TECHNOLOGY BRIEF
Intel® Solid-State Drive 710 Series
Non-Volatile Memory Storage Solutions from Intel

High Endurance Technology in the Intel® Solid-State Drive 710 Series

Optimized solution for applications with high endurance requirements.

With solid-state drives (SSDs) becoming more prominent in data centers, there is an increased focus on SSD endurance – the ability of an SSD to withstand large amounts of data writes. The Intel® SSD 710 Series introduces a solution for these data center environments and endurance-focused applications: High Endurance Technology (HET). This technology, consisting of silicon-level and system- level optimizations, helps the Intel SSD 710 Series extend its endurance beyond what can be achieved using standard multi-level cell (MLC) NAND.

What is SSD endurance?

SSD endurance is the amount of data that can be written to an SSD during its lifetime. More specifically, it is the ability of an SSD to withstand multiple data writes within a set of criteria defined in the JEDEC *Solid State Drive (SSD) Requirements and Endurance Test Method* (JESD218) specification.

In JESD218, SSD endurance for data center applications is specified as the total amount of host data that can be written to an SSD, guaranteeing no greater than a specified error rate (1E-16) and data retention of no less than three months at 40 °C when the SSD is powered off.

What is High Endurance Technology?

After a large amount of program/erase (P/E) cycles on an SSD, data retention issues on individual NAND flash cells may occur. Additionally, increased read errors may happen when retrieving previously written data – wearing out the life of an SSD.

High Endurance Technology (HET), incorporated into the Intel SSD 710 Series, combines Intel® NAND Flash Memory silicon enhancements and SSD system management techniques to help extend the endurance of the SSD.

SSD endurance is dependent on the capacity of the SSD and the amount of P/E cycles its NAND flash cells can support. When a host issues a write command to an SSD, the SSD data management scheme may consume multiple P/E cycles on the NAND flash cell. The ratio of host writes to NAND writes during this operation is known as *write amplification*. For example, when writing 100 GB to an SSD, NAND flash cells may be written two times, resulting in 200 GB of NAND flash cell writes, which results in a write amplification of 2 (200 GB / 100 GB = 2).

HET tries to extend inherent NAND flash cell P/E cycle capability while at the same time lowering write amplification, thereby increasing overall drive-level endurance.

Why is High Endurance Technology Needed?

Unlike typical client environments where a small amount of writes occur over the lifetime of an SSD, some data center applications have higher write requirements. An MLC NAND-based SSD without HET may be insufficient to meet these endurance needs of data center applications because standard MLC NAND flash cells without HET inherently contain less P/E cycle limits. Typical MLC NAND can fail to meet the data retention and bit error rate criteria specified in JESD218 at greater than its typical rated P/E cycle count.

Based on different program pulses, a page of NAND flash cells is programmed at four different levels: L0, L1, L2 and L3. Figure 1 shows the probability distribution of each cell's threshold voltage level for a read after being programmed. The erased state of the cell is L0. The remaining levels are other programmed states. Read references (R1, R2, R3) are placed between these levels (L0, L1, L2, L3). The gap between each level is referred to as *read margin*.

Figure 1: Read Margin Between Programmed States

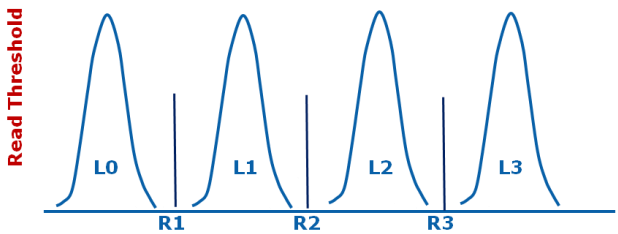

With more P/E cycles, read margin shrinks, resulting in both data retention problems and higher read errors beyond the limit specified by JESD218. This is caused by:

- Loss of charge due to flash cell oxide degradation
- Over-programming caused by erratic program steps
- Programming of adjacent erased cells due to heavy reads or writes in the locality of the NAND flash cell

