WP25-1220-01

vtGuard Technology

Introduction

Employing NAND flash-based storage has become a mainstay in most of today's industrial embedded applications. Solid state drives (SSDs) satisfy the high performance, low power, and strict reliability requirements of a growing list of embedded systems. For these reasons, suppliers have been able to reduce the cost per gigabyte to make these solutions increasingly more viable. This advanced technology, however, is not without its drawbacks.

One key design challenge SSD vendors face is the susceptibility of an SSD to data corruption due to an unexpected power failure. This is a critical concern for designers of most embedded systems that require "five nines" uptime, or 99.999% of a computer's time available to do work, particularly for systems operating in harsh environments. This issue is further compounded by the difficulty in maintaining a steady and uninterrupted power source - especially in light of the proliferation of battery-operated and portable devices in the Internet of Things (IoT), which are particularly sensitive to power loss.

Power failure threats for embedded systems can range from supply power spikes to brown-outs. These conditions leave the system susceptible to data corruption and may cause failures in the field, resulting in the potential loss of revenue from equipment returns.

Because this is a key issue, it is a critical to understand how a given storage solution will operate in a specific application. This white paper will discuss the effects of data corruption on an SSD due to power loss during a write operation. It will also examine the advantages of Virtium's vtGuard and vtGuard+ technology as protection against data loss in embedded and industrial applications.

SSD Structure

To understand the challenges associated with unexpected power loss to the drive, it is important to first understand the basic structures of a solid-state storage device and NAND flash. As seen in Figure 1, a solid-state storage device, commonly referred to as a Solid-State Drive (SSD), consists of a controller, NAND flash, sometimes DRAM, and other electrical and timing components.

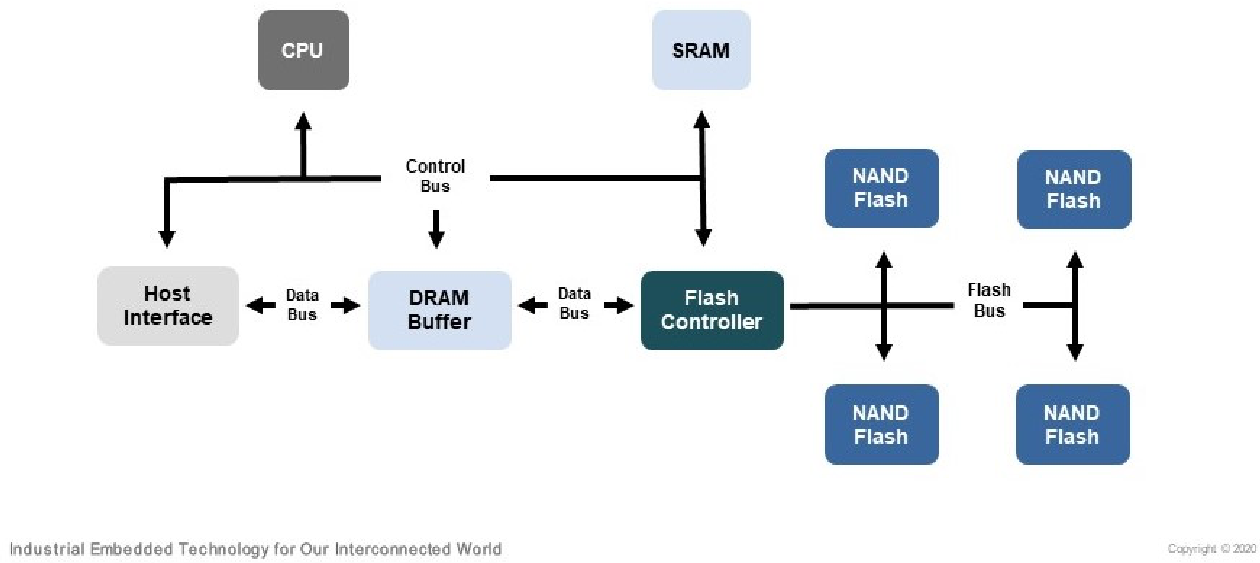

Figure 1: Simple SSD Block Diagram. The diagram above shows the main components of an SSD, which consists of the host interface, the controller, a data buffer, optional SRAM/DRAM, and the flash array.

