

Multi-Tier Caching Technology™

Technology Paper

How Layering an Application's Cache Improves Performance

Modern data storage needs go far beyond just computing. From creative professional environments to desktop systems, Seagate provides solutions for almost any application that requires large volumes of storage. As a leader in NAND, hybrid, SMR and conventional magnetic recording technologies, Seagate® applies different levels of caching and media optimization to benefit performance and capacity. Multi-Tier Caching (MTC) Technology brings the highest performance and areal density to a multitude of applications. Each Seagate product is uniquely tailored to meet the performance requirements of a specific use case with the right memory, NAND, and media type and size. This paper explains how MTC Technology works to optimize hard drive performance.

MTC Technology: Key Advantages

Capacity requirements can vary greatly from business to business. While the fastest performance can be achieved using Dynamic Random Access Memory (DRAM) cache, the data in the DRAM is not persistent through power cycles and DRAM is very expensive compared to other media. NAND flash data survives through power cycles but it is still very expensive compared to a magnetic storage medium. Magnetic storage media cache offers good performance at a very low cost. On the downside, media cache takes away overall disk drive capacity from PMR or SMR main store media.

MTC Technology solves this dilemma by using these diverse media components in combination to offer different levels of performance and capacity at varying price points. By carefully tuning firmware with appropriate cache types and sizes, the end user can experience excellent overall system performance. Since the disk drive on-board controller manages the data flow, the operating system and application work seamlessly in executing read/write operations.

