Western Digital.

TECH BRIEF

# Recording Technologies

The world is currently in the midst of a digital data explosion. IDC predicts that global data creation will rise from 64 zettabytes (ZB) in 2020 to almost 180ZB by 2025. Companies succeed globally by turning data into information, and using that information to make better decisions and profit in the marketplace. As a result, this growing data generation requires significant increases in the world's ability to store data.

Only a small portion of that 180ZB data created will be stored long-term. Much of that data will be used for its intended purpose and immediately discarded. Some will be retained short-term until it is no longer useful, and replaced by more recent data. But certain types of data are necessary or desirable to be retained over years or decades.

This creates enormous pain points for cloud service providers, enterprises, and consumers to store data. The entire data storage industry must increase efficient HDD storage capacity in response. It has created one of the most exciting periods in HDD history, as multiple technologies and recording formats are being introduced to solve this challenge for the world.

## Increasing HDD Capacity

### Areal Density

The driving factor in HDD capacity is, and always has been, areal density. Areal density is the amount of data that can be stored per square inch of a platter's surface, usually expressed in Gb/in$^2$ or Gbits/in$^2$.

There are many ways to increase HDD capacity. Using larger physical drive form factors (e.g., 2.5" vs 3.5" HDD form factors or increased vertical height such as 7mm vs 15mm 2.5" drives) allows for more or larger platters within a drive, increasing the capacity of an HDD without increasing areal density. Increasing the circumference of the platters within a drive, such as from 95mm to 97mm in a 3.5" form factor, creates more physical area on which to store data, and capacity can be increased without increasing areal density. Increasing the number of platters within each drive can increase the capacity of an HDD, but is subject to simple physical and volumetric limits. However, increasing areal density is the most significant driver of improved HDD capacity over time.

Figure 1 demonstrates the two key measurements in determining and increasing areal density: tracks per inch (TPI) and bits per inch (BPI). HDDs comprise concentric magnetic tracks around the media, with the track pitch defined as the spacing between the center of each track. Increasing TPI by squeezing the tracks more closely together, decreasing track pitch, will increase areal density.

Likewise, the bits themselves are collections of grains of material the width of a track and of a length that the read head requires to successfully read a unique value. These are aligned circumferentially along the track, and increased BPI is achieved by making the bits shorter, which increases areal density.

TPI and BPI are increased via recording formats, where the arrangement of the bits on the media within a specific recording technology is used more efficiently, and via recording technologies, where the magnetic properties of the head and media are adjusted to make the physical bits smaller.

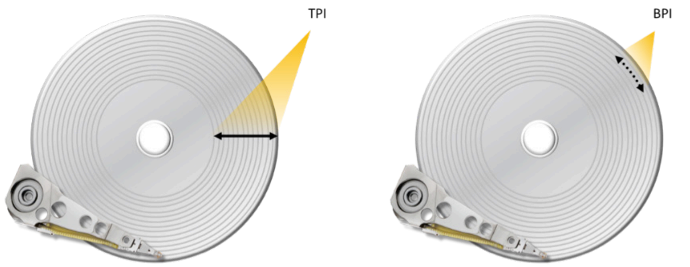


Figure 1: Visualization, tracks per inch (left) and bits per inch (right)

## Recording Formats

### Physical/Logical Sector Size

Early disk drives did not have standardized sector sizes. Each drive was simply a physical device and the drive control logic existed outside the hard disk drive, in the host. The number of bytes per sector was dependent on the HDD manufacturer, the operating system or application accessing the drive on the host, and what sort of error detection/correction mechanism the host used to ensure data integrity. In the 1980s, Western Digital invented the Integrated Drive Electronics (IDE) interface, which moved the drive controller into the physical drive. This created and standardized a command set for host-drive interaction, part of which defined a logical sector as 512 bytes long, with each byte consisting of 8 bits of data.

