

# End-to-End Data Path Protection

Application Note AN004

February 2012



## Table of Contents

1	Introduction .....	2
2	Error Correction and Detection .....	3
2.1	SATA Link CRC .....	3
2.2	32-bit CRC .....	3
2.3	ECC Detection and Correction .....	3
2.4	Data Fail Protection .....	4
3	Soft Errors .....	5
3.1	Types of Particle Upset .....	5
3.2	Classes of Failure due to Soft Error in an SSD.....	5
4	Summary .....	6

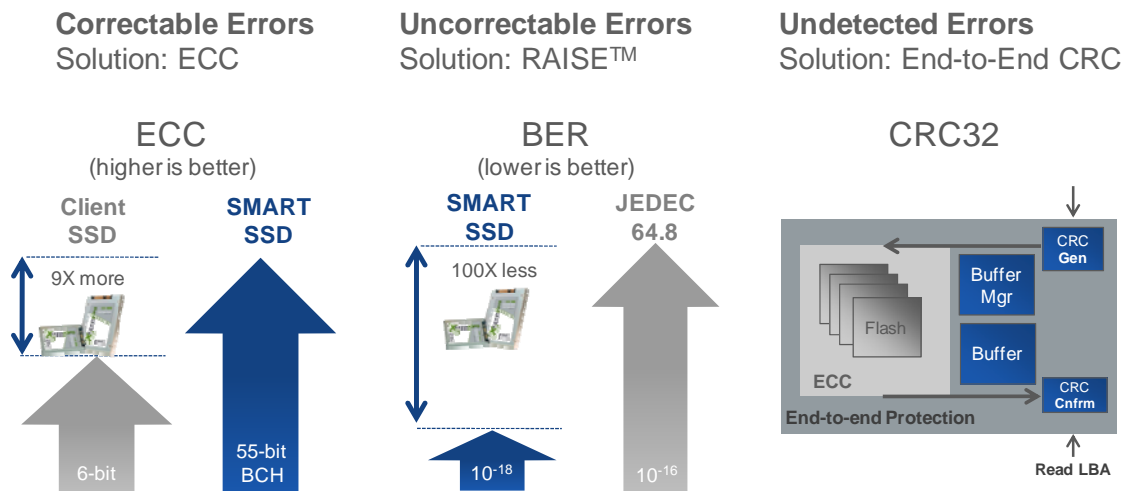
## 1 Introduction

Threats to data integrity in Solid State Drives, as in HDDs, fall in to two basic categories, NAND flash memory errors (analogous to media errors in HDDs) and errors that occur to data in-transit either to or from the flash. Errors in NAND flash consist of both correctable data errors and more catastrophic flash failures.

**Correctable Errors:** Bit errors occur fairly frequently throughout the flash array, and increase as the flash ages. These type of errors are detected and corrected on-the-fly by the controller’s ECC engine (refer to Figure 1). SMART SSDs (XceedIOPS2, XceedStor 500S and Xcel-200) feature a powerful ECC engine that is capable of correcting 55 bits per 512 bytes, using a BCH algorithm. This is substantially better than most Client SSDs on the market today, which are typically capable of detecting and correcting 4 or 6 bits per 512 bytes.

**Uncorrectable errors:** When data bit errors occur in numbers too great for the ECC engine to correct, or when NAND flash pages or blocks fail outright, an SSD’s primary data protection capability breaks down. Without additional protection against these threats, data loss is a real risk. SMART SSDs feature a Data Fail protection scheme based on the RAISE™ capability of the controller (refer to Figure 1). RAISE correction can recover entire flash memory blocks, and reduces the Uncorrected Bit Error Rate (UBER) risk to better than one in  $10^{18}$  bits read, exceeding JEDEC requirements by two orders of magnitude.

