Power Loss Imminent (PLI) Technology

Introduction

In recent years, there have been major shifts in weather patterns. Inclement extremes have become the norm with heavy snow storms, hurricanes, flash flooding, and ominous thunder and lightning storms, all of which can interrupt electrical power. Additionally, our nation's power grid and substations are at risk of terrorist attacks seeking to cause havoc and disruption. Whether an act of nature or of evil, our daily lives and the critical data we rely on to run governments, corporations, banks, local businesses, schools, and households are impacted by power loss events. It's imperative to protect critical data, and to protect the integrity of the data at risk due to power loss events.

Data centers and their administrators are responsible for the thousands of storage drives at risk in a power loss event, and whether the last operation prior to the power outage was completed. One of the newest technologies embedded in a Solid-State Storage device is Power Loss Imminent (PLI) technology. PLI technology significantly reduces the possibility of losing data during a power loss event. The PLI technology feature in a Solid-State Drive (SSD) continuously prepares for unexpected power loss, thereby protecting the data from corruption and destruction.

The Storage Power Loss Problem

During a clean shutdown, most host systems initiate the ATA "STANDBY IMMEDIATE" command to the SSD, signaling it to prepare for shutdown. In preparing for shutdown, the SSD saves any data in transition or in flight (in temporary buffers) to the non-volatile NAND Flash media.

However, during an unsafe or unexpected loss of power, the SSD abruptly loses power before the host system can initiate the ATA "STANDBY IMMEDIATE" command, or mid-way through the command. An unexpected loss of power can mean that the data in the SSD controller's temporary buffers will not be saved to the non-volatile NAND Flash media, which often results in data corruption or loss due to incomplete or partial writes.

PLI Technology Details

Not all SSDs contain PLI technology. SSDs with PLI technology contain energy storing capacitors. The energy storing capacitors provide enough energy (power) to complete any commands in progress and to make sure that any data in the temporary buffers is committed to the non-volatile NAND Flash media. These capacitors act as backup batteries for the drive.

Figure 1: 1 of 2 PLI Capacitors shown from Intel® SSD DC 3700 Series Product

Figure 1 highlights (yellow box) a PLI capacitor in an SSD. The SSD supply voltage detector circuit constantly monitors the drive's supply voltage. If voltage falls below a predefined level, indicating a power loss event is imminent, the SSD will use the backup energy in the capacitors to write any data in temporary buffers to the NAND Flash media. After power is restored, the PLI capacitors will be charged again, ready for a future power loss event.

Figure 2 illustrates the switching sequence that engages the PLI capacitors to the SSD controller during a power loss event. As shown, switch 1 (SW1) is closed, connecting the capacitor energy to the SSD ASIC Controller, and switch 2 (SW2) is opened, disengaging the host power supply source. The energy stored in the PLI capacitors will be used to finish transferring any data in the buffer to the non-volatile NAND Flash media. When power is restored the PLI capacitors will be charged up again and ready for a future power loss event.

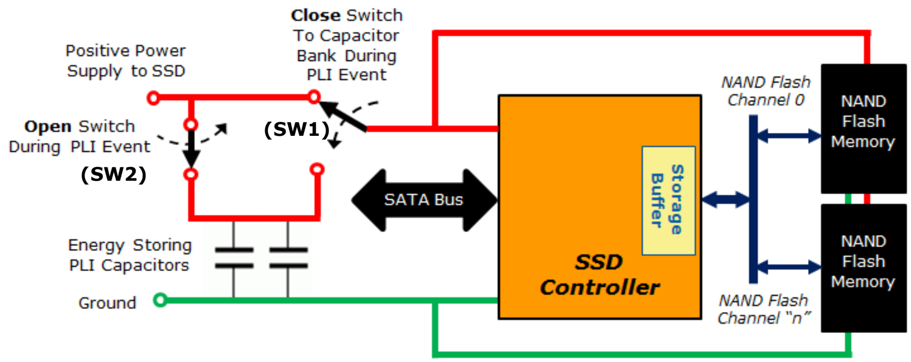

Figure 2: PLI Capacitor Engagement Sequence During Power Loss Event

Architecting and Implementing PLI Technology Solution into a Solid-State Drive

Designing PLI technology into an SSD is quite involved. The design process requires intricate architectural integration into the whole SSD's operation, from enumeration to a proper shut down. A critical PLI feature is the ability to anticipate needs of the system administrators, and to monitor the health, performance and usage count of the PLI circuitry. PLI architecture incorporates the SMART command and functionality necessary to periodically verify the PLI circuitry continues to function properly over the life of the SSD. Additionally, the PLI capacitor chemistry choice needs to be carefully selected for reliability over the lifetime of the product. This capacitor choice is key and, depending on the capacitor type used, will dictate supporting circuits and firmware which provide storage robustness through power loss events.

