Samsung SSD 950 PRO

Upgrade to unmatched performance, power efficiency and endurance

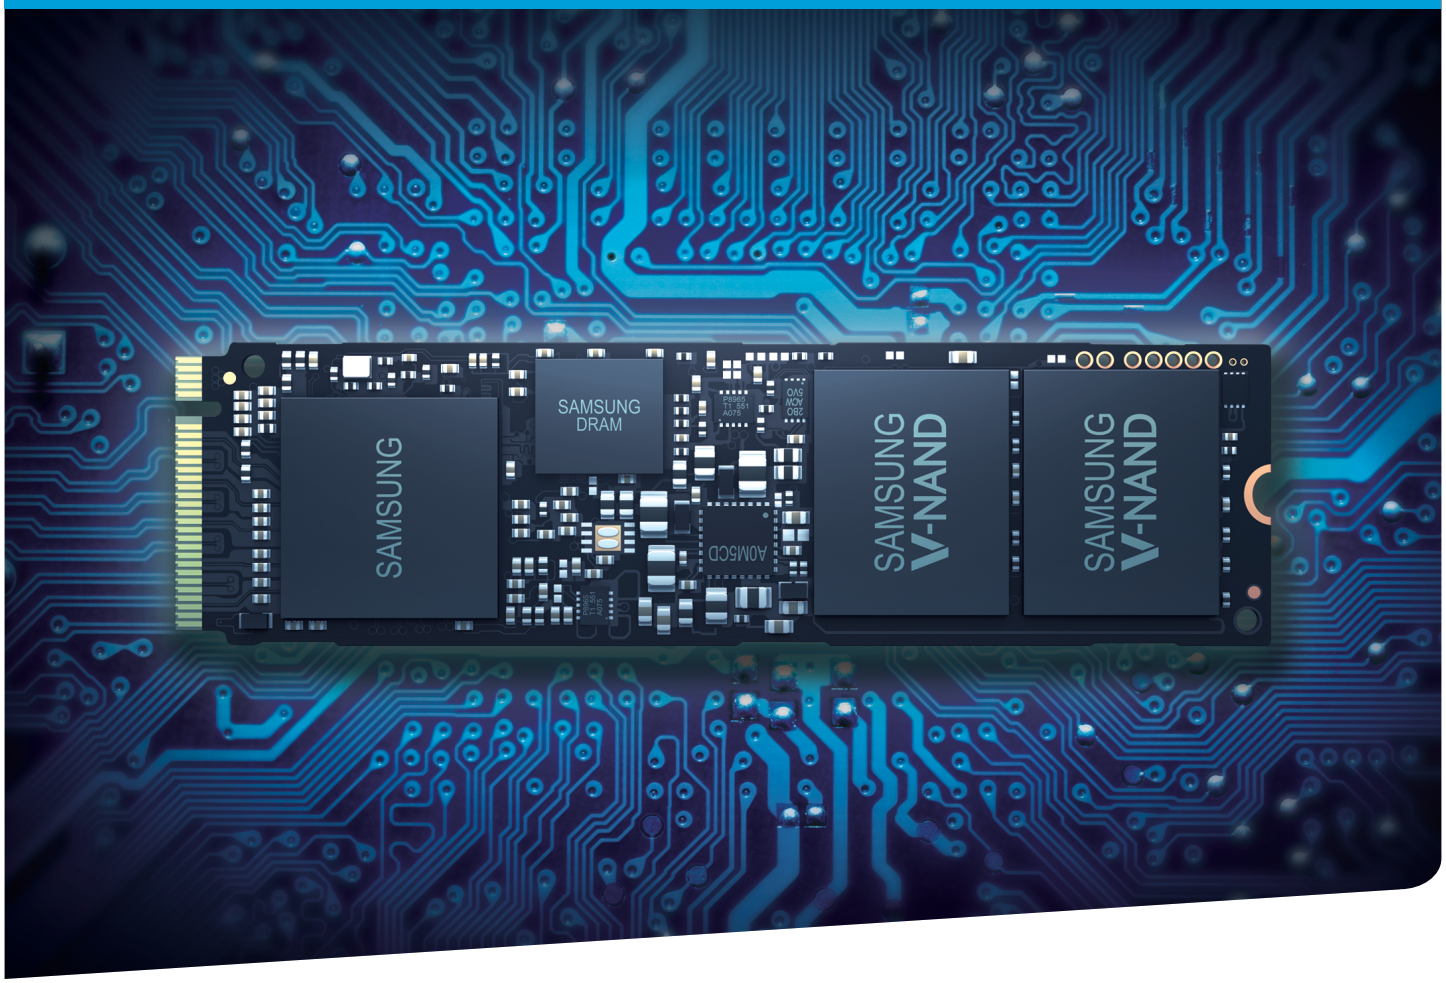

SAMSUNG

Table of Contents

SAMSUNG

Table of Contents

SAMSUNG

Consumers are upgrading to SSDs for enterprise-level performance, power efficiency and endurance

1.Executive summary

Today's consumers are demanding the same level of performance, power efficiency and endurance for their computing needs as typically found only in enterprise-grade equipment. Increasingly they are turning to solid state drives (SSDs) to address their needs.

Samsung 950 PRO is an ultra-fast SSD with first-in-its-class V-NAND technology and a new generation Non-Volatile Memory Express (NVMe) controller, which supports Peripheral Component Interconnect Express (PCI) Express® (PCIe®) Gen 3 x4 lanes. It is designed specifically to cater to the needs of client applications with features equivalent to enterprise-grade SSDs.

1.1. Introduction and hypothesis

The NVMe controller interface used in the Samsung 950 PRO has a variety of enhancements over the traditional Advanced Host Controller Interface (AHCI). Enhancements include:

- Support for 8 queues with 64K command queue depths in each queue
- Low latency for each command completion cycle
- MSI-X interrupts for effective utilization of all the cores in an I/O cycle
- No locking mechanism for parallel operation of multiple I/O threads

Consequently, the 950 PRO is able to achieve sequential read and write (R/W) performances up to 2,500 MB/s and 1,500 MB/s, respectively, while consuming peak power as low as 5.5 W and 5.7 W, respectively. Ultra-low power consumption as low as 2.5 mW, in active as well as idle state, results from the adoption of the latest PCIe L1.2 low power standby mode, finally proving that a perfect balance can be achieved between performance and power consumption.

Not only does the 950 PRO offer high sequential R/W performances, but it also provides consistently high random R/W performances of up to 300K/110K input/output operations per second (IOPS) respectively, with very low latency.

Additionally, the 950 PRO offers the following advantages:

- **Option ROM**, which uses the SSD as a primary boot device in UEFI BIOS
- **Proprietary NVM Express (NVMe) driver** for higher performance
- **Self-Monitoring Analysis and Reporting Technology (S.M.A.R.T.)** for bad block detection and increased endurance levels, providing high lifetime cycles.

This white paper describes how the Samsung 950 PRO provides an exceptional user experience and outclasses competitor PCIe or SATA SSDs in performance, power consumption, form factor, endurance and reliability when used in client systems.

1.2. Industry trends

The SSD is becoming an increasingly popular alternative to traditional hard disk drive (HDD) storage. A hard disk employs mechanical arm movements to read data from a specific location. Conversely, data is saved and accessed electronically in an SSD without involving any mechanical movement, thus allowing it to perform faster and be more durable and reliable. Although initially designed for enterprise and data center storage, SSDs are consistently making their way into client devices, such as PCs, notebooks, tablets, and cameras. Figure 1 shows the difference between HDDs and SSDs, and the important areas in which SSDs bring significant improvements.

