Solid State Storage Initiative
Glossary of Terms

September 2011

The Solid State Storage Initiative

About SNIA

The Storage Networking Industry Association (SNIA) is a not-for-profit global organization made up of some 400-member companies and 7,000 individuals spanning virtually the entire storage industry. SNIA's mission is to lead the storage industry worldwide in developing and promoting standards, technologies, and educational services to empower organizations in the management of information. To this end, SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. For additional information, visit the SNIA web site at http://www.snia.org.

Introduction

Welcome to the SNIA Solid State Storage Initiative's (SSSI) glossary of solid state storage (SSS)-related terms. Though solid state storage has been available for decades, within the past three years many new SSS products, solutions, and technologies have entered the marketplace. With this flood of new SSS technologies has come a wealth of new terms. SNIA has for years maintained a dictionary of storage and networking words and phrases already in common use. The SSSI Glossary of Terms complements the SNIA Dictionary by offering definitions for terms directly related to solid state storage which may not yet be commonly known or used outside the SSS industry. As these relatively new SSS terms become more widely used and accepted, they will be submitted to the SNIA Dictionary for inclusion there. The SSSI Glossary of Terms is a dynamic repository; it can easily grow and evolve. The SSSI Governing Board is actively soliciting comments, suggestions, and questions. Please send them to asksssi@snia.org.

ActiveRange

The range of LBAs that can be used for a given Test Code or Preconditioning Code expressed as a percent of the total addressable LBAs. [SSS PTS term]

ActiveRange Amount

The sum of the capacity referenced by the LBAs that are accessed for a given Test or Preconditioning Code equal to, or less than the capacity referenced by the ActiveRange LBAs. [SSS PTS term]

ActiveRange Segment

A collection of contiguous and equal sized LBA ranges within the ActiveRange where the Test Code and Preconditioning Codes are allowed to access. The starting LBA of each ActiveRange Segment is randomly distributed across the entire ActiveRange. The set of ActiveRange segments spans the ActiveRange. Note that ActiveRange segments shall not touch to form a single segment. [SSS PTS term]

ActiveRange Segment Size

The size of a single ActiveRange Segment is determined by taking the ActiveRange Amount and dividing by the number of ActiveRange Segments as prescribed in the Test Code or Preconditioning Code. [SSS PTS term]

Client SSD
An SSD implemented in a single user non-shared environment, such as a personal computer or desktop workstation.

Enterprise SSD
An SSD for which the typical environments are multi-user servers and other shared implementations. Also known as an EFD (Enterprise Flash Drive)

Flash Memory
Flash is non-volatile read/write semiconductor memory which is used in SSS devices (see Solid State Storage). A unique characteristic of Flash is that a memory location (cell) cannot be written unless that cell has not been previously written or has been erased. Obviously this requirement can slow write performance considerably. However, there are techniques utilized by SSS devices (see Garbage Collection and Trim) to mitigate the impact of this constraint.

Another attribute of Flash is that a cell can degrade after a finite number of Program (Write) / Erase (P/E) operations. The Write Endurance of Flash is specified in P/E cycles. Here too, SSS devices have a method (see Wear Leveling) to alleviate this limitation.

There are two types of Flash: NOR and NAND. NOR Flash has long erase/write times, but allows random access to any location. This makes it suitable for storage of program code that needs to be infrequently updated. NAND Flash has faster erase/write times, higher density, lower cost per bit than NOR Flash, and higher endurance. However it allows only sequential access to data. This makes it suitable for SSS devices. All mentions of Flash memory in this document refer to NAND Flash.

Flash stores data bits in cells. Originally Flash was designed to hold one bit per cell, which is known as Single Level Cell (SLC) Flash. Subsequent generations of Flash products were designed to hold two or more bits per cell. This is called Multi Level Cell (MLC) Flash. Of course MLC Flash is higher density memory than SLC Flash, and thus MLC Flash offers a lower cost per bit than MLC. However, MLC Flash has lower endurance than SLC Flash.

