Western Digital.

WHITE PAPER OCTOBER 2022

Shingled Magnetic Recording (SMR) HDD Technology

Achieve a better Total Cost of Ownership (TCO) with Western Digital Ultrastar® SMR HDDs for data-intensive workloads

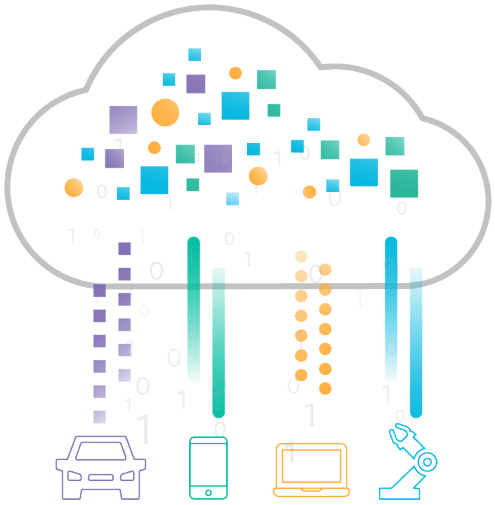

Figure 1: Zettabytes of data

Zoned Storage devices are a class of storage devices—both HDDs and SSDs—that enable host and storage devices to cooperate to achieve higher storage capacities, increased throughput and lower latencies.

As its name suggests, Zoned Storage involves organizing or partitioning the drives into zones. In the case of HDDs, Zoned Storage is implemented using a technology known as shingled magnetic recording (SMR). In the case of SSDs, Zoned Storage is implemented using a technology known as Zoned Namespaces (ZNS).

Learn more about Zoned Storage at http://zonedstorage.io/introduction/

Introduction

Data volumes generated by cloud services, video/video analytics, 5G-connected workloads, and mega digital platforms have created an almost unlimited opportunity in which data that is properly stored and analyzed may be the basis for valuable business insight and digital business transformation possibilities.

From data-in-flight transitions to data-at-rest, data storage becomes one of the key infrastructural components. However, the demand for scalable, cost-effective mass storage at an attractive economic price point continues to challenge the technological capabilities of the storage industry. The issue pressing data storage architectures stems from cases where the cost of storage may exceed the value of the data at hand, leaving data un-stored and untapped.

To meet these data demands, at a cost-effective Total Cost of Ownership (TCO), Western Digital offers its Ultrastar® Shingled Magnetic Recording (SMR) HDDs. SMR HDDs provide for higher areal density, and thus greater capacity, than Conventional Magnetic Recording (CMR) HDDs.

As Cold Data Warms, Data Center Storage Categories Evolve

The emergence of greater data volumes and data generation workloads has changed the method by which cloud and data architects view and approach data. With increased processing power and data analytics technologies, data that was previously in "deep archive" (data sitting offline on removable media) is migrating to an "active archive" (data kept online and accessible), where continuous value can be extracted from data sets. The simplistic notion of archived data as being data written to the cheapest media available and rarely accessed—namely tape drives—is rapidly changing.

In addition, many hyperscale and cloud storage customers are now beginning to realize that their workloads are trending toward data that is written sequentially and rarely updated, as opposed to random and frequent reads. In these instances, customers seek storage solutions with the lowest overall TCO, i.e., power consumption, rack density and dollar per terabyte. These active, sequential workloads have the potential to create a new paradigm opportunity for innovative storage solutions.

Figure 2: Unorganized data written to HDD

How Shingled Magnetic Recording Addresses the Capacity Challenge

Western Digital's SMR technology provides up to an additional 11% increase in areal density, compared with same-generation drives using CMR technology. Physically, this is achieved by writing the data sequentially, then overlapping (or "shingling") it with another track of data. By repeating this process, more data tracks can be placed on each magnetic surface. Conventional magnetic recording places gaps between recording tracks on HDDs to account for track mis-registration (TMR) budget. These separators impact areal density, as portions of the platter are not fully utilized. Shingled magnetic recording removes the gaps between tracks by sequentially writing tracks in an overlapping manner, forming a pattern similar to shingles on a roof.

Data Center Grade SMR: Performance Without Compromise

Some storage vendors achieve higher capacities and lower cost per terabyte by eliminating features from their drives. For example, reducing reliability specifications, product warranties and performance specifications is one way to drive down costs, however, the end result is desktop-class HDDs.

