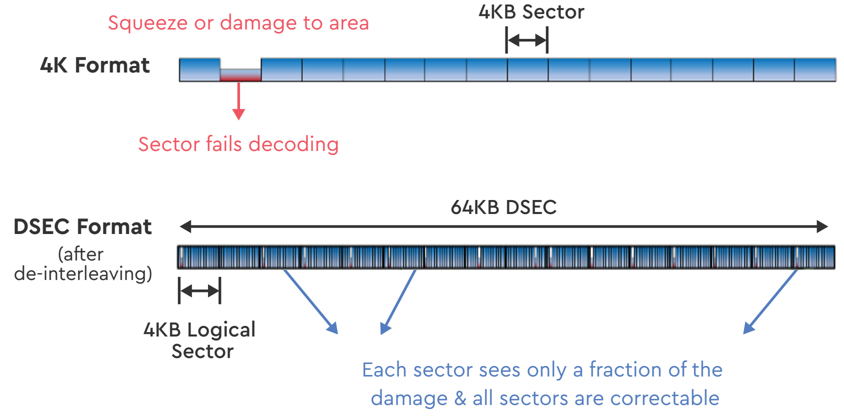


Figure 3: Enhanced Correction Of Damaged Sectors with Distributed Sector Encoding

Distributed sector spreads a logical sector over many physical sectors, improving SNR. Distributed sector, conversely, spreads a physical defect over many logical sectors, improving robustness. This regains the margin that would be lost by increasing TPI, and this margin allows for track overlap to be much tighter with UltraSMR than SMR.

## Soft-Decoded Track ECC

SMR HDDs have traditionally had parity sectors at the end of each track and the parity sectors are decoded logically by the controller by a process known as Hard Track ECC, also known as "erasure" mode. A single sector which exceeds the correction threshold of the sector-level LDPC engine and requires additional correction is assumed "erased" and regenerated by a combination of the other data sectors and the parity sector, as shown in Figure 4. For this, the other data sectors and the parity sector must all be read correctly, and the resultant regeneration is based upon the logical '0' or '1' value of the read bits in all these other sectors. This is a limited capability, as it relies on no more sectors to be degraded than the total number of parity sectors, so if you have 'N' parity sectors, you can only correct 'N' or fewer data sectors (≤'N' data sectors).

