Envoy Data Memory

SSD CONTROLLER Technical Brief

Introduction

Solid State Drive (SSD) is a type of storage device which stores data in its embedded NAND Flash chips. Because of the physical characteristics of NAND Flash, the access latency of SSD is much shorter than traditional Hard Disk Drive (HDD). Besides, NAND Flash is a kind of Non-volatile Memory (NVM) which means the stored data will not be lost even the power supply has been cut off. To date, SSD has become one of the most popular storage devices in the market since its price has been more and more attractive. In other words, consumers are allowed to have a much better user experiences (ex: extremely short data access time) but still with affordable costs.

In the meanwhile, NAND Flash industry has advanced rapidly for several years. More and more multi-level cell (MLC) NAND Flash architectures have been proposed like 2-bit / 3-bit MLC Flash. These brand new technologies surely save lots of cost but they are also accompanied by some trade-offs. One obvious trade-off is that the latency of accessing MLC Flash is much slower than single-level cell (SLC) Flash. However, users always look forward to better user experiences. Therefore, SSD manufacturers utilize many methodologies to enhance the performance of their products. The concept of "Cache" is popularly implemented inside SSD to improve user experiences. After all, the thing user care the most is the transmission rate between Host (ex: Personal Computer) and Device (ex: SSD) instead of the timing their data is programmed into Flash "physically." Based on this idea, a general strategy is to put the received data into an area which is much faster than the main storage area (i.e. MLC NAND Flash) once any data was just transferred from Host. In this way, the processing time of data is able to be shortened significantly even though the data has not been programmed into the main storage area yet.

We know that Random Access Memory (RAM) can be regarded as one of the fastest storage media nowadays and that is the reason RAM has been commonly used to exchange data with Central Processing Unit (CPU). Also, RAM is also very popular to be adopted for cache implementations, no matter Static Random Access Memory (SRAM) or Dynamic Random Access Memory (DRAM). As mentioned previously, for pursuing better user experiences, cache techniques are also widely implemented in current SSD products. With appropriate designs and implementations, consumers are allowed to practically experience the advantage of accessing speed brought by RAM. Fig. 1 illustrates a basic idea of cache in SSD.

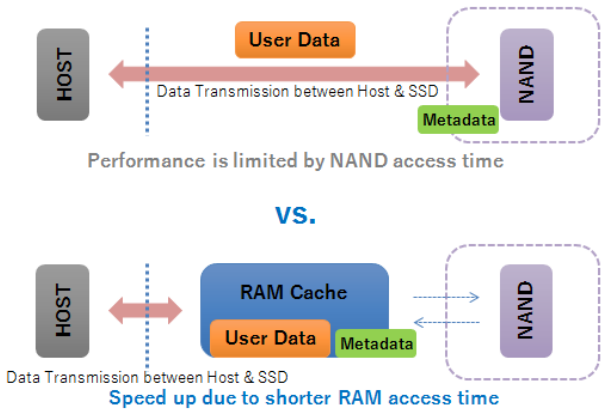

Fig.1 RAM Caching Idea Adopted in SSD

Concerns for Data Loss in RAM Cache

Although RAM can be used as cache in SSD design, there is still one thing we need to think about. Except for the attractive data access speed provided by RAM, another physical characteristic of RAM is volatility. RAM is a well-known volatile memory so the data in RAM cannot be kept once power supply is gone.

Nevertheless, the adoption of RAM in SSD design indeed brings significant advantages of user experiences. If we are able to reduce or even eliminate the risk of data loss in RAM due to the volatility characteristic of RAM, it will definitely strengthen the data integrity capability of SSD end product.

In fact, we can try to deal with this caching data integrity concern through two main ideas. The first idea is putting recoverable data into RAM cache to achieve SSD performance burst purpose. The purpose of DRAM adoption in SSD design is caching some "Hot" data in order to reduce access time. Therefore, we should consider caching some hot data which can be recovered with the original source data stored in flash. In this way, even the caching data is corrupted because of sudden power loss, SSD controller is still able to rebuild the data according to the existing data in flash. Fig. 2 illustrates this idea briefly.

