



Samsung Solid State Drive

White Paper



Preface

This content is designed to help facilitate a general to moderate understanding of SSD technology as well as Samsung's industry leadership and unique product offerings. Much of the content requires at least a basic understanding of SSDs. Thus, it may help to review the glossary and familiarize yourself with some fundamental SSD terminology before delving into this material. Please keep in mind that many concepts have been simplified for the sake of understanding.

While each paper may stand on its own, you will benefit from a more clear understanding if you review all of the content in this booklet in the order in which it is presented. In addition to explaining why SSDs are a worthwhile upgrade from their traditional mechanical counterparts, you will find an introduction to NAND technology (the fundamental building block of SSDs), a discussion of various SSD-specific technologies, an overview of how to understand SSD benchmarks, and an outline of the most important Magician software features to help keep your SSD in top shape.

This booklet is the product of many months of effort and collaboration. The Brand Product Marketing Team sincerely hopes that you find this content both useful and informative.


01 | Why SSDs Are Awesome

An SSD Primer

Solid State Drive (SSD) technology has been receiving a lot of attention recently – and for good reason. It represents the first major advancement in primary storage technology in years. Upgrading to an SSD does not provide minor performance enhancements. Rather, an SSD will transform your computing experience with its revolutionary random access speeds, its multi-tasking prowess, and its outstanding durability and reliability. Let’s explore in more detail some of the things that make an SSD so special – and why it could be the most important computer upgrade you ever make.

Move Over, Mechanical Drives.

A Hard Disk Drive (HDD) works by way of a mechanical drive head that must physically move to access locations on a rapidly-spinning magnetic disk. If you’ve ever seen the inside of an HDD, it looks a bit like a high-tech record player. A Solid State Drive (SSD), on the other hand, has no moving parts and is therefore capable of accessing any location on the drive with equally fast speed and precision. This distinction is behind each and every benefit that will be discussed below.



Access Time and Latency

One of the reasons SSDs so dramatically outperform their traditional Hard Disk Drive (HDD) counterparts is because of their incredibly fast data access times (the delay between the request and completion of a read or write operation). A modern SSD boasts random data access times of 0.1ms or less, whereas mainstream 2.5” consumer HDDs take about 10~12ms or more. As you can see in the below comparison table, SSDs are over 100 times faster than HDDs in accessing data, including data seek time and latency. Moreover, while HDDs are deeply affected by fragmentation, SSDs deliver consistent performance thanks to their use of integrated circuits instead of physical spinning platters – while rotating disks must wait for spindles motors, heads, and arms to physically locate data locations, SSDs can access any location with lightning speed.

840 SSD (500GB)	Category	2.5” SATA HDD (500GB, 7200rpm)	Difference
NAND FLASH	Media	Magnetic Platters	
540 / 330	Seq. R/W Speed (MB/s)	60 / 160 (*140 / 70)	x 3~8 / 2~5
98,000 / 70,000	Ran R/W Speed (IOPS)	450/400	217 / 175
0.1	Data Access Time (ms)	10~12	x100~120
78,700	Benchmark Score (PCMark Vantage)	5,600	x 14
0.127W (active)	Power Consumption (Operation)	1.75W	x 13 ↓
0.046W (idle)	Idle Power	0.8W	x 17 ↓
20G (10~2000Hz)	Vibration	0.5G (22~350Hz)	x 40
1500G/0.5ms	Shock (Operation)	350G/2.0ms	x 4
1.5M hours	Reliability (MTBF*)	700k hours	x 2

*HDD performance may vary by brand and model. The above data is for explanatory purposes only.

Let's continue with an analogy. Imagine you are making a cake and you need to collect all of the ingredients. You have a big house – and a kitchen to match. The flour is in the cabinet on the other side of the room, the eggs and milk are in the refrigerator, but you forgot the sugar downstairs in your car. It can take a considerable amount of time to gather the things you need, since you have to physically walk to the location of each ingredient and bring it to your prep area. Now imagine that, through some miracle, all of the things you need are instantly accessible without moving from the spot you are standing. How much time would that save? This simplified analogy helps to explain one of the primary advantages SSDs enjoy over HDDs. There is no need to wait for a mechanical drive head to swing into place, just as there is no need for you to move from your prep area in the analogy above. The result? Your PC just became a dramatically more responsive companion.

Furthermore, an SSD benefits from massive latency advantages. Think about how quickly you can get up and moving in the morning with and without a cup of coffee. An HDD's speed profile will be more akin to your pre-coffee morning self, whereas an SSD would represent how you feel after a couple shots of espresso. This measurement is important because how long it takes to get ready and actually locate a piece of data directly affects how fluid the user experience is.


Performance

Multi-taskers and speed demons rejoice. The incredible access times just discussed contribute directly to the performance of a machine under typical to heavy loads. The ability to access any drive location without a performance penalty means you can have more applications open at the same time with less lag. SSD users also experience dramatic improvements in boot time, shutdown, application loads, web browsing, application installations, and file copies. The copy dialogue box and the spinning hourglass will be gracing you with their presence much less frequently and for much shorter periods of time. You'll also be able to open and close your laptop without the frustrating delays, allowing you to experience the “sleep” and “wake” features the way they were meant to work – without a “groggy” state in between.

The performance benefits just discussed are a result of the sequential and random performance capabilities of an SSD. Fast sequential speeds allow for quick file copies and smoother performance when working with large files, like videos. However, it is random performance, measured in Input/Output Operations Per Second (IOPS) that is, perhaps, the most important performance metric for SSDs. A large portion of storage activity is made up of 4K random writes, a metric that measures how well a drive will perform when writing small chunks of random data (e.g. changing a small piece of a Word or text file and then saving the changes). Users spend a majority of their time not copying large files or installing applications, but multitasking (e.g. email, web-surfing, listening to music, etc.) and working with various work and media files - tasks influenced by IOPS. An SSD can offer up to a 200x improvement in IOPS over a traditional HDD (results may vary based on HDD model). For this reason, Samsung put a heavy focus on random performance when designing its SSD lineup, offering users industry leading Random Performance of up to 100,000 IOPS. This is performance for the real world; performance you will notice and appreciate every day; performance that represents a dramatic leap forward in usability.

Performance Sustainability

Moreover, SSDs will deliver consistent performance regardless of how much data they contain. HDDs, on the other hand, can suffer a performance decrease of over 50% when full or containing a lot of fragmented data. This is the result of the physicality of the disks, which benefit from being able to sequentially write to the space along the outer area of the platter (where there is more surface area to cover) in the early days of a drive's life. Over time, as the outer sectors fill with data, the drive must write to progressively smaller sectors, which naturally store less data. Thus, additional movement is required to switch to the next available sector if more space is required. Additionally, data becomes fragmented with extended use, forcing the mechanical drive head to jump among inner and outer sections of the platter, negatively affecting performance even further.



* If an HDD is rated for maximum sequential performance of 160MB/s R/W, such performance is guaranteed only for the initial, clean state, after which performance may decline to around 70-80MB/s

Durability

How resistant do you think a record player would be to a few healthy blows? Would the arm and needle be able to withstand any significant force? If you're too young to remember what a record player is, think about

a CD player. What happens when it is bumped around too much? Now think about your digital MP3 player. Comparing those devices, which would you trust most to take on a road trip? Any time the complexity of mechanical parts is involved, there is an increased risk of damage to one of the components. Some things wear out naturally with time, and other things simply cannot take any significant abuse. An SSD can withstand more than 40 times the vibration (20G) than an HDD (0.5G) and up to 1500G of shock (compared to less than 350Gs for a typical HDD). That means your precious data is much more likely to survive a few bumps and bruises as you go about your life. It also means you can watch a video on a bumpy train or plane ride without worrying about skipping frames. Your inner road warrior can celebrate - your laptop just got a lot more travel friendly.

Reliability

Remember cassette tapes? If not, consider yourself lucky. Compared to their digital counterparts (the CD), they were drastically more prone to failure. Their physical tape mechanism was notorious for getting stuck, tangled, and breaking; and the process of finding the song you wanted to listen to was tedious at best. Enter the CD. This digital successor to the cassette represented an advance not only in reliability (no more fragile tape to get caught or broken), but in usability. Much the same can be said about the move from HDDs to SSDs. No mechanical components means fewer parts to wear out and fail. While Samsung SSDs enjoy a Mean Time Between Failures (MTBF) of up to 1.5 million hours for consumer drives, high-quality consumer HDDs offer a MTBF of only 0.5-0.7 million hours. Samsung's SSDs, in particular, enjoy one of the lowest Annual Failure Rates (AFR) in the industry thanks to rigorous quality control and a fully automated advanced bin sorting processes. Samsung's bin sorting process singles out only the highest quality NAND wafers and components for use in SSD products, which require the utmost quality in order to safeguard precious user data. In fact, Samsung's quality assurance procedures have been acknowledged by OEM customers worldwide, each of which must run their own qualification processes before implementing a drive into consumer products.

Efficiency

Imagine if the only way you could get anything done was if you were running at full speed frantically flailing your arms about. Such is the exhausting life of an HDD. When an HDD wants to do any work, it must spin up its magnetic platter to full velocity, consuming a considerable amount of energy in the process. On top of that, the mechanical read head must swing into place continuously to locate the correct data. An SSD, free from the constraints of mechanical hardware, can access any data it wishes while consuming only minimal power. In fact, Samsung's 840 Series SSD draws an average of only 0.127 watts of power in operation, whereas an HDD requires an average of 1.75 watts. Furthermore, because SSDs complete their tasks with such incredible speed, they spend more of their time in idle mode than an HDD (at which an 840 Series consumes only 0.046W compared to an HDD at 0.8W), which translates into noticeable improvements in battery life and means you can spend less time looking for a power outlet when you take your digital life mobile.

Bringing it all together

Upgrading to an SSD will make your computer dramatically faster, more durable, and significantly more reliable. In fact, no other single PC upgrade can contribute so many positive effects to user experience. Free your computer from the limitations of its crusty, old, mechanical storage, and free yourself to be a happier, more productive PC user.

Why Samsung

Samsung has been the number one supplier of SSDs in the preinstalled storage business for more than 6 years. No one has more experience with SSDs or a better track record for reliability. You can trust that Samsung will continue to lead the way in SSD technology, just as you can trust that a Samsung SSD will breathe new life into your PC.

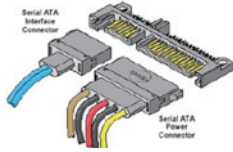


02

Understanding SSD System Requirements

SATA Interface Basics

SATA is the name of the interface standard that allows a storage device, in this case an SSD, to communicate with a host system, in this case a PC or laptop. It facilitates the functioning of all storage features, from basic input/output (or read/write) requests to advanced AHCI-enabled features like Native Command Queuing (NCQ), hot-swapping, power management, and RAID. SATA has gone through several revisions in its lifetime, boasting significant performance gains with each iteration. The SATA interface generation a PC is equipped with will have a direct impact on maximum SSD performance. For this reason, it is crucial to understand the difference between each SATA revision and the capabilities each will afford a modern SSD.

SATA Revisions

Generation	Speed	Special Features	Photograph
SATA 1.0	1.5Gb (150MB/s)		
SATA 2.0	3Gb (300MB/s)	Native Command Queuing (NCQ) support	
SATA 3.0	6Gb (600MB/s)	New NCQ Management feature, Improved power management, support for new connectors to accommodate new device form factors, queued TRIM Command	
SATA 3.1	6Gb (600MB/s)	mSATA support	
SATA 3.2	8Gb/s & 16Gb/s (800MB/s & 1600MB/s)	SATA Express support (not finalized yet)	

SATA 1.0

Released in 2003, SATA I, also known as SATA 1.5Gb/s, is the first generation of the SATA specification. As its name implies, it is able to communicate at up to 1.5Gb/s. This early revision does not support some modern storage features, such as Native Command Queuing. Maximum transfer rates are around 150MB/s, slightly better than the older PATA interface (also called UDMA 133) it replaces, which had a maximum speed of 133MB/s.

SATA 2.0

SATA 2.0, also known as SATA 3Gb/s, is the second generation of the SATA specification. Released in 2004, it is capable of communicating at up to 300MB/s. This revision also introduces Native Command Queuing (NCQ), a feature that improves multitasking performance. SATA 2.0 is more than capable of handling the best hard drives on the market, but it is easily overwhelmed by the onslaught of Flash devices (SSDs) that are now popular. Computers purchased in 2005 or later are likely equipped with SATA 2.0 support.

SATA 3.0

Released in 2009, SATA 3.0, also known as SATA 6Gb/s, is the third generation of the SATA specification. It is capable of communicating at up to 600MB/s, with overhead taken into account. In addition to its speed improvements, this revision also introduces NCQ management, improved power management features, and queued TRIM support (allowing TRIM commands to be queued with I/O requests, which was not possible on earlier implementations). While it is currently capable of supporting the massive speeds of today's SSDs, it is already being outpaced as SSD technology advances. This is the standard that most modern SSDs are built to support, although they are backwards compatible with the earlier standards as well.

mSATA

Building upon the foundations of SATA 3.1, mSATA was designed to address the rising trend of continually shrinking Notebook computers. Smaller profile laptops require smaller SSDs, which in turn require smaller connectors. Notebooks that feature the mSATA interface are becoming more popular but still represent a relatively small portion of the overall market. Maximum transfer speeds are equivalent to the standard SATA 3.0 interface.

SATA Express

SATA Express, once finalized in 2013, will be capable of speeds faster than SATA 6Gb/s thanks to its use of the PCI Express (PCIe) interface. Interface speeds may be increased to 8Gb/s or 16Gb/s (or perhaps even higher speeds later). Future motherboards will offer slots for both SATA Express and traditional SATA devices, and the SATA Express interface will be backward compatible with the standard SATA 3.0 interface.

How SATA Affects SSD Performance

SATA Version

Most SSDs available on the market today are designed for SATA 6Gb/s (SATA 3.x). Many consumers still have older machines, however, which do not have support for the latest SATA revision. Fortunately, SSDs are backwards compatible with older SATA versions. They will, however, be limited by the maximum bandwidth supported by the host machine (e.g. a SATA 6Gb/s SSD connected to a SATA 3Gb/s computer will not be able to reach transfer speeds more than 300MB/s, even though the drive might be rated for performance well over 500MB/s). This is why it is important to understand the capabilities of existing PC hardware before upgrading to an SSD – to avoid any potential disappointment if speeds do not match advertised rates.