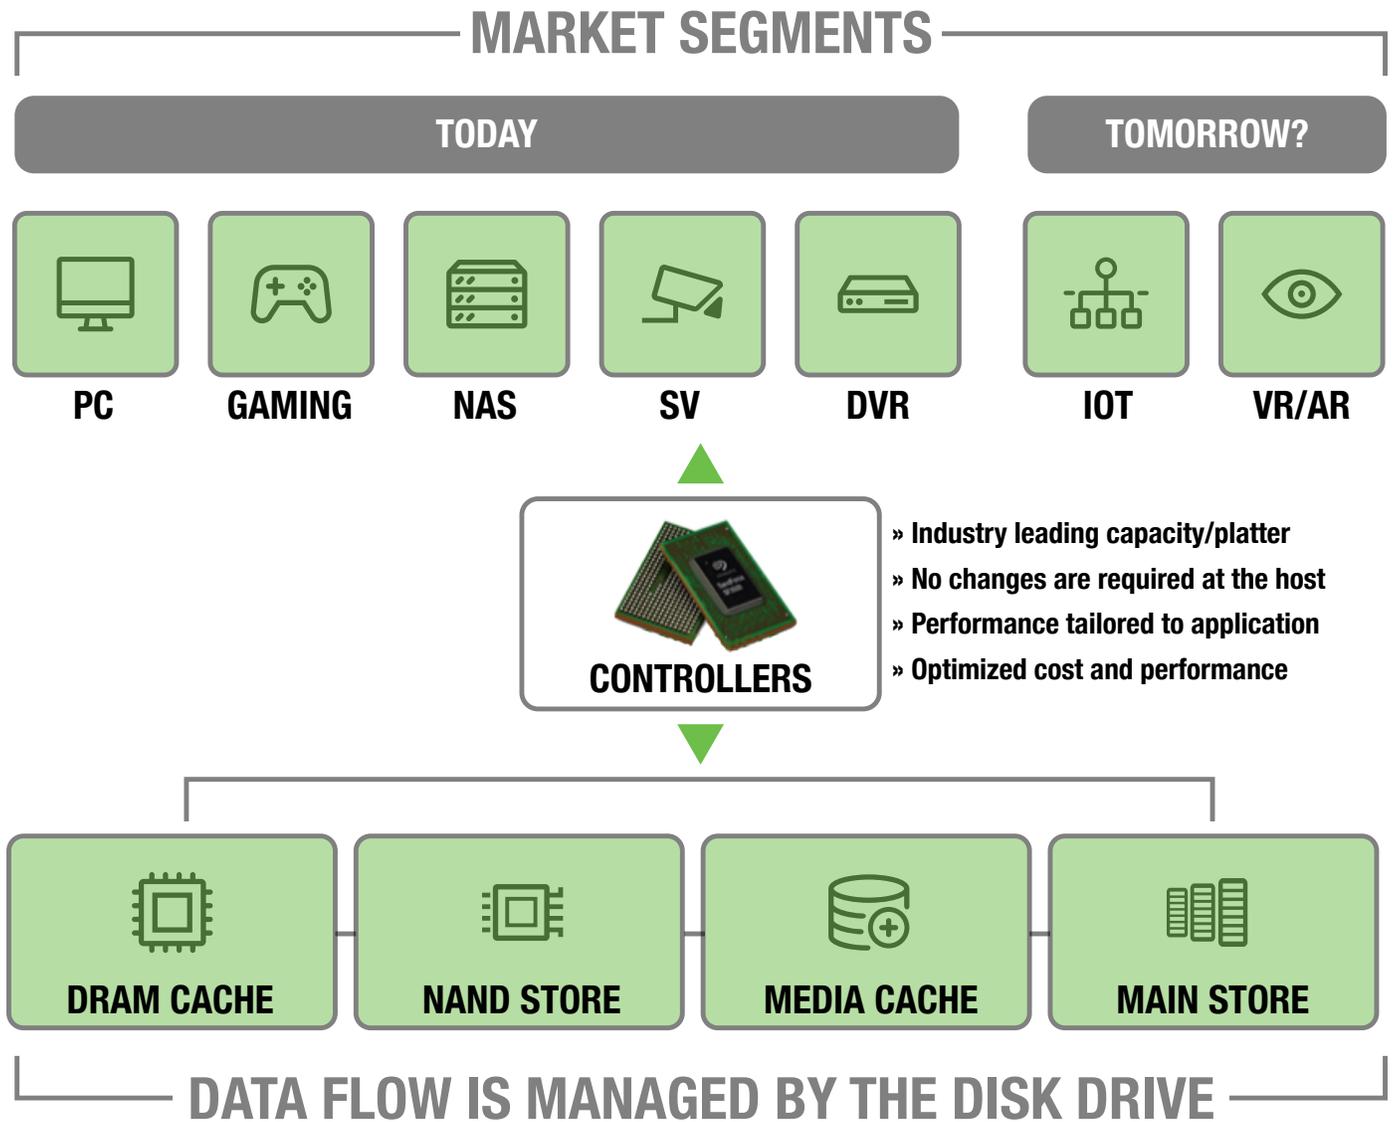


Figure 1. MTC Technology is the performance solution for the best daily end user experience.

Single-Tier vs. Multi-Tier Caching

Traditional disk drives use single-tier cache to speed up data flow from the host to the magnetic media main store—also known as permanent storage. In this scheme, shown in Figure 2, DRAM serves many purposes. A small DRAM buffers randomly written data and reorders them before writing to the main store. Reducing the number of seeks during a write improves random write performance significantly. During sequential write operations, such as streaming, DRAM helps to minimize fragmented writes by merging writes to disk. DRAM also serves as a read look-ahead cache, which can help certain random read operations. However, due to the limited size of the DRAM, the benefit of the read look-ahead cache is limited.

