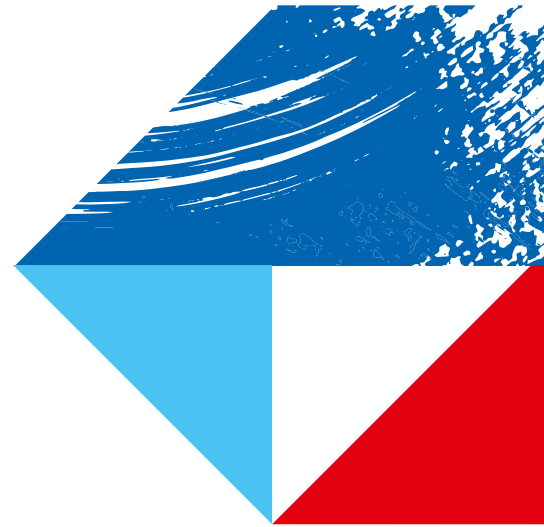


Taking Advantage of SMR based HDDs in Surveillance Camera Systems



Ongoing expansion of the surveillance camera market is leading to more intense demand for large-capacity, high-performance hard disk drives (HDDs). These HDDs are being used in numerous forms of video recorder hardware.

In response to this rising market demand, Toshiba has strengthened its S300 series of surveillance optimized HDDs. Available in 4TB and 6TB capacity options, they are based on shingled magnetic recording (SMR) technology.

These models offer the following key features:

- Higher storage density figures in a compact 3.5-inch form factor.
- High-performance capabilities that allow them to record video data simultaneously sent from up to 64 connected cameras to the SMR area on the disk without using a media cache (MC).
- Access to an additional conventional magnetic recording (CMR) area on the disk to directly record system data of a small size.

1. Introduction

As the resolution levels of video data increases and the image processing technology using artificial intelligence (AI) progresses, surveillance camera system deployments are becoming ever more commonplace. Throughout Europe, they are used to ensure people's safety, keep assets secure and address a broadening array of commercial and smart city analysis functions.

The volume of data generated by surveillance cameras is growing dramatically. This is due to the increase in resolutions that are being witnessed. At the same time, demand is increasing for surveillance camera systems capable of recording multiple video streams simultaneously. Both these dynamics drive the need to increase HDD capacity in the digital video recorders (DVRs) and network video recorders (NVRs) that are deployed in surveillance systems.

So that it can attend to the elevated data storage requirements of its customers, Toshiba has developed 3.5-inch Surveillance HDDs with SMR technology. These application optimized HDDs enable efficient and simultaneous recording of multiple video data

streams from cameras without using an intervening media cache. This is done by reserving part of a disk as a CMR area where system data, which are small in size, are directly written.

The following report describes the technical challenges in achieving the SMR performance required for Surveillance HDDs, as well as the technologies used to overcome these challenges.

2. Challenges for SMR HDDs incorporated into surveillance camera systems

2.1 Firmware specifications

Designed to work with an existing filesystem, the DVRs and NVRs in surveillance camera systems require HDDs that are compatible with the existing command sets. In addition, drive-managed (DM) firmware is needed so the CMR HDDs of existing systems can be replaced by high storage density SMR HDDs. With DM firmware, data streams from the host system are buffered in a MC, an area on a magnetic disk allocated during HDD operation. When the HDD runs short of free MC space, the data in the MC is written to a band, or a collection of overlapped tracks that are sequentially written (as illustrated Figure 1). If this occurs as a result of continuous HDD operation for an extended period, MC-to-band data transfer occurs frequently. This will lead to degradation of random-write performance⁽¹⁾. This is one of the key challenges that has to be tackled.

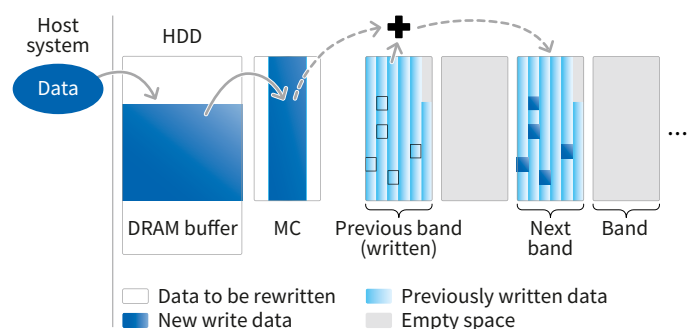


Figure 1. Data rewriting procedure in drive-managed SMR (DM-SMR) HDD – An MC is internally utilized for random writes that require rewriting of existing data.

