Western Digital.

TECHNICAL BRIEF

Reimagining HDDs with OptiNAND™ Technology

Prepared by:
Western Digital

Introduction

The runaway train that is data proliferation shows no signs of slowing. More than twice the amount of data will be created in the years 2021–2026 than has been created since the advent of digital storage.[1] New sources such as AI/ML, blockchain, sensors, 5G networks, connected automobiles, and more are driving this growth—and with it, a rapidly increasing need for robust, reliable storage. Our customers are relying on Western Digital to help them meet this need. Western Digital's technological expertise and leadership, together with a diverse portfolio of high-performance HDD and flash products, make us uniquely positioned to help customers meet this demand at scale. We are now expanding our vast portfolio with an innovative solution created entirely in-house at Western Digital.

OptiNAND Technology

When it comes to cost-effective storage of data at scale, hard disk drives (HDDs) continue to play a central role. IDC projects that in 2025, HDDs will represent 82%[2] of capacity sold to the enterprise market. Investments in HDD technology remain critical to supporting worldwide data growth.

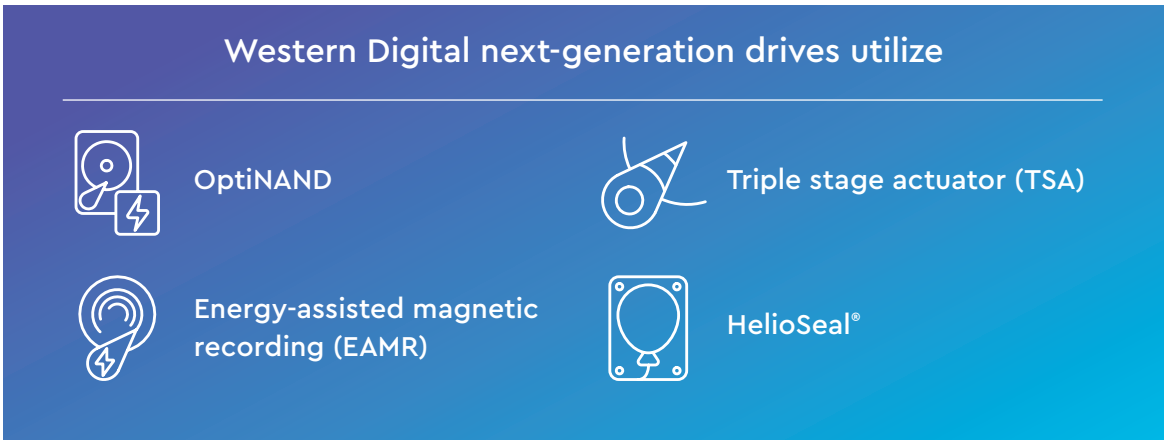

Western Digital has developed OptiNAND technology for flash-enhanced drives by vertically integrating our leading NAND flash with our world-class HDDs. OptiNAND integrates an iNAND® Universal Flash Storage (UFS) Embedded Flash Drive (EFD) with traditional spinning disk media, and incorporates innovative changes to the firmware algorithm and system-on-a-chip (SoC). OptiNAND is not a hybrid technology. This reimagined storage architecture brings together the two fundamental technologies of Western Digital to deliver a solution that will enable new innovations, forming the basis for future capacity, performance, and drive resilience gains.

HDD architectures now contain as much memory and processing power as PCs from the early 2000s. Before OptiNAND, non-volatile memory in an HDD was used primarily for booting and storing tiny amounts of metadata. As HDD storage architecture has become more sophisticated, the addition of a flash layer is a logical step in the system's memory hierarchy. Flash is more cost-effective than DRAM, with data persistence across power cycles. Flash also provides faster access than disk, enabling time-sensitive calculations to be performed while keeping the disk free to perform host operations. The new OptiNAND-enabled memory hierarchy utilizes the drive SoC to control communication with the iNAND EFD. With OptiNAND, key drive housekeeping functions can take advantage of an increase in metadata capability. This can reduce future DRAM needs as well as enable more sophisticated mechanisms to achieve greater capacities, increased performance, and enhanced reliability.