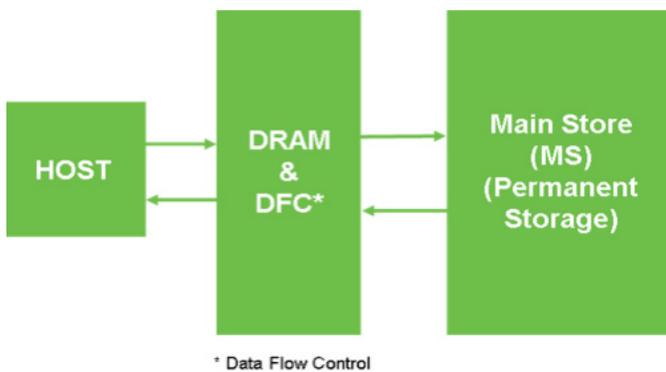


Figure 2. Traditional HDD Caching Data Flow

MTC technology adds various caching layers to further improve both read and write performances. Depending on Seagate's implementation, some or all of the caching tiers can be utilized for a specific product. In data shown in Figure 3, NAND flash and Media cache layers are added to provide additional performance boost. In this example, NAND flash is used to store most-often-read data to increase read performance while media cache is used to merge small block writes. Since the NAND flash size is in the order of gigabytes compared to the hundreds of megabytes found on DRAM, cache hit is higher by order of 100 times. For fragmented write operations, media cache boosts performance as DRAM reaches its limit. Media cache can sustain a near-sequential data rate of up to scores of gigabytes to better handle burst workloads. All user data are protected and moved in and out of the main store storage without host intervention, making the MTC Technology completely transparent to the user.

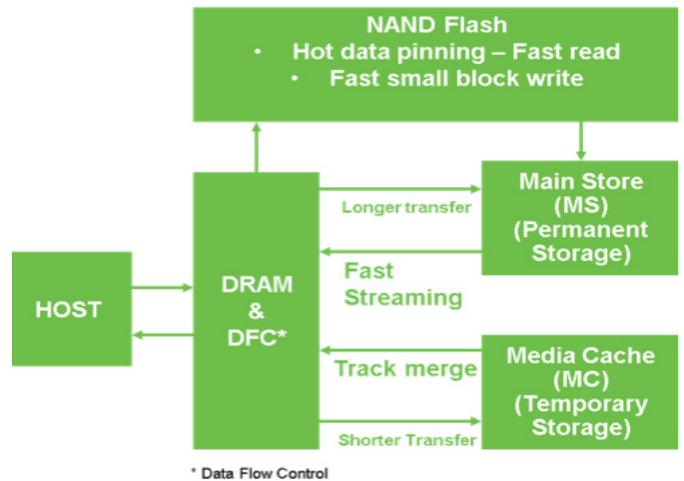


Figure 3. MTC Technology HDD Caching Data Flow

MTC Technology: Understanding the Layers of Performance and Capacity

As each layer of caching has an associated cost per storage unit, it is important to understand the performance metric of each component as well as the practical capacity provided by each component. Figure 4 shows the approximate buffering/storage capacity, performance metric as well as unique performance characteristics of each component. Combination of these components and firmware implementation are used to optimize the performance and capacity requirements for a given application.