2.2 Storage usage specific to surveillance camera systems

A surveillance camera system records video data streams from many cameras simultaneously in a single HDD or a redundant arrays of independent disk (RAID) arrangement. This is generally performed as follows. The disk space is managed as a collection of clusters, i.e., a unit of logical block address (LBA) space of several hundreds of MB to GB. Each cluster is assigned to one camera. A video data stream from a given camera is written sequentially from the beginning of a cluster. When it becomes full, a new cluster close to the current one is allocated to that camera. This process is repeated to record hours worth of video content (see Figure 2). Demand is growing for solutions to increase the number of cameras that can be connected to a surveillance camera system.

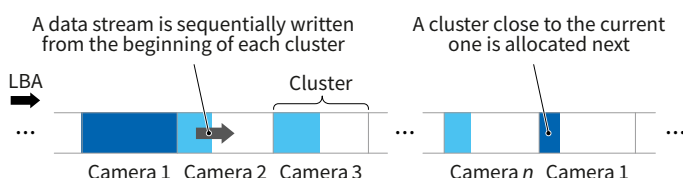


Figure 2. Example of video data arrangement in surveillance camera system
– A video data stream from a given camera is written from the beginning of a cluster. When the current cluster becomes full, a cluster close to it is allocated next.

3. Overview of the S300 series

Table 1 shows the main specifications of Toshiba's S300 3.5-inch SMR Surveillance HDD in 4TB and 6TB. To develop these models, it was necessary to make enhancements to the recording layer of the magnetic disk and the performance of read-write heads, so as to increase the linear recording density while using SMR to increase the track density. These S300 models incorporate DM firmware using existing HDD commands. This firmware is tailored to achieve high multi-stream performance (supporting connections of up to 64 surveillance cameras). The subsequent sections describe the fundamental technologies used to develop the S300 series.

S300 Surveillance HDD	6 TB	4 TB
	HDWT860UZSVA	HDWT840UZSVA
Form factor	3.5-inch	
Interface	SATA 6.0 Gbit/s	
Recording Technology	SMR	
Number of cameras supported	Up to 64	
Drive bays supported	Up to 8	
Rotation speed	5,400 rpm	
Max. data transfer speed	184 sustained MB/s Typ.	
Buffer size	256 MB	
Workloads	180 TB/year	
MTTF (hours)	1 000 000	
Power consumption (active idle)	2.81 W typ.	2.33 W typ.
External dimensions (mm)	147 (L) x 101.85 (W) x 26.1 (H)	
Weight	680 g	650 g

Table 1. Main specifications of S300 4TB and 6TB.

4. Video data recording methodology

4.1 SMR using spare bands and bypass writes

The S300 series records data in SMR format on the units of a band. An unused area is inserted between bands to prevent a write to a given band from corrupting data on an adjacent band. A collection of multiple bands is called a band group. Although all the bands in a band group have the same capacity, the capacity of a band depends on whether it is located in an outer or inner track. Each band group has one spare band, a work space where no valid data exists (as shown in Figure 3).

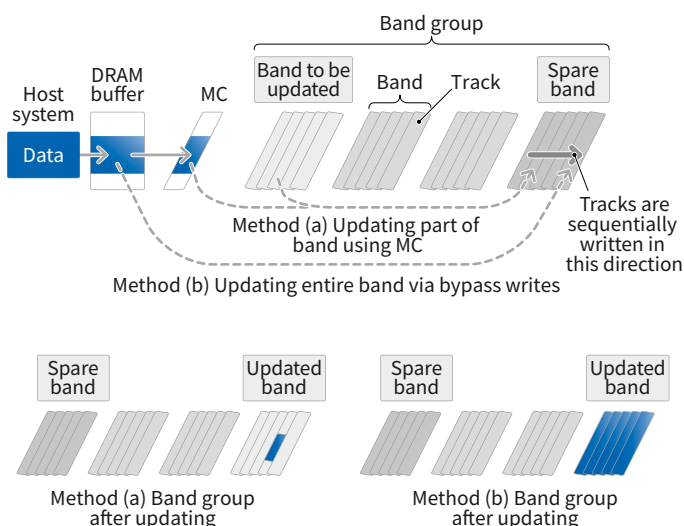


Figure 3. Data updating operation using spare bands – New data are written to a spare band. Data updating occurs safely as the firmware flags the updated band as a spare band and the previous spare as a valid band.

