Envoy Data Memory

SSD Controller Technical Brief:

Power-Loss Data Protection with Flushing Mechanism

Envoy Data Memory • 1310 W. Boxwood Ave Gilbert, AZ 85233 • email: sales@envoydatamemory.com • 1-800-368-6971

Introduction

A volatile memory means that without the power, all data in the memory will vanish. DRAM is a volatile memory with fast access time and frequently used as temporary cache or buffer between host controller and backend storage. In today's SSD designs, the utilization of the DRAM cache (either internal or external) has become a common practice to boost the overall SSD performance, especially on small file transfers. Besides the performance benefits, cache can also improve the endurance of the SSDs by consolidating multiple small transfers before pushing into NAND flash, reducing the amount of block erase during the process. This is the reason why Envoy Data Memory (EDM) controller reference design requires an external cache to be connected. However, there are two sides in every door. The major drawback of cache memory -- it requires power to maintain the stored information. This raises a common concern for the SSD adopters: What happen when SSD loses its power? In this brief we will introduce ***SMART Flush Cache Technology*** (***SFCT***) used by the EDM controller to protect the data while suffering power loss issue.

What is Flush Cache?

During a proper system shutdown, the operation system would typically issue a STANDBY IMMEDIATE command to signal the SSD to prepare for system shutdown. This will provide enough time for SSD controller to flush the data in the cache to NAND Flash, where the data will be preserved after power loss. However, during an unsafe power shutdown, the in-flight user data and system metadata in cache may not have chance to transfer to the NAND flash. Thus, it is critical for SSD firmware to implement intelligent protection scheme to preserve data integrity in an event of unexpected power loss. One of the methods to protect our data is by flush cache -- copy every data in the volatile write cache to the non-volatile media in a very short time.

How does the Controller Employ the Flush Cache?

The EDM SSD controller reference design requires an external cache to be connected. The cache size is typically 256 or 512MB. Within the cache buffer, the use of cache dedicated for user data is kept at maximum 8MB in order to reduce the amount of user data sits in the cache.

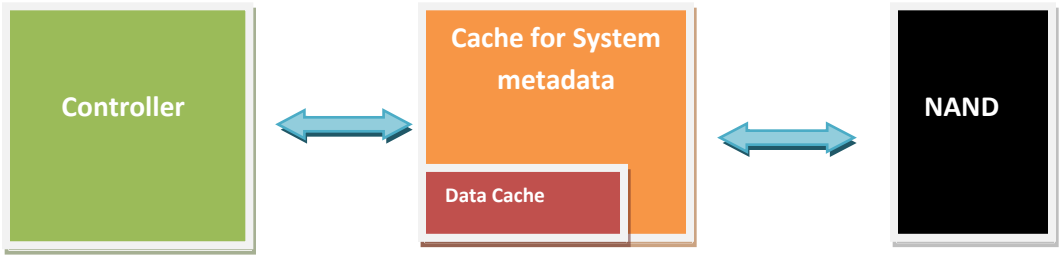

The SFCT allows incoming host data to make a "pit stop" in cache and then pushed into the NAND media immediately, as long as there is no bottleneck on the flash interface. When the flash interface becomes jammed due to particular file sizes (random 4K), cache are being treated as "organizers". The incoming write data will be consolidated into groups before written into flash to improve the write amplification and acknowledge host only when the data is fully committed to the NAND media.

EDM SMART Flush Cache Technology Highlight

- **S**mart Trigger – Smart trigger includes device sleep (DEVSLP) trigger and SATA link loss (SLL) trigger. For DEVSLP trigger, the SSD flushes cache before the device goes to sleep mode. For SLL trigger, once the SSD detects the loss of SATA link, controller will flush cache to NAND to preserve data integrity.

- **M**anagement Pair-Page Data (GuaranteedFlush)– During flush cache, the SSD intelligently manages the pair-page of the NAND. Once the data is committed to the NAND media, the following page writes will not impact the previous committed data.

- **A**rrangement of **R**andom Data – The SSD arranges and combines small random data into a larger packet before flushing to the NAND. This will reduce write amplification factor (WAF) and enhance performance.

- **T**ime-base Idle Trigger – The SSD will flush cache if the host idles for 200ms to guard against contingencies.

Conclusion

This whitepaper delivers fundamental understanding of flush cache. Although there are several host-initiated power-fail protection mechanisms to protect data, SMART Flush Cache Technology (SFCT) not just passively waits SATA-link-loss happen to trigger flush cache but also actively flushes cache when the host idles for a short time to guarantee data integrity. With in-house validation and on-field deployment by our customers, the SSD proves to provide the reliability required by consumer, industrial, and enterprise-level application.

Appendix

FLUSH CACHE CMD is a host issued command defined in ATA specification to guarantee data integrity. Below is an example of how FLUSH CACHE CMD is used by Windows OS to help it recover from power-loss. During a file copy task, the operation system may break down the task into sub-tasks. The following picture describes a possible process flow.

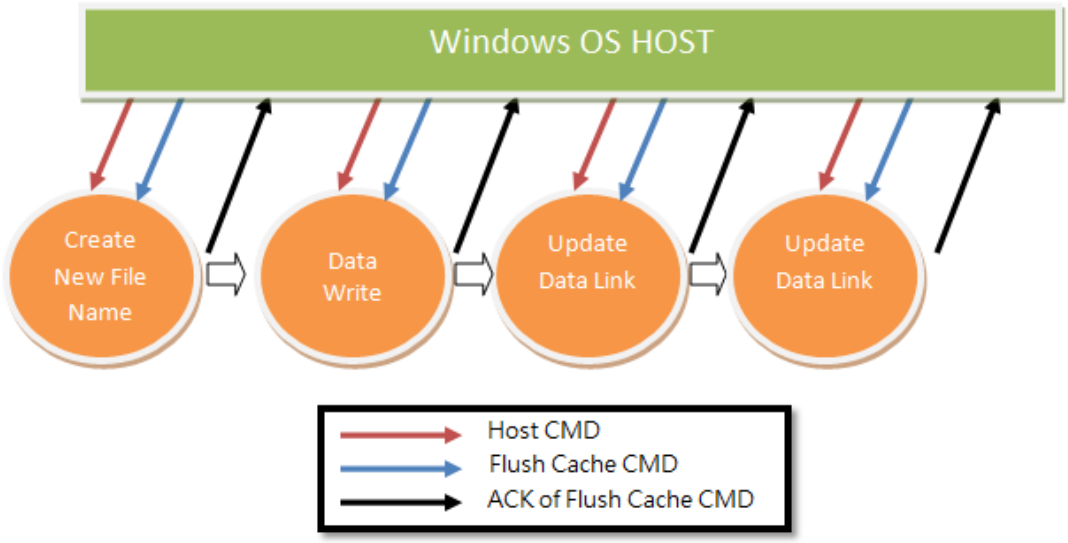

The FLUSH CACHE CMD following each important stages of a given task allows the OS to re-built system table / file system table (For Windows, CHKDSK, for Linux, FSCK) upon return from power fail.