Data is the lifeblood of any business. That's why Western Digital believes cost savings cannot come at the expense of availability and reliability. Ultrastar SMR solutions maintain the industry's highest levels of reliability, quality and performance. Rather than eliminating features to achieve higher densities and cost savings, the Ultrastar solution actually adds features such as write-verify capabilities, design optimizations, and patented HelioSeal® technology to reduce both cost per terabyte of storage and TCO. The latest Ultrastar SMR HDDs utilize innovations such as energy-assisted magnetic recording technology and triple stage actuators to achieve higher capacities.

UltraSMR Technology

UltraSMR is a collection of hardware, controller, and read channel technologies that enable a +4TB capacity advantage over CMR HDDs.

- Two-Dimensional Magnetic Recording (TDMR)
- Distributed Sector Encoding (DSEC)
- Soft-Decoded Track ECC (sTECC)
- OptiNAND™

Additional Western Digital proprietary firmware and servo advances are the glue that bring these technologies together to achieve the capacity gains that our customers need.

The write head designed for SMR drives is wider than required to read a single track of data. 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. Overlapping tracks are grouped into bands (called zones) of fixed capacity for more effective data organization and partial update capability. Recording gaps between bands are laid to prevent data overwrite by the wide write head from one band to another.

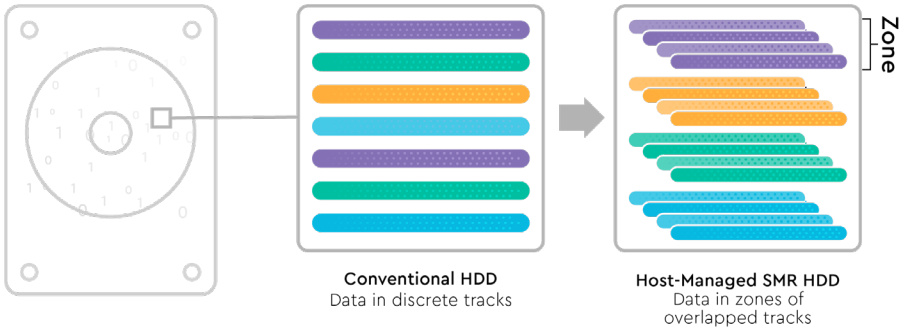

Figure 3: Comparison of conventional and SMR HDDs (Image source: ZonedStorage.io)

Fundamental Implications of SMR

Because of the shingled format of SMR, all data streams must be organized and written sequentially to the media. Consequently, should a particular track need to be modified or re-written, the entire "band" of tracks (zone) must be rewritten. Because the modified data is potentially under another "shingle" of data, direct modification is not permitted, unlike traditional CMR drives. In the case of SMR, the entire row of shingles above the modified track needs to be rewritten in the process.

SMR HDDs still provide true random-read capability allowing rapid data access similar to a traditional CMR drive. This makes SMR an excellent technology candidate for both active archive and higher-performance sequential workloads.

There are two types of SMR, drive-managed SMR and host-managed SMR. Drive-managed SMR emulates a traditional block storage device, ensuring compatibility with legacy systems. However, performance may be unpredictable in certain workloads. Host-managed SMR promotes SMR management to the host to help manage the data streams, read/write operations and zone management. This presents a new drive type to the system, requiring new ATA/SCSI command sets to utilize the drive. Ultrastar DC HC600-series data center HDDs utilize host-managed SMR.

Figure 4: Serialized data written to HDD

Host-Managed SMR

As the name implies, the host manages all write operations to be in sequence by following a write pointer. Once data is written to the zone, the write pointer increments to indicate the starting point of the next write operation in that zone. Any out-of-order writes, or writes to areas not properly indexed for writes by associated counters, will force the drive to abort the operation and flag an error. Recovery from such an error is the responsibility of the controlling host. This enforcement allows host-managed SMR to deliver predictable, consistent performance.

With host-managed SMR, data is organized in a number of SMR zones ranging from one to potentially many thousands. There are two types of SMR zones: a Sequential Write Required Zone and an optional Conventional Zone.

The Conventional Zone, which typically occupies a very small percentage of the overall drive capacity, can accept random writes but is typically used to store metadata. The Sequential Write Required zones occupy the majority of the overall drive capacity where the host enforces the sequential data stream of all write commands. (It should be noted that in host-managed SMR, random read commands are supported and perform comparably to that of standard CMR drives.)

