

Complete Power Failure Protection for SQFlash

Author: Precyan Lee

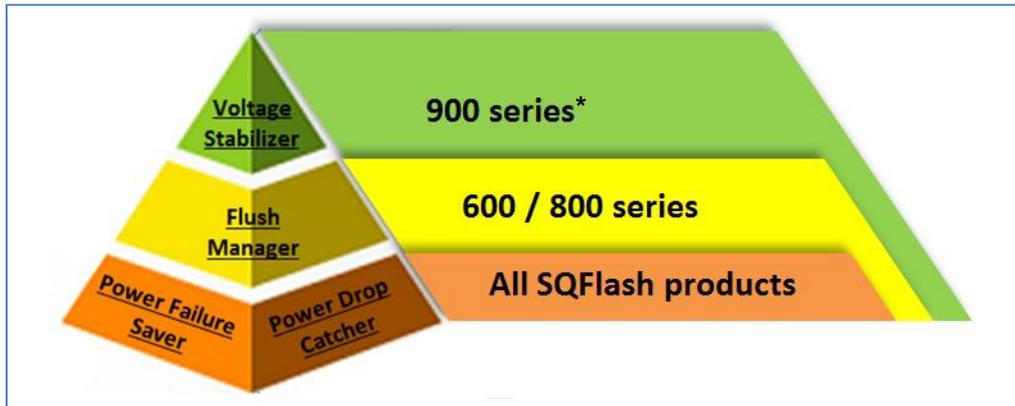
E-mail: precyan.lee@advantech.com.tw

Table of Contents

Introduction.....	1
1. Power Failure Saver.....	2
2. Power Drop Catcher.....	4
3. Flush Manager	6
4. Voltage Stabilizer	9
5. Design Verification of Power Failure Protection.....	11

Introduction

Advantech SQFlash features complete power failure protection in all product SKUs, proper power failure protection mechanisms are implemented in different product series. This complete scheme combines four different functions – Power Failure Saver, Power Drop Catcher, Flush Manager, and Voltage Stabilizer.



* Voltage Stabilizer is also available in 830V series

This white paper discusses specific protection functions in SQFlash product lines. Also, the verification of power failure resistance in each of the SQFlash design / development stages will also be covered the end of this white paper.

1. Power Failure Saver

When using flash products, unavoidable conditions may occur, such as high/low temperatures, unstable voltage and vulnerable power supply, which may affect the reliability of the data. If the power suddenly goes down or surges when the machine is operating or the voltage is lower than the specification, then processing data and information could get lost, even the memory device itself may get damaged. Therefore, in industrial applications, a reliable flash device with stable OS is required for normal operation. Within all Advantech SQFlash memory devices, just such protection exists.

A. How it works...

i. Voltage Detection

The SQFlash controller IC has built in low voltage detection, when the power source drops lower than 2.9 V, the internal RESET will be triggered and the controller will stop all activity in the flash. While the working voltage is 2.9 ~ 3.3 V, the flash will still keep running. This will prevent the invalid data being written or access to the flash during power loss.

ii. Data Integrity Check after Power Resume

When the controller IC is writing data and suddenly the power fails, the controller IC will judge if it has completed one full page of written data before power failure; If not, the data of this incomplete page will be flagged as having failed. The controller IC will then implement a recovery mechanism in FW, and after powering up, it will check the data written in the flash previously, if the data has been detected as invalid it drops the invalid part, then merges the old valid data with the newer valid data to make a new combination of data that are all valid. This will eliminate the chance of getting wrong data caused by non-complete writing to flash.

B. Verification of Power Failure Saver

i. Test setup

Condition for H/W Platform:

- Power control machine x 1
- SOM CPU board x 1
- SOM environmental board x 1

Condition for S/W Tool:

