



White Paper

Eliminating Drive Corruption from Power Disturbances

SiliconDrive PowerArmor Technology



Executive Summary

Many embedded OEMs face operational failures because of less-than-optimal power conditions. Approximately 75% of all field failures are due to power disruptions. Spikes, brownouts, surges, blackouts, and other power problems can cause data loss and/or corruption in applications.

WD has integrated PowerArmor technology into all of its SiliconDrive technology to virtually eliminate the damage power interruptions can create. PowerArmor's voltage detection circuitry alerts the host system of any power fluctuations and prevents the transmission of commands until power levels are normalized. Address lines are latched to prevent data from being written to the wrong sector.

When power goes out, drives can be corrupted and data ruined, resulting in downtime as drives get reformatted, operating systems reinstalled, or products returned. This directly impacts overall reliability, dependability, customer goodwill, and cost of ownership. PowerArmor now offers designers a way to virtually eliminate these costly problems.

Table of Contents

Introduction	2
Background	3
PowerArmor Technology	4
Storage Basics	6
Powerdown Tester Need	9
Conclusion	12

Introduction

Equipment manufactured by embedded OEMs often operates in less-than-ideal power conditions. Power disturbances ranging from spikes to brown-outs can cause a significant amount of data and drive corruption, causing field failures and potential loss of revenue from equipment returns.

WD's SiliconDrive solid state storage technology is specifically designed to meet the high performance, high reliability, and multi-year lifecycle requirements of embedded OEMs in the netcom, military, industrial, interactive kiosk, and medical markets. WD has integrated PowerArmor technology into all of its SiliconDrives to virtually eliminate drive corruption in the event of unexpected power disturbances.

Background

Because SiliconDrive has no moving parts, it is intuitive that it performs much better in environmentally rugged applications than rotating hard disk drives, whose precision mechanics cannot take the rigors and duty cycles of most embedded systems. In most cases, SiliconDrive also offers significant power savings that translates into longer battery life for mobile applications. The following table contrasts these parameters for SiliconDrive and hard drives.

Table 1: SiliconDrive Parameters

Parameter	SiliconDrive	2.5" Rotating Drive
Operating		
Shock	>1000G per MIL-810F	200G
Vibration	16.3G rms per MIL-810F	1G
Temperature	-40° to 85°C	5° to 55°C
Altitude	80,000 feet	50,000 feet
Power Consumption	0.2W	2.5W
Requalification Cycle	3 to 5 years	1 year
R/W Speed (MBps)		
Small File Random	2.5/1.5	0.03/0.03
Large File Sequential	8.0/6.0	31/31
MTBF	>4,000,000 hours	300,000 hours
Duty Cycle	100%	20%–40%

While solid state drives are generally more robust than hard drives, not all solid state storage products are created equal. Even though they physically look the same, SiliconDrives in the CF form factor differ greatly from CompactFlash cards designed for the consumer electronic market. The inherent technology differences are outlined in the following table.

Table 2: SiliconDrive Inherent Technology Differences

Parameter	SiliconDrive	Compact Flash Card
Write/Erase Endurance	>2M cycle per block	<100K cycles per card
Error Correction (ECC)	6 bits	1 bit–2 bits
Wear-leveling Algorithm	Over the entire SiliconDrive	Over free space only
Powerdown Protection	Yes	No
R/W Speed (MBps)		
Single Sector	1 MBps–2 MBps	30 KBps–40 KBps
Large File	6 MBps–8 MBps	3 MBps–5 MBps
Requalification Cycle	3–5 years	1 year
Maximum Capacity	8 GB	4 GB

The term *standard CompactFlash card* can be a bit misleading. Flash cards must pass the test suite defined by the CompactFlash Association (CFA). These tests allow for a relatively wide range in some timing parameters. These parameters can vary with both hardware and firmware changes. Most consumer-oriented products have enough other system overhead to render these changes insignificant. On the other hand, embedded systems usually have strict timing requirements and even the slightest changes to the host interface can cause significant issues.