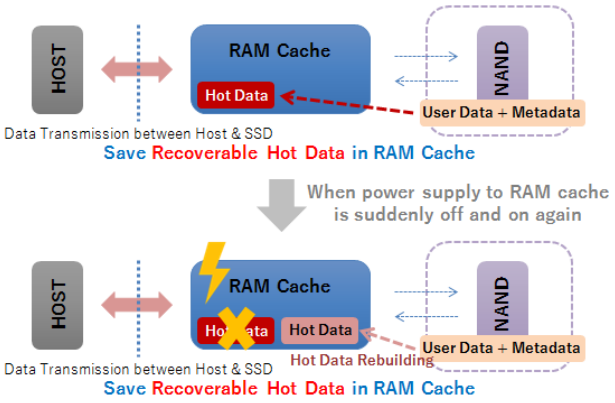

Fig.2 Put "Recoverable" Hot Data in RAM Cache

Unfortunately, there is another scenario which cannot be covered by this idea. If the data programmed by users has just come in from host side to RAM cache of SSD, this kind of data can be regarded as "Unrecoverable" data. Once power loss happens during the period of user data caching, this data is gone anyway. To reduce the risk of caching data loss further, EDM will apply a proprietary technology called SmartFlush to SSD product line.

EDM Technology: SmartFlush

EDM proprietary technology introduced as SmartFlush is based on one fundamental principal: ***Flush caching data from RAM to flash at appropriate timing***.

This strategy is really straightforward. We try to move caching data from RAM (volatile) to flash (non-volatile) prior to unexpected power off happens. Fig.3 illustrates the main idea of SmartFlush mechanism in brief.

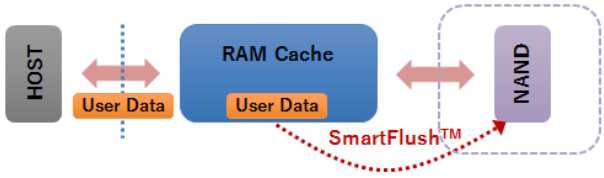

Fig.3 Main Idea of SmartFlush

The critical part of this mechanism is about the timing to flush data from RAM to flash. If the condition of launching data flushing is easy to trigger, SSD possibly keeps staying in busy state even host stops sending commands to device. As a result, SSD is not allowed to enter sleep mode so that the power consumption might not be saved as expectation. Besides, triggering data flushing so frequent is also likely to affect the fluency of SSD operation. On the contrary, if this flushing mechanism is rare to be launched, the possibility of caching data loss is going to be higher of course. Apparently, the balance between the frequency of launching data flushing from RAM to flash and the fluency of SSD operation is the key.

Two suggested timings to launch caching data flushing are listed below.

- ***Timing 1:***
  Data size in RAM cache is larger than a page in flash.

This is a very basic suggestion obviously. Since flash types commonly used in SSD market now are page-based programmable, it makes sense that do not flush caching data to flash until the size of caching data is more than the capacity of one single page in flash.

- ***Timing 2:***
  Host stops sending commands to device.

The main spirit of SmartFlush is preventing caching data loss. From the viewpoint of users, this operation should be performed in background since it is not asked by users. In other words, this prevention mechanism should not cause any significant effect on general user experiences. Consequently, while host stops sending requests to device, it can be regarded as an appropriate timing. Of course, as mentioned earlier, the frequency of launching data flushing should be considered with the overall specification of SSD as well. The real implementation of this technology will vary with the hardware design and firmware design in SSD products, not limited to the contents introduced above.

Summary

In practical cases, SmartFlush is implemented with more complex methodologies and those methodologies could be different between different EDM flash-based products by considering different user scenarios. This article simply introduces the fundamental concepts of SmartFlush instead of going through all the details. No matter how the implementation methodologies vary, the benefit brought by SmartFlush mechanism is the same: With this smart feature offered by EDM, users' confidence in data integrity can grow to the next level.