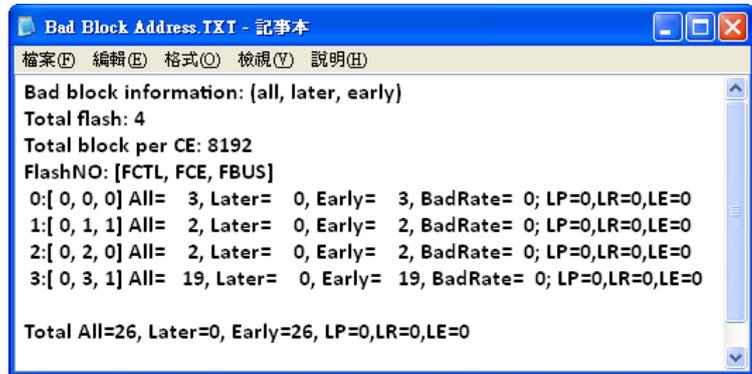
- Windows XP SP2
- Burn-In Test V6.0 Pro
- Advantech power off count tool
- Auto-run agent tool

ii. Test method

- Sudden power down during data writes
- Verify data integrity after power resumes

iii. Test result

Normally works 5496 times with no later Bad Block occurrence.



2. Power Drop Catcher

In some ruggedized environments, the power supply could drop very sharply in a short period of time. Although power might recover within milliseconds (< 1ms), for low power Flash storages like CFast / mSATA, it could cause devices to freeze up and be unable to boot up or read / write correctly. So SQFlash embedded has a power Reset IC command in all product series to prevent this kind of sudden power drop and retain SSD controller functionality.

A. How it works...

The normal reset circuit has a wide range of reset timings from 20ms ~ 60ms, when power fluctuations occur and reset timing is variable, the SSD controller might be at risk. So the embedded Reset IC controller will constrain the reset timing to exactly 55ms. In such cases, even if there's a sudden power drop, the Reset IC will always maintain exactly 55ms to catch up with SSD controller reset timing, so SSD risk will be significantly reduced.

B. Verification of Power Drop Catcher

i. Test setup

Condition for H/W Platform:

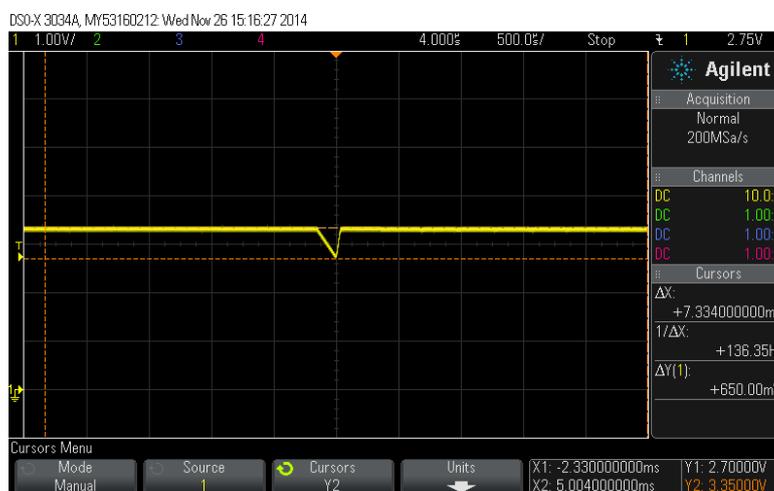
- Power control machine x 1
- MIO-5250 with SATA / PCIe switch IC x 1

Condition for S/W Tool:

- Windows 7 32-bit
- Burn In Test V6.0 Pro
- Advantech Power off count tool

ii. Test method

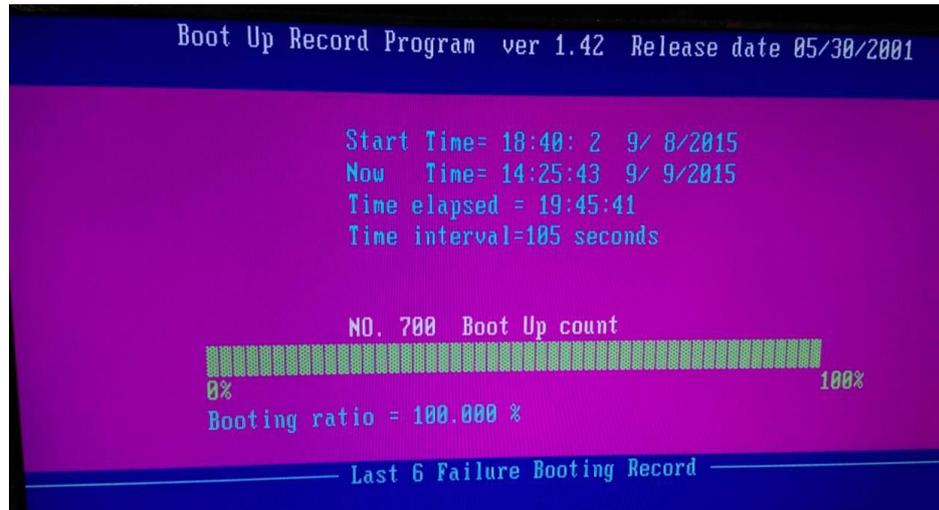
- Power on / off test to verify if the disk can boot up properly when there's a power drop during the boot up process. (1V voltage drop for 0.5 ms)