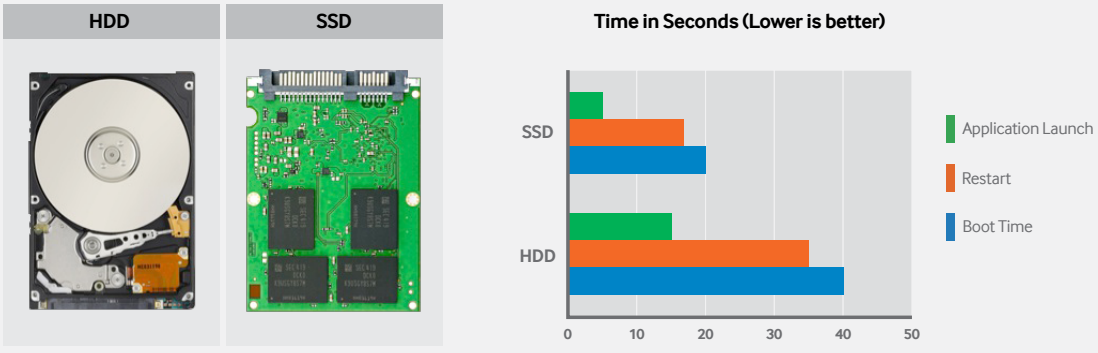

Figure 1. Advantages of the SSD compared with the HDD
(measurements are indicative and not accurate)

SAMSUNG

SATA and PCIe interfaces are gaining popularity for their high bandwidths and robust features

SSD technology continues to advance at a very fast pace and is currently the leading technology in terms of IOPS per dollar, as well as IOPS per watt. SSDs can replace a large number of 7,200/15,000 RPM drives with their significantly low power consumption when compared to spinning drives for a given number of transactions.

1.3. SSD interfaces (SATA versus PCIe)

Today, SSDs are available with a variety of system interfaces based on the performance requirement of the applications. Common interfaces include Serial ATA (SATA), Fibre Channel (FC), Serial Attached SCSI (SAS), Advanced Technology Attachment (ATA/IDE) and PCIe. Of these, the most popular interfaces, SATA and PCIe, are shown in Figure 2.

Figure 2. The SATA versus the PCIe interface

1.3.1. SATA

SATA is an evolution of the Parallel ATA (PATA) interface. SATA offers several advantages over the PATA standard, such as native hot swapping and faster data transfers through an I/O queuing protocol. By replacing the mode set from a slower ATA to a faster AHCI, it is able to introduce advanced features like a queued TRIM command, Native Command Queuing (NCQ) and hot-plugging with the same physical connector. Since its introduction in 2003, it has gone through three major revisions, resulting in a quadrupling of bandwidth from 1.5 Gb/s to 6 Gb/s as shown in Table 1.[1]

Generation	Speed	Special Features
SATA 1.0	1.5 Gb/s (150 MB/s)	
SATA 2.0	3 Gb/s (300 MB/s)	Native Command Queuing (NCQ) support
SATA 3.0	6 Gb/s (600 MB/s)	NCQ management feature, Improved power management, Support for new connectors, TRIM command
SATA 3.1	6 Gb/s (600 MB/s)	mSATA support

Table 1. SATA speeds

1.3.2. PCI Express (PCIe)

PCIe is a high-speed serial computer expansion bus standard with numerous improvements over its predecessors, such as a higher maximum system bus throughput, lower I/O pin count, smaller physical footprint and better performance scaling for bus devices[2]. The PCIe technology is undergoing constant development and improvements and as of 2015, it has reached its fourth version, multiplying its bandwidth to almost 8 times greater than the first version

Generation	Raw Bit Rate	Interconnect Bandwidth	Bandwidth/Lane Per Direction	Total Bandwidth for a x16 Lane
PCIe Gen1	2.5 GT/s	2 Gb/s	~250 MB/s	~8 GB/s
PCIe Gen2	5 GT/s	4 Gb/s	~500 MB/s	~16 GB/s
PCIe Gen3	8 GT/s	8 Gb/s	~1 GB/s	~32 GB/s

Table2. PCIe speeds

NVMe lifts the limitations of bandwidth bottlenecks with its scalable host controller interface for PCIe-based SSDs

A PCIe link between two devices consists of one or more lanes, which are dual-simplex channels composed of two differential signaling pairs. Physical PCIe links may contain from 1 to 32 lanes[3]. The PCIe interface has the ability to aggregate multiple individual lanes to form a single link. For example, two single lanes (x1) can be combined to form a single link capable of transmitting double the bandwidth of a single lane. Likewise, x4 or x8 lanes can be formed. Figure 3 shows examples of PCIe connectors with different lanes and their associated bandwidth.

