White Paper:

POWER LOSS PROTECTION IN SSDs

How SSDs are Protecting Data Integrity

SAMSUNG

INTRODUCTION: POWER LOSS PROTECTION IN SSDs

SSDs are usually categorized as client and enterprise SSDs. Client SSDs are designed for PC applications, which includes laptops, desktops, tablets and other devices purposed for personal computing, whereas enterprise SSDs have been built for more strenuous server and data center usage. It's generally well understood that enterprise SSDs offer higher endurance and better consistent performance, which are often required in enterprise environment. The lower cost per gigabyte and high performance makes client SSDs tempting for enterprise users as well, but it's not always understood that client SSDs don't provide the same guarantee of data integrity.

This white paper explains the differences in power loss protection techniques used in Samsung's enterprise and client SSDs. Power loss protection is a critical element in data integrity because nearly all modern SSDs contain a volatile DRAM cache that is vulnerable to unexpected power losses. Potential data loss or corruption can be a serious risk for an enterprise handling important customer data, hence understanding the significance of power loss protection is important when evaluating drives for enterprise usage.

WHY SSDs NEED POWER LOSS PROTECTION

When the host operating system (OS) sends an IO write command to the SSD, the data is first received and processed by the SSD controller. The data is then moved to the volatile DRAM cache to improve performance and perform write optimization, after which data is eventually written to non-volatile NAND. While NAND is non-volatile and thus data in it is protected against power losses, the data that has yet to reach NAND is vulnerable and subject to corruption in case of an unexpected power loss.
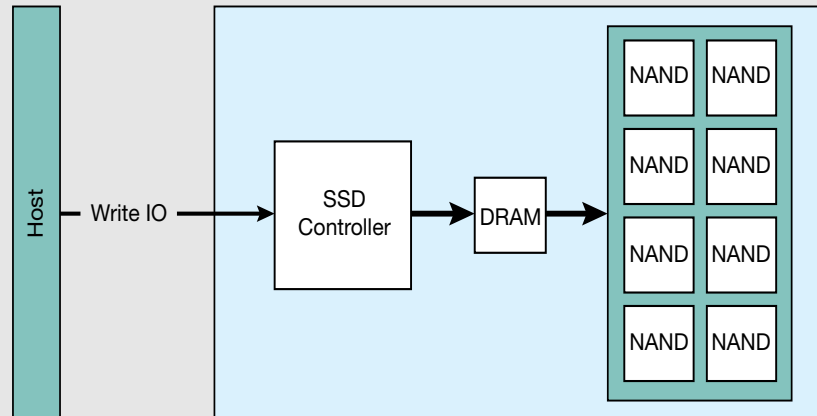
The Journey of Data Inside an SSD

Usually SSDs send an acknowledgement to the OS that the data has been written once it has been committed to DRAM. In other words, the OS considers that the data is now safe, even though it hasn't been written to NAND yet. The reason why SSDs use DRAM for write IO caching is because NAND can only be programmed at the page-level (typically 8KB or 16KB in modern NAND), so write IOs need to be combined for optimal performance and endurance. Otherwise, for example a common 4KB write, would have to be written to NAND as 8KB or 16KB depending on the page size, which would seriously reduce endurance and performance.

FLASH TRANSLATION LAYER: KEEPING DATA ORGANIZED

The other purpose of the DRAM is to cache the Flash Translation Layer (FTL), which is used to map Logical Block Addresses (LBA) used by the operating system to physical NAND addresses. The FTL is constantly changing as new write IOs from the host are written to the drive and the drive performs background tasks such as garbage collection and wear leveling, thus using DRAM for FTL caching is necessary because DRAM has substantially lower latency than NAND, its endurance is practically unlimited and it's bit-addressible. Without the DRAM cache, SSDs would not

perform as well as they do, and there would also be endurance concerns as FTL changes consist of very small writes and NAND can only be programmed at the page-level (typically 8KB or 16KB in modern NAND).

Because of the volatile nature of DRAM (i.e. all data is lost in an event of power loss), SSDs need to incorporate power loss protection (PLP) techniques to prevent against data loss and corruption. During a normal shut down or sleep, an SSD is given enough time by the operating system to write all contents of DRAM to NAND, but in case of an unexpected power loss, the

data in the DRAM may be lost. Since DRAM hosts the most current version of the FTL, an unexpected power loss may corrupt all data in the SSD if proper power loss protections are not in place.

Unexpected power losses may be due to power outages, but also due to power supply glitches that don't impact the rest of the system. Hence an uninterruptible power supply (UPS) or any other form of backup power is not enough to prevent against unexpected power losses, and SSDs need internal mechanisms to guarantee data integrity.