512-byte sectors require additional space beyond the 512 bytes of data stored. Part of the extra information is the error correcting code (ECC) needed to determine if the data that was read from the sector was read correctly. As areal density increases, the physical bits become smaller and harder to read, and read errors become more common. The ECC algorithms used over time have advanced and become more powerful to be able to identify and correct read errors.

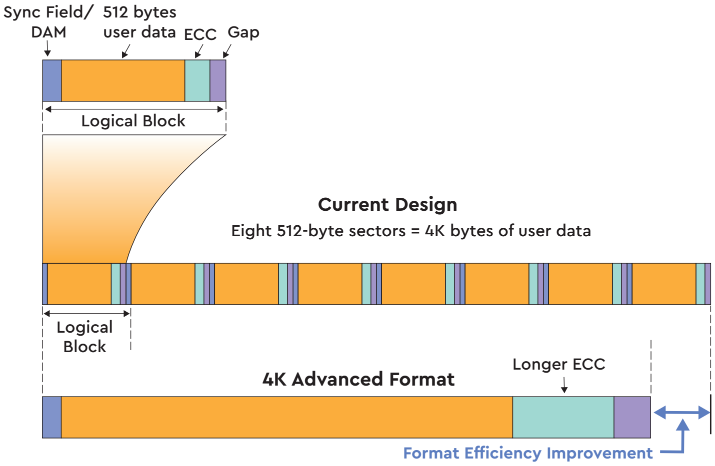


Figure 2: Eight 512-byte sectors compared to one 4096-byte sector

As the ECC algorithms increased in power and complexity, HDD manufacturers determined that it would be more efficient to increase the physical sector size from 512 bytes to 4096 bytes, or 4KB. Figure 2 demonstrates how ECC applied over a 4096-byte sector can be relatively smaller than over eight 512-byte sectors, improving format efficiency—the ratio of total drive capacity to user data. In essence, simply increasing the quantity of physical bits in a sector was a "free" way to increase usable areal density of an HDD. The standard to allow for 4KB physical sectors was completed in 2005, and HDD manufacturers started shipping drives with 4KB physical sectors in 2011.

Of course, nothing is truly free. In an ideal world, all host software would be rewritten to use 4KB logical sectors to match the HDD change, but that is not feasible. HDDs based on 4KB physical sectors needed to be backward-compatible to a host, operating system, and software ecosystem based upon 512B sectors. This requires the drive to emulate 512B logical sectors despite the physical change on the media. If the host does not write data to the drive aligned to the physical 4KB-sector boundaries, the drive is forced to perform a "read-modify-write" operation to integrate the new writes into existing physical sectors, which dramatically hurts drive performance.

Figure 3 outlines three ways in which we describe logical and physical sector formats. Some legacy and lower-capacity HDDs have continued to retain the 512B physical sector size. Because the physical and logical sector sizes are the same, these drives are described as 512B native (512n). Most larger drives have moved to 4096B physical sectors, which created issues as many host applications could not be rewritten to accept 4096B logical sectors. Through a significant effort, the ecosystem made the necessary changes to ensure that a host could know that a drive was emulating 512B logical sectors above a 4KB physical structure, known as 512-byte emulation (512e). The host would then be able to align its writes to the natural 4KB physical boundaries while still using 512B logical sectors, avoiding the read-modify-write operations. Modern hosts are now able to utilize 4KB physical sectors with 512B emulation without a performance penalty. While much of the storage ecosystem could not overhaul itself to switch to 4KB logical sectors, some host applications did make the change. Drives sold for those applications are called 4K-native (4Kn), as both the logical and physical sector size is 4096B. Today, all three drive types, 512n, 512e, and 4Kn, coexist in the market, depending on model and capacity.