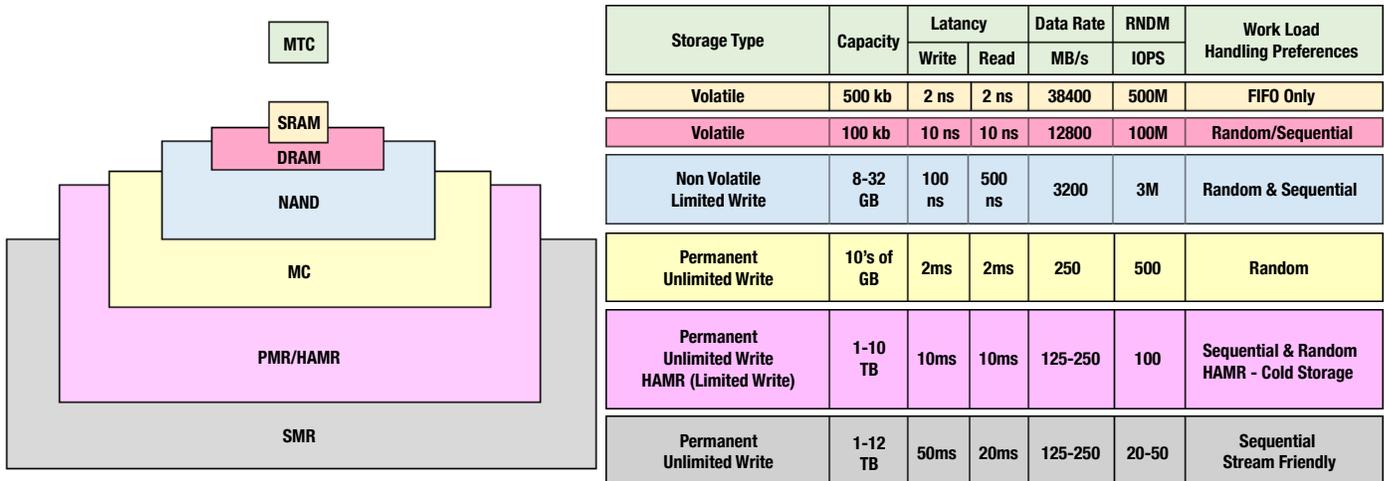


Figure 4. Capacity + Performance + Unique Characteristics + Firmware = High-Functioning Storage

Terms and Definitions			
MTC Technology (Multi-Tier Caching Technology) – Seagate’s trademarked caching scheme which utilizes multiple memory schemes and firmware to optimize data flow to enhance both read and write performance.		PMR – Perpendicular magnetic recording	SMR – Shingled magnetic recording
MC (Media Cache) – A temporary non-user addressable storage on magnetic media used to buffer incoming or outgoing data.	MS (Main Store) – A permanent user addressable storage on magnetic storage. Data can be either storage on PMR or SMR media.	DRAM (Dynamic Random Access Memory) – Volatile memory used to temporarily store data for computing purpose.	NAND Flash – Non-volatile solid state memory used to store bulk data.

MTC Technology Media Cache & Main Store Usage Cases

The usage of media cache and main store based on each application is illustrated in Table 1. Each application has typical data storage usage noted by file size and its inherent workload characteristics by sequential write, sequential read, random write and random read. These file sizes and workload characteristics drive cache size and type requirements. For example, compute environment deals with small files that are written and read randomly on the HDD. This workload is characterized by high likelihood of random write and random read. Media cache and other caching usages can be high, while the main store data storage requirement is fairly low. By contrast, single-drive surveillance applications deal with large files that are highly sequential read and write. This workload doesn't require much media cache to boost performance. It does, however, require lots of main store data storage capacity.

Table 1

	File size	SW	SR	RW	RR	MC usage	MS usage
Compute	MB	L	L	H	H	High	Low
Gaming	GB	M	M	L	L	Low	High
SV	TB	H	H	M	M	Low	High
NAS	GB	L	L	H	H	High	High
DVR	GB	H	M	L	L	Low	High

SW=Sequential Write SR=Sequential Read RW=Random Write RR=Random Read L=Low likelihood M=Somewhat likelihood H=High likelihood

Conclusion

Seagate is committed to delivering storage solutions with the highest performance and capacity available. Unique to the Seagate brand, MTC Technology breaks barriers by leveraging media cache, NAND and DRAM in a variety of hard drives for high-performance workloads. BarraCuda Pro, the essential HDD for creative professionals and high-performance desktops, is an excellent example of a drive that hits a high sequential data rate thanks to MTC Technology. Due to the adaptability and affordability of a MTC Technology approach, individuals and small and medium businesses alike can hone in on the full capabilities of their storage drive—ultimately leading to faster and more efficient workflows.

seagate.com

AMERICAS Seagate Technology LLC 10200 South De Anza Boulevard, Cupertino, California 95014, United States, 408-658-1000
 ASIA/PACIFIC Seagate Singapore International Headquarters Pte. Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888
 EUROPE, MIDDLE EAST AND AFRICA Seagate Technology SAS 16-18, rue du Dôme, 92100 Boulogne-Billancourt, France, 33 1-4186 10 00

© 2017 Seagate Technology LLC. All rights reserved. Printed in USA. Seagate, Seagate Technology and the Spiral logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. Multi-Tier Caching Technology is either a trademark or registered trademark of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners. When referring to drive capacity, one gigabyte, or GB, equals one billion bytes and one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Seagate reserves the right to change, without notice, product offerings or specifications. TP702.2-1705US, May 2017.