Fresh Out of Box (FOB)
The condition of a new/unused solid state storage (SSS) device when first received from the manufacturer. SSS devices that are FOB, or initialized (or Purged) to a near FOB state, exhibit a short period of higher performance which then levels off to a relatively sustained level of performance called Steady State. Typically, the storage cells in an FOB device will have experienced few or no program/erase cycles, except those done at the factory. This device is ready to have data stored with ample pre-erased blocks. [SSS PTS term]

Garbage Collection
Working in the background, Garbage Collection accumulates data blocks previously marked for deletion, performs a whole block erasure on each "garbage" block, and returns the reclaimed space for reuse by subsequent write operations. The controller inside the SSS device performs this bit of housekeeping. Garbage Collection can improve write performance by eliminating the need to perform whole block erasures prior to every write. See also Trim.

Hybrid HDD
A Hard Disk Drive (HDD) that contains some amount of NAND Flash memory which acts as a non-volatile cache, enabling SSS-like access times to data stored in the device. The HDD firmware may "learn" the typical access patterns and optimize performance by placing frequently accessed data, applications, and system files into the Flash cache.

IO Demand
Measured number of OIO executing in the host. [SSS PTS term]

IOPS
The number of Inputs/Outputs per Second (IOPS) provides an indication of the performance of a device in applications generating random reads and/or writes. An SSS device will typically exhibit higher IOPS than HDD because the SSS device contains semiconductor storage technology which does not have the electromechanical latencies of HDD.

Latency
The time between when the workload generator (or host) makes an IO request and when it receives notification of the request's completion

Logical Block Address (LBA)
The address of a logical block, i.e., the offset of the block from the beginning of the logical device that contains it.

MaxUserLBA
The maximum LBA addressable in the User Capacity [SSS PTS term]

Measurement Window
The interval, measured in Rounds, during which test data is collected, bounded by the Round in which the device has been observed to have maintained Steady State for the specified number of Rounds and five Rounds previous. [SSS PTS term]

Nonvolatile Cache
A cache that retains data through power cycles

Outstanding IO (OIO)
The number of IO operations issued by a host (or hosts) awaiting completion. [SSS PTS term]

OIO/Thread
The number of OIO allowed per Thread [SSS PTS term]

Over Provisioning
The practice of incorporating extra Flash storage capacity reserved for controller use. To improve the probability that a write operation arriving from the host has immediate access to a pre-erased block, many Flash SSS products contain extra Flash storage capacity reserved for use by the controller and not visible to the host as available storage. Over provisioning is common because it can help SSS designers mitigate various performance challenges resulting from garbage collection and wear leveling solutions, among other Flash management activities.

Performance Test Specification (SSS PTS)

A SNIA SSS Technical Work Group (SSS TWG)-generated standard that defines a set of device level tests and methodologies to enable comparative testing of solid state storage devices. There are two separate SSS PTS documents:

1. Enterprise SSS PTS, to be used for Enterprise SSDs

2. Client SSS PTS, to be used for Client SSDs. The Client SSS PTS differs from the Enterprise SSS PTS in the preparation of the Device Under Test (DUT) for steady state performance measurement and in the amount and type of test stimulus applied to the DUT. For example, preconditioning LBA ranges may be limited in the Client PTS to less than 100% of the available LBAs while the test stimulus Active Range may be limited to a reduced number of uniquely touched LBAs. The use of limited preconditioning and test active ranges are meant to provide a test stimulus that shares more characteristics with a typical Client user workload.

The SSS PTS documents may be found at http://www.snia.org/pts.

Preconditioning

A technique applied to a Flash-based storage device to achieve Steady State performance. Typically, preconditioning is done prior to performance testing and consists of iterative writes in a known, fixed pattern. In the SSS PTS, Preconditioning is performed as follows:

1. Workload Independent Preconditioning: The first Preconditioning step comprised of a prescribed workload, unrelated to the test workload, as a means to facilitate convergence to Steady State.

2. Workload Dependent Preconditioning: The second Preconditioning step comprised of running the test workload itself, after Workload Independent Preconditioning, as a means to put the device in a Steady State relative to the dependent variable being tested.

Preconditioning Code

The Preconditioning steps defined in the SSS PTS. [SSS PTS term]

Purge

The process of returning an SSS device to a state (see FOB) in which subsequent writes execute, as closely as possible, as if the device had never been used and does not contain any valid data.

Queue Depth (QD)

The depth of a queue in Threads [SSS PTS term] Note that QD is defined as the number of OIO in IOmeter.