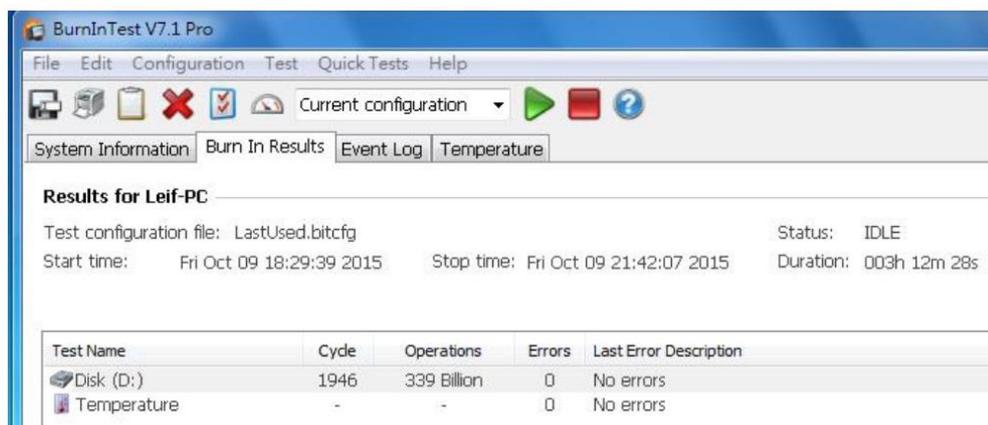
- Burn-in test to verify if the disk can function stably after booting up.

iii. Test result

100% successful boot up over 700 test cycles.



3 hour burn-in after booting into OS without any errors



3. Flush Manager

DRAM (dynamic random-access memory) is a volatile memory with fast access time, frequently used as a temporary cache or buffer between a host controller and backend storage. In today's SSD (solid-state drive) designs, the utilization of DRAM cache (either internal or external) has become a common practice to boost overall SSD performance, especially on small file transfers. Apart from performance benefits, caching can also improve the endurance of SSDs by consolidating multiple small transfers before pushing them to NAND Flash, reducing the amount of block erase during the process. However, there is one major drawback of cache memory—it requires power to maintain the stored information in the DRAM. This raises concerns about data integrity under situations such as unstable power supply or even power failures. Thus, it is critical for SSD firmware to implement intelligent protection schemes to preserve data integrity in the event of unexpected power loss. Several new cache mechanism technologies have been implemented in the latest controller for SQFlash SATA Flash drives to better handle power-loss situations. These technologies will be introduced in following sections.

A. How it works...

i. Smart Flush

SQFlash SATA controllers utilize smart algorithms to reduce the amount of the data which resides in the external cache. Smart Flush technology allows incoming host data to make a “pit stop” in cache and then be pushed into NAND storage immediately, barring bottlenecks on the Flash interface. If the interface backs up, the cache is organized and consolidated into groups before being written to improve the write amplification.

ii. Cache-off

An ACK will only be delivered to the host when the data is fully committed to the NAND media. Other solutions on the market may send an ACK to a flush command when the data is not yet committed to the NAND. Such an implementation gives a false-positive result of the data's integrity and runs the risk of power failure issues. Once the data is committed to the NAND media, the following page writes will not impact previously committed data. This is made possible by intelligently managing the pair-page of the MLC Flash.

B. GuaranteedFlush (Pair Page Management)

When each flash cell is programmed, neighboring word-line cells might be accidentally corrupted if unexpected events occur. In multi-level cell (MLC) flash, each word-line page consists of a number of fast and slow pages. Fast pages have better endurance and SLC-like program times compared with slow pages. The figure below is an example of a flash pair page table. In this case, we consider page (F)#0 and page (S)#4 as a pair, and page (F)#1 and page (S)#5. NAND Flash programming always follows two rules: 1. the program order always follows the page sequence and 2, page (F) is always programmed prior to page (S). If interruption occurs during slow page programming, then it becomes high risk resulting in corresponding page corruption. See figure 1 as below.