Unlike drive-managed SMR, host-managed SMR is not backwards-compatible with legacy host storage stacks. However, host-hanaged SMR allows enterprises to maintain control and management of storage at the host level.

How to Implement Host-Managed SMR: High-level Considerations

Host-managed SMR is the first SMR implementation to gain enterprise market adoption, even though it requires customers to modify their storage software stack.

ZBC/ZAC Command Sets

First and foremost, a new specification of commands has been defined for host-managed SMR. These new command protocols are all standards-based and developed by the INCITS T10 and T13 committees. (There is no industry standard for Drive Managed SMR because it is purely transparent to hosts.) Please refer to the latest draft of the ZBC/ZAC when designing to host-managed SMR.

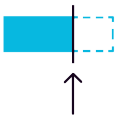

Optimizing Write

Data should be written at a zone write pointer position and in portions aligned to the disk physical sector size, and must be delivered in sequence without gaps within a zone. This pointer will be incremented at the completion of every write command. Consecutive writes must start from the last pointer address—otherwise the drive will flag an error and abort the operation.

Multiple SMR zones may be opened for writes at any time, where the exact number is a function of available disk internal resources. The maximum number is documented in each drive's identify data.

Read Is as Good as CMR

Reading data from an SMR drive is no different from traditional CMR drives. However, the reduced track width may impact settling times and re-read counts, which depend on the actual track width and the drive's tracking capability.

Performance Expectations

To ensure write operations do not corrupt previously written tracks, some host-managed SMR HDDs perform read verification of the previous track under certain conditions. This read-verification process ensures delivery of maximum capacity at maximum bit integrity, while delivering write throughput near but slightly lower than CMR HDDs.

Ecosystem Considerations

Due to the nature of host-managed SMR design, it is not a plug-and-play implementation with legacy systems. There are three ways to design for host-managed SMR depending on the customer's system structure and ability to modify their application layer.

1. Customers that own software applications and do not depend on the operating system and file system to control the HDD can start with new ZBC/ZAC command specifications that are newly defined for host-managed and host-aware SMR. These new command sets are based on T10/T13 committee standards. The applications will need to be rewritten with the new command sets as well as ensuring that all data streams are sequential.
2. Customers that run on open-source Linux®, which relies on its own kernel to control the HDD, can use recent Linux kernels and file systems which include ZBC/ZAC support. ZBC/ZAC support was included in the Linux kernel since version 4.10.0. Support for SMR within the btrfs file system was included since kernel version 5.12.0. Different off-the-shelf Linux distributions support SMR to varying levels. It is recommended to visit zonedstorage.io or work with the individual distribution maintainers to determine support in each distribution.
3. If a customer lacks the software control at the application level while retaining control over Linux kernel choices, an SMR device abstraction driver can be used to present SMR disks as regular CMR disks to the application layer. This feature has been provided by the dm-zoned device mapper driver since kernel version 4.13.0. With this solution, existing file systems and applications can be used without any modification.
4. If a customer lacks the software control at both the application and kernel levels, specialized hardware support to help manage SMR access constraints can be provided by components such as host-bus adaptors or intelligent disk enclosures. However, this customer must work directly with the HBA or enclosure vendors for SMR support.

Conclusion

Data and storage continue to grow at an unprecedented rate, hence the capacity-optimized HDD market has started to address new storage segments that are predominately sequential, highly accessible and at unprecedented capacity points. To effectively service these new segments, a purpose-built solution that leverages innovation, such as host-managed SMR, addresses the economical means for capacity data storage.

To capitalize on the capacity advantages, customers do need to make certain changes to accommodate for host-managed SMR – i.e., the host software (or end application) needs to be modified, all data streams need to be sequential and new command sets are required. Because host-based types of SMR will become more pervasive and ubiquitous in the future, the investment is fully leveraged for subsequent generations of SMR products and even recording technologies.

Western Digital's Ultrastar SMR HDDs are built with reliable HelioSeal and industry innovations to achieve the highest shipping capacities. Ultrastar SMR HDDs address the critical need for true enterprise-grade mass storage at a compelling TCO.

Learn More

zonedstorage.io

https://www.westerndigital.com/products/internal-drives/data-center-drives/ultrastar-dc-hc670-hdd#ultrastar-dc-hc670-26-tb

westerndigital.com/company/innovations/zoned-storage

documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/collateral/tech-brief/tech-brief-ultrasmr-technology.pdf

Western Digital.

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