SSD APPLICATION CLASSIFICATIONS

Application Note

Document #AN0029 – Viking SSD Application Classifications | Rev A

Purpose of this Document

This application note was prepared to help OEM system designers evaluate the performance of Viking solid state drive solutions by using the same benchmarking methodology that Viking performs in it's SSD test facility. The SSD performance stated in the Viking SSD datasheets can be achieved by following the same Viking approach to SSD benchmarking which has been outlined in this document.

Table of Contents

Table of Tables

Table of Figures

No table of figures entries found.

1 Introduction

SSDs can be found in nearly every computer system by the end of 2017, from workstations, desktops and laptops to mobile computing devices and cloud storage. The system designer today, has a myriad of SSD choices, each with varying attributes such as cost per gigabyte, NAND flash endurance, performance characteristics, and various form factors, that optimize their storage options to best fit an application.

The storage industry is quickly evolving and adopting new technologies to satisfy the increasing need for more data storage. These new technologies, however fit some application better than others, and SSD manufacturers offer various types of SSDs that are designed for certain applications.

This white paper was written to help system designers understand the various types of SSD classifications or “classes” so they can make informed decisions based on key SSD attributes that affect feature/benefits. In the storage industry today, the more common application for SSDs fit into the following general market segments: client (end-user), enterprise (business), datacenter (Cloud), industrial/embedded and military/aerospace. But, not all these market segments use an SSD the same way, so getting an understanding about the SSD application and usage model or workload, would be useful in selecting the right feature/benefits in the SSD.

2 SSD Characteristics: Features and Attributes

An understanding of the following key SSD characteristics would be helpful:

- What affects Cost/Gigabyte($/GB) ?
  - Larger capacity SSDs have a lower $/GB since the NAND cost is much much greater the rest of the SSD cost (controller, PCB, case, etc..)
  - Usable capacity is less the total raw capacity because of the amount of over-provisioning needed for spare blocks that will be used up as the SSD wears-out
  - NAND type (The more bits per NAND cell , the lower the NAND cost)
  - Higher controller performance, more firmware features, the amount of DRAM cache, hold-up circuits for power-loss may add more cost

- What affects Performance ?
  - “Fresh Out of the Box” (FOB) performance can be 2x the sustained performance (after the SSD has been totally filled for the 1st time)
  - Performance varies based on workload (random access (ie, frequent access to small files) vs sequential access (i.e. streaming, access to very large files)
  - Performance varies with data characteristics:

  - compressed data (i.e. ZIP, JPG, PDF, MPEG) with random patterns is slower than uncompressed data (ie .doc, .xls, .txt, .dat) with repeatable patterns
  - Performance varies with the type of interface
  - Performance is occasionally interrupted/paused if SSD garbage collection and housekeeping is not implemented during SSD idle mode
  - Performance and latency varies with type of NAND, DRAM cache and que depth.
  - Performance is higher at larger capacities and with utilized capacity and the amount of free space available on the drive, including over-provisioning. (SSD over-provisioning works like HDD short stroking).
  - The SSD interface to the host will determine the absolute maximum performance capability for reading/writing to an SSD (i.e. SATA v2.6 is 1.5/3 gigabits per second and SATA v3.0 is 6 Gbit/sec, etc.)
  - The SSD architecture (i.e. RAID, multiple flash channels, native vs. bridging) and the complexity/efficiency of the controller (data intelligence, garbage collection, recycling, wear leveling, etc.) also have a major impact on performance.
  - Applications that involve frequent transfers of small random amounts of data to the drive will benefit from SSD firmware that is optimized for high IOPS numbers which provide a quick response time to the host (i.e. operating systems, databases, OLTP etc.). Applications that involve moving large sequential amounts of data (i.e. video streaming, data acquisition, data backup/restore, image processing etc.) will benefit from SSD firmware optimized for high throughput (i.e. MB/s).