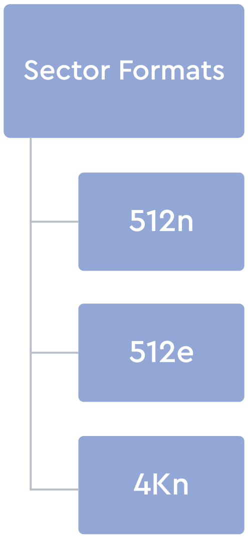


Figure 3: Three types of sector formats

## Track Layout

Recording data onto a magnetic media and reading the data off thereafter require different mechanisms and physical structures. Crucially, the reader head is physically narrower than the writer head, which means that tracks oriented based upon the width of the writer are sacrificing potential areal density.

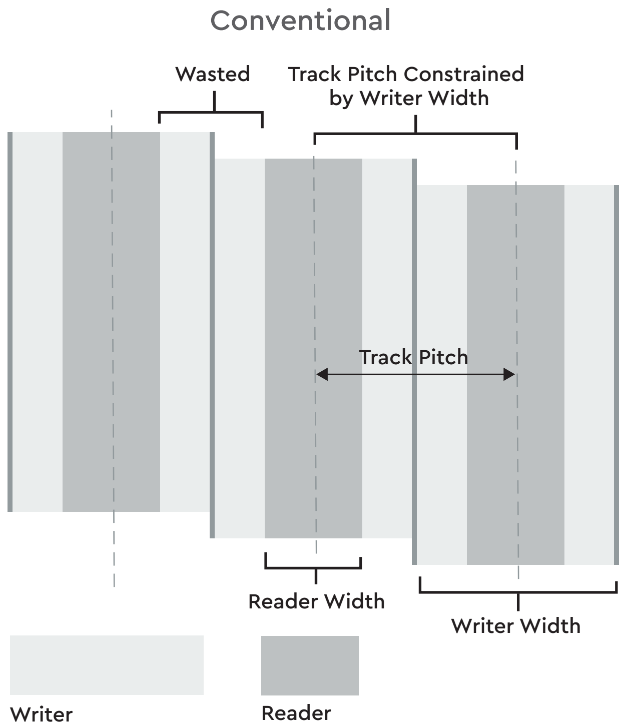


Figure 4: Writer head size and reader head size compared to track pitch in conventional recording format

HDDs historically have been architected based upon the writer width being the controlling factor in track pitch. Figure 4 demonstrates tracks that have been aligned on the circumference to be the width of the writer plus some guard band between tracks to avoid magnetic interference. This recording format has traditionally been described as PMR—more on that later—but with the introduction of the below shingled format, this is now referred to as conventional magnetic recording (CMR). The downside of CMR is that due to the guard band, and the natural difference in writer vs reader width, it is wasting space. CMR is not the most efficient use of space on an HDD platter that can be found. The upside is that having this guard band between tracks allows you to overwrite sectors individually. Thus, all data can be written and/or updated in place. For each piece of data, it has a logical block address (LBA), and that logical block address maps to a predictable physical location, with some unique exceptions. This has been the format for HDDs for decades, and thus the performance envelope of the CMR HDD is well known and predictable. Operating systems, software, and benchmarking utilities have been written with a certain expectation of performance that is consistent with a CMR format.

In order to maximize areal density, it is possible to take advantage of this disparity between the writer and reader widths. Crucially, if one can overlap the tracks such that the track is only slightly wider than the reader plus guard band, the track pitch can be tighter than is available in CMR. This is referred to as shingled magnetic recording (SMR), which alludes to the way shingles are placed onto a roof. Shingles on a roof are overlapped, which means that to repair an individual shingle, the shingles above it must be pulled up to make a repair. The same concept occurs in an SMR HDD, as illustrated in Figure 5. An existing piece of data cannot be overwritten without corrupting the overlapping tracks. The drive is organized into zones of several hundred megabytes, and to overwrite a single sector in a written zone requires that the entire zone be overwritten from the starting sector with the changed data. Because this is not tenable from a performance standpoint, the typical management for a single sector overwrite will be to write it to a new location and marking the previous location as obsolete. Thus, data organization must be handled in such a way that the LBA for a piece of data and its physical location may not have any predictable relationship. This is very similar to the way that SSDs are architected, as blocks must be erased before new data is written, and requires many of the same management techniques, such as garbage collection.

