

Dropbox Magic Pocket Achieves Exabyte Scale with SMR HDDs



Challenge

- Dropbox needed to define and build a cost-effective, proprietary hybrid cloud model
- The new infrastructure needed to be dependable enough to fulfill Dropbox's value of being worthy of their customer's trust

Solution

- Magic Pocket – a multi-regional, exabyte-scale proprietary cloud capable of handling Dropbox's explosive growth in business and personal solutions
- Western Digital Ultrastar® DC HC600-series SMR HDDs for highest capacity, cost-effective storage

Key Results

- Up to 20% Cost Efficiency
- >99.9999999999% annual data durability and 99.99% data availability
- First wide-scale deployment of a cloud with SMR HDDs

The Explosive Growth of Dropbox

Dropbox is a global collaboration platform founded in 2007 with the simple premise that anyone should be able to access, edit, collaborate and share their files at any time and from anywhere. To appreciate how visionary this was, we can look back at the technology landscape in 2007. In 2007 the first smartphone was released with up to 8GB of storage, the first mainstream eReader was introduced, and HGST (later acquired by Western Digital) announced the first 1TB HDD. Files, images and videos were typically shared on USB sticks and the term 'cloud' was not part of our common vernacular.

Dropbox's success in creating a smart workspace that synchronizes files to the cloud has turned into over 700 million users across a spectrum of business and personal customers. Dropbox supports PCs, smartphones and tablets.



"One of Dropbox's key values is to be worthy of their customer's trust. This means they can never lose a user's data and never have an outage such that a user cannot access their data. Millions of customers access their files on Dropbox every day and they need the service to work every time," states Preslav Le, Principal Engineer.

Defining Magic Pocket

To deliver on this promise of trust and to accommodate their explosive growth, Dropbox needed to define and build its own file storage infrastructure. Magic Pocket is Dropbox's solution for an exabyte scale proprietary cloud to accommodate and personalize their products while driving efficiency and performance.

Magic Pocket is an immutable, content-addressable, block storage system. When a user uploads a file to Dropbox, the file is split into 4MB (maximum) immutable blocks. Immutable means that these blocks never change. This allows blocks to be moved around to increase efficiency and ensure reliability through redundancy. The system is content-addressable with the index/name of each block being a hash of the content of each block. A file system sits on top of this data to provide the interface to millions of users.

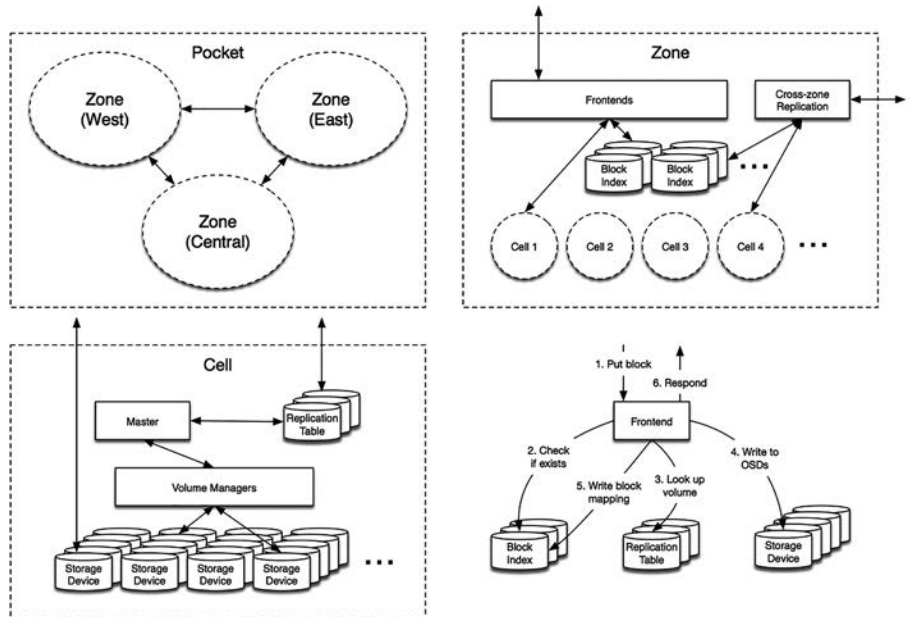
"In 2018, we became the first large technology company to adopt SMR at scale. We adopted these high-density drives to add hundreds of petabytes of capacity, resulting in significant cost savings over conventional PMR drives. Partnering with Western Digital has been a seamless experience, resulting in expedited adoption timelines."

Ali Zafar, Sr. Director of Platform Strategy and Operations

"For Dropbox, SMR is more about density, and more density means a lot to us."

Preslav Le, Principal Engineer

Magic Pocket extends across 3+ geographic regions with annual data durability >99.999999999% and data availability >99.99%. When it first launched, Dropbox transferred 500PB of customer data from a public cloud to Magic Pocket.



Shingled Magnetic Recording (SMR) HDDs Enable Magic Pocket

When evaluating the cost of infrastructure, storage can be one of the largest expenses. The question Dropbox needed to solve was how to pack more data into a fixed number of disks – Magic Pocket allows for up to 102 HDDs in a single chassis – while keeping all their data accessible with low latency. Even if a user wants to see an image stored several years ago, they want to access that image immediately.