Samsung knows that many people still use older computers with SATA 2.0 at home or work, and these people can still greatly benefit from an SSD upgrade. During its design process, Samsung optimized the 840 Series SSD for maximum performance on all SATA interfaces, current and legacy. An 840 Series SSD will outperform both an 830 Series (SATA 6Gb/s) and a 470 Series (SATA 3Gb/s) SSD connected to the same SATA 3Gb/s port, as shown in the table below.

SATA 3Gb/s Performance Improvements (470 Series as Baseline)

	Sequential		Random (4KB, Q32)	
	Read	Write	Read	Write
470 (256GB)	100%	100%	100%	100%
830 (256GB)	+4%	+0%	+52%	+95%
840 (250GB)	+8%	-2%	+68%	+179%
840 PRO (256GB)	+8%	+6%	+68%	+179%

* Based on internal test results. Sequential Performance results via CrystalDiskMark. Random Performance results via Iometer 2010. System Configuration (SATA 3Gbps & 6Gb/s): Win7 Ultimate x64 SP1; Intel Core i7-3770 @3.4Ghz; 4GB DDR3, Intel 7-Series Chipset Z77

In fact, an 840 Series SSD will outperform any SATA 3Gb/s SSD on the market on the same system setup. An investment in an 840 Series SSD is, therefore, future-proof in that an upgrade now will benefit a SATA 3Gb/s system, but the SSD will also be ready to offer even better performance in the event that one decides to upgrade to a SATA 6Gb/s system in the future. Drives designed specifically for an older SATA revision are limited by their aging hardware and will be unable to saturate the available bandwidth on any SATA interface.

Speaking of computer upgrades, SATA expansion cards are widely available in the market. These cards promise to add and/or upgrade SATA ports on an existing system. They are limited, however, by a certain amount of overhead and will often not be capable of providing a modern SSD with the throughput it needs to reach maximum performance. Thus, Samsung recommends connecting your 840 Series SSD to a native SATA 3.0 (6Gb/s) port to enjoy its full potential.

Locating a native SATA 3.0 (6Gbp/s) port is relatively simple, but it requires some attention to detail. Most motherboards include both SATA 3.0 and SATA 2.0 ports, and they are often located side-by-side. While they are usually color-coded, there is no industry standard defining which color represents which SATA revision. Hence, it is important to carefully read the labeling to determine which is which.



Maximize SATA Capabilities with AHCI

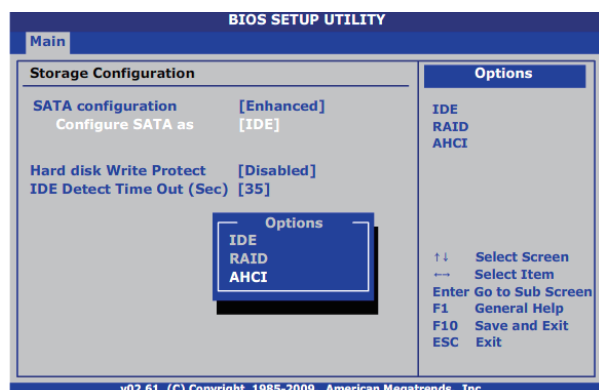
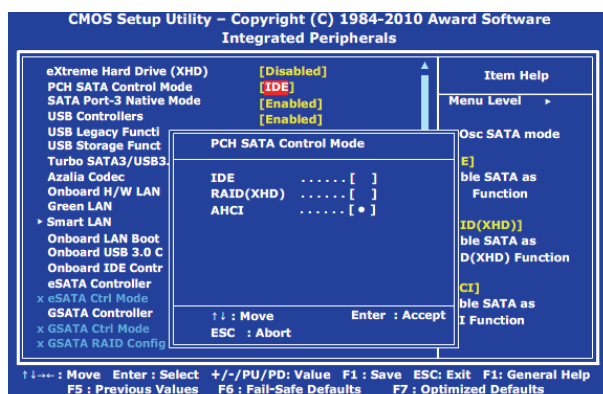
AHCI, Advanced Host Controller Interface, is an open interface championed by Intel to allow the use of advanced SATA features (e.g. NCQ, hot plugging, power management). Basically, it defines a standard method for storage devices from various vendors to communicate with the host system, allowing software engineers to take advantage of specialized functionality. In order to enjoy the full performance of your Samsung 840 or 840 PRO Series SSD, AHCI mode must be enabled on your system through the BIOS.

If AHCI is not properly configured on your system, the Random Read/Write performance of your SSD will be limited to a Queue Depth of 1 (QD1), severely limiting the performance improvements you will notice over a conventional Hard Disk Drive (HDD) while multi-tasking (by 80-90%). Since the latest versions of Windows, including Windows Vista and Windows 7, include standard AHCI drivers, there is no need to install AHCI drivers manually. Some older systems, however, may not support AHCI even if they are equipped with a SATA 3 (6Gb/s) interface. Samsung's Magician software can help you determine whether or not AHCI is supported and/or enabled on your PC.

How do I enable AHCI?

Most modern PC systems support AHCI mode and enable it by default. If it is not enabled on your system for some reason, however, you can manually enable it via your computer's BIOS. The procedure will vary slightly depending on the BIOS version on each system, but the general steps are the same: Restart the computer and press the appropriate key to enter the BIOS (this is usually DELETE, F2, or F10); Find the section for "Integrated Peripherals" or "Storage configuration;" Change "Configure SATA as" or "PCH SATA Control Mode" to "AHCI." Consult your PC's user manual for instructions for your specific machine.

Ideally, AHCI mode should be enabled via the BIOS BEFORE installing the Operating System (OS). On newer systems, this is generally the case. If it is not, and AHCI is enabled via the BIOS for a Windows Installation that was not originally configured for AHCI, it is not uncommon for stability problems to arise. Windows XP does not include native AHCI drivers, making it more complicated to enable AHCI mode because it is the user's responsibility to locate and install the correct AHCI drivers for the OS. In general, PC chipset vendors (e.g. Intel, AMD, nVIDIA) maintain their own AHCI drivers for Windows XP users. It is best to consult the website of your motherboard vendor or PC manufacturer to be sure.



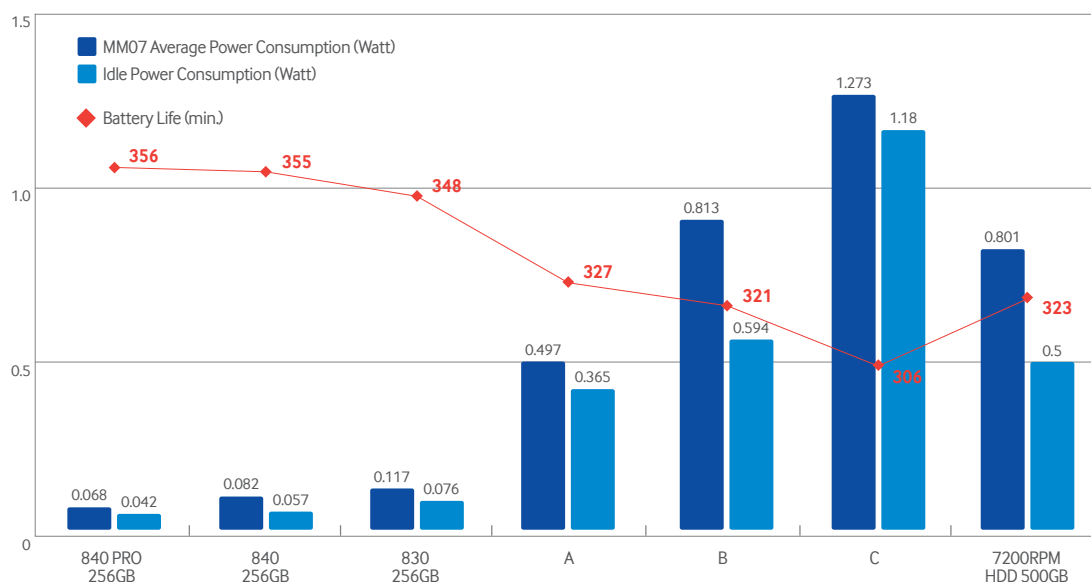
What specific advantages does AHCI provide to Samsung SSDs?

Significantly Lower Power Consumption

AHCI mode is required for Samsung's SSDs to demonstrate their superior power consumption. This is because, without AHCI, the SSD cannot communicate with the host system to take advantage of advanced power management features like Host-Initiated Power Management (HIPM) and Device-Initiated Power Management (DIPM). Maintaining the physical SATA interface (called the PHY) consumes a considerable amount of power, and HIPM and DIPM allow for the PHY to be put into a reduced power mode. After some predetermined period of inactivity, either the host or the device can signal the PHY to enter its reduced power state. Samsung's 840 and 840 Pro Series SSDs make aggressive use of DIPM to put

the PHY into reduced power consumption mode, contributing to the drive's reputation for having the industry's lowest idle power consumption.

Without HIPM or DIPM, power consumption will increase dramatically, as shown in the graph below via drive "C," which does not use these power management features. In general, such power management features are enabled by default on laptop systems. Desktops, which are free from battery constraints, generally operate with these features disabled – thus increasing drive power consumption in favor of slight performance advantages.



NCQ

AHCI also affects the ability of Samsung's SSDs to reach their maximum Random QD32 performance. This is because AHCI is required for Native Command Queuing (NCQ) to function. Because SSDs enjoy extremely fast speeds, storage is no longer the bottleneck of a PC system. NCQ helps an SSD deal with situations where it finds latency on the host, a phenomenon unheard of with traditional HDDs. Additionally, NCQ allows the SSD controller to process commands in parallel, improving performance and reducing latency.

Hot plugging

AHCI is also required in order to enable "hot plugging," a kind of plug-and-play functionality that allows an SSD to be plugged into or removed from a system without the need to shut the computer down. This is convenient for multi-drive systems where an SSD is in use as a secondary drive. Rather than having to quit all running programs and shut down the computer, a new secondary SSD can be inserted to complement or replace an existing drive without adding unnecessary steps to the workflow.

What if my system doesn't support AHCI?

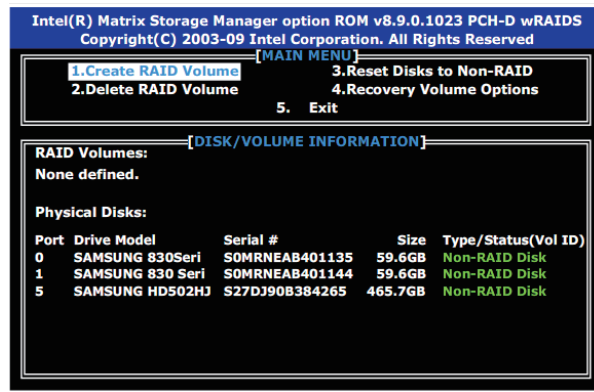
If you have a system that does not support AHCI, or if you fail to enable it properly, your system will run in "IDE emulation" mode, a kind of compatibility mode that mimics the functionality of the older "Integrated Drive Electronics" ATA interface. While your SSD will not display its full performance potential or be able to take advantage of all of the latest SATA functionality, it will still be able to function on your system.

AHCI Wrap-up

Enabling AHCI is crucial to get the most performance out of your Samsung SSD. If you notice that performance is lower than expected, one of the first things to check is whether or not AHCI is properly configured. Samsung's latest Magician 4.0 software can assist in determining if AHCI is supported and/or enabled. Once set up, AHCI will enable your Samsung SSD to achieve superior performance, lower power consumption, and improved convenience.

Enhancing SATA Performance with RAID

Today's SSDs are maximizing the performance potential of the current SATA 3.0 interface generation. One way to increase performance beyond standard single drive SATA configurations is to pair 2 or more drives together using SATA's native RAID support.

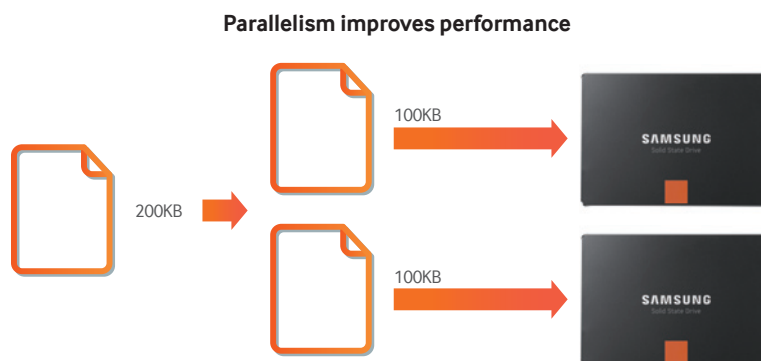


RAID, which stands for Redundant Array of Independent/Inexpensive Disks, is a type of storage system in which a number of drives (at least 2) are combined into one logical unit. RAID is used to improve performance, improve reliability, or some combination of these two. Data can be distributed among the drives in a RAID array in one of several ways (called RAID levels). The most common RAID levels are RAID 0 and RAID 1. With the introduction of its 7 Series Chipsets and the latest Intel Rapid Storage Technology (IRST) drivers (11.0 or later), Intel is now fully supporting SSD technology, including the TRIM maintenance command, in RAID 0 arrays. In the past, the lack of TRIM for RAID 0 was a source of frustration, as the performance improvements initially gained through the RAID array were mitigated by the performance deficits caused by the lack of TRIM. Thus, with the addition of TRIM support for RAID 0, it is useful to understand RAID technology and who (and why) an individual might choose to use it.

RAID Levels

RAID 0


RAID 0, which requires a minimum of two drives and whose primary purpose is speed, divides files into chunks, or stripes, which are split among the available drives and written in parallel. By doing this, it is also possible to read smaller sections of the original file in parallel. This parallelism is what allows for the drastic performance improvements that RAID 0 offers. If one drive fails, however, all data (having been broken into chunks) is destroyed; and the likelihood of failure increases as more drives are added to the array.



RAID 1

RAID 1, which also requires a minimum of two drives, writes data identically to multiple drives and offers performance

similar to a single SSD. Its purpose is redundancy. As such, the array will continue to operate as long as at least one drive is still functioning. Capacity is limited to the size of the smallest drive in the array. The primary goal is to reduce downtime in the event of a hardware failure. However, this solution, despite its redundancy, is not a replacement for a good backup regimen, as it cannot protect from data corruption or security breaches.