PLI Technology Architecture

Monitoring SSD health and reliability, meeting and anticipating system administrator needs, all while balancing the financial and performance cost of restoring an entire data center after a power loss event, provides additional challenges in defining a good architecture. The architecture of the Intel® SSD DC S3700 Series and S3500 Series has taken these needs into account. Figure 3 shows an operational architectural diagram depicting the key PLI components, the SSD voltage detector, the charge storing capacitor bank, and the SMART (<u>S</u>elf-<u>M</u>onitoring <u>A</u>nalysis and <u>M</u>onitoring <u>T</u>echnology) attribute extraction module.

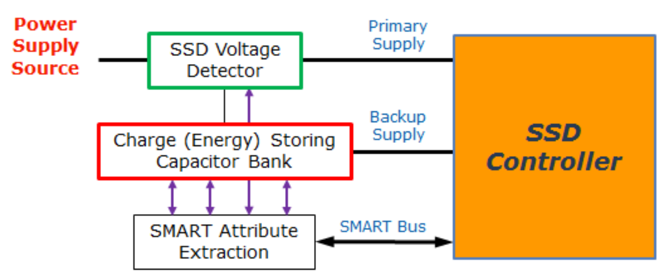

Figure 3: Intel® Solid-State Drive PLI Architecture

The SSD voltage detector circuit constantly monitors the drive's supply voltage. If the voltage falls below a predefined level, indicating a power loss event is imminent, the SSD will use the backup energy supply from the capacitor bank to write any data in temporary buffers to the NAND Flash media. After the power is restored, the PLI capacitors will be charged up again and ready for a future power loss event. The SMART attribute extraction module provides system administrators with specific health information and status of the PLI capacitors, which is the heart of the PLI technology. Since the Intel® SSD DC S3700 Series and S3500 Series support SMART attributes and have specific PLI SMART attributes, the system administrator can retrieve the drive's PLI SMART attribute reported values.

Charge (Energy) Storing Capacitors for PLI

The selection of PLI energy storing capacitors is a very critical component choice. Accurate selection of components is crucial to the PLI technology and the architectural implementation essentially revolves around this choice. There are many critical decision points and trade-offs regarding this selection including cost, capacitance offerings, energy storing efficiency, form factor, operating temperature range, capacitor use life, and its endurance.

Examining the capacitor family tree in Figure 4, there are a variety of choices for each solution and application. For the Intel® SSD DC S3700 Series and S3500 Series, Aluminum Electrolytic (Wet) capacitors were chosen for their specific attributes:

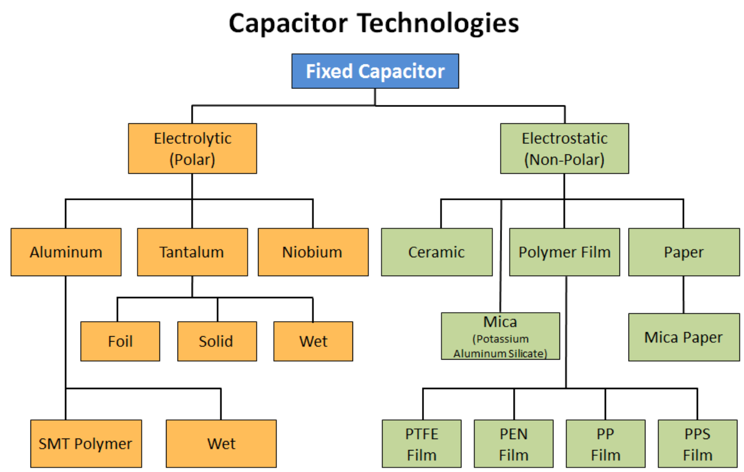

Figure 4: Fixed Capacitor Technology Tree

- Wide Capacitive Size Offerings: 0.1 to 68,000 μF
- Suitable form factor
- Economical cost
- Rated Voltage Range of 6.3 Volts to 450 Volts
- Good Endurance: Up to 5 years at 80 °C
- Wide Operating Temperature Range: -40 to 105 °C