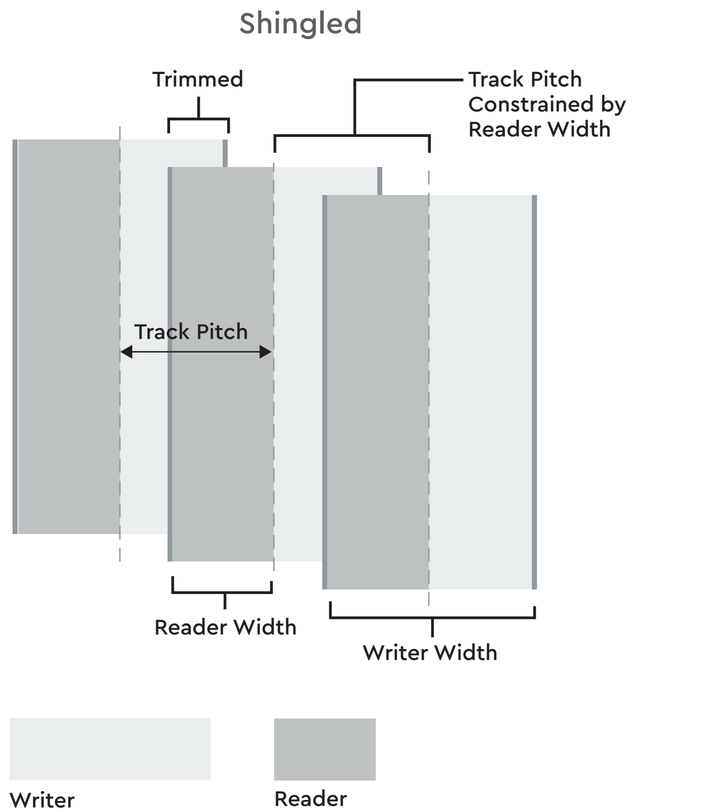


Figure 5: Track overlap in shingled magnetic recording format

There are two types of SMR HDDs, as shown in Figure 6. The first, drive-managed SMR (DM-SMR) appears to the host as a typical HDD block device, and the drive firmware manages the background translation table and garbage collection activities internally. DM-SMR drives are backward compatible with hosts expecting CMR. The second type are Zoned Storage HDDs, also known as host-managed SMR (HM-SMR). Zoned Storage designates a new device type to be reported to the host, and involves a unique command set, the zone ATA command set (ZAC) for SATA HDDs and zoned block command set (ZBC) for SAS HDDs. Zoned Storage is also used in SSDs, with products using the Zoned Namespace (ZNS) NVMe™ standard. In this case the host must be aware of the SMR structure and issue commands consistent with the rules for access of Zoned Storage. To use Zoned Storage, the operating system, file system, and software applications must be aware and architected to be SMR-friendly. DM-SMR is most suitable for client/PC systems, and other systems where the workload allows sufficient idle time for the drive to handle its management. Zoned Storage devices, due to the significant software impact of host management, are most suitable for enterprise and data center applications where the software stack and application are managed and written/tuned specifically for Zoned Storage.

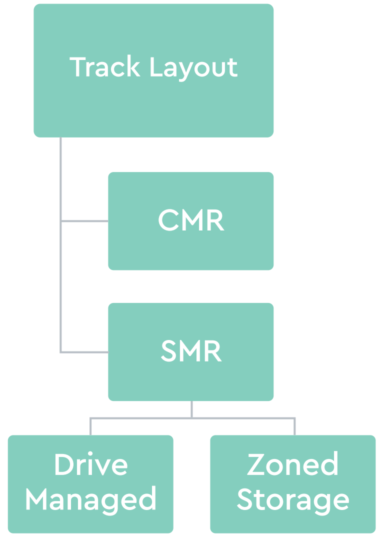


Figure 6: Types of track layout and SMR data management

As drive capacities continue to rise, SMR will be a valuable extension of areal density by taking advantage of the relative widths of the writer and reader. Much like the change to 4KB physical sectors, the transition to SMR—particularly Zoned Storage HDDs—will require significant software effort to ensure that drive performance can be maintained with a shingled track layout. As this software work occurs, it will open additional capacity needed for many applications.