Figure 2: Read Margins Reduce at High Write Cycles

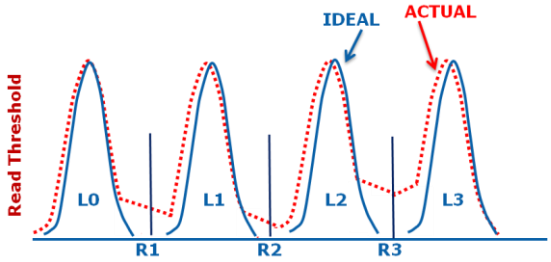

An SSD targeted for client applications may try to overcome these issues using an on-board Error Correction Code (ECC) engine. However, once beyond a certain level of P/E cycles, these SSDs can fail to recover data near the end of life of the SSD. HET uses "beyond ECC" error recovery steps to cross this hurdle of ECC limitation and extend the MLC NAND capability to a higher P/E cycle count.

How is High Endurance Technology Achieved?

To extend the endurance of an SSD, the NAND P/E cycle limit needs to increase while system overhead during writes – which reflects on higher write amplification – needs to decrease. The Intel SSD 710 Series with HET contains NAND components that demonstrate high read margins and improved retention quality. At the SSD system level, special programming sequences are used to mitigate some of the program state disturb issues that may occur.

By using a scheme called *background data refresh*, the SSD moves data around during periods of inactivity to re-allocate areas that have incurred heavy reads. This is in addition to the wear-leveling scheme that already exists in other Intel® SSDs.

NAND characterization and NAND program state disturb management schemes extend P/E cycle capability. Additionally, the SSD comes with extra spare area that lowers write amplification. The combined effect of these items allows the SSD to maintain endurance and retention required by data center applications.

What Does High Endurance Technology Deliver?

Table 1 shows the endurance rating specifications for the Intel SSD 710 Series with HET.

Table 1: Intel SSD 710 Series Endurance Rating Specifications

	100 GB SSD	200 GB SSD	300 GB SSD
4KB Random Write Endurance (up to)	500/900[1] terabytes	1,000/1,500[1] terabytes	1,100/1,500[1] terabytes
8KB Random Write Endurance (up to)	700/2,000[1] terabytes	1,000/1,900[1] terabytes	1,500/3,000[1] terabytes

1. Value based on 20% over-provisioning

Figure 3 shows a plot of normalized read margin (the gap between levels) and terabytes written. To reach the uncorrectable bit error rate target and data retention target, normalized read margin needs to be at least above 0. Figure 3 shows how the Intel SSD 710 Series (with system assist and three months power-off retention) can extend the terabyte written limit. The figure also shows a clear distinction between the Intel® SSD 320 Series and Intel SSD 710 Series.

Figure 3: Normalized Read Margin for Intel® SSD 710 Series

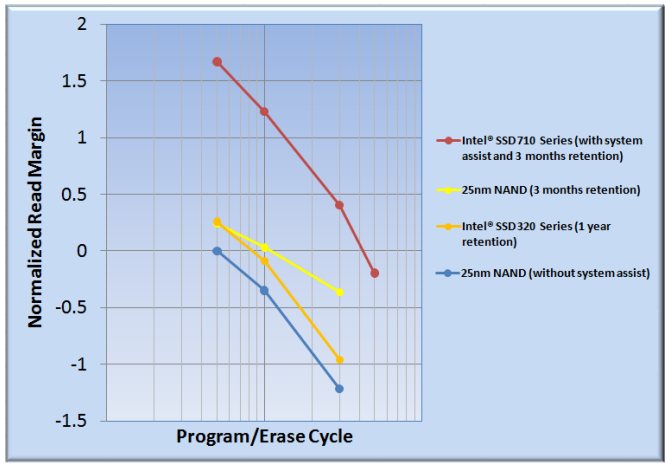

Measuring Estimated Life of Intel® SSD 710 Series

To monitor endurance on the Intel® SSD 710 Series, the SSD supports three SMART attributes: E2, E3, and E4. See the *Intel® Solid-State Drive 710 Series Product Specification* for more information on the SSD and these attributes.

Solid-State Computing Starts with Intel Inside.®
For more information, visit **www.intel.com/go/ssd**

Printed in USA

Please Recycle 325778-001US