The family of electrostatic capacitors was not considered because of their DC bias characteristic phenomenon wherein the effective electrostatic capacitance decreases when a DC voltage is applied. With the need for PLI capacitors to constantly be ready to discharge, this capacitance degrade is not acceptable nor is the resulting varied energy source caused by the degrading capacitance. The other good electrolytic capacitors, Tantalum and Niobium, were not chosen because they are more costly, as shown in Table 1, below.

Table 1: Sample Fixed Capacitor Pricing

Capacitor Type	Supplier	Part Number	Capacitance	Rated Voltage	Price# (Quantity of 100)
Electrolytic Aluminum Wet	Nichion	UKL1V101KPDANA	100 μF	35 V	$0.225
Electrolytic Tantalum Wet	Vishay	138D107X9035F2	100 μF	35 V	$0.401
Niobium Oxide	AVX	NOJD1D7M010RWB	100 μF	10 V	$1.27

SMART (Self-Monitoring Analysis, and Reporting Technology)

Integrating SMART technology into the PLI architecture is essential. Providing the ability for system administrators to periodically check the health of the SSD and the health of the PLI technology, particularly after a power loss event, is crucial.

Intel® SSD DC S3700 Series and S3500 Series have SMART technology incorporated into their PLI architecture. The Intel SSD DC S3700 Series and S3500 Series provide two convenient SMART attributes for the system administrators with key data on power loss events as well as the SSD PLI circuitry health. The SMART power loss and PLI attributes are summarized in the table below.

Table 2: Intel SSD DC S3700 Series and S3500 Series PLI SMART Attribute Descriptions

<table>
<tr><th>SMART ID</th><th>SMART Attribute</th></tr>
<tr><td rowspan="2">AEh</td><td>Unexpected Power Loss</td></tr>
<tr><td>Raw value reports number of unclean shutdowns cumulative over the life of the SSD. An unclean shutdown is the removal of the SSD’s power <u>without</u> the host issuing a “STANDBY IMMEDIATE” as the last command (regardless of PLI activity using capacitor power). The normalized value, which is a percentage value, always is a 100.</td></tr>
<tr><td rowspan="3">AFh</td><td>Power Loss Protection Failure</td></tr>
<tr><td>Last test result of PLI (Provides 3 Health Check Report Out Values)</td></tr>
<tr><td>1. “Microseconds to Discharge Capacitors” (Bytes 0-1): The SSD provides a self-test where the backup energy storing PLI capacitors are partially discharged to make sure they are able to release and sustain their energy as designed; to write any in flight, committed data in the SSD controller storage buffers to NAND Flash memory. In the case of the Intel® SSD S3700 Series and S3500 Series, it should meet its minimum 25 µS.
2. “Minutes since last test saturates at max value” (Bytes 2-3): SSD also keeps track in minutes since the last PLI capacitor test was done by the SSD. So if the system administrator wanted to know when the last test was done, that value, in minutes will be reported up to the maximum recording time of 2 bytes of data (1111,1111, 1111,1111) recording bits, in minutes, or 65,535 minutes or 2,730 days.
3. Lifetime number of tests, not incremented on power cycle, saturates at max value. (Normalized value: set to a “1” on test failure or a “11” if the capacitor has been tested in an excessive temperature condition, otherwise 100 (Bytes 4-5):
a) Cumulatively tracks the number PLI test, not counting power on cycling of the drive.
b) Reports a “1” if there is a PLI test failure that signals some more investigative or diagnostic work is required.
c) Reports a “11” if the PLI capacitor was done in an excessive temperature, beyond SSD specified operating temperatures in the product specification.</td></tr>
</table>

The red arrows in Figure 5 point to examples of reading the SMART attributes through the Intel® SSD Toolbox. The report from the Intel SSD Toolbox clearly show power loss, "AEh", result of 8, indicating that 8 unclean shutdown have occurred since the drive was put into service. Additionally, the SSD Toolbox shows the PLI Loss Protection Failure, "AFh"; the last PLI capacitor discharge test result was 678 which exceeds the minimum 25 µS specification; the last PLI test was done 1 minute ago; and the number of tests done was 1. To the system administrator, this is valuable statistical data to determine if any action is needed. As an added capability, the Intel® SSD DC S3700 Series and DC S3500 Series also support the SMART Command Transport (SCT). The SCT allows the administrator to issue commands with which they can manually invoke a new capacitor test, or set the capacitor test intervals as needed.