Who should use RAID?

Any desktop user can implement RAID on their system with relative ease, but only a certain subset of users will be able to truly justify the expense and benefit from what it has to offer. As outlined above, each implementation has a different purpose.

RAID 0, thanks to its incredibly fast read and write speeds, is ideal for an individual who must work with large files that need to be modified or written quickly, such as a graphic designer who spends a lot of time in Adobe Photoshop or a videographer who uses Avid to edit digital video files. Some gamers also enjoy the performance advantages that RAID 0 provides. Because it lacks redundancy, however, RAID 0 is not suited for mission-critical tasks and should be supplemented by a robust backup system.

RAID 1, which writes all data identically to two or more drives, is ideal for mission-critical workloads, such as small servers, that cannot experience significant downtime. The speed is the same as it would be with a single disk, so there is no performance advantage or disadvantage. The sole benefit is its redundancy, which will make it easier to get up and running in the event that one SSD fails. As mentioned above, RAID 1, despite its redundancy, is not an alternative to a solid backup. It cannot protect against user error (e.g. inadvertently deleting files) or data corruption.

Why Samsung

Samsung, as the largest supplier of SSDs to the preinstalled market, must qualify its SSDs on an endless variety of hardware configurations, including all current and legacy SATA interfaces. Its SSDs are designed to outperform the competition on all SATA implementations, making a Samsung SSD the smartest choice for any PC upgrade, regardless of whether it is an aging or a state-of-the-art system.

03


NAND Basics

Understanding the Technology Behind Your SSD

Although it may all look the same, all NAND is not created equal: SLC, 2-bit MLC, 3-bit MLC (also called TLC), synchronous, asynchronous, ONFI 1.0, ONFI 2.0, Toggle 1.0, Toggle 2.0. To the uninitiated, all of this looks like nonsense. As the SSD market gets more complex, however, so does your buying decision. This paper will help to lay down some of the differences between various types of NAND technology, with the ultimate goal of making you a savvy SSD shopper.

A Brief Introduction to NAND Flash

NAND Flash memory stores data in an array of memory cells made from floating-gate transistors. Insulated by an oxide layer are two gates, the Control Gate (CG, top) and the Floating Gate (FG, bottom). Electrons flow freely between the CG and the Channel (see diagram to the right) when a voltage is applied to either entity, attracted in the direction to which the voltage is applied. To program a cell, a voltage is applied at the CG, attracting electrons upwards. The floating gate, which is electrically isolated by an insulating layer, traps electrons as they pass through on their way to the CG. They can remain there for up to years at a time under normal operating conditions. To erase a cell, a voltage is applied at the opposite side (the Channel) while the CG is grounded, attracting electrons away from the floating gate and into the Channel.



To check the status of a cell, a high voltage is applied to the CG. If the floating gate holds a charge (electrons are trapped there), the threshold voltage of the cell is altered, affecting the signal emanating from the CG as it travels through to the Channel. The precise amount of current required to complete the circuit determines the state of the cell. All of this electrical activity effectively wears out the physical structure of the cell over time. Thus, each cell has a finite lifetime, measured in terms of Program/Erase (P/E) cycles and affected by both process geometry (manufacturing technique) and the number of bits stored in each cell. The complexity of NAND storage necessitates some extra management processes, including bad block management, wear leveling, garbage collection (GC), and Error Correcting Code (ECC), all of which is managed by the device firmware through the SSD controller.

SLC vs. 2-bit MLC vs. 3-bit MLC NAND

NAND technology has been naturally progressing with the needs and expertise of the industry. In the simplest terms, the data stored in NAND flash is represented by electrical charges that are stored in each NAND cell. The difference between Single-Level Cell (SLC) and Multi-Level Cell (MLC) NAND is in how many bits each NAND cell can store at one time. SLC NAND stores only 1 bit of data per cell. As their names imply, 2-bit MLC NAND stores 2 bits of data per cell and 3-bit MLC NAND stores 3 bits of data per cell.

Advantages of MLC NAND

The more bits a cell stores at one time, the more capacity that fits in one place, thus reducing manufacturing costs and increasing NAND manufacturing capacity – a phenomenon called “bit growth.” This phenomenon has allowed NAND technology to penetrate a continually greater number of usage applications at increasingly higher capacities over the years.

NAND technology’s first home was in external storage devices (e.g. USB memory devices) at very modest capacities. As the technology matured, NAND found applications in a host of digital devices, including digital cameras, MP3 players, and mobile phones. Having proven itself a strong and durable performer, the technology made its way into consumer and finally enterprise solid state storage devices (SSDs). NAND’s rise in popularity and usefulness was directly a result of the

advances in semiconductor technology that allowed engineers to squeeze more bits into each cell. Capacities ballooned from megabytes (MBs) to gigabytes (GBs) as manufacturers were able to produce the same NAND bit capacity with less capital investment. A self-reinforcing cycle of increasing demand and decreasing pricing helped manufacturers to continue to meet industry needs without increasing supply costs – benefiting both consumers and device makers.

Cost vs. Capacity

	SLC	2-bit MLC	3-bit MLC
Price/GB	\$\$\$\$	\$\$	\$
External Storage	Mid 90's~	Mid 2000s~	'09~
USB Flash Drive*	4MB~	1GB~	2GB~
Memory card*	2MB~	1GB~	?
Digital Devices	Mid 90's~	Mid 2000's	'11~
Digital Camera	4MB~	128MB~	
MP3 Player	16MB~	512MB~	2GB~
Client SSD (2.5")	'06~	'08~	'12~
	32GB~	64GB~	120GB~
Enterprise SSD	'08~	'12~	'13~
	50GB~	128GB	?

* Less than 5% of Memory Cards and USB Flash Drives use SLC Memory as of 2012

Limitations of MLC NAND

Of course, adding more bits to each cell makes it more difficult to distinguish between states, reducing reliability, endurance, and performance. Indeed, determining whether a container is either full or empty (SLC) is much simpler than determining whether it is one quarter full, one half full, three-quarters full, or entirely full (MLC). This is why it can take up to 4 times as long to write and up to 2.5 times as long to read 3-bit MLC NAND than its SLC predecessor.

Another side effect of storing more bits per cell is an increase in the rate at which the NAND cells degrade. The state of a NAND cell is determined by the number of electrons present on the floating gate. The Oxide layers that trap electrons on the floating gate wear out with every program and erase operation. As they wear out, electrons become trapped, which affects the overall electrical properties of the cell and, consequentially, subsequent program and erase operations. With the oxide weakened, charges sometimes leak from the floating gate. While this is not a huge problem with SLC NAND because there are only two states to distinguish between, it can be a huge problem for 3-bit MLC NAND because there are 8 states to distinguish and very little room for differentiation – just a few electrons can make the difference between one state or another. Compounding matters is the fact that the oxide layer gets smaller with each advance in process geometry – as we shrink the size of the cell, the oxide becomes thinner. A thinner oxide layer will wear out faster, meaning the cell will have a lower lifespan.

MLC NAND Today

Manufacturers have made massive strides in managing the limitations of NAND technology. Sophisticated “bin sorting” algorithms allow 1st-tier manufacturers like Samsung to select only the highest-quality NAND chips for use in SSD devices. Advanced wear-leveling code ensures that NAND cells wear out evenly (to prevent early drive failure and maintain consistent performance), while garbage collection algorithms preemptively prepare fresh storage space (to improve write performance). Improved Error-Correcting Code (ECC) is able to detect and recover from errors at the bit level caused by the natural wear out of individual NAND cells. Supporting all of these maintenance features is over-provisioning, which ensures that the SSD controller always has swap space available to accomplish its tasks. As a precautionary measure, some vendors choose to pre-set a certain amount of mandatory over-provisioning at the factory, and there is always the option to manually set aside additional space for even further-improved performance (e.g. under demanding workloads).

NAND Characteristics

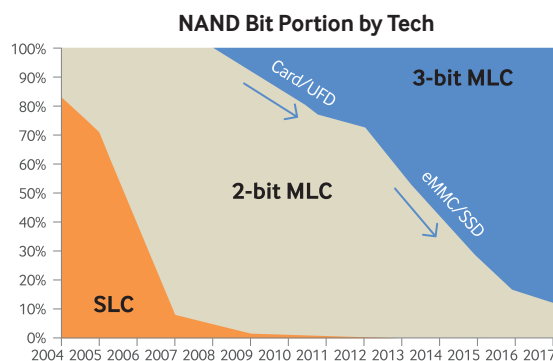
	SLC	2-bit MLC	3-bit MLC
Bits per Cell	1	2	3
Performance	★★★★	★★★	★★
Endurance	★★★★	★★★	★★
ECC complexity	★★	★★★	★★★★

* Today's SSD controllers are capable of minimizing the risks associated with 2-bit MLC and 3-bit MLC NAND through techniques like advanced Error Correcting Code (ECC).

On top of all of the above advancements, the modern trend towards increasing storage densities has a convenient side benefit – with higher capacities also comes higher performance, because increasing the number of NAND chips in an SSD allows for more parallelization, helping to overcome the inherently slow program times that MLC NAND suffers compared to its SLC predecessor (this same phenomenon also explains why the same or similar NAND flash can deliver very different performance, lifespan, and reliability among various NAND-based devices). The additional free space has the added benefit of providing the controller with unused capacity to use as a kind of non-official over-provisioning space. Together, all of the progress SSD vendors have made in NAND technology has led to lower cost drives with performance, endurance, and reliability that can satisfy the demands of an ever-increasing audience.

Industry Trends

History is set to repeat itself thanks to the above and other continuous advances in NAND technology. The entire Information Technology and Consumer Electronics industries, which voraciously consume NAND chips for devices like MP3 players, smart phones, memory cards, USB flash drives, car navigation systems, and the other devices that make up this digital age, have benefited greatly from increased capacity at more favorable pricing. It is now time for PC storage, which is rapidly moving away from traditional Hard Disk Drive (HDD) technology in favor of SSDs, to enjoy similar benefits. The “bit growth” phenomenon discussed above is pushing the industry towards wider



3-bit MLC NAND adoption. In fact, we may soon see 3-bit MLC NAND-based products begin to dominate the SSD industry. In reality, 2-bit MLC NAND’s performance and lifetime characteristics are such that they far exceed the requirements of most computing tasks. There are even fewer applications that require SLC-level endurance and reliability these days. Thus, the real-world benefits are well worth any tradeoffs in performance and endurance we must make as we move towards denser memory technology. Today’s NAND-based storage will still far outlive the useful lifespan of the devices it powers. Choosing to sacrifice a bit of excess lifetime and performance in favor of dramatic cost benefits will allow manufacturers to deliver NAND technology and its many advantages over traditional storage technology to a far wider audience.

Samsung has taken the first step by introducing 3-bit MLC NAND to the SSD market with its 840 Series SSD, made possible by its fully integrated design approach, proprietary controller architecture and firmware algorithms, and superior NAND quality. As mentioned previously, increased SSD capacity leads to increased performance and endurance, so SSDs using this technology will only improve as densities grow. Thus, 3-bit MLC may represent the beginning of the next personal storage revolution.

Asynchronous vs. Synchronous NAND

At the heart of all electronic devices is the concept of a regulatory signal, which coordinates the rate at which instructions are executed and acts as a kind of conductor to keep everything in sync. Historically, NAND devices used two regulatory signals: the “RE signal” (for Read Events) and the “WE signal” (for Write Events). As NAND evolved and increased in speed, however, it became necessary to introduce a new signal, called a “strobe.” Present in all modern DDR NAND implementations, the strobe helps the controller to handle read and write data at high speed. It starts after receiving a signal from the host indicating a read or write event – think of a stopwatch, which can start and stop on demand. Depending on its current task, the strobe plays slightly different roles. During a write event, it is directly managed by the controller, whereas it plays a more supporting role during a read event.

With ONFI 2.0, synchronous logic was introduced to the ONFI NAND specification. Synchronous devices operate a “free-running” clock, which means that, as long they are receiving power, the clock will continuously run – much like a standard wall clock. This clock is used as a reference for the strobe. Modern ONFI NAND implementations eschew the use of a clock once again in favor of using only the strobe. Every generation of Toggle NAND was asynchronous.

Asynchronous




Synchronous




Traditional vs. DDR NAND

The advent of DDR NAND was a key breakthrough in increasing NAND speed. Traditional NAND, also known as “Single Data Rate (SDR)” NAND, was only capable of processing data on “one edge” of its regulatory signal. ONFI 1.0 and pre-Toggle Samsung and Toshiba NAND belong to this category. The figure below shows a Read operation for illustrative purposes.



In contrast, modern NAND can handle data on “both edges” of its regulatory signal (or strobe, as it is known). This concept is similar to how DDR RAM works (DDR stands for “Double Data Rate”).



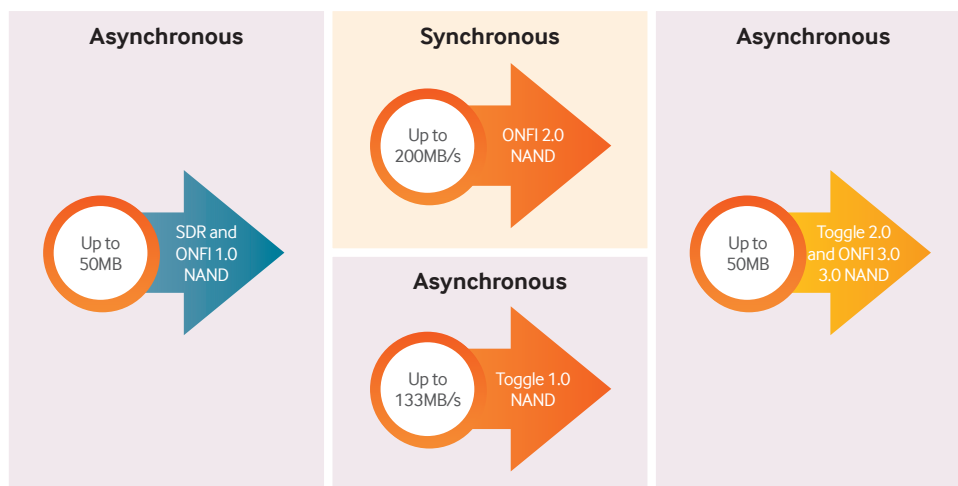
Toggle DDR NAND works much the same way as the DDR NAND diagram above. The strobe that is used to coordinate the NAND operations has rhythmic rises and falls, and Toggle DDR NAND is capable of processing data on both the rise and fall.