The most industry-recognized differentiator between SiliconDrive technology and commercial flash cards is in write/erase cycle endurance. SiliconDrive products have superior error correction capabilities and wear-leveling algorithms that can significantly extend the life of standard storage components. The challenges of solid state storage management and the methodology for determining the operating life of solid state storage are discussed in the *NAND Evolution and its Effects on Solid State Drive (SSD) Useable Life White Paper* located at http://www.wdc.com/WDProducts/SSD/whitepapers/en/NAND_Evolution_0812.pdf.

PowerArmor Technology

The power issues described in this section are significantly less well-known, but are much more prevalent in embedded applications. Approximately 75% of field failures are the results of power-related corruption. SiliconDrive is specifically designed for the embedded OEM market and integrates various techniques for mitigating power-related issues. These techniques are not implemented in most flash cards because they are not economically viable for consumer-based applications, but they are essential to eliminate unscheduled downtime in embedded systems.

WD patented its PowerArmor technology to eliminate drive corruption due to power disturbances. **Figure 1** shows how PowerArmor integrates voltage detection circuitry to provide an “early warning” of a possible power anomaly. When a voltage threshold has been reached, the SiliconDrive sends a busy signal to the host so that no more commands are received until the power level stabilizes.

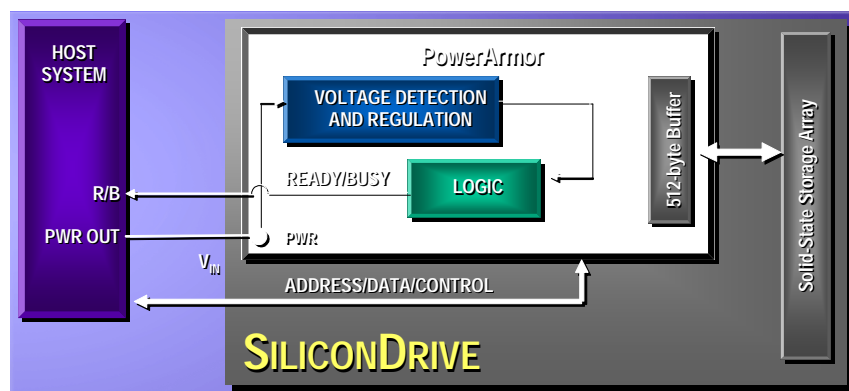


Figure 1: PowerArmor Voltage Detection Circuitry Integration

Address lines are then latched, as shown in [Figure 2](#), to ensure that data is written to the proper location. In contrast, most microcontroller-based flash cards' address lines can float to undetermined states if input power drops below the specified minimum operating level.

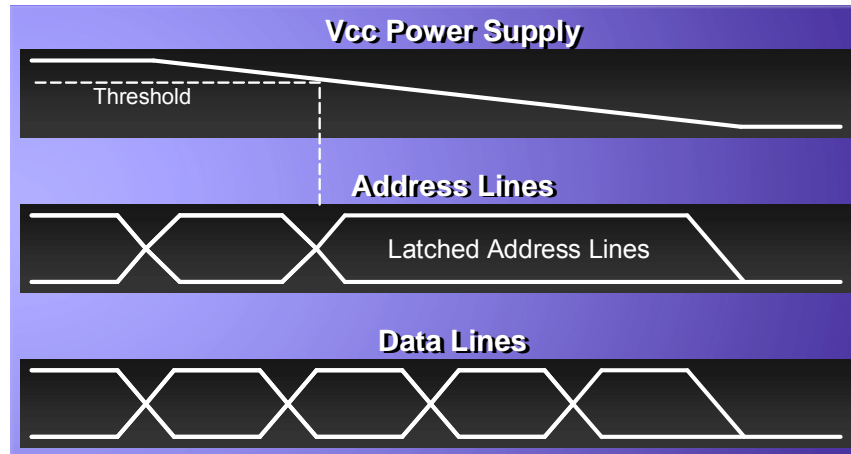


Figure 2: Latched Address Lines

SiliconDrive technology integrates a RISC-based DSP that allows for a 512-byte buffer size. This is 1/4 to 1/8 the size of a microcontroller-based flash card. The smaller buffer size limits the time that the data is in volatile memory, which reduces the power and time required to clear the buffer, updates the control bytes, and completes the data transaction.