# Recording Technologies

## Bit Orientation

In the mid-2000s, a major advance in HDD recording technology took place. This was the evolution from Longitudinal Magnetic Recording (LMR) to Perpendicular Magnetic Recording (PMR).

Prior to this time, LMR bits were written to the media along the surface of the track, akin to bar magnets being laid end-to-end with their magnetic north and south aligned circumferentially in the direction of the track. This process takes up significant area on the platter, and thus the technology had an upper limit roughly in the 100Gb/in$^2$ areal density range.

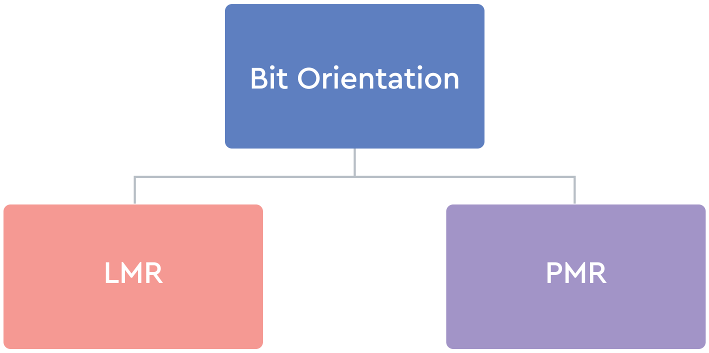


Figure 7: Two types of bit orientation

With the invention of PMR, the bits are written perpendicular to the platter surface, with the north or south pole being aligned vertically, akin to bar magnets being stacked together vertically like dominoes. Exposing only one end of the magnet saved a significant amount of space, which dramatically improved the capability for areal density. PMR allows for significantly higher BPI.

In the fast-moving world of data storage, a recording technology introduced in 2006 may seem ancient. However, it is important for one reason: all additional recording technologies discussed in this document—an alphabet soup of acronyms—are all PMR. Recording formats such as CMR or SMR, and energy-assisted recording technologies such as ePMR, MAMR, and HAMR, are all fundamentally underpinned by PMR in that every one of them uses the perpendicular recording structure inherent in PMR.

## The Trilemma

The push to higher areal density involves balancing multiple competing forces, known as the trilemma:

- To increase areal density, the size of the written bits must shrink. In order to maintain an acceptable signal-to-noise ratio (SNR), the size of the magnetic grains in the media must shrink. As this happens, the amount of energy needed to flip the magnetization drops.
- In order to keep the magnetization from flipping due to thermal energy, media materials with higher anisotropy (magnetic resistance to flipping) must be used to avoid unwanted changes.
- As the writer defines the size of the written bit, then smaller written bits require smaller writers, and smaller writers generally generate less field. However, to write data on high anisotropy media, the magnetic field generated by the write head must increase to overcome the anisotropy. This adjustment of the writer's field is done through changes to the geometry of the write head, increasing the moment of the write head material, and flying the write pole closer to the media.

With conventional recording, the trilemma is becoming more difficult to solve. The fly height of the head cannot be reduced much more, the smaller geometry of the writer is difficult to optimize further, and it already uses materials with the highest known magnetic moment. Thus, further increasing the media anisotropy doesn't seem to be an option, because the SNR degrades without the additional writer field.