Processing data on “both edges” of the strobe obviously results in significant speed and efficiency gains. The different generations of Toggle DDR and ONFI NAND are distinguished by how fast the strobe (or clock in the case of Synchronous ONFI NAND) itself runs (how fast it can complete one cycle).

Toggle NAND vs. ONFI NAND

In the NAND flash industry, there is some variation in implementation depending on manufacturer. Samsung and Toshiba make Toggle NAND, whereas everyone else makes what is called ONFI NAND.

Each generation of NAND, whether Toggle or ONFI, differs in the speeds it is capable of delivering: ONFI 1.0 and pre-Toggle NAND from Samsung and Toshiba (both asynchronous) are capable of speeds up to 50MB/s, ONFI 2.x (synchronous) and Toggle 1.0 NAND (asynchronous) are capable of delivering up to 200MB/s and 133MB/s respectively, and ONFI 3.0 and Toggle 2.0 (both asynchronous) are capable of delivering up to 400MB/s. Today’s newest NAND chips are either Toggle 2.0 or ONFI 3.x.



So why choose one standard over the other? The latest ONFI NAND chips, because ONFI has alternated between synchronous and asynchronous methods, must be built with backwards compatibility for both synchronous and asynchronous logic. Because Toggle NAND has always been asynchronous, Toggle 2.0 NAND benefits from reduced complexity and, in turn, increased flexibility when it comes to designing smaller chip sizes and bringing next-generation products to market faster. Furthermore, because of Samsung and Toshiba's positions as market leaders in the NAND industry, a majority of the market has standardized on Toggle NAND.

NAND Type & Performance

How do speed differences among the various generations of individual NAND chips affect the overall speed of an SSD? And what does this mean for you in terms of real-world impact?

In practice, the use of synchronous versus asynchronous NAND has no impact on SSD performance. What actually affects performance is the generation of the NAND (e.g. SDR/ONFI 1.0, Toggle 1.0/ONFI 2.0, etc.). When discussing NAND type, therefore, it is inaccurate to refer to it simply as either "synchronous" or "asynchronous" when referring to performance.

Why Samsung?

As the #1 player in the worldwide memory market for over 20 years, no one has more experience manufacturing NAND than Samsung. Furthermore, Samsung's fully-integrated manufacturing approach means that it has intimate knowledge of every nuance of each component of your SSD. While generic controller manufacturers must optimize their chips to work with both ONFI and Toggle NAND, Samsung can focus all of its design efforts on making Toggle NAND work perfectly with its proprietary controller technology. The end result is product characterized by awesome performance and unrivaled reliability – a product only Samsung could build.

04

Understanding SSDs

A Peek Behind the Curtain



While an SSD might look like a simple device on the surface, there is a great deal of complexity behind the scenes. In order to ensure that your SSD stays in prime working condition, the SSD controller must manage complex performance and lifetime optimization algorithms. To facilitate an understanding of why the design and implementation of these algorithms is so crucial to your SSD experience, it is necessary to understand the limitations of NAND technology and what exactly is going on behind the scenes to compensate for them.

Writing and Erasing NAND

Before we can understand the various maintenance algorithms that the controller employs to keep your SSD neat and tidy, we need to understand a few basics about how we read and write data to a NAND chip. Data is stored in a unit called a “page,” which is finite in size and can only be written to when it is empty. Therefore, in order to write to a page that already contains data, it must first be erased. This would be a simple process, except each page belongs to a group of pages collectively known as a “block.” While data is written in pages, it can only be erased in blocks. To help illustrate the concept of pages and blocks, let’s use an analogy. Think of the Etch-A-Sketch toys many of us used as children. You could continue drawing on the screen until it was full – think of each artistic flourish you add to the canvas as a separate “page” and the total canvas as the “block” that all of the pages exist in – but when you wanted to erase something, you had to shake the toy and erase everything at once. Extending this thought process, there would be many Etch-A-Sketches inside each NAND chip. This restriction obviously has great implications for how the controller manages the NAND flash that stores your data.



Performance Over Time

With the above limitations in mind, let’s look at what would happen if we wanted to re-use a page that contained old data. If the SSD wanted to reuse an existing, no longer valid page, the other (valid) pages in the block would have to be copied into an empty NAND cell while the entire block was erased. The SSD would then rewrite the entire set of valid data, old and new, to the intended cell. This process is complex and time-consuming. Thus, the SSD controller avoids directly “overwriting” existing data in favor of working some magic through the Flash Translation Layer (FTL), a special mechanism that helps coordinate communication between the NAND flash and the host system. Utilizing the FTL, SSDs implement a logical to physical mapping system called Logical Block Addressing (LBA). Physical NAND writes are not required to correspond directly to the space the host system requests. Instead of performing all of the unnecessary copies described above to overwrite a piece of old data, the SSD writes the new data to the next available page and simply marks the old data as “invalid.”

In reality, the technique just described only delays the inevitable – at some point the SSD will suffer some sort of performance deficit regardless of whether or not it has to perform the complex overwrite procedure above. Let’s think of what happens when the drive is full. As an SSD is filled with more and more data, there will naturally be fewer free blocks readily available. The SSD is then forced to actively consolidate valid data and prepare free blocks in order to write new data and perform maintenance. This process of moving and consolidating data takes time, which is perceived as decreased performance, and requires free space. This is why Over Provisioning, which guarantees a certain amount of free swap space to use for Garbage Collection and other maintenance activities, is so important for SSD performance – it allows the Garbage Collection algorithm to prepare free space in advance through data consolidation.

Garbage Collection

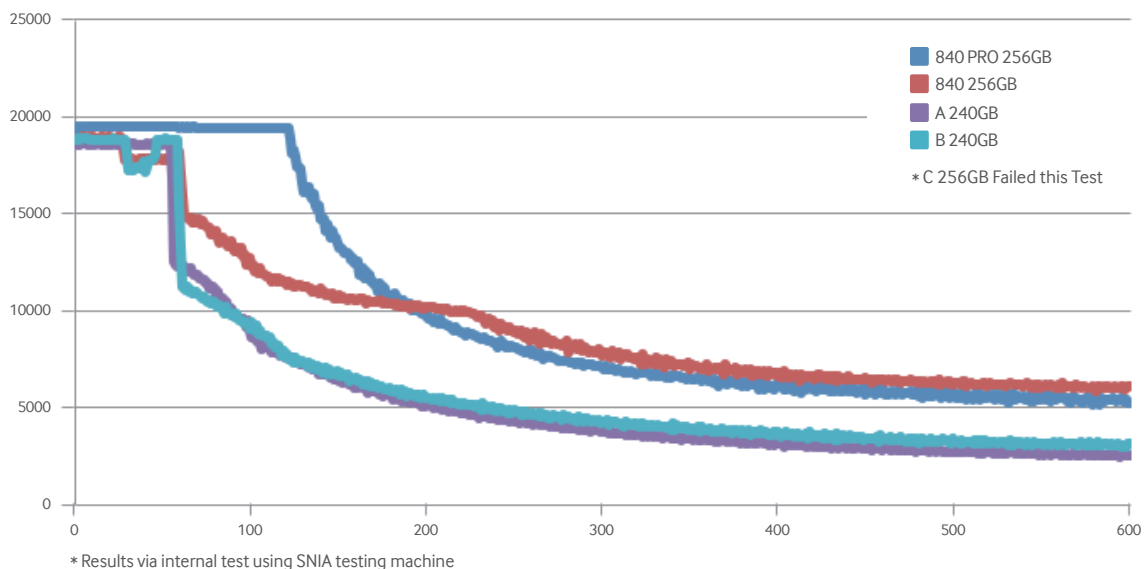


Now that we know a bit about why the SSD needs free blocks to be readily available, we can talk about one of the ways the SSD controller prepares free blocks. One of those processes is called “garbage collection.” This name is a bit misleading, as the process actually involves the collection of good data, not garbage. The concept, however, is rather simple.

By nature, because of the nuances we discussed above, SSDs are a bit obsessive about how they organize their valid data. They prefer to have it neatly stacked together, and they run more efficiently when they can keep it that way because it is easier to find free space to write to. Unfortunately, everyday use makes this type of neatness difficult, because the OS is constantly writing and deleting data of various sizes, leaving data strewn haphazardly throughout the SSD. Garbage Collection remedies this by combing through the Swiss cheese of data left behind, collecting any valid data and carefully placing it together. By doing this, invalid data is left separate and can be erased to make more free space, which means no waiting is necessary when you try to write new data to a page that may have previously been filled with “garbage” data.



There are two methods of Garbage Collection, one that takes place in the background (when the SSD sits idle) and one that takes place in the foreground (when the SSD is working with data). Samsung’s SSDs use a proprietary implementation of the latter, as this allows them to actively manage data and create free blocks on demand. Thus, in an extended workflow where the drive may not have a chance to sit idle and prepare free blocks, Samsung’s SSDs are always prepared for what’s next. The chart below shows what happens to performance when an SSD is hammered with a non-stop data flow for 6 hours. The 840 and 840 PRO Series SSDs level out at a significantly higher performance level, and the 840 PRO is able to hold on to its peak performance longer, than many other drives.



TRIM

Modern operating systems have enabled another form of SSD maintenance, TRIM. TRIM is a facility by which the OS can notify the SSD when data is either marked for erase or no longer valid. TRIM helps to make Garbage Collection more efficient by preparing invalid data for deletion. Remember, SSDs are new technology, so computers were built to interface with traditional hard disk technology. Hard disks are not subject to the same write/erase limitations that SSDs are – they can easily overwrite data in an existing location without erasing it first. Therefore, when the OS “deletes” data, the data does not actually go anywhere. The space in which it resides is simply marked as “free space” that may be used later. By default, because it doesn’t know it’s not working with its longtime HDD companion, the OS doesn’t let the SSD know that a particular piece of data is no longer valid and that its corresponding memory location is now free – after all, there is no reason to do so. With the introduction of SSDs, however, there is now a compelling reason to increase communication about file validity between the OS and the storage device. Enter TRIM. TRIM allows the OS to inform the SSD which data are no longer valid, allowing the SSD to skip over invalid data when performing Garbage Collection instead of moving

around old data. Once a block is full of pages that all contain invalid data, that block is considered free and may be erased. The TRIM command is sent to the SSD controller automatically by the OS every time it deletes a file. As it requires OS support, not all users will be able to use native TRIM functionality. On PCs, TRIM is supported in Windows 7 or later. On Macs, TRIM is only supported for Apple's OEM SSDs and is not supported for Samsung's (or any other manufacturers') aftermarket SSDs. Users of older Windows operating Systems (Windows XP, Windows Vista) may use Magician's built-in "Performance Optimization" feature to manually pass the TRIM command to the SSD on demand (or via user-specified schedule).

Bad Block Management & Error Correcting Code (ECC)

In addition to maintenance at the drive level, the SSD must also perform maintenance at the chip level. In every NAND cell, each page contains a few extra bytes of extra capacity that the SSD controller uses to store a "parity bit." Error-Correcting Code (ECC) uses this parity bit to compensate for other bits that may fail during normal operation of the drive. When the controller detects a read failure, it will invoke ECC to try and recover from it. If recovery is not possible, the firmware's bad block management feature will retire the block and replace it with one of several free "reserved blocks." "Bad blocks" can be made during read, program, or erase operations and are actively managed to guarantee expected SSD performance.

Wear Leveling

NAND flash memory suffers from one final limitation: each cell has a finite lifespan and can only withstand a limited number of program/erase cycles (called P/E cycles). The specific amount of P/E cycles depends on the process technology (e.g. 27nm, 21nm, 19 nm, etc.) and on the program mechanism (e.g. SLC, MLC). In order to overcome this limitation, the SSD firmware employs a wear-leveling algorithm that guarantees that write operations are spread evenly among all NAND cells. Using this technique, no single cell should be unduly stressed and prematurely fail. If too many cells were to fail, the entire block would have to be retired as just discussed above. There are only a limited number of reserved blocks, however, so this event should be avoided to prolong overall drive life.

Wrap Up

Fortunately, all of the above procedures (with the exception of TRIM if you're using an older Windows OS) happen transparently and without action on behalf of the user. While specific implementation will vary, most modern SSDs include all of these features. In fact, without features like Wear Leveling and ECC (to extend drive life and protect data integrity) and TRIM and Garbage Collection (to maintain SSD performance), SSD quality and user experience would suffer.

Why Samsung?

Maintenance procedures like wear-leveling and Garbage collection, which are created to overcome the unique properties of NAND flash memory, work together to help ensure that your SSD performs well over extended use. Together, these algorithms actually increase write activities to the NAND, which reduces overall lifespan. Thus, the key in designing a great SSD is finding the optimum balance among lifespan, performance, and reliability. As the #1 player in the memory business for over 20 years and the largest global supplier of SSDs in the preinstalled storage business, Samsung has unrivaled knowledge of and experience with SSD technology. Samsung's unique, integrated approach to SSD manufacturing affords it full control of every component. You can trust that Samsung's expertise is safeguarding your precious data, and your productivity, when you purchase a Samsung SSD.

05

Maximize SSD Lifetime and Performance With Over-Provisioning

Over-Provisioning (OP), the practice of allocating a specific, permanent amount of free space on an SSD, is a widely-used method for improving both SSD Performance and Endurance. Historically, Samsung has not implemented mandatory OP on any of its SSDs. With the introduction of the 840 Series and the reality of increasingly complex NAND fabrication processes, however, Samsung has chosen to implement a minimum amount of OP in its mainstream drives (the 840 PRO will not feature mandatory OP).

What is OP?

An SSD controller is responsible for the massive task of managing all data traffic and storage for the drive. NAND technology's intrinsic complexities require a lot of extra work behind the scenes. A data write is not as simple as placing data into an empty memory bank. Each NAND cell has a limited lifespan – it can only endure a specific number of data reads/writes. An additional layer of complexity is added by the fact that overwriting old data on NAND requires an erase of the entire NAND block (this same block may contain other data that is still valid). As a result, the controller is constantly moving data around to ensure that the cells wear evenly and to preemptively prepare “free blocks” to use for future data writes.