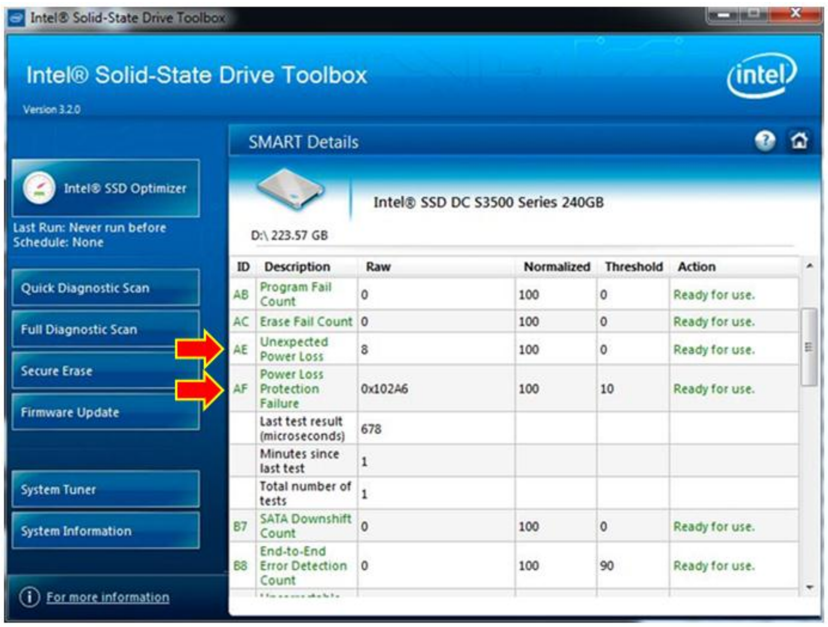

Figure 5: Intel® SSD Toolbox Reporting Sample PLI SMART Attributes on the Intel® SSD DC S3500 Series

The SMART attribute data report from the Intel SSD Toolbox provides power loss event statistics and the PLI technology health log data, which is very useful to the system administrator. With the normalized values from AF SMART Attribute, the system administrator can take necessary actions to replace the drive before data is compromised by a power loss.

Validating Solid-State Drives with PLI

Validating the PLI technology in an SSD involves verifying that the PLI circuitry and its switch timing are working properly; meaning always ready to supply the energy required to complete the write of the SSD controller's temporary buffer data into the NAND Flash in the SSD. Additional power loss validation cases include some rare but plausible scenarios, such as a power loss event during the drive's firmware update, or a secure erase that will not brick (kill) the drive. One of the key validation deliverables of any PLI-enabled SSD is that the drive is stable upon power up, enumerates after 15 seconds, and any commands acknowledged as complete encountered no data loss.

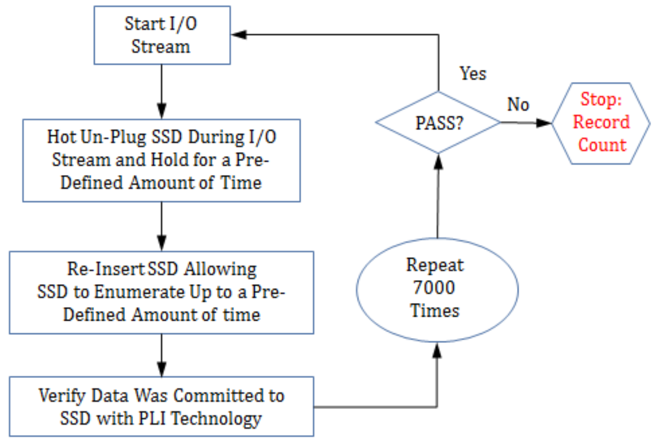

Figure 6 : PLI Validation Process Flow

Performing data integrity checks is absolutely essential but challenging because of the many scenarios of in-flight data. The two power loss write scenarios of typical concern are the shorn writes of unaligned data and unserialized writes. Intel takes these power loss scenarios into consideration when validating the Intel® SSD DC S3700 Series and DC S3500 Series of products.