## Energy-Assisted Magnetic Recording

To push through the challenges posed by the Trilemma, there are two options. The first is to find ways to apply additional energy to change the behavior of the magnetic write head such that the field produced either becomes stronger or more consistent. If the field is made stronger or more consistent, it allows higher anisotropy media to be used, and the bits can be made smaller. Western Digital's energy-assisted PMR (ePMR) would fall into this category. The second is to apply additional energy to temporarily modify attributes of the media, making it easier to write with conventional field strength. Two methods have been proposed to do this, microwave-assisted magnetic recording (MAMR) and heat-assisted magnetic recording (HAMR). Either allows for higher-anisotropy media, and a corresponding reduction in minimum bit size and increase in areal density.

## ePMR

In Western Digital's implementation of ePMR today, a DC bias current is applied through the write pole. As shown in figure 8, the bias current generates a magnetic field which creates a preferred path for the magnetization reversal in the write head. This increases the consistency of the write field from pass to pass over the media, thus reducing jitter and increasing signal-to-noise ratio. The additional consistency and predictability of the magnetic field allows bits to be written more cleanly and tracks packed more closely together, and thus increases areal density.

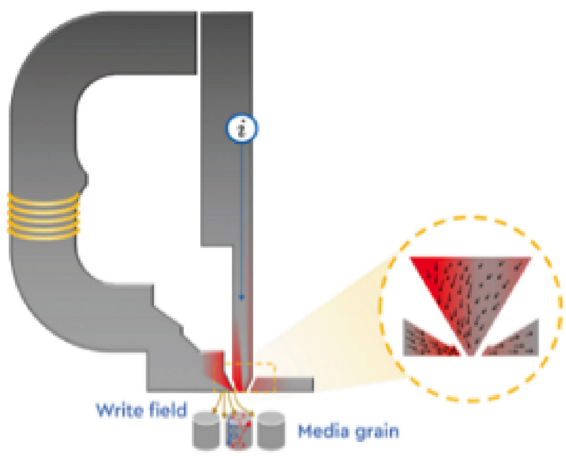


Figure 8: Recording head. Red = relative field coming from ePMR current

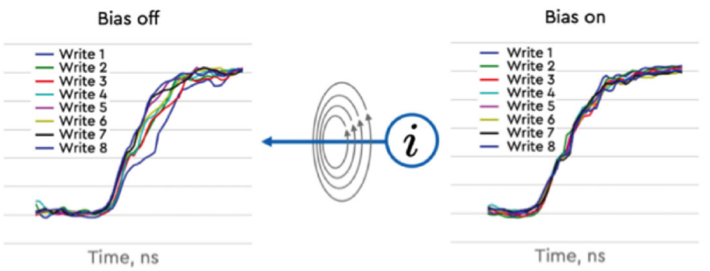


Figure 9: Jitter with and without ePMR

Additional development and design continue on ePMR, and future versions may increase write field strength in addition to providing a more consistent field. This will allow the use of higher-anisotropy media.

## HAMR

The goal of HAMR is to temporarily reduce the coercivity (i.e., the field required to switch the media) of the specific area intended to be written. By applying energy to the media itself, the grains can be temporarily made easier to flip, but retain the high-anisotropy grain's coercivity once the energy is removed.