All of this management work requires the SSD to have a kind of “swap space” available to use as temporary storage while the controller goes about its business. The controller will use any available free space for this, but free space becomes a premium commodity as we fill our drives with data. OP is a way to set aside a minimum amount of free space, inaccessible to the user or the OS, which the SSD controller can utilize as a kind of “work bench.”

How do I set OP?

Samsung Magician can assist users in setting up OP on their SSDs. Magician's simple, graphical OP tool will recommend an optimum amount of OP space (for a consumer workload) and then configure the SSD with just one click. This decision can always be altered or reversed at a later time.

Some drives, however, like most SandForce-driven SSDs and the new SSD 840 Series, set aside a permanent amount of OP. This minimum amount of OP is not user-configurable and will vary from vendor to vendor. There is no “right” amount of OP to set aside, and it is best to vary it by capacity and situation.

If you don't want to use Magician's recommendation, or if you are a professional or more advanced user with different requirements, you are free to set aside your own OP space. To decide how much to set aside, you must consider your specific usage patterns as well as the size of your SSD. Users with large SSDs (250GB+) will likely not use all of the available space on their drives as long as they are not storing large amounts of photos, video or other media files. Thus, a casual user with a large-capacity SSD may not need to set aside any extra space for OP. The SSD will naturally use any available free space to perform its maintenance algorithms. If you have a small SSD, on the other hand, it is recommended to set aside some OP (between 6.7 and 10% of total drive space) to minimize the risk of accidentally filling the drive to capacity. While filling a drive with data isn't harmful, it will have a severe impact on performance. Users with smaller drives are much more likely to fill them with data, effectively taking away the SSD controller's working space and making it more difficult to prepare free blocks as well as accomplish basic maintenance tasks.

Also, different OP ratio is recommended by usage applications and workload. One might wonder why the recommendations should be variable. Data access patterns are the determining factor. Activities that produce frequent Read/Write requests, especially Random Read/Writes, put extra stress on the SSD, which in turn increases Write Amplification (a phenomenon by which physical NAND writes outnumber logical write requests from the host). Heavy workloads wear out NAND cells faster, which increases the need for Bad Block management so that the controller can retire worn out cells. Retiring a NAND cell requires the controller to copy all of the cell's valid data to a new block, taken from a limited number of “reserved” blocks. Without empty blocks ready to accommodate this process, performance will suffer. Working in tandem with bad block management, and reducing the need to use it altogether, is Wear-Leveling,

which ensures that no single cell is written to more than others. This process also requires data to be copied from one block to another, which further increases the necessity for free blocks to be available. To ensure free blocks are available, Garbage Collection algorithms consolidate good data and erase blocks of invalid data. This process too, however, requires free blocks to use for temporary storage while the invalid blocks are cleared.

OP helps with garbage collection, wear-leveling and bad block management by effectively increasing the size of the controller's "work bench," thus giving it extra free space to use while it consolidates and moves data or retires worn out cells.

All of the above tasks are the reason WAF exists – writing data to an SSD is not a simple one-to-one operation. There are complex processes taking place behind the scenes at all times.

Why Use OP?

OP has a direct effect on SSD performance under sustained workloads and as the drive is filled with data. Guaranteeing free space to accomplish the NAND management tasks discussed above (Garbage Collection, Wear-Leveling, Bad Block Management) means the SSD does not have to waste time preparing space on demand, a process that requires additional time as data is copied, erased, and recopied. An added benefit is that OP makes all of the SSD maintenance procedures more efficient, reducing the WAF by ensuring there's room to work. Consider this scenario: You are a chef with very limited counter space (say one arm's length across and half an arm's length deep). Preparing a 5-course meal in such a limited area would mean you have to waste a lot of time moving things around or putting them away to make room for other tasks. Now, imagine you quadruple your work space. You can leave everything you need at hand, reduce repeating steps (like getting the salt from the pantry and putting it back away), and increase your working speed. The same goes for SSDs. Give them more room to work, and they can do it more quickly and efficiently, reducing WAF and increasing performance at the same time.

OP and the 840 Series SSD

As NAND process technology advances, the chips themselves become increasingly smaller. As they shrink, the chips also become less reliable at holding data. There are a number of ways to mitigate this problem, including using Error-Correcting Code (ECC). As explained above, any time we give the controller more work to do, it requires more room on its "work bench" to do it efficiently.

The 840 Series represents the first consumer SSD to implement 3-bit/cell MLC (also called TLC) technology. This technology, as its name suggests, stores 3-bits of data per cell, as opposed to the 2-bits per cell that today's more common Multi-Level Cell (MLC) NAND holds. Because, as just mentioned above, the size of the chips themselves is also shrinking, in effect we are squeezing even more information into a smaller space.

This is nothing a good firmware algorithm can't handle, however. Samsung's 3-bit/cell MLC-based SSD 840 Series, equipped with mandatory OP, will still far outlast the useful life of the hardware it powers.

Closing Thoughts

OP, while already popular among SSD enthusiasts, will continue to become an important safeguard for SSD performance and endurance as the NAND industry continues to shrink chips to save on cost, increase efficiency, and expand capacities. Samsung sets a lower value for its mandatory OP on the 840 Series than any other competitor, a feat that is possible because of its top-tier NAND chips and superior maintenance algorithms. As the NAND industry continues to evolve, Samsung will continue to be on the cutting edge, and optimizing features like OP is just one of the many ways it will continue to provide SSD users with the best storage devices on the market.

06

Protect Your Privacy

Security & Encryption Basics

In today's mobile world, security is a major concern for both individuals and businesses. With the 840 and 840 Pro Series SSDs, Samsung is adding peace-of-mind through the implementation of hardware-based AES Full Drive Encryption.

Terminology

In order to understand the encryption technology built into the latest Samsung SSDs, it is necessary to understand some basic security terminology.

AES

The Advanced Encryption Standard (AES) is an encryption standard approved by the National Institute of Standards and Technology (NIST) for the safeguarding of electronic data. After being adopted by the US government, the standard is now used worldwide. This cipher, usually implemented with either 128-bit or 256-bit encryption keys, is widely used to protect sensitive information and is found integrated at both the hardware and software level. All 840 and 840 PRO Series SSDs are equipped with a high-performance hardware accelerator that implements AES encryption with a 256-bit key.

FDE

Full Drive Encryption refers to a storage device in which nearly everything is encrypted rather than encrypting only certain files or folders. This solution is attractive for high-security environments because it makes it simple to destroy all data on the drive by destroying and replacing the cryptographic key(s) that protect it. With this technology, the swap space and temporary files are also encrypted, and, when implemented through hardware rather than software, even the bootstrapping code is encrypted. By using a Trusted Platform Module (TPM), standardized by the Trusted Computing Group, in conjunction with FDE, the integrity of the boot environment can also be verified.

SED

Self-Encrypting Drive (SED) is a term that refers to a storage device that implements hardware-based FDE. Therefore, an SED is a special case of FDE. SEDs boast better performance, security, and manageability compared to software-based FDE implementations, which commonly suffer severe performance degradation as a result of the encryption overhead. Also, because the encryption key exists only inside the SED itself, it is impossible to access it via the host (operating system). Software-based solutions are vulnerable to several types of attack because they must store the encryption key in main memory. Finally, because SEDs provide drive-level encryption that is independent of the operating system and any other data management tools (e.g. compression utilities, data loss prevention, de-duplication, etc.), users can easily install an SED into any system without worrying about operating system or application interference.

OPAL

OPAL is the name for an SED storage specification developed by the Trusted Computing Group, the same group responsible for the TPM microchip mentioned above. It defines a means by which to place an SED storage device under policy control. Its goal is to protect the confidentiality of user data and prevent unauthorized access to the drive while still maintaining compatibility with multiple storage vendors through a standardized management interface. Most systems require 3rd party software to utilize the OPAL Storage Specification, although Windows 8's BitLocker feature supports this functionality natively.

Choosing the Right Option

With the introduction of the SSD 840 and 840 Pro Series SSDs, Samsung has added AES hardware-based SED technology to its consumer SSD lineup. Simply enabling the ATA password via the BIOS will automatically render all data on the drive unintelligible without the proper password. Because it is implemented at the hardware level, there is no performance penalty like there is with a software-based FDE implementation. This feature is a valuable privacy tool for anyone who uses a portable computing device (e.g., laptop), especially frequent travelers.

In addition to addressing personal security concerns, there are many industries that either require or would benefit from SED technology, including healthcare, insurance, government, law enforcement, and finance, among others. SED technology helps protect sensitive information from physical attack on a lost or stolen laptop, making it nearly impossible to access data stored on a drive even if the drive is removed from its original system and installed in another PC or the NAND chips themselves are removed from the SSD.

While they do feature SED technology, the 840 and 840 Pro Series SSDs do not support the OPAL storage specification management interface. OPAL drives are geared towards enterprises that need to manage security protocols and want to have advanced control over authentication. With third-party software support, IT managers can set detailed security provisions to restrict access by partition, physical location of the laptop, etc. Anyone interested in this level of security management should research enterprise-class TCG/OPAL SED options.

Someone who wants to manage a personal machine or an SMB that depends on its employees to handle most of their own IT support, however, will find that the SED feature of Samsung's 840 and 840 Pro Series SSDs is well-suited to their needs. These SSDs offer basic, yet robust, security with minimal effort and expense.

Enabling AES Encryption

AES encryption is always active on an 840 or 840 Pro Series SSD. In order to benefit from the encryption feature, however, the user must enable an ATA password to limit access to the data. Failure to do so will render AES-encryption ineffective – akin to having a safe but leaving the door wide open. To set an ATA password, simply access the BIOS, navigate to the “Security” menu, enable “Password on boot” and set an “HDD Password.” Administrators also have the option of setting a “Master Password,” which can allow a lost user password (“HDD Password”) to be recovered. The “Master Password” may also be used to unlock and/or erase the drive (depending on the settings), effectively destroying, and thus protecting, the data but allowing the drive to be reused. The setup procedure may differ slightly depending on the BIOS version installed on a particular machine. It is best to consult the user manual if there is any confusion.

07

Communicating With Your SSD

Understanding SMART Attributes

SMART (also written S.M.A.R.T.), which stands for Self-Monitoring, Analysis and Reporting Technology, is an industry standard reliability prediction indicator for both IDE/ATA and SCSI storage drives. When analyzing SMART attributes, it is very important to remember that they vary in meaning and interpretation by manufacturer. SMART simply refers to a signaling method between sensors in the drive and the host computer – the communication protocol is standardized but what it reports is not.

Technology Overview

SMART monitors computer drives to detect and report on various reliability indicators. The technology aims to anticipate failures and warn users of impending drive failure, allowing the user to replace an ailing drive to avoid data loss and/or unexpected outages. Of course, SMART can only warn of predictable errors, which result from slow processes like mechanical wear and can be predicted by analyzing certain indicators (such mechanical problems accounted for 60% of HDD failures). Unpredictable failures, like a sudden mechanical failure resulting from an electrical surge, have no measurable variables to track and analyze. Modern SMART implementations (in HDDs) also try to prevent failures by attempting to detect and repair sector errors. All data and all sectors are tested to confirm the drive's health during periods of inactivity.

In addition to the functions discussed above and the individual SMART attributes outlined in the next section, SMART-enabled drives are also capable of reporting a SMART status. This status represents one of two values, usually “drive OK” and “drive fail” or “threshold not exceeded” and “threshold exceeded.” A “drive fail” or “threshold exceeded” value indicates there is a high probability the drive will fail in the future; however, the failure may not be catastrophic – the SMART Status simply indicates that the drive will not perform within the manufacturer's declared specifications. So, for example, rather than complete data loss, the drive may simply begin to run slower. As with any technology, the SMART status is not infallible and may not necessarily indicate past or present reliability. The SMART sensors may malfunction, for instance, or a serious mechanical failure may destroy access to the SMART status.

Finally, it is important to remember that SMART attributes vary in both meaning and interpretation by manufacturer. Some attributes are considered trade secrets, and not all drives report the same SMART attributes. A manufacturer, in theory, could report only one SMART value and advertise its drive as SMART-enabled. The SMART standard simply refers to a signaling method between sensors in the drive and the host computer, not a standardization of the attributes themselves.

SMART Attributes

Device manufacturers who implement SMART technology enable a set of attributes and corresponding thresholds. Please note that names and descriptions may vary by OEM. Also, many attributes were defined for use with traditional HDDs. As a result, some attributes are used with modified meaning by SSD vendors since their names are not applicable to SSD technology. Below are the SMART Attributes associated with the Samsung 840 and 840 PRO Series SSDs, which are displayed in decimal format.

ID # 5 Reallocated Sector Count

The raw value of this attribute represents the number of sectors that have been moved as a result of a read error, write error, or a verification error. If the firmware detects any of these types of errors, all valid data in the block the error originates from must be transferred to a new block. This number should be low because a high number would indicate a large number of failures.

ID # 9 Power-On Hours

The raw value of this attribute shows the total count of hours the drive has spent in the power-on state. When the system or SSD is in Hibernation Mode, the Power-On Hours value does not increment. Samsung's SSDs support the

DIPM (Device Initiated Power Management) feature. Thus, with this feature enabled, this attribute excludes any time the device spends in a “sleep” state. With DIPM off, the recorded value will include all three device power states: active, idle, and sleep.

ID # 12 Power-On Count

The raw value of this attribute reports the cumulative number of power on/off cycles. This includes both sudden power off and normal power off cases.

ID # 177 Wear Leveling Count

This attribute represents the number of media program and erase operations (the number of times a block has been erased). This value is directly related to the lifetime of the SSD. The raw value of this attribute shows the total count of P/E Cycles.

ID # 179 Used Reserved Block Count (total)

This attribute represents the number of reserved blocks that have been used as a result of a read, program or erase failure. This value is related to attribute 5 (Reallocated Sector Count) and will vary based on SSD density.

ID # 181 Program Fail Count (total)

This attribute represents a total count of the number of failed program requests (failed writes).

ID # 182 Erase Fail Count (total)

This attribute represents a total count of the number of failed erase requests.

ID # 183 Runtime Bad Count (total)

Equal to the sum of the Program Fail count (attribute 181), the Program Erase Fail count (attribute 182), and the Read Fail count. This summary value represents the total count of all read/program/erase failures.