Data from the host system is buffered in a DRAM buffer first. The buffered data then gets transferred to a spare band in one of two ways. Which way will depend on the data size. The two options are:

(a) By updating part of a band - Here data in the DRAM buffer has first been written to the MC. Such data can be combined with the existing data in the band to be updated and moved to a spare band. Subsequently, the firmware will flag the original band as a spare band and the previous spare band as a valid band.

(b) By updating all the data in a band - When a large sequential data stream comes from the host system, data in the DRAM buffer is written directly to a spare band. When the end of the spare band is reached, the firmware flags that the band should be updated as a spare band and the previous spare band as a valid band. This process is then repeated.

Method (b), which transfers data directly to a band without using the MC, is called a bypass write. Although bypass writes can only be used to update all the data in a band, they are completed with minimal operation. Since most of the video data from a surveillance camera system is sequential, the key to improving the write performance lies in recording as much data as possible using bypass writes.

4.2 Bypass writes using free bands

As detailed in Section 2.2, a surveillance camera system records video data streams from many cameras in nearby bands simultaneously. To process such data efficiently, it is necessary to bypass write nearby bands (i.e. multiple bands in the same band group) simultaneously. Section 4.1 describes an example of a bypass write using a spare band.

A downside of this method is that only one stream of bypass writes can be performed for a given band group. Therefore, instead of spare bands, the Toshiba S300 series is designed to allow the use of free bands (i.e. the bands that hold obsolete data and can be overwritten) so that multiple bypass writes to a given band group can proceed simultaneously (refer to Figure 4)⁽²⁾.

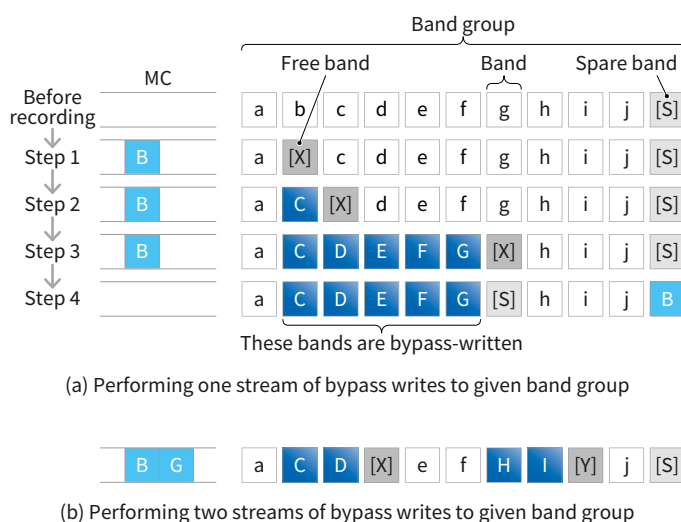


Figure 4. Bypass writing operations in band group using free bands – Free bands are generated when data are written to the MC. Bypass writes are performed using free bands.

First, let us look at how a stream of bypass writes is performed for a band group (Figure 4 (a)). The letters *a* to *j* denote existing data stored in each band, whereas *[S]* and *[X]* signify spare and free bands respectively.

The following paragraphs describe the procedure for sequentially updating the consecutive data of *b* to *g* with *B* to *G*.

- **Step 1.** B is written to the MC. At this point, *b* becomes obsolete, making its retention unnecessary. Therefore, the band in which *b* was held is flagged as a free band [X] so that it can be used to record other data.
- **Step 2.** C is written to [X]. This is performed as a bypass write, while leaving existing data *c* intact. So, the existing data is not lost even if the bypass write is terminated prematurely. When the writing of C is completed, the band in which *c* was held is flagged as a free band [X].
- **Step 3.** D, E, F, and G are sequentially written in the same manner as Step 2.
- **Step 4.** B in the MC is transferred to [S], and [X] is flagged as a spare band [S]. This means that [X] is deleted simultaneously.

Next, let us look at how multiple bypass writes are performed for the same band group (shown in Figure 4 (b)). In this example, two streams of bypass writes (B-C-D and G-H-I) are performed simultaneously. This process uses two free bands, [X] and [Y], generated as a result of writing B and G to the MC, in order to bypass-write C and H simultaneously.

Since any number of free bands can be generated when a band of data is written to the MC, three or more streams of bypass writes can proceed at the same time. Bypass writes are performed band by band. Since the band size has no bearing on the cluster size of a surveillance camera system, the leading and trailing ends of video data are not large enough to fill a band. Therefore, these fractional data elements are combined before being written to a band, in the same manner as Method (a) described in Section 4.1.