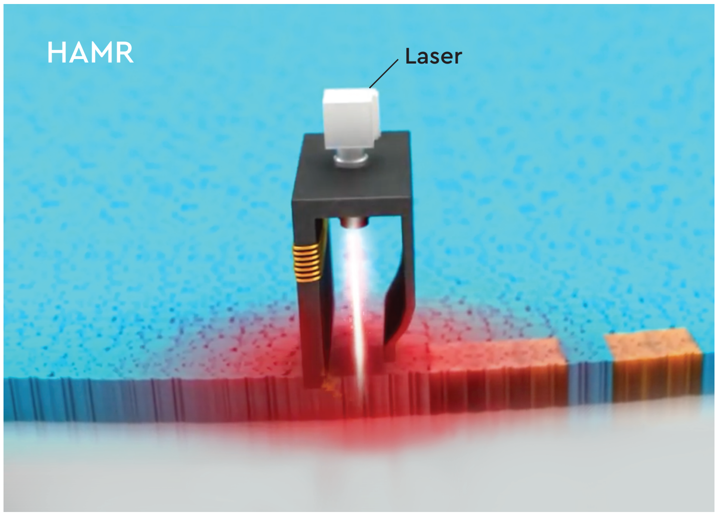


Figure 10: Heat assisted magnetic recording using a laser in the recording head

HAMR accomplishes this via heat. When the media temperature increases locally, the anisotropy goes down, and that region of media is easier to write. When it cools down, the anisotropy goes back up, and the media is harder to write and more thermally stable. Thus, by heating only the area needed to be written, a conventional magnetic write head can be used and still effectively write higher-anisotropy media than that used for conventional PMR. HAMR accomplishes this using a laser and an optical transducer to heat the local area of the media above its Curie temperature so that it loses its magnetic moment; the write field being applied during cooling then stabilizing the grains in the desired magnetic polarity.

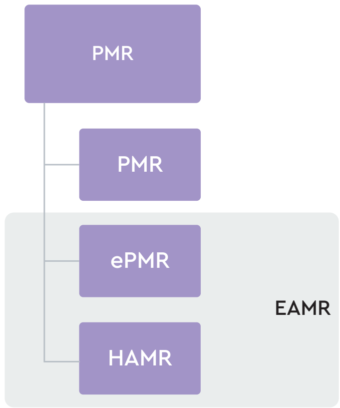


Figure 11: Three types of perpendicular magnetic recording. ePMR and HAMR are considered energy-assisted magnetic recording (EAMR).

MAMR continues to play an important part in Western Digital's overall HDD strategy. The ePMR innovations that we ship today in our leading HDD products result from years of research and development in energy-assisted methodologies, including MAMR and HAMR. The innovations in ePMR and those still to come, along with corollary technologies such as OptiNAND™, allow ePMR to be the bridge between PMR and the transition to HAMR.

## The Recording Taxonomy

The rise of SMR adoption and the introduction of new energy-assisted magnetic recording technologies are occurring independently, but they are not in competition with each other. Recording technologies can be used in combination with recording formats, as shown in Figure 12, and technologies like ePMR, MAMR, and HAMR will drive areal density forward for the future of HDD storage. SMR extends those areal density increases even further for the applications and workloads that work well within the performance characteristics of SMR drives.

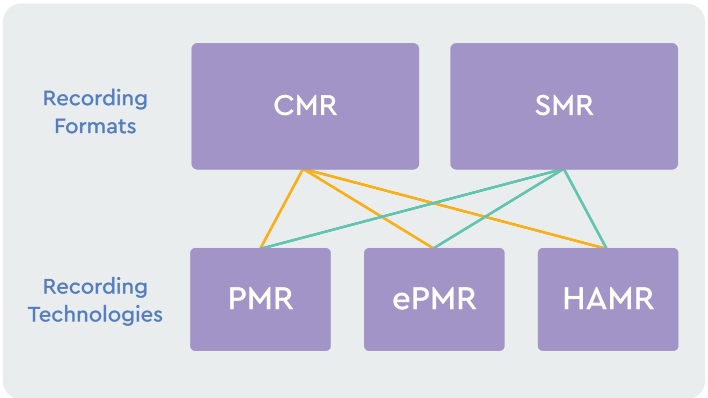


Figure 12: Recording format and recording technology combination options

Going forward, particularly in capacity enterprise, CMR and SMR will coexist regardless of the underlying recording technology used. Each application must balance the needs of capacity, performance, and cost.

The areal density benefits of SMR on top of an energy-assisted recording technology will be compelling for many use cases.

The combination of SMR and EAMR will drive HDD capacity up for the next decade and beyond. HDD technology has changed significantly over past years and continues at a pace of rapid development. New inventions and developments need to ensure that the world's rapidly expanding data will be able to be stored reliably and cost-effectively.