- What affects SSD life?
  (a.k.a Endurance; unlike HDD, flash-based SSDs wear out)
  - Limited life based on amount of usage, (maximum TeraBytes Written(TBW)
  - Performance throttling Drive Write Per Day (DWPD)
  - Thermal throttling
  - Amount of Over-Provisioning (OP). The more OP, the longer the SSD life, but usable capacity will be lower than total raw capacity.

  Note: DRAM-based SSDs (such as NVDIMMs) do not wear-out, but have higher Cost/Gigabyte.

Understanding these SSD characteristics will lead to a better choice of SSD for it's intended application, since some SSD attributes are more important than others in certain usage classification, as show in the following table and also described below.

Table 2-1: SSD Features/Attribute by Application Class

SSD Features/Attribute	Client (End-user)	Enterprise (Business)	Datacenter (Cloud)	Embedded, Industrial, Telecom	Military, Aerospace
Feature-Rich firmware that is field upgradeable	no	yes	yes	no	no
DRAM Cache	no	yes	yes	no	no
Value cost basis	yes	no	maybe	no	no
Feature and Performance cost basis	no	yes	maybe	yes	yes
Low Latency	no	yes	maybe	yes	yes
SSD Life, Endurance, DWPD, TBW	low	high	low	high	
Performance throttling to increase SSD life	no	yes	no	no	no
Thermal throttling to increase SSD life	no	yes	no	no	no
Amount of Over-Provisioning for optimized endurance	low	high	high	high	high
NAND Flash Type	TLC,QLC	TLC, MLC	TLC, MLC	SLC	SLC
TLC NAND configured as psuedoMLC = 2/3 capacity TLC capacity	no	yes	yes	no	no
TLC NAND configured as psuedoSLC = 1/3 TLC capacity	no	no	no	yes	no
3D Xpoint Flash with PCIe interface	no	yes	yes	no	yes
SSD Preconditioned for consistent performance	no	yes	yes	no	no
Firmware optimized for IOPS (random read/writes)	yes	maybe	maybe	maybe	no
Firmware optimized for MB/sec(sequential read/writes)	no	maybe	yes	maybe	no
SSD firmware optimized for uncompressed data (write amplification and Compression enabled)	no	yes	no	no	no
Power fail (PFAIL) management, data loss protection, supercaps and hold up circuits	no	yes	yes	no	yes
Capacity Goals	low	high	high	low	high
Hot Swap, Hot Plug	no	yes	yes	no	yes
DualPort	no	yes	yes	no	maybe
NVMe Interface	no	yes	yes	no	yes
SATA Interface	yes	yes	yes	yes	yes
SAS Interface	no	yes	no	no	maybe
eMMC Interface	yes	no	no	maybe	maybe
Small Form Factors (M.2 , slim SATA, mSATA, eMMC, USB, CFAST, SD, microSD)	yes	no	no	yes	yes
U.2 form factor (2.5", 3.5", 1.8")	yes	yes	yes	maybe	yes
Power Consumption	low				low
Hardware Encryption, Self Encrypting Drive (SED)	yes	yes	yes	yes	yes
Full Data Path ECC	no	yes	yes	no	maybe
Advanced SSD Specific SMART support	no	yes	yes	yes	yes
Firmware Stored Redundantly (eliminates SSD bricking)	no	yes	yes	yes	yes
Industrial Temperature	no	no	no	yes	yes
Software Write-Protect (initiated at end of SSD life)	no	yes	yes	yes	no

SSD Features/Attribute	Client (End-user)	Enterprise (Business)	Datacenter (Cloud)	Embedded, Industrial, Telecom	Military, Aerospace
Highest UBER possible	no	yes	yes	yes	yes
Military Purge Routines, Sanitization	no	no	no	no	yes
Conformal Coat Sealant	no	no	no	yes	yes
User Initiated Self-Destruct	no	no	no	no	yes