Power Loss Scenario of Shorn Writes with Unaligned Data

When attempting to complete a file or data write, depending on which operating system was used to deploy the Solid-State Drive, the associated file system may not align its data to certain transfer size(s). Misalignment could cause data fragmentation during transfer across the SATA bus to the SSD. In this fragmented state, the data is unaligned, so if a shorn write were to occur during this stream of data, a data integrity transfer failure could possibly happen. The best way to mitigate and minimize this fragmented state of data transfers is to use an operating system such as the newer Windows* server operating system, which will lay down the file system aligned to 4KB. For users of Linux* OS, there are commands to confirm 4KB alignment.

Validation flows can include the shorn writes with both aligned and unaligned data and compare the success rates. Intel utilizes its own software tool to analyze and track the stream of data transmitted across the SATA bus so when power is removed (by unplugging the drive) during the test, Intel SSDs can record, then verify that upon drive re-insertion, all the right LBA's of data were indeed written.

Power Loss Scenario of Unserialized Writes

Unserialized writes occur in an SSD when the SSD writes cached data in order to optimize write throughput and efficiency. When unserialized writes caching occurs, either at the drive or file system level, this data is vulnerable to being lost, depending on where it is temporarily stored. So, how susceptible are SSDs to unserialized writes during a power loss event? Is it possible to determine if a PLI-enabled SSD is able to solve this particular scenario? Unserialized write failures are difficult to validate because file systems may reorder write sequences. This makes it difficult to validate and predict susceptibilities of unserialized writes. Regardless, because Intel SSDs do not consolidate writes, therefore unserialized writes will not occur.

What Do PLI-Enabled SSDs Do for Me?

PLI technology significantly improves data integrity during power loss events whether from acts of nature, careless mistakes, or even inadvertent hot removal of the SSD storage device. Improved data integrity can help speed the data center recovery after a power loss event with fewer data path checks and validation. Bringing the data center back online as fast as possible will minimize users' delayed access to their data, media, or web materials.

Do I Need PLI in My Storage Device Platforms?

PLI technology offers many advantages by adding higher levels of integrity to data storage environments. Before deciding, it is critical to evaluate the geographical location and its susceptibility to power outages. In addition, evaluate the robustness of the data center regarding power outage risks and mitigating the loss of critical data. Another critical factor in deciding if PLI is the right data center choice is evaluating the workload. If it is more write than read intensive, what is the mix? Is the nature of the data being read or written more random or sequential? If it data is long sequential writes, how long are they? What file system is deployed and does it coalesce the writes? Obviously there is a lot to analyze and consider, including whether the cost of time to rebuild backup data is less than the cost of buying a drive with PLI technology. Down time is lost time, and will impact the total cost of ownership of the entire data center.

Featured Intel® SSD Products with PLI Technology

The Intel® SSD DC S3700 Series and S3500 Series with SATA 3.0 interface contain enabled PLI Technology hardware and the firmware to manage the SSD through the power loss event.

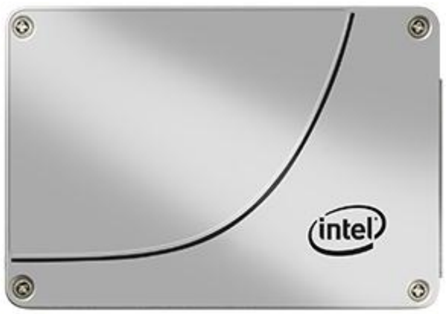

#References:

Sample Pricing: As of February 28, 2014

Nichion Electrolytic Aluminum Wet:
http://www.mouser.com/ProductDetail/Nichicon/UKL1V101KPDANA/?qs=d7wrpMys%2fBuKyOHCoT5fmw%3d%3d

Vishay Tantalum Wet:
http://www.mouser.com/ProductDetail/Vishay/138D107X9035F2/?qs=pEtYqTnH3INJqCnQ3SzcqQ%3d%3d

AVX Niobium Oxide:
http://www.mouser.com/Passive-Components/Capacitors/Niobium-Oxide-Capacitors/_/N-5g8e?P=1z0wrkmZ1z0x8hf