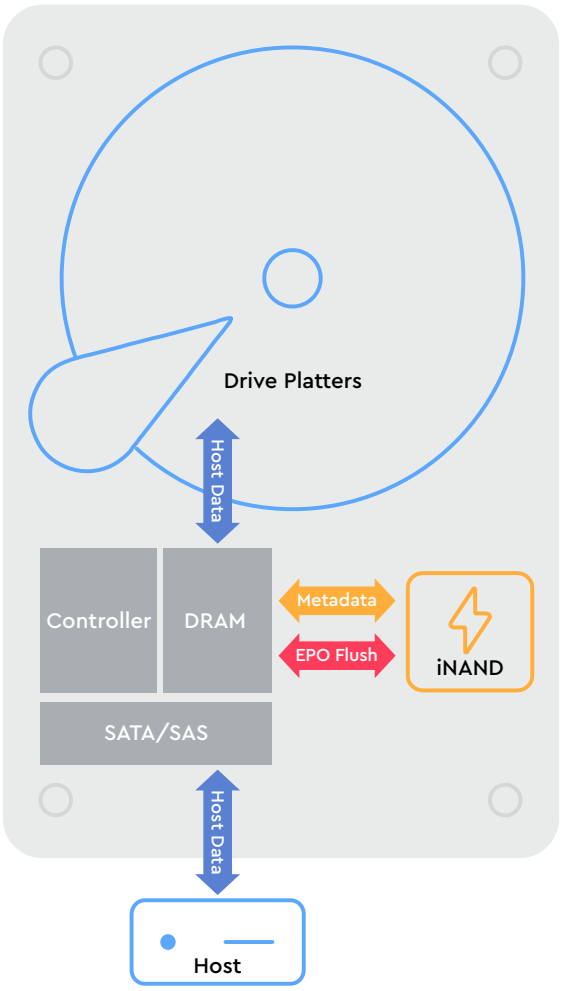

Capacity

OptiNAND, combined with triple stage actuator (TSA) technology, enables higher areal density through increased tracks per inch (TPI) to provide the highest capacities. HDDs generate gigabytes of metadata that can be utilized to increase areal density. This data is too large to be cost-effectively maintained in DRAM, while retrieving this data on demand from disk interferes with host operations and performance. OptiNAND enables cost-efficient storage and fast access to this massive quantity of metadata that can be stored and accessed in real time, freeing up valuable space on the rotating media for user data.

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 iNAND, freeing up disk space for the customer while enabling faster data access. 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 iNAND 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.

OptiNAND enables a broad portfolio of 20TB[3]+ CMR HDDs and 26TB+ SMR HDDs.

Performance

Many systems are intolerant of data loss on power failure, and those systems put the drive into Write Cache Disable (WCD) mode. The drive in WCD must be able to guarantee that all data for which the drive has issued a command completion is safe. As a result, the performance of those drives is limited to caching only the quantity of data that it knows can be safely saved during an Emergency Power Off (EPO).