3 Finding the right SSD for an application

3.1 Client-class SSDs for End-Users

Client applications are fairly straightforward because there are many well-known use cases and metrics associated with desktops, laptop/notebooks, mobile workstations, and tablets. In these applications, SSDs are used for storing the operating systems and user data that is generated or downloaded by an individual (end-user). Performance is largely subjective based on the person's needs with the most demanded SSD features being, low cost, instant-on and quick application response time. This means that client SSDs are typically optimized for read speed. Write speed is not as important, if cloud storage is being used or if the computer is NOT media intensive (i.e, movies, large image processing). Client SSD speeds are often limited by outside connectivity (USB, wireless Internet speeds) or human interaction (typing/reading/watching at much slower rate than the SSD is capable of). The SSD is often in idle mode with client applications – enough time for the SSD to take care of any internal housekeeping tasks (garbage collection, wear leveling etc..) that will help the client SSD maintain performance and last longer. The more common interfaces for Client SSD for the year 2017 and beyond, are eMMC, USB or SATA.

3.2 Enterprise-class SSDs for Business-Users

Enterprise class SSDs were originally developed to replace enterprise class hard disk drives (HDD). The interface for these drives was Fibre Channel and SAS, with its dual port modes, DIF, data integrity enhancements and speeds up to 12Gb/s , that offers benefits of higher reliability than SATA, but the more recent trend is PCIe/NVMe for reliable high-speed storage of mission critical enterprise data.

Table 3-1: Enterprise SSD Features

	SATA	SAS	PCIe
Capacity	1TB to 4TB	1TB to 16TB	365 GB to 10.24 TB
Form	2.5", M.2,	2.5", 3.5”	2.5” and PCIe half

Factor	mSATA, SlimSATA		size/PCIe full size
Interface	6Gb/sec	6Gb/sec and 12 Gb/sec	
Read/Write	Up to 1200/750 MB/s	Up to 1200/750 MB/s	Up to 6.7/4.4 GB/s
Read/Write IOPS	Up to 145,000/40,000	Up to 200,000/40,000	Up to 1,300,000/1,240,000

Although NAND flash PCIe SSD are currently mainstream for enterprise class applications, 3D XPoint or Optane flash SSDs are now being introduced in the year 2017 for 2.5" PCIe drive configurations. 3D XPoint or Optane flash SSDs move the speed of SSDs closer to DRAM speeds (7-8 times faster than NAND SSDs) for high speed “hot” data caching while slower NAND flash SSDs or fast HDDs (HDD short stroking) are for “warm” or "cold" data storage. 3D XPoint or Optane flash has a much lower latency (~20usec) than NAND flash (~90 usec), higher IOPS and higher MB/s bandwidth. Enterprise class 3D XPoint or Optane flash SSDs are characterized by full data path protection and consistent IOPS/latency curves regardless of workload and capacity constraints. This makes these SSDs ideal for VMware-based applications servicing hundreds or thousands of users. The more common interfaces for Enterprise SSD for the year 2017 and beyond, are high performance SATA (6Gb/s), SAS (6 and 12Gb/s) and NVMe/PCIe.

3.3 Datacenter-class SSDs for the Cloud

Datacenter SSDs have been designed as the main storage building block for application-specific servers and storage appliances for the Cloud. Internet search, social media and on-line shopping websites have created this application class for SSDs. These SSDs are generally lower cost 6Gbps SATA SSDs as well as PCIe SSDs in high capacities that have a useful life of about 3 years, or the tax-life a fully depreciated/amortized server. Data center SSDs generally feature the read/write speeds of a 6Gb/sec SSD: ~550MB/s and ~50K+ IOPS. Data center SSDs are targeted to applications that require a lower cost per gigabyte while maintaining adequate IOPS with low latencies. Datacenter SSDs need to support a wide range of mixed workloads, both reading and writing of small and large files.