ID # 187 Uncorrectable Error Count

The total number of errors that could not be recovered using ECC.

ID # 190 Air Flow temperature

The current temperature of the area surrounding the NAND chips inside of the SSD.

ID # 195 ECC Error Rate

The percentage of ECC correctable errors

ID # 199 CRC Error Count

The number of Cycle Redundancy Check (CRC) errors. If there is a problem between the host and the DRAM or NAND flash, the CRC engine will tally the error and store it in this attribute.

ID # 235 Power Recovery Count

A count of the number of sudden power off cases. If there is a sudden power off, the firmware must recover all of the mapping and user data during the next power on. This is a count of the number of times this has happened.

ID # 241 Total LBAs Written

Represents the total size of all LBAs (Logical Block Address) required for all of the write requests sent to the SSD from the OS. To calculate the total size (in Bytes), multiply the raw value of this attribute by 512B. Alternatively, users may simply consult the Total Bytes Written indicator in Magician 4.0.

08

Benchmarking Utilities


What You Should Know

With the abundance of Benchmarking Utilities available and the varying results they can report, it can be confusing to decide which tests to pay attention to and which to question. Sometimes, two blogs using the same tools present slightly differing results. In fact, there are many variables that can affect test results, including PC configuration and the options set by the user in the utility itself. Further complicating matters, the results may look cryptic to novices. In this paper, we will explore some of the most common benchmarking utilities, their configuration options, and how to interpret some common test results.

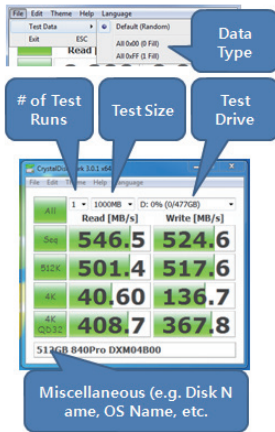
Types of Benchmarking Tools

Categories

There are 3 main types of Benchmarking Tools: Synthetic, Trace-Based, and Application-Based. This paper will explore some of the most common tools, CrystalDiskMark (CDM), Iometer, PCMark Vantage, and PCMark 7. In addition to delving into the details of each application, it is also helpful to understand the wider application categories to which the tools belong.



Synthetic Benchmarking Tools



Synthetic Benchmarking Tools mimic the type of stress a drive would experience during various types of workloads using patterns specifically designed to behave like real-world processes.

Test data may be either Non-Random (Repeating) or Random (Non-Repeating). Whether the data is Random or Non-Random may have an effect on SSD performance for some drives because of differences in controller technology, firmware design, NAND components, and DRAM cache size and performance. For example, repeating data is compressible, which means that an SSD that utilizes a compression algorithm will be able to shrink the amount of data it must actually write to the NAND flash. Samsung SSDs perform equally well with Repeating and Non-Repeating data because the Samsung MDX controller does not use compression. Thus, when there is a large discrepancy between competing drives, it is a good idea to check whether or not the benchmarking utility was sending Random or Non-Random data and whether or not the tested drives utilize compression.

Some Synthetic Benchmarking tools also allow the user to specify Sequential or Random data access patterns. Randomly accessed data, as its name suggests, is stored and accessed in a random order (non-neighboring memory locations), whereas sequentially accessed data is stored and accessed sequentially (in neighboring memory locations). Files accessed randomly are typically small in size (less than 64KB), while files accessed sequentially are usually over 128KB.

Synthetic tests provide the most basic, fundamental information about drive performance, but they do not represent real-world application performance.

Crystal Disk Mark (CDM)

CrystalDiskMark, freely available via the web, is a very popular storage benchmarking program. Its simple interface allows the user to specify whether to test sequential or random data, the size of the test range, and the number of repeat tests. CDM, which Samsung uses when measuring the sequential performance of its SSDs, reports the maximum sequential and random read/write performance in Megabytes per second (MB/s). Random performance is reported for three separate workloads: 512KB data, 4KB data at a Queue Depth (QD) of 1, and 4KB data at a QD of 32. The most important of these are the latter two, as most consumer workloads will be similar to 4KB data at QD 1 and the SSD's

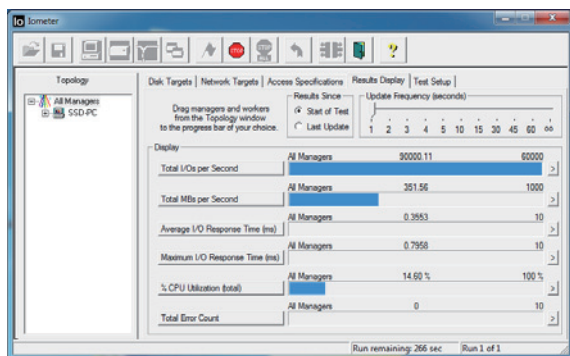
maximum random performance is represented by 4KB at QD 32.

Random performance is most often reported in Input/Output Operations per Second (IOPS), which can be calculated from CDM's results by dividing the total Bytes per second by the data transfer size (4KB). For example, in the image to the left, if we divide the Random Read speed of 408.7MB/s by 4KB (4,000MB), we get 0.102175MB/s or 102,175 IOPS.

While the majority of client PC workloads will not exceed a QD of 1, some usage scenarios may generate a QD of 2-6 or even up to 10 (in limited applications). Data center applications, on the other hand, may generate massive numbers of Input/Output (I/O) requests, creating a QD of 32, 64, or even 128 in some cases (depending on the number of access requests per second). This is why it is useful to look at multiple performance metrics when evaluating the overall performance of a drive and why it is important to know which metrics are more appropriate for the intended usage application.

Iometer

Iometer, also freely available via the web, is another popular synthetic benchmarking tool. It is commonly used to measure IOPS (random performance) and is the program Samsung uses to measure the random speeds of its SSDs. Its interface is significantly more complex than CDM, but it offers more control over the testing scenarios. Users can



customize what types and sizes of data to use (or mix) as well as specify the test length, the number of workers (threads), or the queue depth, among other options.

When choosing variables for a Random performance benchmark, remember that IOPS performance will vary greatly depending on the parameters of the test data (e.g. the number of threads, queue depth, data block size, etc.) and system configuration (e.g. storage drivers, OS background processes, RAM, etc.). IOPS numbers are directly affected by the NAND components, storage controller, and NAND interface speeds, which may all influence random data

seek time. For the vast majority of users, the most meaningful Iometer scores will be those of 4K random Read and Write performance at a Queue Depth of 1-32. The example screenshot above shows Random Write performance for the 840 PRO SSD.

Trace-Based Benchmarking Tools

Trace-based Benchmarking tools are essentially a recording of actual Input/Output (IO) requests from a set of real-world processes (e.g. running Microsoft Word or Windows Media Center or moving image files from a USB disk to the SSD and subsequently deleting those files). Think of a cassette tape or a CD-R, which, once they contain a recording, can be played back on any number of compatible devices. A trace-based tool will record the time it takes to finish a complete trace (playback) on a particular system configuration, allowing a user to compare results among SSD vendors. Trace-based tools are more realistic than synthetic tools and are much faster to complete than an application-based benchmark.



PCMark Vantage and PCMark 7

PCMark Vantage, designed for Windows Vista, and PCMark 7, designed for Windows 7, consist of pre-recorded traces of real-world user scenarios (actual applications are not run, but emulated) designed to test PC performance at the system and component level. These tools include a variety of test scenarios that test all system components, but storage reviewers select only data relevant to SSDs to report in their reviews. Tasks could include: viewing and editing photos, video, music and other media, gaming, communications, productivity, and general storage.

Application-Based Benchmarking Tools

An Application-Based Benchmarking tool differs from both of the above types of tool in that actual applications (e.g. Microsoft Word or Adobe Photoshop) are installed and run while performance is recorded. Applications are chosen

based on what system component the user wishes to stress. The problem with this approach, however, is that most applications are not very storage-centric. Since all SSDs are so fast in everyday tasks, these types of benchmarks are less useful because they will show minimal differences between SSDs.

Interpreting Benchmark Results

Measurements

Among the many measurements you will come across, it is best to familiarize yourself with some of the most common:

- **Sequential Read** speeds are reported in megabytes per second (MB/s) and indicate how fast the SSD will be at completing tasks like accessing large multimedia files, transcoding, game level loading, some types of game play, watching and editing video. This is the speed at which the drive can read data from contiguous memory spaces.
- **Sequential Write** speeds, also reported in megabytes per second (MB/s) indicate how fast the SSD will be at tasks like application installation and document backup. This is the speed at which the drive can write data to contiguous memory spaces.
- **Random Read** speeds, reported in Input/Output Operations Per Second (IOPS) indicate how fast the SSD will be at completing tasks like antivirus scans, searching for email in Outlook, web browsing, application loading, PC booting, or working in Microsoft Word. This is the speed at which the drive can read data from non-contiguous memory spaces.
- **Random Write** speeds, also reported in IOPS, indicate how fast the drive will be able to complete tasks like downloading email, compressing files, or browsing the web. This is the speed at which the drive can write data to non-contiguous memory spaces.

Some other common terminology you will encounter includes:

- **Idle Time**, which is the period of time when an SSD is not handling requests to read or write data.
- **Latency**, which is the response time of the SSD, or the time between when you initiate a request to read or write data and when the SSD completes the request.
- **Queue Depth/Concurrent IOs**, which is the outstanding (pending) number of IO requests that must be handled by the drive. Multiple requests are placed in a queue of a specific length, where they will wait for their turn to be addressed by the SSD if another task is already underway. The most common queue depths to test are a Queue Depth of 1, which is typical of light consumer workloads, and a Queue Depth of 32, which is representative of a heavy workload as might be seen on a server (e.g. web server, database server, etc.). Where in that spectrum an SSD will perform best varies based on the drive's firmware algorithm. Being optimized for high or low queue depths or vice versa is not necessarily a bad thing, but rather a design decision.

Your Results May Vary

Remember that each of the measurements discussed above will vary from system to system. For example, the number of IOPS possible on a given system will vary greatly depending on: the balance of read and write operations, the mix of sequential and random patterns, the number of worker threads, the queue depth, and the data block sizes – and this is without taking into account the effect of basic hardware configurations (microprocessor speed, DRAM size and speed, chipset type and drivers, BIOS settings, etc.).

When working with Synthetic Benchmarking tools, keep in mind that Random Write performance will differ based on the test range. For example, write tests with very small ranges may not be able to fully utilize the internal parallelism of the NAND chips, resulting in low performance result.

Another major factor is whether or not an SSD uses compression technology. SSDs that use compression will suffer significant performance drops (up to 90% of claimed maximum performance) when working with incompressible data (e.g. JPEG photos, digital audio or video files, zipped files, etc.). Because everyday home and office PC users have and use a considerable amount of this type of data in real-world usage scenarios, Samsung chooses not to use compression technology in any of its current SSDs in favor of delivering industry-leading read and write performance with all data types. This is also why Samsung uses incompressible (or Random) data when testing its SSDs.

The best approach is always to look at a variety of test results together, as one particular test in isolation can never show the whole picture of a drive's performance profile.

Samsung's Approach

Benchmarking tools are a great way to show off the raw performance capabilities of an SSD, but they do not, and cannot, truly represent a real world usage scenario. Samsung's SSDs do not utilize compression, but many benchmarking tests rely on heavily compressible data. Therefore, in some mainstream tests, Samsung's SSDs will appear to perform below a drive that uses compression. Today's users work with a lot of multimedia files, however, which are largely incompressible (JPEGs, Movie files, MP3s). An 840 Series will deliver consistently excellent performance regardless of data type, rather than sacrificing performance with one type of file for the sake of another. This is an important distinction, and a point that may be glossed over if the data is not analyzed carefully.

Furthermore, peak speeds are not a good indication of everyday performance. Users are typically not installing applications or copying massive files on a regular basis. Many manufacturers like to brag about peak performance, and Samsung could have designed its SSD to have a similar performance profile. Instead, it chose to focus on real-world users as its baseline rather than topping a particular benchmark. That's not to say the 840 Series won't top the charts (it will) – it just wasn't the design focus. The best picture of a Samsung SSD's performance will come from a system-level, trace-based performance test like PCMark Vantage or PCMark 7 because these tests closely approximate actual consumer usage scenarios. If one must rely on benchmarking tests rather than hands-on testing, these tests are where to focus.

When making a purchase decision, benchmarking test results can be an important factor, but they should be just one of many. Reliability, the drive's ability to sustain performance over extended use and over its entire lifespan, product quality, software solutions (like good data migration and SSD management tools), customer care, and technical support are also very important factors to consider. As the number one SSD supplier to the pre-installed SSD market (PC OEMs), Samsung has an unrivaled reputation in all of these criteria, making the 840 Series an ideal choice for almost any PC user looking to take their computing experience to a higher level.

09

Why Integration Matters

What Samsung's Vertical Integration Means to You

The three most crucial components of any SSD are the NAND flash memory, the controller, and the firmware. The first does the critical task of storing your precious data; the second two work together to accomplish the complex and equally essential tasks of managing data storage and maintaining the performance and lifespan of the SSD. Using the control program contained in the drive firmware, the controller (an embedded microchip) executes automated signal processing, wear-leveling, Error Correcting Code (ECC), bad block management, and garbage collection algorithms, communicates with the host device (e.g. a PC), and facilitates data encryption, among other tasks. Redundantly stored in NAND flash, the firmware may be updated manually to improve or extend functionality when the SSD manufacturer releases an update.

The NAND flash itself is a kind of digital repository, while the controller functions as the command center for everything the SSD does – from actually reading and writing data to executing the Garbage Collection and wear-leveling algorithms that keep the drive clean and speedy. Thus, the firmware, which drives the controller, is arguably the most important component of a good SSD because it delivers the largest difference among branded SSDs in terms of performance, reliability, and endurance (lifespan). Developing good firmware requires precise engineering skills and painstaking integration among the NAND flash, controller, DRAM and other SSD components. Intimate knowledge of NAND characteristics, semiconductor process technologies, and controller characteristics must go hand-in-hand. Samsung represents one of the few manufacturers that designs all of these components in-house; and, among those who do, Samsung enjoys the most experience (with over 6 years as the number one supplier in the preinstalled SSD market).