Larger buffer sizes require additional capacitance to hold up the power longer. This extra capacitance is physically large and generally does not scale to form factors smaller than 2.5" drives or 3.5" drives. Because SiliconDrive technology does not require this added capacitance, it can scale to virtually any industry-standard form factor.

Storage Basics

To fully appreciate the benefits of PowerArmor technology, it is important to understand the basics of storage. This includes how the storage address space is segmented, as shown in [Figure 3](#), and the possible ways to corrupt data.

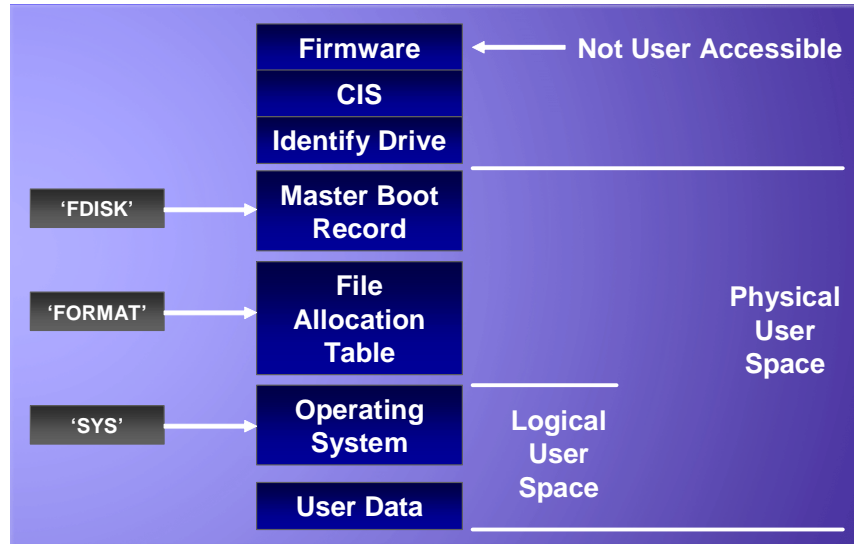


Figure 3: Storage Address Space Segmentation

Table 3: Storage Address Space Segmentation

Item	Description
Firmware	Storage device internal run-time firmware.
CIS	Identification file for removable devices such as CF and PC Card form factor products.
Identify Drive	Identification file for all drives identified as fixed IDE devices. Note: Because these files are set by the storage vendor, they can be read but not written.
Master Boot Record (MBR)	Record that contains partition and boot information that the computer must read and interpret to communicate properly with the drive. The MBR is created using the Fdisk or similar utility.
File Allocation Table (FAT)	A table that the operating system uses to locate files on a drive. Due to fragmentation, a file may be divided into many sections that are scattered around the drive. The FAT keeps track of all these pieces and is created when the drive is formatted.

Data Corruption

Data is read and written by the host system in minimum 512-byte increments called *sectors*. Data corruption occurs when there is a power disturbance during the sector write operation.

On subsequent power-up, a read sector error occurs when the sector is read. The read sector error occurs because there is a mismatch between the data in the sector and the error checking information for that sector. Typically in solid state drives, the sector is replaced by a spare on the next write to that sector. The solid state drive ceases to operate when all spare sectors are depleted.

File System Analogy

A read sector error is analogous to a bad cluster after running Scandisk. When a cluster is labeled bad, it is unusable in the file system. Similarly a solid state drive views sectors with read errors as a sector that must be replaced with spares. The difference in these two scenarios is that if there are bad clusters after running Scandisk, the drive becomes smaller. In a solid state drive, spares get used. When all spares are exhausted, the drive ceases to operate on the next read error.