3.4 Industrial/Embedded-class SSDs

SSDs for industrial and embedded systems generally support the manufacturing and communication infrastructures. Some examples of infrastructure applications include routers, switches and base stations for the networking and communications industry; security and monitoring devices for enterprise networks; medical and gaming equipment; factory automation and digital signage.

While Client SSDs are employed for read intensive applications and enterprise SSDs for write-intensive workloads, the datacenter and industrial/embedded SSDs need to support a wide range of mixed workloads. As an example, casino gaming slot machine SSDs might only be written to once and then write protected, but read from as the games are played. A cell phone base station SSD may need to be continuously written with traffic log information and read periodically. So, infrastructure equipment data patterns can span the spectrum from 99% read/ 1 % write to 1% read, 99% write.

Like the enterprise-class SSD's, infrastructure SSD applications are mission-critical and must be designed for 24/7 operation - many times in harsh, extended temperature environments. SSDs for infrastructure-based embedded and industrial systems often use small form factors such as M.2, SlimSATA, mini-SlimSATA, mSATA, CompactFlash, 10-pin eUSB and SD/MicroSD cards with capacities of 100GB or less. Real time OS-based systems and Linux varieties have OS requirements less than 4GB.

Although SLC NAND, is the predominant NAND for industrial/embedded SSDs for mission-critical scenarios requiring high endurance, reliability and long life, there is a higher Cost/Gigabyte with SLC, although psuedoSLC NAND looks like a promising alternative. For high SSD capacity requirements, MLC or TLC may be the preferred device, but there is a higher cost of product re-qualification on NAND flash die revision with lower cost MLC or TLC since these devices may go through several die iterations for every one die iteration of SLC. If the capacity requirement is low, SLC still be the better NAND choice for the SSD, from a total cost of ownership (TCO) and performance point of view.

4 SSD Application Class Definitions

SSD application classes are defined by usage model and the workload associated with an application. These definitions provide a common set of guidelines around which to specify SSDs. Not all SSD suppliers follow these guidelines however, and it is not mandatory. At the moment, the JEDEC JC-64.8 SSD committee defines application classes only for client and enterprise SSDs in document JESD218. The workloads associated with these application classes are explained in JESD219. Viking recommends a review these document to get a better understanding on benchmarking and comparing SSD performance numbers from different vendors. JEDEC definitions are helpful in specifying client and enterprise SSDs, but they don't cover all of the considerations for datacenter, embedded or military SSDs. Therefore, it is important for designers to always look at SSD datasheets to fully understand the assumptions and conditions under which the product performance and endurance numbers were specified. For example,

Endurance workload testing might assume the following:

- Active use (power on) time and temperature
- Retention use (power off) time and temperature
- Functional failure and uncorrectable bit error rate requirements

and

Performance workload testing might assume the following:

- Preconditioning
- Testing with min and max entropy (compressed and uncompressed data)
- Varying que depth, various block sizes, LBA boundary (4G, 8G LBA or Full LBA)
- Use of different workload models (i.e. workstation, server, database)

The table below outlines these elements for the two JEDEC application classes.

Table 4-1: JEDEC Application Classes

Class	Workload	Active Use (Power ON)	Data Retention[1] (Power OFF)	Functional Failure Requirement	UBER Requirement
JEDEC Client	JE5D219 client	40°C (8 hours per day)	30°C, for 1 year	≤3%	< 1 sector in 10^{15} bits read
JEDEC Enterprise	JE5D219 Enterprise	55°C (24 hours per day)	40°C, for 3 months	≤3%	< 1 sector in 10^{16} bits read

NOTE:
1: After endurance requirement has been met

All of these metrics are interrelated when it comes to endurance and changes in assumptions for one parameter can lead to changes in another.