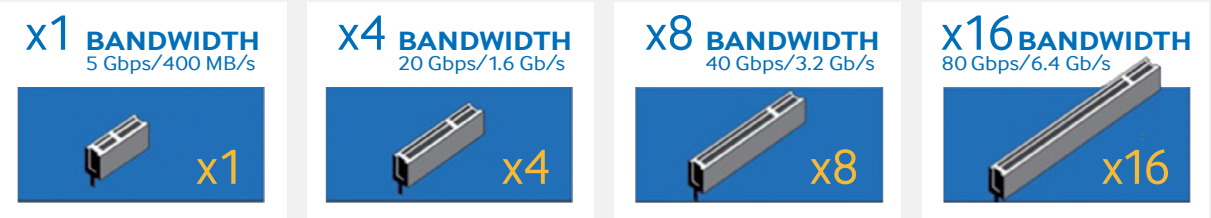

Figure 3. PCIe connector examples

1.4. AHCI/SATA interface limitations

Although the AHCI mode set for SATA fulfilled its intended architecture and design goals quite well for devices such as HDDs and optical drives, it performs inefficiently when applied to SSD technology. The main reason is that the storage of data in an SSD is different from that of spinning media; an SSD bears more resemblance to system memory (DRAM). To understand this clearly we will dig a little deeper into the design and storage method of an SSD.

In an SSD, the information that is stored on the device actually resides on NAND chips. There is a primary SSD controller that branches out to NAND chips via multiple lanes. The NAND chips have a rated speed at which they perform and there can be several of these chips for increased capacity. More chips connected to the controller via multiple lanes allow faster data transfer from the device. Figure 4 illustrates the SSD architecture at a very high level.

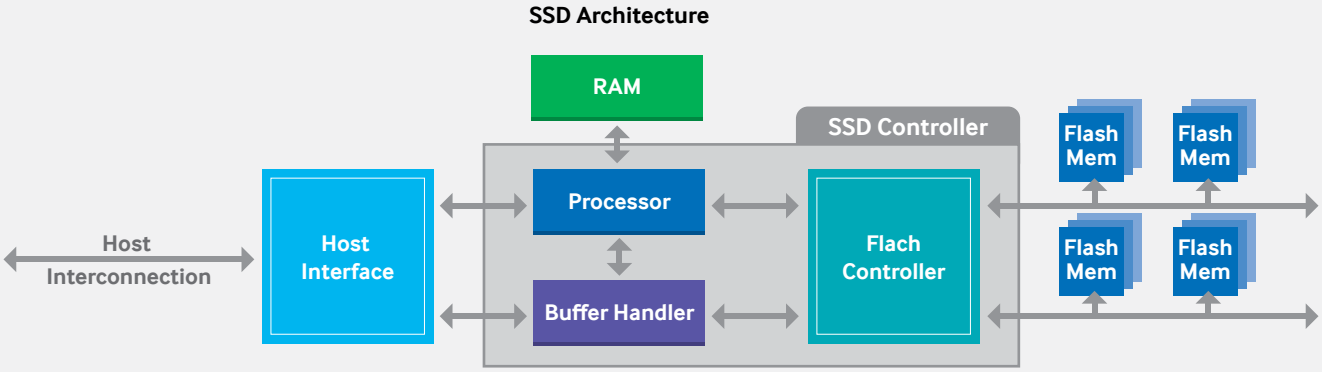

Figure 4. SSD architecture

Earlier NAND chips had speeds of approximately 50 MB/s per chip, with usually 4 to 8 chips on a typical SSD, accounting for speeds up to 400 MB/s (8 x 50 MB/s). Since the current SATA 3.0 has speeds up to 6 Gb/s, it allows SSDs to reach a maximum speed of up to 600 MB/s, which is quite sufficient for the older NAND chips (with max speed up to 400 MB/s).

However, current NAND chips in consumer SSDs usually run at around 200 MB/s[4]. With the advent of the latest NAND chips like AT&T® Toggle® 2.x and Open NAND Flash Interface (ONFI) 3.x, PCIe SSDs are capable of delivering speeds up to 400 MB/s[4] per chip, resulting in an overall speed up to 3.2 GB/s. This means that the device is capable of delivering a higher bandwidth, but the interface itself is imposing a limitation on the available bandwidth. One can easily see how a 6 Gb/s SATA has become a bottleneck in deriving the maximum possible throughput from current SSDs.