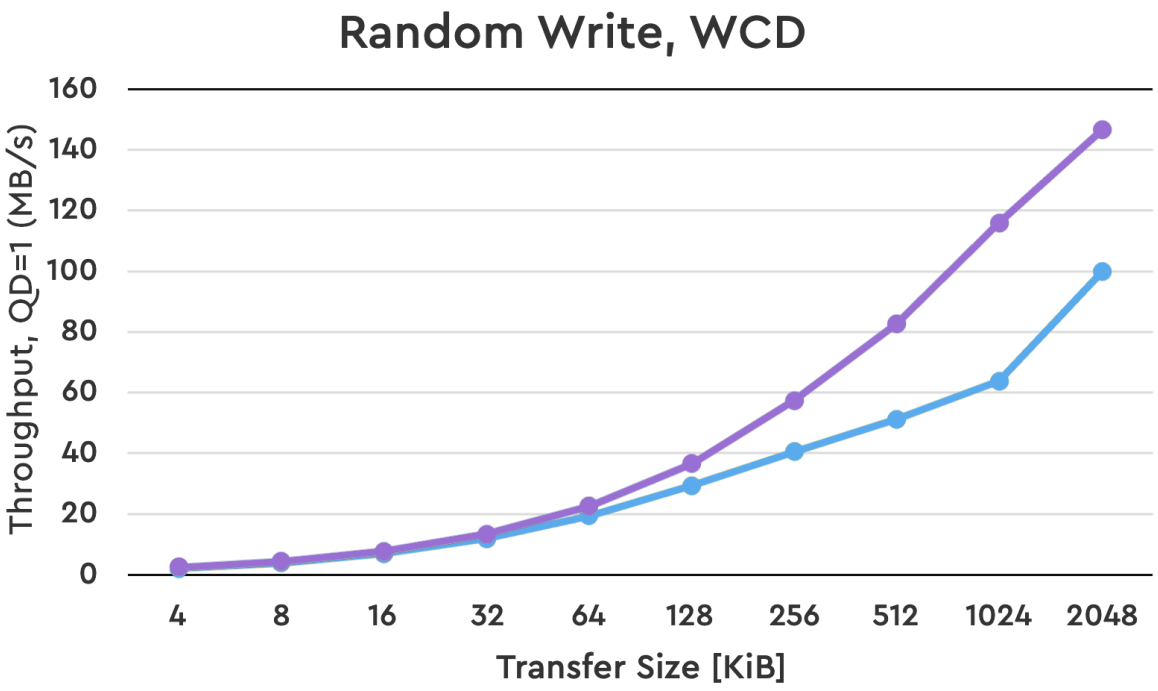

The ArmorCache™ feature of OptiNAND allocates 128MiB of NVM for EPO flush of customer data, with the result that more internally queued commands and ones of greater transfer length can be safely written than on drives without OptiNAND. Random write IOPS in WCD mode increases to the same level as Write Cache Enabled (WCE) mode. At the most common transfer lengths of 256KB or greater, IOPS improvement is even more significant compared to a drive without ArmorCache. The performance advantage of these OptiNAND drives is greatest in the most difficult and aggressive write workloads.

	WCE Mode		WCD Mode	
	Performance	Power Loss	Performance	Power Loss
Prior drives	High	Data at risk	Lower	Data is safe
HDDs with ArmorCache	High	Data is safe	High	Data is safe

Western Digital iNAND Technology

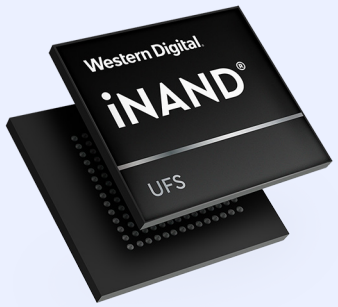

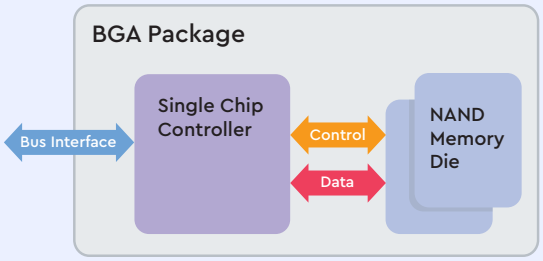

iNAND EFDs are built on Western Digital's proven 3D NAND flash and can be optimized to deliver the right combination of program/erase (P/E) cycles, performance, data retention, and temperature requirements.

iNAND products consist of:

- Single Chip Controller compliant with JEDEC specs
- Multiple NAND die, depending on configuration
- Single BGA package

NAND flash can store multiple bits within a single NAND cell. Single Level Cell (SLC) stores one bit/cell, while Triple Level Cell (TLC) can store 3 bits/cell. SLC offers higher performance, higher P/E cycles, and higher data retention margins compared to TLC, which enables 3X storage in the same footprint as SLC. iNAND technology can enable both SLC for write-heavy operations and TLC for read-heavy operations in the same device.