Figure 1: SMART SSD Reliability Features



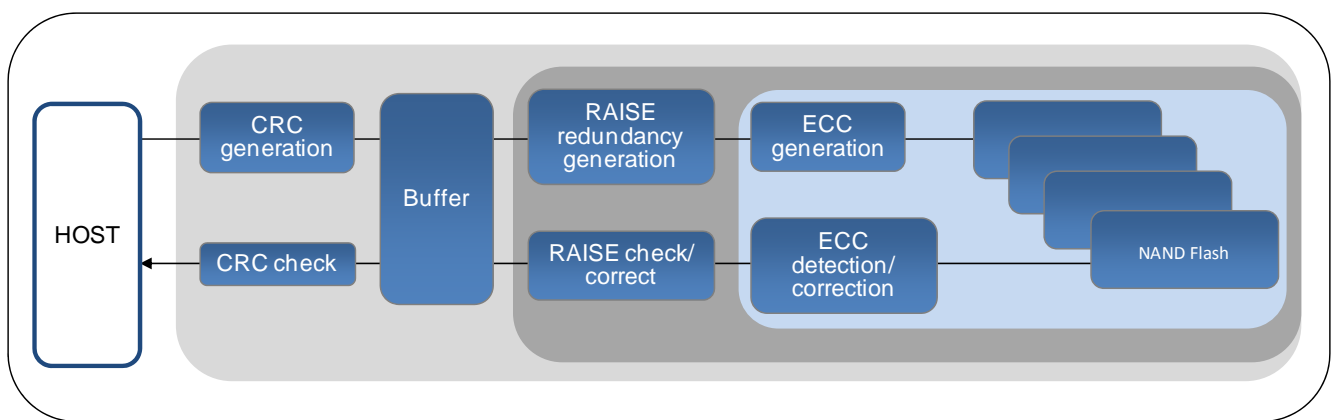
**Undetected Errors:** While exceedingly rare, single bit errors can escape detection by the ECC engine. Undetected errors in this subcategory result will result in bad data returned being to the host as good. SMART SSDs include 32-bit CRC Data Path error detection (refer to Figure 1) to protect against this subcategory of error. No more than one in  $10^{10}$  errors that escape the ECC engine will also escape CRC detection. Furthermore, data in transit to/from flash memory are susceptible to external events such as electrical noise and soft errors due to particle induced bit flip. Unless an SSD is also able to detect these errors, in spite of any other error protection subsystems, the integrity of mission-critical data will still be at risk. SMART SSD’s 32-bit Data Path CRC error detection provides protection here as well.

While all SSDs designs include ECC engines, and some include CRC, few include protection for catastrophic flash memory failures. An SSD that lacks the capability to recover from an uncorrectable error, such as a page or block failure is unlikely to meet the requirements of high reliability applications. SMART SSDs implement all three, and are highly suited for the most demanding Enterprise and Industrial applications.

## 2 Error Correction and Detection

Figure 2 shows a block diagram of the various data protection domains within the SMART SSDs. The protection domains overlap to ensure seamless protection.

Figure 2: Data Protection Domains



### 2.1 SATA Link CRC

Data on the SATA link is protected by industry-standard link CRC. Each Data FIS is determined to be either intact or in error (for example, due to noise on the link) as received by the SMART SSD SATA Link layer.

### 2.2 32-bit CRC

Once data is written from the host to the drive, the controller calculates a 32-bit round-trip CRC redundancy for the data. The data traverses along the normal high-speed data path, including the embedded SRAM write buffer, to be written to the flash. Errors incurred along this path are detectable via the hardware CRC detection capability and via single-bit parity of the SRAM buffer. The round-trip data path CRC protects data from the moment the CRC is generated (just inside the SATA link), to the flash memory (where the CRC redundancy is stored with the data), and back out to where the CRC is checked before the data is delivered to the host system. It guards against undetected errors arising from operating condition anomalies, particle bombardment, and in extreme cases limitations in ECC correction capability.

### 2.3 ECC Detection and Correction

Once the data passes the internal SRAM write, ECC redundancy is generated (see Figure 2). Within the ECC data correction domain, data errors are not only detectable (by the ECC engine), but are also correctable. The

ECC engine used by the SMART SSDs is capable of correcting up to 55 bits per LBA, using the BCH algorithm. This correction power is best-in-class for both Enterprise and Client SSDs.