Dropbox partnered with Western Digital to enable a leading-edge storage solution for Magic Pocket. SMR HDDs were chosen as the storage technology because they offered greater storage density compared to Conventional Magnetic Recording (CMR) HDDs. SMR HDDs enable greater bit density than CMR HDDs by overlaying data tracks. More data on a spinning disk produces higher capacity HDDs. Higher capacity HDDs mean that more data can be put into the fixed area of a server chassis, requiring fewer racks. Fewer racks saves on CPU, RAM, housing, and floor space and provides other benefits that all contribute to a better cost structure and lower total cost of ownership. Because all data needs to be accessible all of the time, HDD disks are always spinning. Higher data density with SMR requires fewer mechanical disks which ultimately helps with reliability.

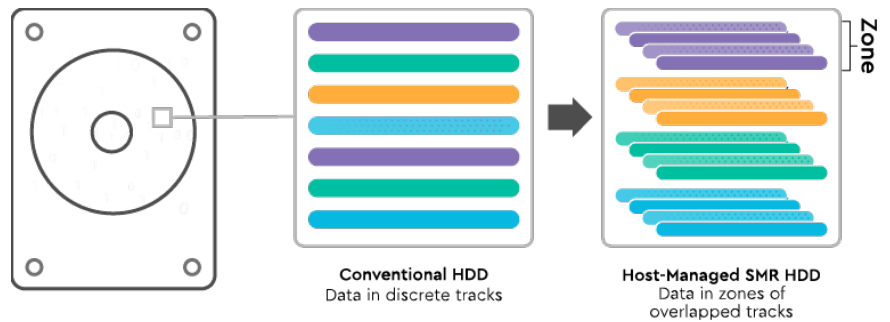


"For Dropbox, the result is more cost-efficient storage with a smaller energy footprint without sacrificing reliability or performance. The simplicity of our infrastructure has also set us up to take advantage of future innovations in data storage technology. Going forward, our efforts will pay dividends as we'll be able to use the same software infrastructure to improve densities with whatever emerging HDD technologies are coming."

Ali Zafar, Sr. Director of Platform Strategy and Operations

SMR HDDs replace random writes with sequential writes. Host-Managed SMR drives require the host to manage the sequential zones on its own. Hosts open, fill and close sequential zones. Host Managed SMR offers the most control over the way data is stored on the drive and is consistent with how Dropbox stores data. Caching random writes can either be done to a small CMR area at the outside diameter of the disk or by using a caching SSD.

Building the new Magic Pocket state-of-the-art storage system meant re-imagining and optimizing the hardware infrastructure and software stack.



Magic Pocket Hardware Optimizations

To accommodate SMR, Dropbox made several hardware optimizations

- They removed RAID from their architecture in favor of Host Bus Adapters (HBA). This decreased complexity and reduced vendors in their supply chain.
- Higher capacity HDDs meant more data per chassis, leading to more data being moved around. To accommodate this data, Dropbox implemented faster network access which had the benefit of improving overall access times.
- Each chassis utilizes an SSD for data caching.

Magic Pocket Software Optimizations

Dropbox realized their workloads are ideal for sequential writes and invested in re-writing their software to implement host-managed SMR

- Data traffic is divided into live writes and background data movement
 - Writes are sequential.
- Immutable block data is kept in sequential zones on the HDD.
 - Blocks of data are organized into 1GB extents, which are written in an append-only fashion. SMR write zones are 256MB, meaning that each Magic Pocket extent fits neatly into 4 zones.
- Metadata is kept in the CMR area of the HDD
 - The ratio of Dropbox's metadata size to file size is similar to the ratio of CMR to SMR tracks on the HDD.
- SSDs are used as the staging area for live writes while flushing data to the HDD in the background
- Proprietary erasure codes are customized for workloads
 - Attaining >99.999999999% durability.
- Optimized garbage collection overhead algorithms and eliminated the file system
 - Reduced cost by achieving over 90% disk utilization.



The Future Is Magic Pocket

Dropbox defined, built, tested and optimized Magic Pocket, an exabyte scale, proprietary hybrid cloud. They transferred 500PB of existing customer data from a public cloud to Magic Pocket with no downtime, no data loss and no performance impact. Dropbox is the first company to deploy SMR HDDs at scale in a data center environment. At one point, they were installing 30 to 40 racks/day into their data center every day with no downtime, no data loss and no performance impact. Dropbox partnered with Western Digital to qualify and deploy the Ultrastar DC HC600-series SMR HDDs.

Dropbox estimates they realized ~20% cost efficiency compared to a similar data center with CMR Drives. In 2019 Dropbox opened their Portland data center with all SMR drives, and in 2021, Dropbox will be deploying 20TB SMR drives to continue to expand their storage capabilities. Magic Pocket is an active project and is still being optimized. Dropbox is using this SMR architecture to more efficiently service tiers that are optimized for less frequently accessed data.

A data center architecture using SMR HDDs helps Dropbox fulfill their mission of customer trust, while increasing data density, delivering better cost structure and offering new opportunities.

¹ One gigabyte (GB) is equal to 1,000MB (one billion bytes) and one terabyte is equal to one trillion bytes. Actual user capacity may be less due to operating environment.

Western Digital.