Overview of SSD Software Architecture
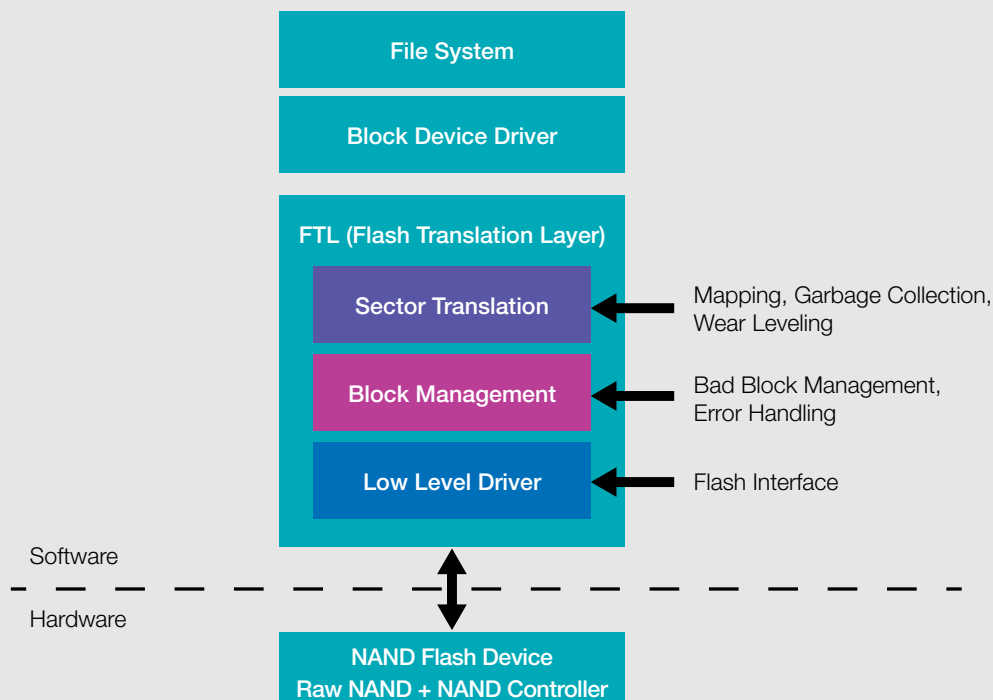
POWER LOSS PROTECTION IN SAMSUNG ENTERPRISE SSDs

Because data centers are used to store customers' data, data integrity and availability are a very high priority. Any data loss or delay in availability will immediately degrade the customer's experience and it's likely that the customer will not continue to use the service due to lack of trust. Imagine that a customer is using an image hosting service to store photos of his or her children for example, and one day the customer finds out that some of the most precious memories have gone missing because of a poorly designed storage architecture. Not only is the customer furious, but the customer will also make sure that others hear his or her experience through social media and word of mouth, which would damage the reputation of the service.

To guarantee data integrity, Samsung's enterprise SSDs feature full power loss protection with backup power circuitry in
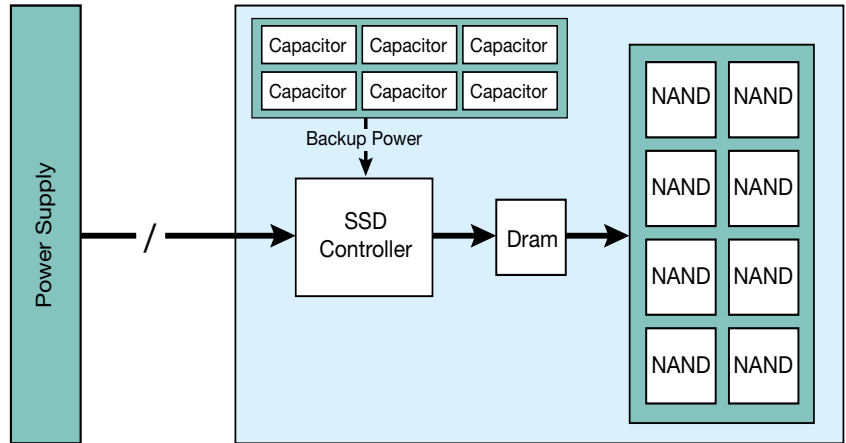
Power Loss Event in a Samsung Enterprise SSD

the form of tantalum capacitors. When the SSD is powered on, the tantalum capacitors are charged and if external power is suddenly lost for any reason, the charged current in the capacitors is discharged to provide additional power to the SSD. This additional power is used to provide power to the SSD controller, DRAM and NAND

to flush all the data in the DRAM to the NAND in order to avoid loss or corruption of data. With the help of the capacitors, Samsung enterprise SSDs are guaranteed to provide full power loss protection for all data that has been acknowledged to the host as written.

$1.7 trillion
annual cost of data loss and resulting downtime

29%
of data loss is due to corruption

53%
of enterprise storage disruptions are caused by hardware failures

# POWER LOSS PROTECTION IN SAMSUNG CLIENT SSDs

Due to the high cost of capacitors and the price sensitivity of the client SSD market, client SSDs do not normally include full power loss protection in the form of backup power circuitry. While the FTL is frequently flushed from the DRAM to the NAND, there is still a window of opportunity where the DRAM contains a newer version of the FTL in the event of an unexpected power loss. Fortunately, there are alternative and more cost efficient mechanisms to prevent the FTL from corrupting.