NAND Flash Technology

Within these devices, NAND serves as the primary storage component of the drive. Figure 2 shows a block of NAND flash. NAND is characterized by page size – the minimum write or program unit – and block size – the minimum erase unit.

	Page size	Spare area	
Block	16,384	1,280	Page 0
	16,384	1,280	Page 1
	⋮	⋮	
	16,384	1,280	Page 255

Figure 2: Block and Page Structure of a 16K Page NAND Flash Drivemponents of an SSD.

For example, a 16K page NAND has (32) 512-byte sectors per page and 256 pages per block. Each page has an associated spare area where metadata, such as Error Correction Code (ECC), is stored. SLC NAND typically have 2KB and 4KB page sizes, while MLC NAND have 8KB and 16KB page sizes and TLC has 16KB page size.

Thus, programming times for SLC are typically shorter than either MLC and TLC. In general, larger NAND page size and longer programming time often indicate a drive's greater susceptibility to power disturbances. In addition, MLC and TLC use dual plane, or paired pages, during write operations, thereby contributing to a longer programming time. For these reasons, SLC SSDs are much more robust than MLC SSDs.

DRAM Technology

Some SSD may also include Dynamic Random-Access Memory (DRAM). A small portion of the DRAM may be used as a data buffer while the remainder would be used as a mapping table. Inclusion of the DRAM would improve the performance of the SSD drive by increasing the amount of mapping data that is immediately available to the SSD controller.

While NAND is the primary nonvolatile memory component in SSD, DRAM is a type of volatile memory that requires power to store data. Thus, data from the DRAM cache would occasionally be pushed to the NAND for physical storage. This provides the most efficient method of data access: when the SSD controller needs to locate data, it checks the DRAM mapping table to find the corresponding address on the NAND.

Writing to an SSD

In most cases, power cycling issues occur while writing to the flash array on the SSD. For this paper, we will focus on the writing of user data.

Writing to the SSD follows this basic sequence:

1. When the host requests a data transfer, the SSD prepares to receive the data.

2. Once the host transmits the data, it is written to the data buffer on the SSD. The size of this buffer can vary by SSD vendor.

3. Following the write to the data buffer, the data is then transferred to the active block in the flash. The active block is a pre-erased block used by the SSD controller to quickly save data.

4. The active block is given a mapping address and added to the map table. This completes the data transfer.

Under normal circumstances, this sequence is properly executed, and each step is implemented in this order. However, a sudden power disruption can occur at any point during this sequence and can disrupt the write to the SSD. Without the proper mitigative measures, disruption may result in data loss. To prevent data loss in the event of a sudden power anomaly, Virtium employs vtGuard technology in all SSD products.

Virtium vtGuard® Technology

Standard vtGuard

Virtium’s StorFly® and TuffDrive® SSDs with integrated vtGuard technology provides voltage detection circuitry that acts as an “early warning” system for potential power anomalies to protect data in the NAND. In addition to the voltage detection circuitry, the vtGuard technology also includes a hold-up capacitor and firmware technology to protect against metadata loss and to ensure the integrity of the SSD once it is powered up again.

It is important to understand the following key items prior to reviewing the individual scenarios in the flowchart.

1. The SSD must have voltage detection circuitry to recognize the impending power loss.

2. The SSD must have a mechanism to hold up the power long enough to "flush" the data that may be in the internal buffer of the SSD or the DRAM before the lower threshold is reached.

3. Once the lower threshold is reached, the SSD controller to assert write protect to the flash to prevent data corruption of existing data.

While some SSD vendors use a capacitive bank, others may use a super-cap or a battery. However, super-caps and batteries often introduce more reliability issues than they solve, especially during system operation over a wide temperature range. Thus, the vtGuard circuit relies solely on the capacitive bank to hold up power. The standard vtGuard technology provides protection for data in the NAND and is suitable for most operations. For additional protection for data in-flight, Virtium also offers products that utilize vtGuard+.

vtGuard+

While most operations are protected under the standard vtGuard technology, vtGuard+ provides extra protection for data in-flight, or data in the process of being transferred from the host to the controller or the DRAM to the NAND. These operations are especially important for critical applications in banking, e-commerce, etc.