ECC correction occurs when data is read back from the flash and compared to the ECC checksum that was originally written together with the data. Errors caused en route to flash memory, but after the ECC checksum has been generated, are stored with the data. Those errors will be corrected (if within the maximum data correction capability of the controller) when that data is subsequently read.

The ECC engine used in the SMART SSDs is one of the most powerful in the industry, providing substantially higher Uncorrectable Bit Error Rate (UBER) than most competing SSD solutions in the market. SMART SSDs uncorrectable error rates are 1 error per  $10^{18}$  bits read, which is better than that of other Client SSDs by two orders of magnitude.

The JEDEC JESD218<sup>1</sup> specification calls out for a minimal UBER of one in  $10^{16}$  bits read for Enterprise-class SSDs. The SMART SSDs are rated for 1 uncorrectable error per  $10^{18}$  bits read, a factor 100x better than the specification requires and 100x better than competing SSD solutions.

At a read rate of 200 MB/sec, one error in  $10^{16}$  bits read translates to an uncorrectable sector every 72 days (2.3 months). The two additional orders of magnitude provided by the SMART SSDs extend that average to 7,233 days, or nearly 20 years between uncorrectable errors (15 years beyond the drive's warranted lifetime).

## 2.4 Data Fail Protection

Before data is written to the NAND flash memory, additional redundancy is generated. This redundancy is stored to enable a RAID-like rebuild of flash memory pages or blocks in the event of catastrophic NAND flash block failure or large quantities of bit errors. High bit errors on reads and block failures can put large amounts of data at risk, not to mention the reliability of the SSD itself (potential risk to firmware images, configuration records, etc.).

Unlike ECC redundancy, which protect a single LBA's worth of data and focuses on bit error correction, the highly efficient RAISE™ redundancy and correction mechanism of the SF-1500 controller is tailored toward rebuilding larger flash memory structures. RAISE™ stands for "Redundant Array of Independent Silicon Elements", this correction mechanism can rebuild and recover data that has suffered even hundreds of bit errors. The RAISE rebuild protection domain is highlighted in dark grey in Figure 2. RAISE is invoked only if the ECC engine is not capable of correcting the error and reduces the probability of undetected read errors from the flash.

The Data Fail Protection capability of the SMART SSDs is unique in the market. Very few other SSDs available today incorporate this capability. When catastrophic NAND failures occur on these SSDs, the data is lost and the drive itself can become unusable. As a result, the actual overall probability of uncorrectable error can be substantially worse than specified.

---

<sup>1</sup> JEDEC JESD218 specification is available for download <http://www.jedec.org/>

### 3 Soft Errors

All Solid State Drives are subject to soft errors, which include the following errors:

1. Data errors in the data path caused by alpha particle (deep space sub-atomic particle) bombardment;
2. Data errors in the data path caused by neutron (sub-atomic particle emanating from the package material of the flash controller) bombardment;
3. Data errors in the data path caused by other external factors (noise, thermal issues, etc);
4. Data errors in the data path caused by hardware or firmware anomalies;

The SMART SSDs exclude the errors described in item 3 and 4, due to a solid hardware design, verification and ongoing monitoring of thermal parameters during operation. Soft Errors are generally thought of in terms of internal storage node excitation due to particle upset, as described in 1 and 2.

Particle upset is calculated in FIT (Failures In Time) and standardized for particle occurrence at sea level at the latitude of New York City. One FIT is normalized to one error for every  $10^9$  device operating hours

Note: At high altitudes or latitudes, results change due to the different incidence of alpha particles from deep space.

#### 3.1 Types of Particle Upset

Particles impacting the following circuit “storage” nodes can flip the state of the node:

- RAM (DRAM or SRAM) cells
- Internal Flip Flop or Latch cells

In the SMART SSDs, Flip Flop upset is insignificant. By far the lion’s share of soft errors is caused by SRAM cell upset. The XceedIOPS SSD employs parity error detection with firmware correction capability on internal SRAM memories

#### 3.2 Classes of Failure due to Soft Error in an SSD

Soft errors can cause a variety of error types. In order of severity, these error types are:

1. No failure (RAM cell or flip flop did not hold critical data at time of particle upset);
2. Data can be corrupted on write and can be corrected on subsequent read;
3. A transaction can fail (due to data being corrupted during read) and causes the host to perform a retry;
4. Data can be corrupted during a write operation and can be rendered irrecoverable;
5. A drive can hang or freeze without data loss, requiring a hard reset or power cycle to recover;
6. A drive can hang or freeze with data loss, requiring a hard reset or power cycle to return to operation;
7. A drive can hang or freeze and not return to operation (file system or firmware corruption);
8. Data corruption on read can go undetected (silent error), with invalid data being returned to the host;

In the SMART SSDs, these classes of failure have different probabilities. Due to the overlapping protection domains and the manner in which errors are detected and handled, silent corruption probability is extremely low:

## 4 Summary

Solid State Drives must protect data along all data paths and within all storage cells of the NAND flash. This includes protection from flash bit errors, catastrophic flash failures, sub-atomic particle-induced soft errors within internal SRAMs and along control and data paths, and across interfaces including the host interface and the flash memory channels.

SMART Storage Systems SSDs provide state-of-the-art error detection and correction, Data Fail Protection and guarantees full integrity of stored firmware, file system data, and meta-data.

Providing overlapping data protection domains, SMART SSDs deliver best-in-class data path protection that exceeds the JEDEC JESD218 specification for Enterprise-class SSDs by two orders of magnitude.

**Disclaimer:**

No part of this document may be copied or reproduced in any form or by any means, or transferred to any third party, without the prior written consent of an authorized representative of SMART Storage Systems. ("SMART"). The information in this document is subject to change without notice. SMART assumes no responsibility for any errors or omissions that may appear in this document, and disclaims responsibility for any consequences resulting from the use of the information set forth herein. SMART makes no commitments to update or to keep current information contained in this document. The products listed in this document are not suitable for use in applications such as, but not limited to, aircraft control systems, aerospace equipment, submarine cables, nuclear reactor control systems and life support systems. Moreover, SMART does not recommend or approve the use of any of its products in life support devices or systems or in any application where failure could result in injury or death. If a customer wishes to use SMART products in applications not intended by SMART, said customer must contact an authorized SMART representative to determine SMART's willingness to support a given application. The information set forth in this document does not convey any license under the copyrights, patent rights, trademarks or other intellectual property rights claimed and owned by SMART. The information set forth in this document is considered to be "Proprietary" and "Confidential" property owned by SMART.

ALL PRODUCTS SOLD BY SMART ARE COVERED BY THE PROVISIONS APPEARING IN SMART'S TERMS AND CONDITIONS OF SALE ONLY, INCLUDING THE LIMITATIONS OF LIABILITY, WARRANTY AND INFRINGEMENT PROVISIONS. SMART MAKES NO WARRANTIES OF ANY KIND, EXPRESS, STATUTORY, IMPLIED OR OTHERWISE, REGARDING INFORMATION SET FORTH HEREIN OR REGARDING THE FREEDOM OF THE DESCRIBED PRODUCTS FROM INTELLECTUAL PROPERTY INFRINGEMENT, AND EXPRESSLY DISCLAIMS ANY SUCH WARRANTIES INCLUDING WITHOUT LIMITATION ANY EXPRESS, STATUTORY OR IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

©2012 SMART Storage Systems. All rights reserved.

**Corporate Headquarters:** 39870 Eureka Dr., Newark, CA 94560, USA ♦ Tel:(510) 623-1231 ♦ Fax:(510) 623-1434 ♦ E-mail: info@smartstoragesys.com

**Flash Design Center:** 2 Robbins Road, Westford, MA 01886, USA ♦ Tel:(978) 303-8500 ♦ Fax:(978) 303-8757

**Flash Design Center:** 2600 W. Geronimo, Chandler, AZ 85244, USA ♦ Tel:(480) 792-8900 ♦ Fax:(480) 792-8901

**Asia:** Plot 18, Lrg Jelawat 4, Kawasan Perindustrian Seberang Jaya 13700, Prai, Penang, Malaysia ♦ Tel:+604-3992909 ♦ Fax:+604-3992903