Figure 4 illustrates the relationship between sectors and files. When a power disturbance occurs during a large file write, the file may be truncated, but it should be truncated at a sector boundary for the storage device to be considered “robust.” File corruption may have occurred in this type of situation, but no spares are used. Significant problems described in “[Failure Examples](#)” on page 7 occur if the corruption occurs in the middle of a 512-byte sector or during the FAT update.



Figure 4: Sector and File Relationship

Failure Examples

- An unexpected power event could occur during the middle of a write operation, resulting in a read sector error.
- Many applications that encounter a read sector error automatically produce a system-level error, which results in system downtime until the error is corrected and the cause of the error is defined and eliminated.
- In addition to losing the data, the drive may thereafter interpret the read sector error as a defective sector, and unnecessarily replace it with a spare sector. When the number of factory-defined spare sectors decreases to zero, the drive needs to be replaced.

- A brown-out or low-power condition could cause the address lines of the storage device to reach an unstable state. If this occurs and write commands are still being accepted by the drive, there is a possibility that data could inadvertently be written to the wrong location and corrupt critical system files. This issue could result in a critical failure requiring the unit to be replaced.

Criticality of Errors Due to Power Disturbances

Data and drive corruption as the result of power disturbances can affect host systems differently depending on where and when the disturbance occurs. The following section discusses different types of files that could be corrupted and their impact on the host.

Master Boot Record

The MBR is a single-sector file that is always located at logical block zero. The MBR must be repaired if it becomes corrupted, or it will be impossible to access any volumes on that drive. Corruption of the MBR renders the drive inoperable. To re-use the drive, it needs to be repartitioned and reformatted.

File Allocation Table

Corruption of the FAT causes access to the files on the drive to be lost and requires that the drive be reformatted. All data is lost.

Critical System Files

When using an operating system, the loss of one or more of the critical system files (config.sys, system.ini, and so on depending on the operating system used) creates system errors or renders the drive unbootable. This situation causes unscheduled downtime, and may require that the drive be reformatted and the operating system be reinstalled.

Corruption of User Data Files

User data files can be also be corrupted. This type of corruption can often occur without the user being aware of it until it is too late. After a power disturbance has occurred, the entire drive should be checked and any sector errors repaired.

Possible causes of file corruption:

- Address lines that float to undetermined states — if the input power drops below the specified minimum operating level but there is still enough power to program the solid state storage component, data could be written to the wrong address.
- Inadvertent operating system overwrite — this corruption is due to system level instability brought on by a loss of power.

It is critical that the solid state drive is not physically damaged or “worn out.” In the case of a critical file overwrite, the host can repartition and reformat the drive. In the scenario of running out of spares, the product must be returned to the vendor.

The return scenario occurs as follows:

1. The drive fails by running out of spares.
2. The OEM sends the drive back to the vendor for failure analysis.
3. The vendor confirms the failure, but then must re-initialize the drive as part of its failure analysis process.
4. After re-initialization, no problems are found and the units are sent back.

In both cases, the data is lost and downtime occurs, but the drive can be used again.

Powerdown Tester Need

Today's embedded systems require technology to reduce the overall impact of an abrupt system powerdown. The varied applications expose their storage products to many forms of power disruption and interference that can cause voltage spikes, brown outs, and complete loss of system power. For this reason, many companies have included powerdown testing in their qualification process for any new storage media they are considering. The ability to handle power problems has an impact on overall reliability, customer goodwill, and total cost of ownership.

Powerdown Tester Design Considerations

This section describes the many factors that must be considered when designing a power down tester.

Powerdown Ramp Rate on Power and All I/O Pins

The steeper the ramp, the faster the voltage falls below the storage components' minimum programming voltage — which results in a greater chance for error. The ramp rates of various systems are shown in [Figure 5](#). The rates for the standard desktop PC and the embedded system emulate the loss of system power. The powerdown tester should also take into account the steeper ramp rate that emulates pulling out the drive while writing.




Figure 5: Ramp Rates of Various Systems

Ramp times are highly application-specific, and a powerdown tester should be designed to validate the worst-case scenario.