The world is creating incredible amounts of data annually and has put demands on the storage industry to meet the storage requirements of that data. From these new demands sprung a wide range of recording innovation, as outlined in the full recording taxonomy tree in Figure 13. The growing adoption of SMR and continued innovation in EAMR will help ensure that HDDs can meet those demands for many years to come.

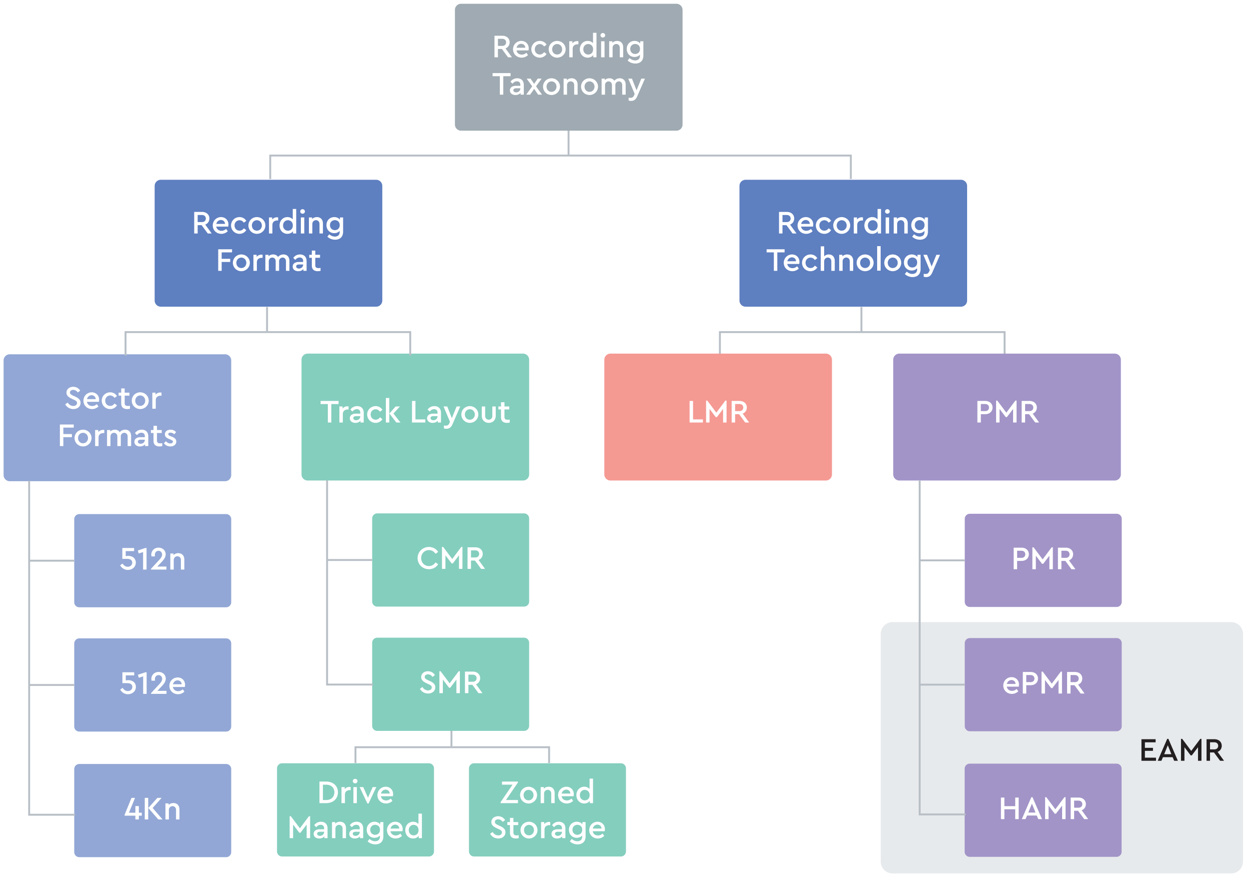


Figure 13: Full recording taxonomy tree



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



May 2022