Data Protection and Drive Resilience

When operating in WCE mode, HDDs without ArmorCache require that the host system gracefully shut down before power is removed. This means that the host will typically send an ATA "Flush Cache" command or a SCSI "Synchronize Cache" command to inform the drive that any cached data needs to immediately be committed to media. The host would then not send any new data for a short amount of time in order to allow the drive to complete the writes, and then would finally power off.

With ArmorCache technology, host flush cache commands are no longer necessary to protect data, as the drive can guarantee the entire user cache can be flushed to iNAND upon EPO. Most HDDs ship into the market with a default of WCE. Unless a host system manually disables this to go into WCD mode, the drives will remain in WCE mode. ArmorCache is available on Western Digital 22TB+ HDDs and is the first and only HDD feature available that provides enterprise power loss protection (PLP) from the risk of data loss due to EPO in WCE mode.

In the event of an EPO, OptiNAND can also securely flush and retain nearly 50x more critical drive metadata than prior generation HDDs. This reduces the likelihood of long-tail failure modes that can occur on EPO, and make drive failure or data corruption less common.

Embedded Flash

Western Digital's vertical integration includes HDD and iNAND manufacturing (using NAND from a joint venture with Kioxia). This provides for adequate supply allocation and stable flash nodes. With Western Digital's unique capabilities in design, development, testing, and qualification of flash-enhanced drives, customers can count on the drive's reliability.

Summary

Western Digital has reimagined HDDs with OptiNAND technology, which integrates an iNAND EFD with traditional spinning disk drives and incorporates innovative changes to the firmware algorithm and SoC. These flash-enhanced drives feature a pioneering storage architecture that brings together the two fundamental technologies of Western Digital to deliver a solution that will enable new innovations and form the basis for future capacity, performance, and drive resilience gains. OptiNAND has broken through the conventional boundaries of storage, adding to Western Digital's legacy of industry-first technologies like EAMR, TSA, and HelioSeal, enabling customers to navigate the phenomenal worldwide growth in data.

Learn More

Energy-Assisted Magnetic Recording Technology for Higher HDD Capacities

Essential Mechanical Innovations to Drive Industry Leading Capacities

ArmorCache™ Write Cache Data Protection

NAND Glossary

SLC—Single Level Cell: one bit per storage cell

TLC—Triple Level Cell: three bits per storage cell

P/E cycles: the number of times a NAND device can be programmed and erased while maintaining its data retention requirement

Data Retention: the amount of time that a NAND device will retain its data

HDD Glossary

ATI—Adjacent Track Interference: interference from the proximity of tracks written next to each other on the media

EPO—Emergency Power Off: a sudden loss of power during which power stored inside the HDD during spin down is used to flush the write cache non-volatile cache to prevent data loss

RRO—Repeatable Run Out: portion of the position error signal that is repeatable for every spindle revolution; the irregularity of the servo track during servo track write can be caused by various disturbances, but some coming from sources like spindle motor are repeatable

TPI—Tracks per Inch: a measurement of the density of tracks on a hard drive spinning disk

ePMR—Energy-assisted PMR: an EAMR technology that applies an electrical current to the main pole of the write head during the write operation

[1] IDC press release, March 24, 2021: "Data Creation and Replication Will Grow at a Faster Rate than Installed Storage Capacity, according to the IDC Global DataSphere and StorageSphere Forecasts"

[2] IDC Worldwide Hard Disk Drive Forecast Update, 2021–2025 – Doc #US47633120; and IDC Worldwide Solid State Drive Forecast Update, 2021–2025, May 2021, Doc #US46412021

[3] One gigabyte (GB) is equal to one billion bytes and one terabyte is equal to one trillion bytes. Actual user capacity may be less due to operating environment.

Western Digital.

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