Powerdown Timing

A complete powerdown test requires power to be removed at various intervals of the write cycle. The ability to strictly control and repeat this variable exposes any weaknesses in the write cycle and accelerates any failures.




Figure 6: Powerdown Timing

Cut Power to the I/O and Vcc at the Same Time

It is important to eliminate any Vcc feedback that may occur if power is cut to Vcc and not to the I/O. Many test platforms may only cut power to Vcc, resulting in invalid test data and potential damage to the drive.

Validation and Verification

Read sector errors may occur at any point in time, yet the drive may continue to operate. The worst-case scenario is that the first read sector error occurs in a critical system file. Designers may choose to run the test until the drive becomes inoperable. The amount of verification performed directly impacts the test time, and testing after every cycle could be prohibitive.

Test Flow

The following figure illustrates a sample powerdown test program flow.



Figure 7: Sample Powerdown Test Program Flow

Table 4: Test Program Flow

Item	Description
Write to Unit Under Test (UUT)	Write a single sector using the ANSI ATA write command. This write isolates the error down to the rawest form.
Timer to Shutdown	A shutdown granularity of 1 μ s allows for variance of the shutdown sequence to powerdown at varying points in the write cycle, as shown in Figure 6 on page 10 .
Shutdown UUT	Ensures all proper power and signals are shut off to simulate the worst case scenario.
Delay	Plan for at least 10 ms to allow any stray capacitance to discharge.

Table 4: Test Program Flow (Continued)

Item	Description
Check for Read Errors	Scan previously-written sector using the ANSI ATA read command and confirm the read sector status. A status of 0x51 constitutes a read sector error. The test code should track and monitor the number of cycles to the first failure, and the number of cycles versus the number of failures.
Validate Data in Other Sectors	Ensure that data was not written to another location by mistake. The amount of validation required significantly impacts test time.

Conclusion

Solid state storage continues to replace rotating hard disk drives in embedded systems as the cost per usable gigabyte (i.e., the cost for the capacity required by the application) continues its rapid decline. The immense success of the iPod nano from Apple, which uses solid state storage instead of hard drives, is currently the most visible example of things to come.

The engineering trade-offs between storage component reliability (lower write/erase cycle endurance and higher bit error rates) and cost accelerates the need for PowerArmor, enhanced error correction, and wear-leveling capabilities. In addition, new technologies like WD's SiSMART, which monitors and predicts the SiliconDrive's remaining usable lifespan, are required for embedded systems with multi-year deployments. This technology drives an even larger differentiation between solid state drives designed specifically for embedded systems and flash cards designed for consumer electronic applications.




Figure 8: Solid State Storage Technology

For more information, visit WD's website at <http://www.wdc.com>.

© 2009 Western Digital Technologies, Inc.
All Rights Reserved

Information furnished by WD is believed to be accurate and reliable. No license is granted by implication or otherwise under any patent or patent rights of WD. WD reserves the right to change specifications at any time without notice.

Western Digital, Inc., the WD logo, SiliconDrive, PowerArmor, SiSMART, and SolidStor are registered trademarks in the U.S. and other countries; SiSecure, SiKey, SiZone, SiProtect, SiSweep, SiPurge, SiScrub, SiSTOR, and LifeEST are trademarks of Western Digital Technologies. Other marks may be mentioned herein that belong to other companies. Data reflects products in production as of June 2009. Not all products are available in all regions of the world. © Copyright 2009 Western Digital Technologies, Inc. All rights reserved.

1 Megabyte (MB) equals 1 Million Bytes; 1 Gigabyte (GB) equals 1 Billion Bytes. Accessible capacity may vary depending on the operating environment.

Contact Us

Western Digital
Solid State Storage Business Unit
26840 Aliso Viejo Parkway
Aliso Viejo, CA 92656

Tel: 949.900.9400

Fax: 949.900.9500

Website: <http://www.westerndigital.com>

For technical questions, contact: appeng@wdc.com

For Marketing/Sales, contact: salesorders@wdc.com