If the detection hardware senses a power anomaly, SSDs using vtGuard+ will stop accepting data from the host system. While the interface is halted, the data in-flight is flushed from the buffer to NAND. The additional capacitors used in vtGuard+ discharge more power than the standard technology, which extends the time to write the data in-flight to the NAND. In doing so, the vtGuard+ technology is able to provide extra protection against corruption of the drive and of the data being transferred.

Power Loss

Graceful or Normal Power Shutdown

In an ideal scenario, the host system would always approach a power-down event with a systematic power-down sequence. Sending the SSD a shut-down notification is the safest method of removing power to ensure data integrity. Both SATA and NVMe interfaces support shut-down notifications that are used by all operating systems (OS) to prepare for power removal or system power down.

When a formal shut-down command is issued to the OS, the system discontinues communication to the SSD and then issues the shut-down notification to the SSD. Once the SSD receives this notification, it knows not to expect any more incoming commands. The system then begins to clean up and discontinue any background processes before going into an idle state and powering down. This guarantees that the SSD has no data movement in the Flash when power is removed and no opportunity for data loss. However, for many industrial embedded systems, which often face unexpected power loss, notifying the SSD to shut down may be difficult, if not impossible or impractical.

Ungraceful or Unexpected Power Loss

Unexpected power disruptions may happen at any time and are especially concerning when it occurs during a write operation. As seen in Figure 3 below, these disruptions may occur by one of the three following scenarios:

1. Power loss during a write operation before the SSD has acknowledged receipt of data.

2. Power loss after the SSD has acknowledged receipt of data but before it is written to the active block.

3. Power loss after the data has been written to the active block but before the mapping table gets updated.

Data Loss Mitigation

Figure 3 below shows a block diagram of the vtGuard technology and its actions in mitigating data loss in each of the three scenarios above. As seen in this figure, vtGuard and vtGuard+ can provide protection against data loss during an unexpected power loss during the following three conditions.