- Workload - consists of the types of data, file sizes, whether that data is sequential or random, and the read and write requirements of the application.
- Active Use - defines the assumed case temperature inside the host system, generally on the SSD case, at which the SSD is written and read. It also defines how often the SSD is used.
- Retention - defines the storage temperature and the length of time the SSD can be powered off while still keeping the data intact after the SSD has reached its endurance specification.
- Data Retention Time - is an important metric point for industrial SSDs. If the SSD has barely been written, the retention time is significantly longer than an SSD that has been in use for a long time.
- Functional Failure Requirement - outlines the number of "acceptable" failures for a given sample size subject to specifically defined conditions.
- UBER - measures the number of sectors that return an Uncorrectable Bit Error Rate based on the number of bits that have been read.

Knowing the use case around which an SSD is specified is vital to understanding its applicability and effectiveness in certain situations. OEMs are cautioned that if SSD specifications do not provide use case data, these specs may be for reference only, and more use case data should be requested from the manufacturer.

5 Comparison of SSD Classes

Table 5-1: Comparison of SSD Application Classes

Definitions	Client/Consumer	Enterprise/Datacenter	Embedded/Industrial
Platforms	PCs, Client, Mobile Devices, Portables, Tablets, etc.	Servers, Storage, Arrays, Cloud, rack-mounted sleds compute and storage	"Fixed Function"; Factory, Auto Compute Systems, Gaming
Performance	Sequential Sensitive	Random Sensitive	
Lifetime	500 TBW	1-5 DWPD	
Capacity	64GB-1TB	0.6TB-8TB	
Endurance(UBER)	10*E-15	10*E-16	
Power Consumption (controller)	<5mW Idle <2.5W Active	5-10W	
Use Case	Mostly Read (80/20), 8hr Duty cycle, 0 to 70.C 1 –3Yr Service Life	Read & Write Intensive,1-5x DWPD, 24/7 Duty Cycle, 0 to 70C, 5Yr Service Life	Wide range of mixed Workloads, 24/7 Duty cycle, -40 to 85 .C, 8+Yr Service Life
Bottom Line	Price & Performance "Low Expectations"	Performance, Capacity, Green & Endurance Levels of Redundancy	Reliability, Endurance, LCM & TCO "Mission Critical"

6 Reference Documents

- Viking whitepaper on performance: AN0006 - SSD Benchmarking Guide
- Viking whitepaper: AN0009 - SSD AES Encryption
- Viking whitepaper: AN0010 – Secure Erase and Military Purge Routines
- Viking whitepaper: AN0011 - Flash Data Retention
- Viking whitepaper: AN0013 - SSD Primer
- Viking whitepaper: AN0025 - SSD Power Fail Protection
- Viking whitepaper: AN0026 - TCG- OPAL
- Viking whitepaper: AN0027 - Wear Leveling
- Viking SSD Product Datasheets
  http://www.vikingtechnology.com/products/ssd/ssd.html
- JEDEC Document JESD218: Application classes for client & enterprise SSDs
- JEDEC Document JESD219: Client & enterprise SSDs workloads

7 About Viking Technology

Viking Technology develops and delivers innovative high-technology products that optimize the value and performance of our customers' applications. Founded in 1989, Viking Technology has been providing Original Equipment Manufacturers (OEMs) with industry leading designs, engineering, product support and customer service for 20 years. For more information visit http://www.vikingtechnology.com.

8 Revision History

9/5/17		Initial release
9/21/17		Add more reference documents. Add more detail throughout. Add TOC

Global Locations				
US Headquarters	Canada Office	Texas Office	India Office	Singapore Office
2950 Red Hill Ave. Costa Mesa, CA 92626 Main: +1 714 913 2200 Fax: +1 714 913 2202	500 March Road Ottawa, ON K2K 0J9 Canada	1201 W. Crosby Road Carrollton, TX 75006 USA	A 3, Phase II, MEPZ- Special Economic Zone NH 45, Tambaram, Chennai-600045 India	No 2 Chai Chee Drive Singapore, 109840
For all of our global locations, visit our website under global locations. For sales information, email us at sales@vikingtechnology.com				

viking TECHNOLOGY

A RF, Optical, Microelectronics and Memory Company