In-House, Proprietary Firmware

Driving the controller and pulling everything together is the SSD's firmware. The more control a vendor has over the firmware, the more precise and effective the total package. Samsung is not obligated to use off-the-shelf firmware like many of its competitors, allowing it to design its own unique performance profile and execute more robust reliability and bug testing. In fact, customers may notice that Samsung releases very few firmware updates. This is a testament to the quality of the original testing procedures and the superiority of the firmware code – fewer problems to fix is a good thing.

In the rare case that Samsung does issue a firmware update, SSD Magician makes it easy to upgrade your SSD. Magician will notify you of any available updates, and most updates can be completed right within the software itself. When a DOS-based firmware update is required, the Magician software will walk you through the creation of a bootable USB drive or CD/DVD, eliminating much of the confusion surrounding the process.

Stability and upgradeability are only part of the story, however. Being the sole steward of its firmware means Samsung can draft customized garbage collection and wear-leveling algorithms, which have a direct effect on SSD endurance (or lifetime). Because of some of the underlying design limitations of NAND memory, a write to an SSD is not as simple as placing data in a storage location. Rather, the SSD is constantly monitoring the use of each individual cell, ensuring that each one is written to an even number of times to prevent uneven wear and, consequently, premature drive failure. This process is called “wear-leveling.” In tandem, because SSDs cannot overwrite old data without first erasing it, the controller is also preemptively preparing “free blocks” of space to save time when a write request arrives, a process that also requires the moving and re-writing of data. Thus, the SSD is constantly moving around data and a simple write procedure is not a “one-to-one” operation. Since each cell can only endure a certain number of read/write cycles, the efficiency and effectiveness of the wear-leveling and garbage collection algorithms directly affect the endurance of the drive. There is no room to waste precious read/write cycles with unnecessary reads or writes – the more efficient the maintenance algorithms, the longer the SSD will last. Only a company like Samsung, with complete control of both the NAND and the controller/firmware combination, can squeeze out every last bit of efficiency to maximize both SSD performance and endurance.

If you are interested in learning more about the jobs SSD maintenance algorithms are responsible for, please refer to the separate “SSD Maintenance” paper.

Integration Matters

Having full control over the NAND, controller and firmware has major implications both for performance and reliability. Samsung has intimate knowledge of every nuance of these components and can tweak them along each step in the development process to ensure that they work perfectly together. Generic SSD controllers have to worry about supporting multiple NAND specifications (from varying manufacturers who each have different manufacturing processes), while Samsung's proprietary SSD controller is engineered solely to work with its own NAND – which means engineers can focus all of their efforts towards one common specification.

As the most-integrated SSD manufacturer in the industry, Samsung controls all of the most crucial design elements of an SSD: NAND, Controller, DRAM, and Firmware. Working together, these components are responsible for a crucial task – storing and protecting your precious data. When it comes to SSDs, no one knows more about building a high-performance, reliable drive than Samsung.

10

The Samsung Advantage

Why You Should Choose a Samsung SSD

In such a crowded SSD market, it can be confusing to figure out which brand to trust or which drive to buy. Many vendors use the same components, repackaged under competing brand names, and feature extensive lineups that can be difficult to navigate. Samsung is unique in the SSD industry, as it represents one of the only fully-integrated SSD manufacturers. Additionally, capitalizing on its years of industry experience and leadership, Samsung strives to maintain a streamlined lineup of user-friendly products. With its unparalleled understanding of SSD technology and its focus on usability, Samsung is the brand you can trust to store your precious data.

End-to-End Integration

Samsung represents one of the few manufacturers that designs and manufactures every major SSD component. In addition to manufacturing both the NAND and the proprietary MDX SSD controller, Samsung also manufactures the DRAM cache that its MDX controller uses as a performance buffer. Additionally, having full control over the firmware means Samsung is free to design its own unique performance profile and execute more robust reliability and bug testing. Each and every piece that makes up a Samsung SSD is fully optimized to work seamlessly with its counterparts - a design luxury that produces a product that not only outperforms, but outlasts, most of the competition.

The Best NAND

As a fully integrated manufacturer that controls its own NAND fabrication lines and produces a wide range of memory products, Samsung has the luxury of choosing only the best quality NAND to use in its SSD lineup. A NAND wafer does not contain NAND chips of uniform quality. Samsung takes those chips with the highest quality, performance, and potential longevity and uses them for the most demanding applications – SSDs. The regular quality chips can be used for less demanding applications, like USB storage devices, that do not experience consistent, heavy use.

The Most Experience

Samsung has been a number one player in the semiconductor industry for over 20 years. For 6 of those years, it has been the largest supplier of SSDs to the preinstalled PC business. Designing a storage device like an SSD is no trivial task – experience matters. Samsung's SSD technology has been going through (and passing) the rigorous testing of the world's most demanding clients for years, and its engineers have been learning each and every step of the way. Very few companies can pull together such vast knowledge and experience across every major SSD component – and Samsung's expertise is embodied by its Average Failure Rate (AFR), which is among the industry's lowest at 0.05%. Samsung's SSD lineup is a culmination of everything it has learned and stands for, and customers love the results.

Simple Upgrade Solutions

With the introduction of the 840 Series SSD, Samsung is expanding its focus on ease-of-use. The full upgrade kit includes everything users need to bring their computing experience into the 21st century: a 3.5" adapter bracket for desktop PC users, a SATA to USB adapter cable for Notebook users, and a comprehensive software suite. In just 3 clicks, the new migration software, powered by Clonix, will migrate all of your precious data and your operating system to your new SSD. The fully-redesigned Samsung Magician software has been engineered from the ground up for better usability and brings everything you need to manage your SSD into one coherent interface. With regards to the SSD itself, the performance profile has been designed to deliver next-generation random performance, ensuring that your PC is more responsive than you ever imagined possible. To top it all off, you can rest assured knowing that your SSD will continue

to offer excellent performance throughout its useful lifespan. With its simple upgrade solution and sustained, industry-leading performance, the Samsung SSD 840 is the single best upgrade you can make to your PC.

Unrivaled Performance

For users seeking the most extreme performance available, Samsung offers the 840 Pro SSD. Featuring the industry's fastest Random Performance (up to 100,000 IOPS), users can expect up to about 15 times increase in performance over a typical SATA HDD. (Based-on traced-based benchmark score comparison) Coupled with the industry's most rigorous qualification and reliability testing procedures and an advanced global support network, the 840 Pro is a welcome addition to any power user or professional workflow. Samsung Magician has been enhanced for advanced users as well, with the addition of superior, easy-to-use disk diagnostic and management features with a simple, graphical drive health indicator.

Trust Samsung

Samsung has a long track record of excellence in the memory industry. Its fully-integrated manufacturing approach means its SSDs are among the best performing and most reliable on the market. With great technology, a proven history of performance and reliability, and a strong commitment to putting it all together in a simple package, Samsung is the SSD brand you can trust for your next PC upgrade.



10

The Samsung Advantage

Why a Samsung SSD is Best for Business

With the SSD marketplace becoming increasingly crowded, it can be difficult to decide which vendor to trust with your business' precious data. With 20 years of experience in the semiconductor industry, a unique, fully-integrated design approach, and industry-leading performance and reliability, Samsung is the brand to choose for business PC upgrades.

Industry-Leading Longevity & Reliability

Long-lasting, reliable SSDs mean reduced downtime and lower repair costs

Samsung has been a number one player in the semiconductor industry for over 20 years. For 6 of those years, it has been the largest supplier of SSDs to the preinstalled PC business. Designing an SSD is no trivial task – experience matters. Samsung's SSD technology has been going through (and passing) the rigorous testing of the world's most demanding clients for years, and its engineers have been learning each and every step of the way. As a result, Samsung is known in the industry for its relentless focus on quality, which is directly enabled by one of its most unique competitive advantages – control.

Unlike most SSD vendors, Samsung designs and manufactures every major SSD component. As a fully integrated manufacturer that controls its own NAND fabrication lines, Samsung has the luxury of choosing only the best quality, highest-performing NAND to use in its SSD lineup.

In addition to manufacturing both the NAND and the proprietary MDX SSD controller, Samsung also manufactures the DRAM cache that its MDX controller uses as a performance buffer. Furthermore, having full control over the firmware means Samsung can execute more robust reliability and bug testing. Each and every component that makes up a Samsung SSD is fully optimized to work seamlessly with its counterparts - a design luxury that produces a product that not only outperforms, but outlasts, most of the competition.

Very few companies can pull together such vast knowledge and experience across every major SSD component – and Samsung's expertise is embodied by its Average Failure Rate (AFR), which is among the industry's lowest at 0.05%. Thus, businesses that choose Samsung SSDs benefit from reduced down time and repair costs, fewer data loss incidents, and increased productivity since fewer employees must deal with inconvenient hardware repairs.

Unrivaled Performance with Business in Mind

Extend hardware lifetime with next-generation performance and reduce IT costs with must-have maintenance and security features.

Featuring the industry's fastest Random Performance (up to 100,000 IOPS), users can expect up to about 15times increase in performance over a typical SATA HDD (Based-on traced-based benchmark score comparison). Even older hardware will benefit from dramatic speed improvements, which means savvy businesses can extend the lifetime of their existing hardware rather than making an expensive investment in new machines, saving precious IT investment dollars in a stubbornly volatile economic climate.

With the 840 Series, Samsung has added AES 256-bit Full Disk Encryption (FDE), a feature many business will find a welcome addition. Simply enabling the Hard Disk Password via the BIOS renders data stored on the drives inaccessible to anyone without the proper password. And, because this feature is hardware-based encryption implemented at the controller level, there is no impact on overall drive performance as there is with software-based FDE solutions.

In addition to the performance and security enhancements an 840 Series SSD offers, the 840 was designed to include special, business-specific features. Worldwide Name (WWN) and LED Status Indicator support enable easier management of the drives as part of a storage network or RAID array. Samsung's Magician software has been also enhanced to include superior, easy-to-use disk management and diagnostic features, including simple Usage Statistics with a graphical indicator of Total Bytes Written (TBW). It is now much simpler for IT managers to monitor

drive health and prepare for replacement before a catastrophic drive failure, saving countless hours of frustration and significant repair expenses.

Trust Samsung

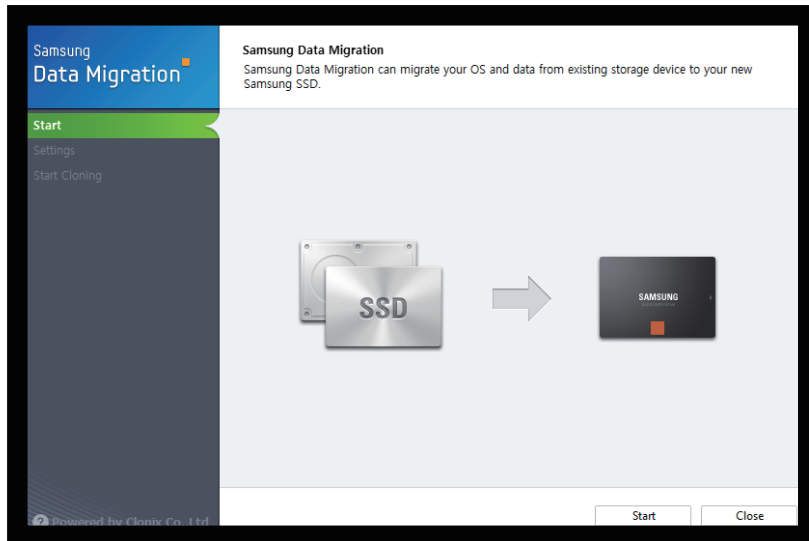
Samsung has a long track record of excellence in the memory industry. With new business-friendly features, redesigned disk management software, and next-generation performance capable of extending the life of aging hardware, Samsung's 840 and 840 Pro Series SSDs represent a solid investment for any business looking to maximize the effectiveness of its IT resources.



11

Samsung Data Migration Software

The simplest way to get your new SSD up and running



Migrate Your Data in 3 Steps


One of the biggest pain points for new users when upgrading from their Hard Disk Drive (HDD) to their new Solid State Drive (SSD) is the process of moving over all of their existing files and programs. Samsung's Data Migration software was designed with novice users in mind to simplify this tedious process. Rather than inundate users with an endless array of configuration options and menus, the Samsung Data Migration (SDM) Software completes this task in 3 simple steps.

Automatic Disk Scan and Diagnosis

The SDM software will scan your system for any connected disks and present your current system status with easy-to-read graphics. SDM hides unnecessary features and partitions from the user, automatically moving the crucial "System Partition" to the SSD and leaving the user free to concentrate on the partitions that contain their operating system, programs, and data.

Recommendation and User Options

Next, SDM will recommend a plan of action and allow the user to customize how the migration will proceed. The first partition, which generally contains the Windows operating system installation, is added automatically. Users can add an additional partition to the migration by simply clicking on its picture. While SDM will automatically partition the drive proportionally, the user can adjust partition sizes to their personal preferences by simply dragging the slide bar in between the two partitions.



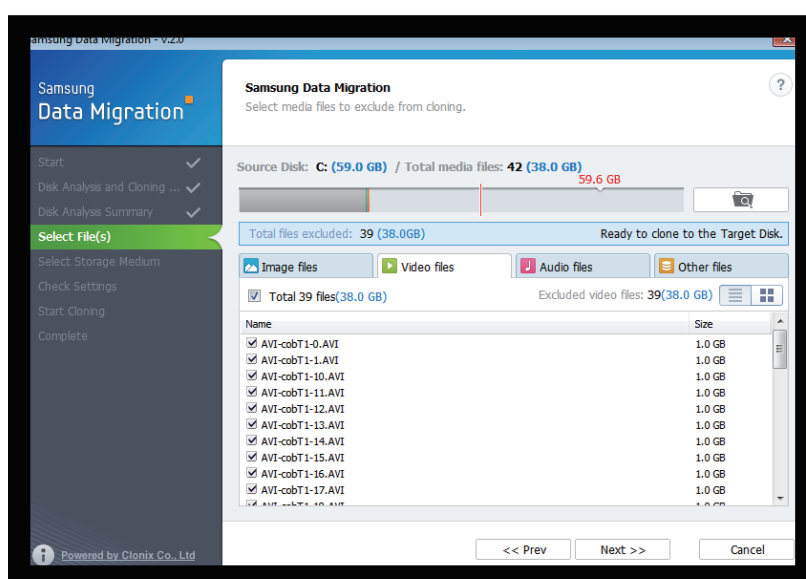
Data Migration

Next, sit back and relax while SDM takes care of the rest. All of the contents from the partition(s) you selected will be

automatically moved to your new SSD. The process time will vary depending on how much data is being copied, and whether you're transferring over SATA or USB, but it will generally complete in less than 1 hour.