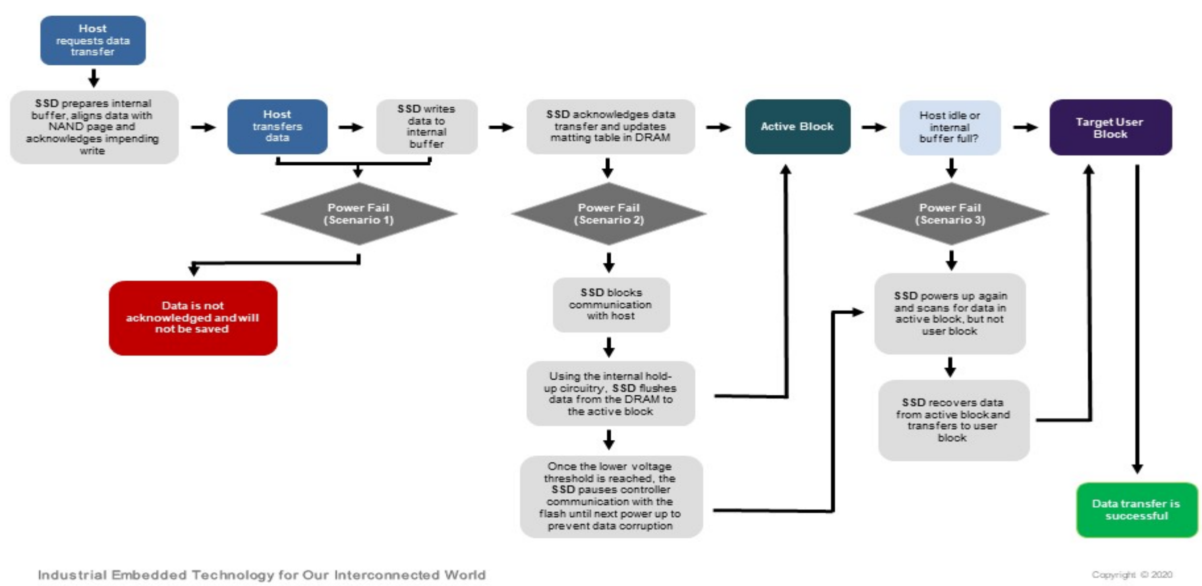

Figure 3: Power Failure Scenarios and vtGuard Protection Mechanisms

Power Fail Scenario 1

In this scenario, a power loss occurs during a write operation before the SSD has acknowledged receipt of the data. Since no user data was acknowledged, no data will be saved. With vtGuard technology, the mapping table is saved, but it may need to be rebuilt on the next power cycle. With vtGuard+, the mapping table will not need to be rebuilt, because of the higher reserve power.

Power Fail Scenario 2

In this scenario, the power loss occurs after the SSD has received and acknowledged the data but before it can write it to the active block.

With vtGuard, once a power failure occurs, the SSD controller and firmware immediately cut off communication with the host and the hold-up circuit is activated. This will allow the SSD to transfer the write cache (metadata, mapping tables, etc.) contents from the internal buffer or DRAM to the active block. This initial behavior confirms that the SSD will be recoverable after sudden power failure events with minimal delay as the previous mapping table is restored from the NAND flash. User data which has been acknowledged to the host may be lost.

As described above, vtGuard+ provides additional protection for data in-flight, or data which has been acknowledged but not yet saved to the flash memory. With vtGuard+, the drive sends a command to the host system to stop sending data while the data in-flight is flushed from the buffer to NAND. Since the SSD has higher on-board capacitance than the standard vtGuard technology, all the data in-flight will also be saved. In this case, the hold-up circuit provides enough power to save both the metadata and the mapping tables as well as the data in-flight.

All of this data will be flushed from the internal RAM buffer to the active block. On the next power-up cycle, as with vtGuard, all the data stored in the active block is remapped to the user space prior to the host's recognition of the drive.

Power Fail Scenario 3

In this scenario, the power loss occurs after the data-in flight has been written to the active block but before it has been written to the correct LBA.

In this state, the data is already saved in the active block, so no acknowledged data is lost. On the next power-up cycle, all the data stored in the active block is remapped to the user space prior to the host's recognition of the drive.

Under each power loss scenario, with vtGuard and vtGuard+, the SSD is able to maintain integrity of the data buffer and the internal mapping table to prevent data corruption and validate usability of SSD when the computer regains power. However, with vtGuard+, the additional detection hardware and extra capacitance that protect against loss of data in-flight would result in a larger and more expensive drive. Therefore, vtGuard+ is best suited for critical applications in which data in-flight preservation is imperative.

Power-Cycle Testing and Validation

For each vtGuard product, Virtium examines the effects of sudden power-down on the drive. Virtium uses proprietary process to validate the vtGuard technology to 5,000 power-down cycles during initial product development. This validation process examines the controller and firmware platform that uses SLC, MLC, and TLC NAND flash. Virtium also performs an additional 1,500-cycle report for each form factor and capacity point at initial release.

Power cycle test outputs are available under NDA. Please contact your Virtium representative for more information.

Conclusion

In the event of a power anomaly, a system is left susceptible to data loss or corruption that may result in system failure in the field. Virtium's vtGuard and vtGuard+ technology protects against data loss and corruption of the drive as a result of this data loss. Each vtGuard product utilizes voltage detection circuitry, hold-up capacitors, and proprietary firmware technology to serve as an "early warning" system and to protect against data loss and ensure the integrity of the SSD once it is powered up again.

vtGuard primarily protects against loss of data in the NAND and is suitable for most operations, whereas vtGuard+ provides additional protection for more critical applications that require preservation of data in-flight. For each vtGuard product, Virtium validates the drive based on several factors that affect data and drive integrity after a power down event.

With a wealth of experience, Virtium's sole mission is to supply its partners and customers with the most reliable and capable products. Visit www.virtium.com or contact your local Virtium sales, manufacturer's representative or distributor for more information, including how to request samples and pricing.

30052 Tomas | Rancho Santa Margarita, CA 92688
Phone: 949-888-2444 | Fax: 949-888-2445
www.virtium.com

WP25-1220-01