The most important mechanism is journaling, which is used at both the FTL and file system level. A journal is a separate log that keeps track of upcoming changes that have yet to be committed to main file system or FTL. In case of an SSD, the journal basically keeps track of all data that has yet to be written to NAND in full in order to avoid the FTL from corruption. If the FTL in the DRAM is newer than the version in NAND in an event of an unexpected power loss, the controller only has to check the data integrity of the journal rather than the whole FTL, which typically only takes a few seconds compared to at least several minutes for a full FTL checkup. In a client usage, a rare few second delay is a non-issue, but in an enterprise deployment with thousands of drives such delays would add up and potentially cause serious problems.

Furthermore, the small amount of pending write data that is stored in DRAM is lost during an unexpected power loss. In PC environment, such data loss isn't critical because modern file systems use journaling in order to avoid corruption, and it's important to acknowledge that hard disk drives (HDDs) also contain a DRAM cache, thus the power loss phenomenon is nothing new in the PC space. However, in the enterprise even small data losses can be critical because they may result in corruption of customer's data, which would degrade the reputation of the service.

Journaling Process Inside an SSD
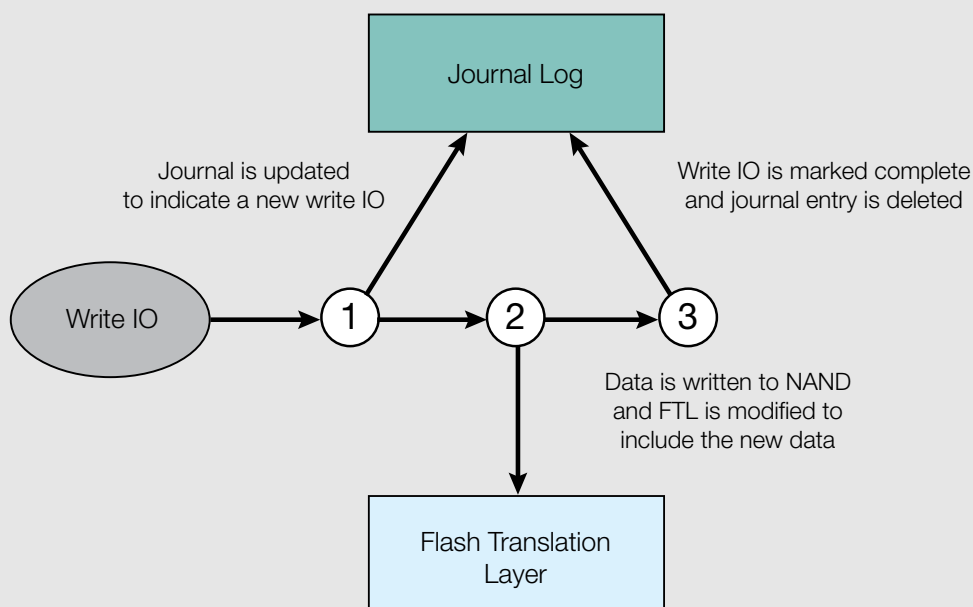
CONCLUSION

As we live in a data driven world, there is nothing more important for an enterprise than to protect its own or its customers' data. Using client SSDs in enterprise may seem like a good investment at first due to the lower cost per gigabyte, but the cost of data losses overtime may considerably outweigh the initial savings in acquisition cost. Thus enterprise SSDs are a better, risk-free investment for customers looking to deploy SSDs in an enterprise environment.

SAMSUNG ENTERPRISE SSD PORTFOLIO

With Built-in Power Loss Protection

PM863 Series Data Center SSDs
- 3 bit MLC NAND
- Designed for read-intensive applications
- SATA 6Gb/s Interface
- Form-factors: 2.5"

SM863 Series Data Center SSDs
- 2 bit MLC NAND
- Designed for write-intensive applications
- SATA 6 Gb/s Interface
- Form-factors: 2.5"

SAMSUNG WORKSTATION SSD PORTFOLIO

Without Built-in Power Loss Protection

950 Pro Series Client PC SSDs
- 2-bit MLC V-NAND
- Designed for high-end PCs
- PCIe Interface
- NVMe protocol
- Form-factors: M.2

850 Pro Series Client PC SSDs
- 2-bit MLC V-NAND
- SATA 6Gb/s Interface
- Form-factors: 2.5"

About the Author

Kristian Vättö is a technical marketing specialist and started his career as a news editor at AnandTech.com in 2011. He later became the site's SSD editor and was responsible for producing highly-detailed and professional SSD reviews. In addition to his work with Samsung, Kristian is currently studying economics at the University of Tampere in Finland.

Learn more: samsung.com/enterprisessd | insights.samsung.com | 1-866-SAM4BIZ

Follow us: ▶ youtube.com/samsungbizusa | 🐦 @SamsungBizUSA

SAMSUNG

WHP-SSD-POWERLOSSPROTECTION-FEB16J