New Feature of SDM 2.0 – Custom Cloning

In most cases, PC users use 500GB, 1TB and even more capacity of HDDs while SSD is less than that. If you have a laptop and you want to migrate your data from your 1TB HDD(500GB used) to a 250GB SSD, you'll find that the Target Disk(SSD) is not large enough to accommodate all files on the Source Disk(HDD). "Custom Cloning" feature will help you select unnecessary media files on the Source Disk to exclude from the cloning process.



Using SDM

All you need to get started is your existing computer, your new SSD, and, if you're using a laptop PC, a SATA to USB adapter cable (included if you purchased an SSD with Installation Kit packaging and available online and at major electronics retailers). The SDM software is included on the CD that came with your SSD and available for download at www.samsung.com/samsungssd.

Once the SDM software is installed, connect your new SSD to your PC either via SATA (for desktop users) or via the adapter cable (for laptop users) and launch the software. As it loads, SDM automatically scans the system for the presence of any connected storage devices, so there may be a noticeable delay as the software gets ready.

Click "Start" to begin the disk scan and diagnosis step. On the next screen, click on up to one (1) additional partition that you would like to migrate to the SSD in addition to the "C:" drive that contains your Windows installation. Adjust the partition sizes as necessary, and click "Next."

The SDM software will transfer all of your existing programs, applications, and user settings from the selected drive partitions. When it's finished, all you have to do is shut down your system, swap out the old HDD for the new SSD, and boot up (be sure to enable AHCI in your BIOS if it's not already set up).

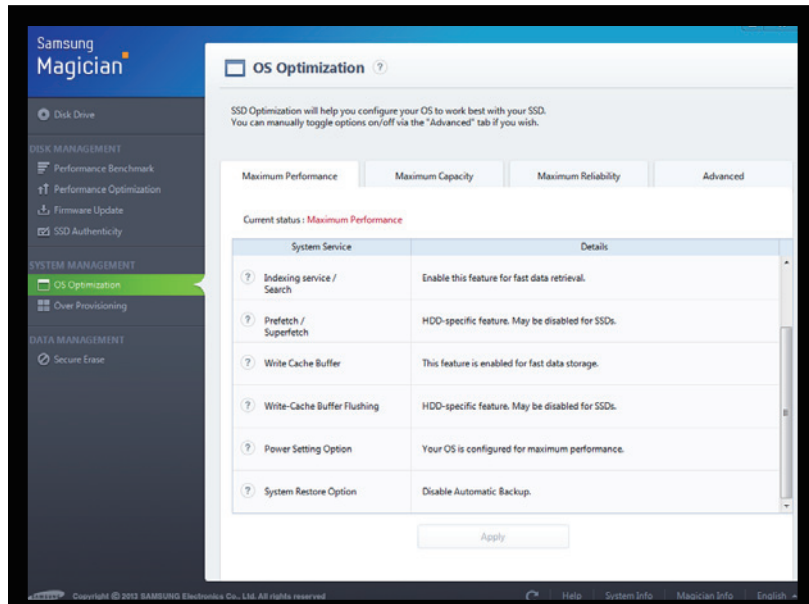
The Single Best PC Upgrade

As simple as that, you've completely transformed your PC into a modern marvel of speed and efficiency. In fact, you may not believe it's the same computer when you see how fast the system boots up. With the help of Samsung's Data Migration software, you're now ready to enjoy industry leading speed, endurance, and reliability from the #1 name in memory with your Samsung SSD.

12

Samsung Magician Software

OS Optimization Feature Overview




Introduction

With Magician 4.0, Samsung is introducing a host of improvements designed to enhance user experience and to make Magician's features more user-friendly and accessible to novices. The totally redesigned User Interface (UI) features simple, graphical indicators to show SSD health status and Total Bytes Written (TBW) at a glance. Users can now determine SATA and AHCI compatibility and status without consulting their user's manuals. The updated Benchmarking feature lets users test other storage devices so that they can clearly see the dramatic performance improvements their new Samsung SSD delivers.

Particularly dramatic improvements have been made to the OS Optimization feature with a simpler, tabbed interface providing access to 3 optimization profiles (Maximum Performance, Maximum Capacity, & Maximum Reliability), along with detailed descriptions of each OS setting and why they are turned on or off for each profile. This feature should reduce user confusion about OS configuration, a common question among new SSD users, and the Magician team is excited to outline what's changed.

Optimization Profiles

Magician's new Optimization Profiles were chosen specifically to address the 3 most common types of user: those seeking to maximize performance, those seeking to maximize storage capacity, and those seeking to maximize drive lifespan (reliability). Each tab will display all relevant OS settings and whether or not they will be toggled on or off for the chosen profile. Additionally, users can see a brief description of a particular OS setting and Samsung's recommended configuration by clicking on the "Help" button (?) to the left of each line item.



The "Maximum Performance" configuration profile sacrifices a small amount of SSD lifespan (or endurance) in favor of

the fastest possible performance. Features like “Hibernation Mode,” “Virtual Memory,” and “indexing,” which require extra writes to the SSD in order to function, are enabled for added speed.

The “Maximum Capacity” profile disables features that require a large amount of extra storage space to function, including “Hibernation Mode,” “Virtual Memory,” and “Automatic Backup,” all of which reserve rather large portions of SSD storage. This configuration is particularly useful for those who purchased a low capacity SSD and want to ensure they have as much space available as possible for their own files without the OS claiming precious space for itself.

“The Maximum Reliability” configuration profile disables all features that produce extra writes to the SSD in order to maximize potential SSD lifespan. It is a well-known fact that SSDs have a limited number of write/erase cycles. While this limit is far beyond any usage level a regular user might encounter, some still prefer to take extra cautionary measures. Features like “Indexing,” “Hibernation Mode,” “Virtual Memory” and “Automatic Backup,” which make frequent writes to the SSD, will all be switched off.

While there may be a slight performance decrease when using the “Maximum Capacity” or “Maximum Reliability” profiles, especially when waking from sleep, SSD technology’s incredible read/write speeds ensure users will still enjoy a noticeable performance and usability improvement over traditional hard disk drive (HDD) technology.

Of course, there is always the option to choose a custom configuration via the “Advanced” tab. Even power users will find this tab useful, as it provides quick access to all relevant OS features either via simple toggle or via a shortcut to the appropriate Windows dialog box.

Explanation of OS Settings

Hibernation Mode

Hibernation Mode enables quick restart times (particularly important for Laptop PCs) by saving the current system status to the hiberfil.sys file (hidden) when the PC power is turned off. Windows 7 or later will automatically reserve up to 75% of the available system memory (DRAM) for the hiberfil.sys file, even if the entire reserved capacity may not be fully utilized. When the system goes into hibernation, the contents in DRAM are copied to NAND flash. Because this produces extra writes and may require a significant amount of storage space to function, this feature is turned OFF for “Maximum Capacity” and “Maximum Reliability” modes. Laptop or Desktop users who frequently leave their machine for a period of time but demand a quick resume time may wish to enable this feature via the “Advanced” tab in Magician, especially if they have a larger SSD.

Virtual Memory

In order to address any potential lack of memory capacity, the Windows OS automatically generates a block of virtual memory (stored in a hidden pagefile.sys file) on the “C:” drive. In the past, before PC Memory (DRAM) was available in high volume, PCs needed to utilize some HDD space to address any memory shortcomings. The Windows OS automatically reserves storage space for Virtual Memory equal to 100% of physical DRAM capacity (e.g. Windows will reserve 4GB of Virtual Memory for a system with 4GB of DRAM or 32GB of capacity for systems using a 64-bit OS and featuring 32GB of DRAM). Reserved space will be dynamically allocated based on resource needs, which often never exceed the available physical memory (DRAM) in recent PCs that feature 4GB of memory or more. Thus, it is possible to reduce or even eliminate the use of virtual memory (with the additional fringe benefit of improving overall system performance and reliability).

The “Maximum Performance” configuration profile uses Windows’ default Virtual Memory settings. Magician sets a minimum of 100MB and a maximum of 1GB of Virtual Memory for the “Maximum Capacity” and “Maximum Reliability” profiles. A minimum amount of Virtual Memory capacity is necessary because some software applications require at least a small amount of Virtual Memory, which is why Magician sets 100MB rather than eliminating it altogether. Additionally, systems running for extended periods (days or weeks at a time) without a restart may run out of Virtual Memory if the capacity is set too low, so Magician sets a maximum of 1GB of Virtual Memory to accommodate this scenario. If your particular system needs require more Virtual Memory, you may adjust settings manually via the “Advanced” tab in Magician.

Indexing Service/Search

Windows automatically indexes the content of data files in order to perform very fast Keyword searches. Users who

frequently perform data searches on their PCs (e.g. business PC users working with Microsoft Outlook or regularly searching for documents) may benefit from turning this feature ON via the “Advanced” tab in Magician to ensure optimal system responsiveness. User who do not frequently search for data on their PCs may turn this feature OFF.

Because it generates a significant amount of random read/write activity, this feature is turned OFF for “Maximum Reliability” mode to reduce SSD writes. It is left on for “Maximum Capacity” mode because the amount of physical space it uses is negligible and outweighed by the potential user benefit.

Users interested in how much capacity their index file is using may view the hidden “windows.edb” file at:

Windows Vista / 7

C:\ProgramData\Microsoft\Search\Data\Applications\Windows\Windows.edb

Windows XP

C:\Documents and Settings\Application Data\All Users\Microsoft\Search\Data\Applications\Windows\Windows.edb

Prefetch/Superfetch

Windows stores common device drivers and frequently used applications in main memory (DRAM) in order to improve performance. By disabling these features, the OS can reduce system memory (DRAM) use. Furthermore, since SSDs have very fast data access times, these features are no longer really necessary. Thus, they are disabled for all profiles.

Write-Cache Buffer

In general, “caching” refers to the pairing of a high-speed storage medium with a lower-speed storage medium to accelerate overall system performance. Applications or data storage devices (SSDs/HDDs) may send a command to the Windows OS to temporarily place data in the faster Cache Memory (DRAM on the SSD’s or HDD’s main board) rather than directly writing it to the slower permanent storage (NAND chips in the case of SSDs or magnetic platters in the case of HDDs). In order to maximize both system and SSD performance for Client PCs, Samsung recommends enabling this feature. If this feature is disabled, there will be significant decrease in Random Read/Write performance. The use of this feature does not affect the life or capacity of the SSD and is thus enabled for all profiles.

Enabling “Write-Cache” does, however, increase the risk of data loss if the system power fails or the drive is removed suddenly. This is because DRAM is a volatile storage medium that loses its contents in the absence of power, which means that data stored in the DRAM cache may not have a chance to be flushed (saved) to storage (NAND) in the event of sudden power loss.

In general, this is an acceptable risk for Client PC systems because they enjoy a lot of idle time and because SSDs write data to NAND very quickly. However, for enterprise systems, which typically run under high load 24/7 and place premium importance upon data integrity and security, this feature is often disabled.

Write-Cache Buffer Flushing

While using the Windows Write-Cache Buffer feature, any data stored in the Cache Memory is lost or damaged in the event of a device failure or sudden loss of power. The Windows Write-Cache Buffer Flushing feature should be enabled to ensure data integrity. In particular, this feature is important if your SSD is being used in a server application. When using an SSD and HDD together, it is always recommended to have this feature ON. Because all configuration profiles use the “Write-Cache Buffer” feature, this system setting is toggled ON for all profiles. Samsung strongly recommends that Write-Cache Buffer Flushing be enabled at all the times. However, if you are more concerned with performance than data integrity, you are welcome to disable this feature to eliminate extra commands being sent to the SSD to flush the data.

Power Setting Options

Samsung SSDs support DIPM (Device Initiated Power Management). Using this feature, SSDs enter a power saving mode when they are idle (i.e. when there is no data to read or write). SSDs installed in Home PCs spend most of their time idle thanks to their extremely fast read and write speeds. Users can manually configure their preferred power setting options via Windows’ “Power Options” dialogue, available via the “Advanced” tab in Magician.

Select either power-saving mode or maximum performance mode, depending on your individual preferences. Laptop PC users typically choose power-saving mode, while Desktop PC users typically choose maximum performance

mode. By default, Magician chooses “High Performance” mode for all configuration profiles because SSDs consume considerably less power than their HDD counterparts and because Samsung SSDs support the DIPM feature discussed above, thus offsetting any extra energy consumption.

System Restore Options

Windows’ default Drive backup and recovery functionality may be used to backup the system, making it possible to restore the system to its original state if problems arise. This feature requires a significant amount of extra storage capacity and produces extra writes, thus it is disabled for all configuration profiles. Users should implement a secondary backup solution using an external storage device to protect themselves from unexpected system failure. Alternatively, users with multiple internal storage devices in their system can redirect the storage location of “System Restore” via the Windows System Restore Dialogue Box, which is accessible via the “Advanced” tab in Magician.

Wrap Up

The “OS Optimization” feature of Magician 4.0 represents only one of many important improvements to Samsung’s Software. Placing advanced OS configuration options in one, convenient, easy-to-use interface helps remove another layer of complexity from the upgrade process. Users can switch between profiles or customize settings at any time, making this one of the most powerful additions to Magician 4.0.



Contributions

Brand Product Marketing Team

Sung-Hee Lee, Director
Woody Yang, Senior Manager
Justin Erfort, Assistant Manager
Kyung-Min Lee, Assistant Manager

Flash Software Development Team

Sungchul Kim, Principal Engineer
Jisoo Kim, Senior Engineer
Jong Tae Park, Senior Engineer
Pilsung Kang, Senior Engineer

Application Engineering Group

JB-JongBaek Chang, Principal Engineer
Jeff Kim, Senior Engineer

Flash Memory Planning Team

Eric Hong-seok Chang, Senior Engineer

Controller Development Team

Myeongeun Hwang, Principal Engineer



www.samsung.com/ssd