Page(F)	Page(S)
0	4
1	5
2	8
3	9
6	12
7	13
10	16
11	17
14	20
15	21
18	24
19	25

Figure1 : NAND Flash programming

In order to ensure data integrity, firmware automatically fills up with dummy data sequentially till its corresponding Page (S) receives a flush cache command, which effectively protects Page (F) data from the impact of non-fully programmed Page (S) during a sudden power loss. The following example shows how GuaranteedFlush protects initial pages (F) 0-3, when unexpected power losses occur.

Page(F)	Page(S)
0	Dummy 4
1	Dummy 5
2	Dummy 8
3	Dummy 9
Dummy 6	12
Dummy 7	13
10	
11	

Figure2 : GuaranteedFlush

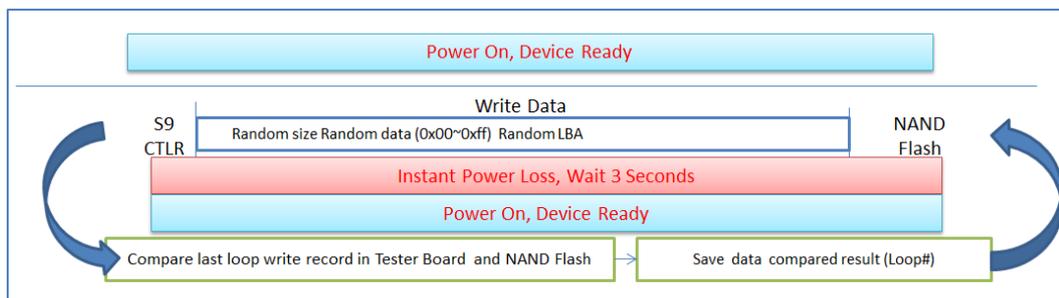
During NAND Flash programming, firmware will fill up pages 4-9 with dummy data automatically once the program page (F) 0-3 has finished programming. The next write starts from page (F) 10. Even if there's a sudden power loss, page (S) 12, and page (F) 0-3 data won't be impacted due to the already programmed pages (S) 4, 5, 8, 9, which helps eliminate the risk of data corruption to page (F) 0-3.

C. Verification test of Flush Manager

i. Test flow

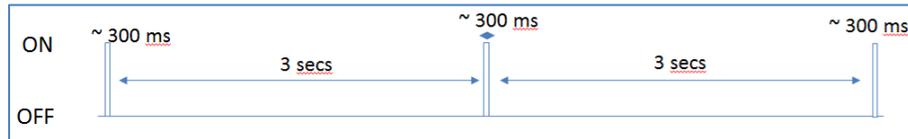
Cycle: Start → Write CMD → Flush CMD → Power Cycle → Read CMD → Compare Data

Testing File Size: Random data



Testing Loops: to simulate instant power interrupt (abnormal power supply)

Power Cycling:



ii. Test result

35,000 cycles passed

```
* SATA SSD (port: 2 ) G3*
*****
SPOR Pattern Testing...
Start Fill Up Disk...
Progress : Fill Up Disk Completed
Progress : Test Round[4000]
Progress : Compare Data

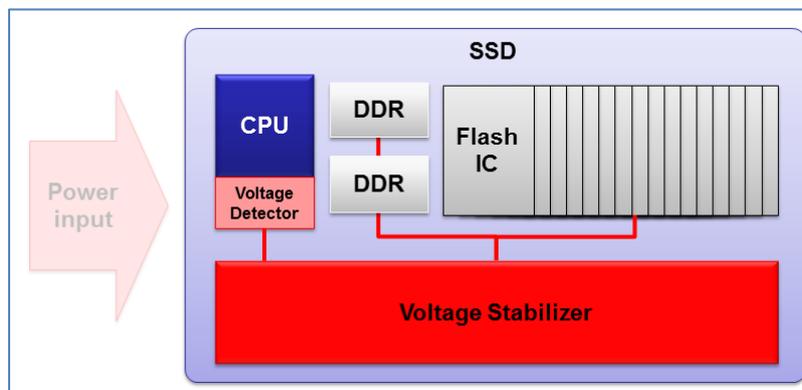
Total Compare Pass Times:35000
Total Compare Fail Times:0
Test Complete!!
Press Any Key to Continue...
```

4. Voltage Stabilizer

SQFlash 900 series product line has built-in Voltage Stabilizer to ensure the SSD internally can always provide a stable power supply to the DDR chip and the Flash IC even when the power input drops below the lower limits of the power input.