Round

A complete pass through all the prescribed test points for any given test [SSS PTS term]

Solid State Card (SSC)

Solid state storage that resides on a printed circuit board and can utilize a standard card form factor such as a PCI card. A Solid State Card would typically use an interface such as PCIe. Compared to a Solid State Module, an SSC normally has a larger physical size, more capacity, and higher performance.

Solid State Drive (SSD)
Solid state storage that may utilize traditional HDD form factors such as 3.5-inch, 2.5-inch or 1.8-inch. Solid State Drives typically use storage interfaces such as SATA, SAS, or Fibre Channel.

Solid State Module (SSM)
Solid state storage that resides in a Dual In-line Memory Module (DIMM) or similar form factor and may use a standard HDD interface such as SATA. A Solid State Module is typically smaller in physical size and has lower capacity and performance than a Solid State Card.

Solid State Storage (SSS)
Any storage capability that is provided by non-moving memory technology rather than moving magnetic or optical media. Solid state storage typically possesses the property of non-volatility and may take various forms such as Solid State Drives, Solid State Cards, or Solid State Modules. Typical interfaces used include SATA, SAS, Fibre Channel, or PCIe. See www.snia.org/forums/sssi/knowledge/standards for more details about SSS interfaces.

Solid State Storage Initiative (SSSI)
An initiative within the Storage Networking Industry Association (SNIA) formed to foster the growth and success of the market for solid state storage in both commercial and consumer environments. (http://www.snia.org/forums/sssi)

Steady State
Operational state of a Flash-based storage device in which the performance is relatively consistent under a prescribed workload. The SSS PTS (see Performance Test Specification) defines Steady State as when, after a Purge and Preconditioning: 1) there has been no more than a 20% variation from the average performance for five consecutive measurements, and 2) there is a maximum of 10% difference between the minimum and maximum measurements on a curve fitted to those five measurements. [SSS PTS term]

Storage Pairing
Low capacity, low cost SSS and high capacity Hard Disk Drive (HDD) in the same system, where SSS provides fast boot and quick access to frequently used files, and HDD provides bulk storage. Some level of file management may be used to optimize performance.

Test Code
The measurement steps defined in the SSS PTS. [SSS PTS term]

Thread
Execution context defined by host operating system/CPU. May be defined as number of processes [SSS PTS term]. Note that IOmeter (www.iometer.org) uses the term Workers for this purpose.

Thread Count (TC)
Number of Threads specified by a test. [SSS PTS term]

Total OIO
Total outstanding IO Operations specified by a test. May be represented formulaically as: (OIO/Thread) * (Thread Count). [SSS PTS term] Not to be confused with OIO as used by IOmeter.

Throughput

A measure of the amount of data that can be transferred from a device (reads) or transferred to a device (writes) within a specified time period, typically measured in MegaBytes per second (MB/s). Throughput is indicative of the performance of a device in an application generating sequential reads or writes.

Trim Command

A method by which the host operating system may inform a NAND Flash-based SSS device about which blocks of data are no longer in use and can be erased. Such blocks may then be written without having to erase them first, enhancing SSS device write performance.

User Capacity

LBA range directly accessible by the file system, operating system and applications, not including over-provisioned capacity

Wear Leveling

A set of algorithms utilized by a Flash controller to distribute writes and erases across Flash cells. The purpose of Wear Leveling is to delay individual cell wear-out and prolong the useful life of the Flash-based storage device.

Workload

A set of operations, such as reading and writing, that is defined so as to exercise a SSS device in a specific manner. The variables defining the workload may include the ratio of reads to writes, and the block size (Access Pattern), and the data pattern.

Workload Dependent Preconditioning (WDPC)

The technique of running the test workload itself, typically after Workload Independent Preconditioning, as a means to put the device in a Steady State relative to the dependent variable being tested. [SSS PTS term]

Workload Independent Preconditioning (WIPC)

The technique of running a prescribed workload, unrelated, except by possible coincidence, to the test workload, as a means to facilitate convergence to Steady State. [SSS PTS term]

Write Amplification

Because a previously written NAND Flash memory location must be erased before it can be re-written, the number of write operations within Flash SSS typically exceeds the number of writes issued by the host. This "write amplification" can be represented in equation form: Write Amplification = (Data Written to Flash) / (Data Written by Host.)

Storage Networking Industry Association
425 Market Street, Suite 1020 • San Francisco, CA 94105 • Phone: 415-402-0006 • Fax: 415-402-0009 • www.snia.org