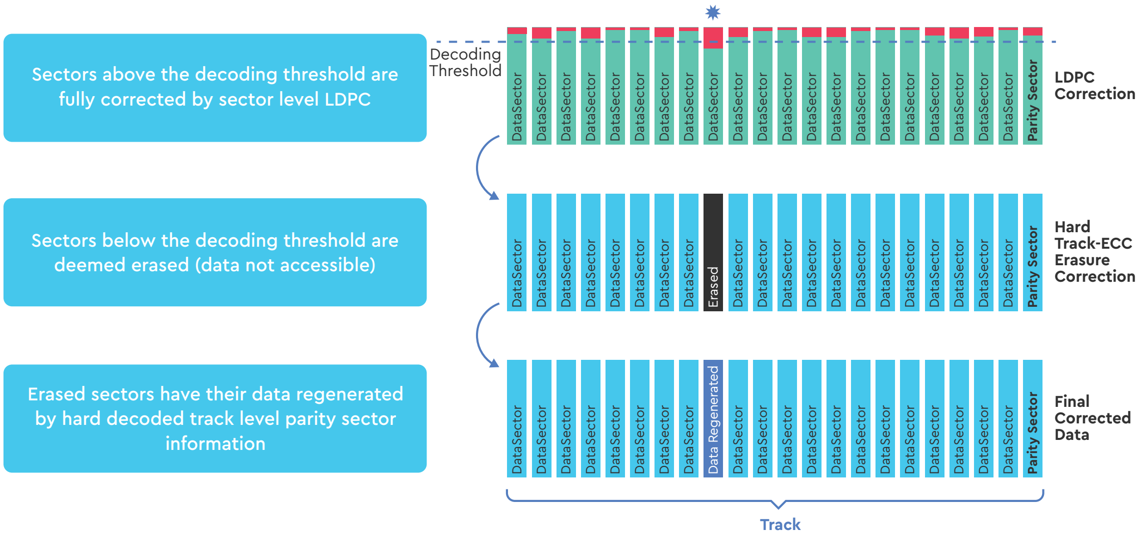


Figure 4: Hard Track ECC Decoding

UltraSMR introduces soft-decoded track ECC (sTECC), which is a correction scheme where the read channel has capability to utilize likelihood analysis of the uncorrectable sector(s) based on the strength of the read signal rather than assuming the sector itself is 100% bad. In the case of multiple unreadable sectors per parity sector, this can be an iterative process where bits that have a strong likelihood of '1' or '0' can be tested against the other sectors and parity sector and, where successful, can then be passed back through the sector-level LDPC engine. Sectors that have been substantially improved by the sTECC process now become correctable, and this increases the ability to correct the remaining sectors, as shown in Figure 5. As a result, a single parity sector can correct multiple uncorrectable data sectors, and the overall efficiency of the error correction capability can be improved by an order of magnitude or more.

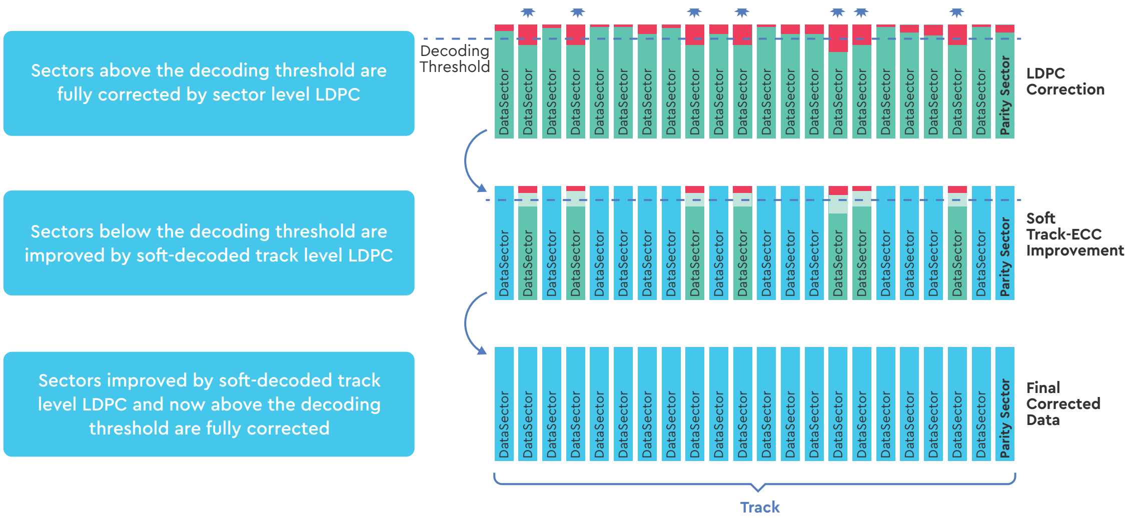


Figure 5: Improved Soft Track ECC Correction

Like the distributed sector, sTECC is ideal for SMR. The parity sectors are generated during track write by a bitwise XOR operation across all the data sectors on the track. Because an SMR write is sequential-only and update-in-place is not allowed, the parity sectors are only generated and written when the zone is written, and do not need to be updated or rewritten until the entire zone is rewritten.

sTECC provides significant error correction capability beyond the recording format advantage of distributed sector. Combined, the two add significant margin to the robustness of read operations in an UltraSMR HDD, which allows for more aggressive track spacing, improved areal density, and increased drive capacity.

## OptiNAND

OptiNAND technology is a significant enabler to UltraSMR due to the addition of an iNAND® embedded flash device (EFD) to the HDD. SMR HDDs allow for multiple active, or "open", write zones. Each write zone can be in a partial write state at any given time. To aggregate and hold data for distributed sector writes or sTECC parity sectors on each track of an open zone requires that if data is being held for a future write, it will be retained if power is lost before it is committed to the media.

While this data may regularly be held in DRAM during normal operation, in some cases it may require more storage or the ability to destage the data to non-volatile media (NVM), particularly upon emergency power loss. OptiNAND provides a large fast iNAND EFD that can handle all the transient data needed to enable UltraSMR.

## Conclusion

UltraSMR is a collection of hardware, controller, and read channel technologies that significantly expand the capacity advantage that SMR provides over CMR HDDs. Additional Western Digital proprietary firmware and servo advances are the glue that bring TDMR, distributed sector, sTECC, and OptiNAND together to achieve the capacity gains that our customers need.

Customers are increasingly driven by the need for higher capacities to meet the TCO requirements of their large data center deployments. With UltraSMR, Western Digital is providing the HDD capacities that satisfy these needs.

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