A. How it works...

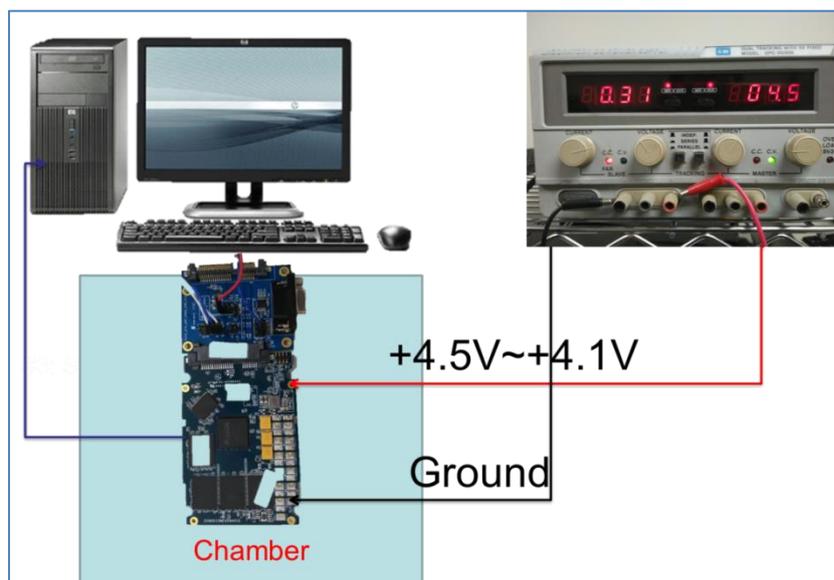
While the built-in voltage detector detects an unstable power input ($< 4.75\text{ V}$ or $> 5.25\text{ V}$), the controller will issue a power failure interrupt and force a Flush CMD first. At the same time, the whole internal power supply will be switched to Voltage Stabilizer (tantalum capacitor array) immediately to ensure stable power is supplied throughout the whole drive. This ensures the Flash IC and DDR IC will not operate with unstable power which could lead to data errors or bad data integrity.



B. Verification test of Voltage Stabilizer

This test demonstrates real low power operation under high temperature to verify the Voltage Stabilizer design.

i. Test setup



ii. Test conditions

Chamber with > 85 °Celsius ambient temperature

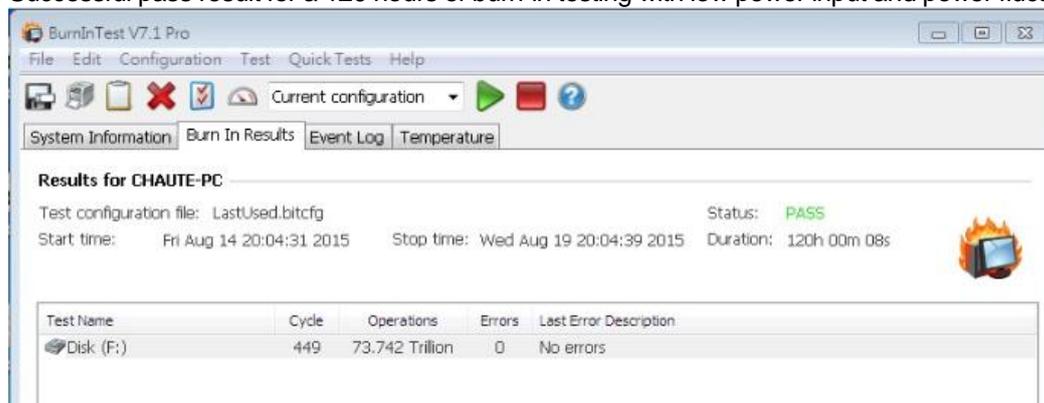


iii. Test flow

- Enable burn-in test with +4.5V power input for 100 hours
- Drop input power to +4.1V power input for 5 minutes
- Recover power to +4.5V for 5 minutes
- Repeat b) and c) for additional 5 cycles
- Stay on +4.5V for another 19 hours

