

Power Loss Data Protection in Enterprise Solid State Drives

Technology Brief: OCZ DataWrite Assurance

Introduction

Though datacenter power management systems feature the highest reliability and redundancy there is, it is critical that system architects take further precaution to preserve the integrity of the data in case of sudden power loss. Because enterprise-class SSDs typically contain mission-critical data, it can be detrimental to a datacenter if in-flight data is lost or corrupted.

To address this challenge OCZ has designed its line of enterprise-class PCI Express SSDs (Z-Drive R Series) with the risk of power interruption in mind, and offers innovative solutions that enable data to survive power drops, outages, or inadvertent disengagement of a live drive. However, it is important to keep in mind the Z-Drive C Series, though without power loss data protection, implements a number of intelligent reliability methods to preserve data integrity.

What Happens During an Unexpected Power Loss?

A NAND Flash-based storage solution must fully complete an in-progress program (write) operation to ensure the data is stored properly within the NAND Flash cell; however, an unexpected loss of power can stop in-flight data in its tracks resulting in malfunction of the datacenter or loss of mission-critical information. This is especially essential for multi-level cell (MLC) NAND, including enterprise-grade MLC (eMLC), which stores two logical pages of data on the same physical page. In this case, a sudden power outage during a program operation may result in loss of the entire page making the power failure even more devastating. Although these bits of lost data may be manageable in a client environment, enterprise applications demand greater reliability.

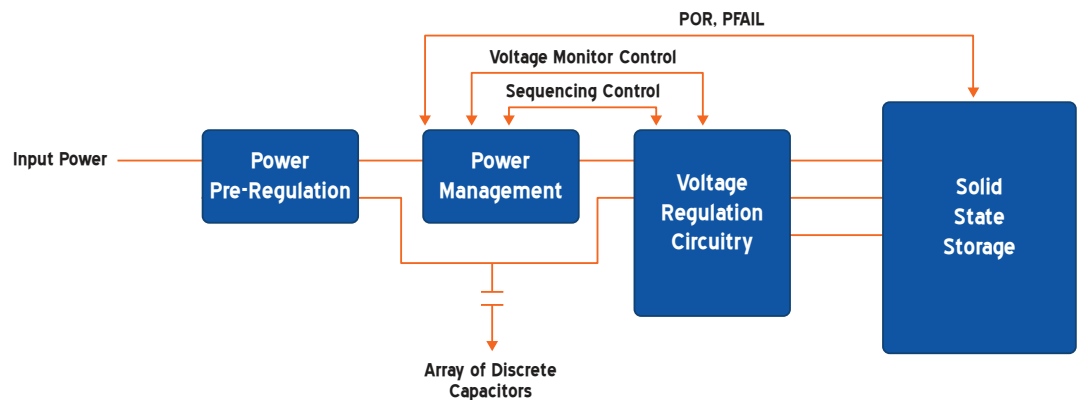
For this reason, power loss data protection is necessary for datacenters and mission-critical IT infrastructures, especially ones that rely on volatile memory/storage sources that manage logical block address (LBA) mapping information, the drive's sole source of locating the physical block addresses. Many purported "enterprise" PCIe-based SSDs store this vital mapping information (metadata) in the host DRAM and an unexpected power failure can lead to a loss of the system's metadata causing severe downtime or loss of access to indispensable data.

How Does an OCZ SSD Handle a Power Failure?

OCZ enterprise-class storage solutions implement both sophisticated firmware and hardware-based power failure protection. In the case an OCZ SSD's primary power source drops below a predefined threshold, the SSD will automatically not accept any new commands from the operating system. The power loss backup circuitry, a self-contained secondary power source, is then activated ensuring that any in-flight data is safely transferred and stored in the NAND flash.

Figure 1 below shows a simplified block diagram of the OCZ DataWrite Assurance design which utilizes discrete capacitors as a power source.

Figure1 - Power Fail Data Protection Block Diagram



Conclusion

Power loss data protection is a key component to ensuring superior data integrity and implementing highly reliable power fail protection at the storage level is second to none in the prevention of data loss. Simply inserting backup power mechanisms into the system architecture provides no instrumental advantage if the vital mapping data is stored in volatile DRAM, a common practice of PCIe-based SSDs. OCZ's superior combination of in-flight data protection as well as LBA mapping (metadata) protection truly differentiates OCZ SSDs from other drives in the market and makes OCZ the only trusted PCIe-based solution for enterprise datacenters.