In this way, S300 series HDDs perform bypass writes using free bands to cater to the characteristics of video data streams from a surveillance camera system.

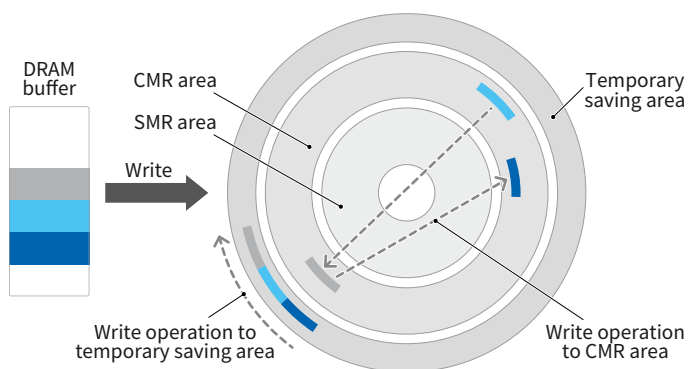


Figure 5. CMR and temporary saving areas on disk – The CMR area is suitable for recording systems data. The S300 series provides an area where they are temporarily saved in order to speed up flush-writes.

5. Recording system data in CMR format

In addition to video data, a surveillance camera system records management and analysis data (herein after collectively referred to as ‘system data’) simultaneously. System data consist of a multitude of small data blocks. Since system data is random in nature, it is unsuitable for DM-SMR. However, the amount of system data is growing because of the ever-increasing performance of surveillance camera systems.

To record the system data of a surveillance camera system efficiently, the S300 series provides a CMR area, which is located in a fixed LBA range. Although random data can be written directly to the CMR area, it is more important to guarantee the integrity of system data than that of video data. In order to do so, the surveillance camera system often requests an HDD to flush-write all the unwritten data, using the Flush Cache and Flush Cache EXT commands. Since typical HDDs cache the write data in a DRAM buffer, much of the write data remains in the DRAM buffer, waiting to be written to a disk. As there is a large amount of data to be flush-written, the flush-write operation must be performed at high speed, so as to record video data streams with adequate stability.

In response to a flush-write request, the S300 series writes all the unwritten data to a disk immediately. At this time, the write operation destined for placement in the SMR area can be completed within a short period, because the write data is cached in the MC. In contrast, the writes to the CMR area require a lot of time because the data destined for the CMR area will result in many seek operations and rotational delays. This is inevitable even if these writes are optimized in the order of write operations. To resolve this issue, S300 series HDDs each provide an area where the data destined for the CMR area may be temporarily saved (see Figure 5). This data is written to the temporary saving area in response to a flush-write request, so as to speed up the flush-write operation.

6. Conclusion

The S300 series of 3.5-inch SMR HDDs in 4TB and 6TB for surveillance camera systems provide customers with high density storage of 2TB per platter. In addition, these models incorporate firmware that is optimized for multi-stream surveillance environments to support connections of up to 64 cameras. With the growing prevalence of 4K (3,840 × 2,160 pixels) and 8K (7,680 × 4,320 pixels) ultra-high-definition (UHD) cameras, the volume of data generated by surveillance systems is expected to keep increasing, and the accompanying data storage resources must keep up.

Surveillance camera systems are also incorporating various innovations to take advantage of the characteristics of SMR HDDs. Therefore, demand for SMR HDDs in surveillance camera systems is expected to increase. In response to the growth of the Surveillance HDD market, Toshiba will continue to extend the capacity of its Surveillance HDDs to satisfy customer requirements.

References

Toshiba Corporation. "Large-Capacity, High-Performance 3.5-inch HDDs for Surveillance Camera Systems Applying SMR Technologies" Toshiba Review: 76 (6), November 2021
chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/https://toshiba.semicon-storage.com/content/dam/toshiba-ss-v3/master/en/company/technical-review/pdf/3_5-inch-hdd-smr_202303_en.pdf

⁽¹⁾ Shimomura, K. 2019. "Large-Capacity HDDs Applying SMR Technology for Data Centers." Toshiba Review: 74(6): 12–16. Accessed August 2, 2021.
https://www.global.toshiba/content/dam/toshiba/migration/corp/techReviewAssets/tech/review/2019/06/74_06pdf/a04.pdf.

⁽²⁾ Toshiba Corporation, Haga, T., Tanaka, H. 2016. "Disk storage device and method." Japanese Patent No. 5951472. Published July 13, 2016.