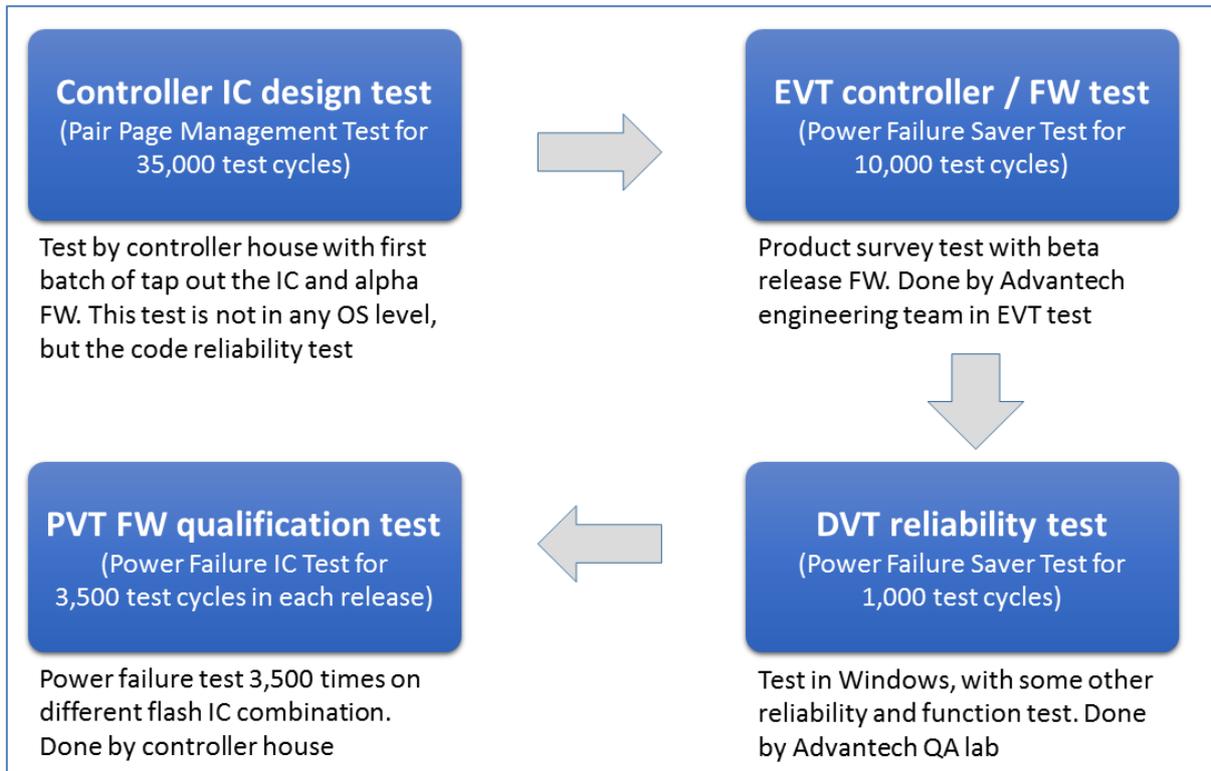
iv. Test result

Successful pass result for a 120 hours of burn-in testing with low power input and power fluctuations.



5. Design Verification of Power Failure Protection

The following flow diagram shows how SQFlash is verified in all design phases.



Advantech Co., Ltd. –Founded in 1983, Advantech delivers visionary and trustworthy industrial computing solutions that empower businesses. We cooperate closely with solution partners to provide complete solutions for a wide array of applications in diverse industries, offering products and solutions in three business categories: Embedded ePlatform, eServices & Applied Computing, and Industrial Automation groups. With more than 3,400 dedicated employees, Advantech operates an extensive support, sales and marketing network in 18 countries and 39 major cities to deliver fast time-to-market services to our worldwide customers. Advantech is a Premier Member of the Intel® Embedded and Communications Alliance, a community of embedded and communications developers and solution providers. (Corporate Website: www.advantech.com).
Copyright© 2013 Advantech Co., Ltd. All rights reserved.