1.5. NVMe as a long-term solution

NVMe is a scalable host controller interface designed to address the needs of enterprise data centers and client systems that utilize PCIe-based SSDs. Prior to NVMe,

The NVMe structure enables the establishment of multiple submission and completion queues to avoid locks

various SSD vendors used different implementations and their own unique driver designs for PCIe SSDs, thereby hindering their wide-scale adoption. Hence, to enable faster adoption and interoperability of PCIe SSDs, industry leaders have defined the NVMe standard[5]. NVMe capitalizes on the low latency and parallelism of PCIe SSDs, mirroring the parallelism of contemporary CPUs, platforms and applications. Figure 5 illustrates the different interconnection methodologies of AHCI and NVMe interfaces.

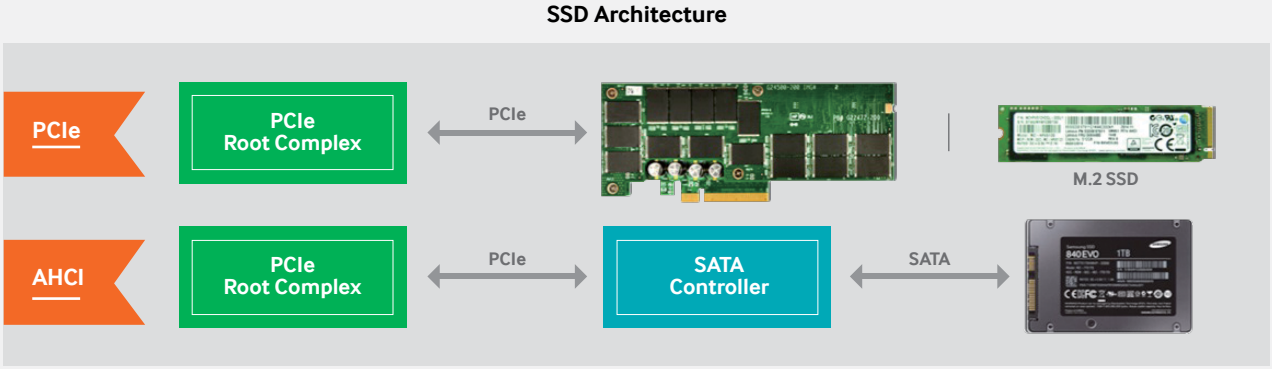

Figure 5. The AHCI versus the PCIe interface

A SSATA SSD is connected to a SATA/AHCI controller, which is then connected to the PCIe root complex. But PCIe-based SSDs are connected directly to the PCIe root complex or CPU as shown in Figure 5, thereby eliminating any intermediate protocols and allowing for faster data flow and processing.

Table 3 lists the important differences between AHCI and NVMe standards.

	AHCI	NVMe
Latency	6 us	2.8 us
Maximum queue & queue depth	1 command queue 32 commands per queue	64K queues 64K commands per queue
Uncacheable register access (each consumes 2K cycles)	6 per non-queued command 9 per queued command	2 per command
MSI-X and interrupt steering	Single Interrupt	2K MSI-X Interrupts
Parallelism and multi-thread support	Requires synchronization lock to issue command	No locking, doorbell register per queue
4 KB command efficiency (4 KB critical in client PCs)	Two serialized host DRAM fetches required	One 64 B fetch

Table 3. Differences between AHCI and NVMe standards

While both AHCI and NVMe interfaces support parallelism, the way each of them provides it is highly influenced by their differing design goals. NVMe was designed to harness the parallelism of PCIe SSDs through multiple cores using threads and I/O paths of the operating systems. The communication between a host and an NVMe controller is based on a series of paired submission and completion queues built by the driver and shared between the driver running on the host and the device. A pair of admin queues is used for the controller management operations, and submission/completion queues are set up for I/O processing per core.

NVMe supports up to 64K command submission/completion queue pairs, as well as multiple command submission queues where command status is placed on a common command completion queue. This structure of NVMe allows the platform itself to establish multiple independent I/O channels directly to the device itself as shown in Figure 6. Here, common applications on client systems can benefit by establishing parallel connection on a per core basis, resulting in a smaller number of interrupts to core 0. The benefits of NVMe architecture compared to AHCI are explained in the following sections.

The NVMe structure enables the establishment of multiple submission and completion queues to avoid locks

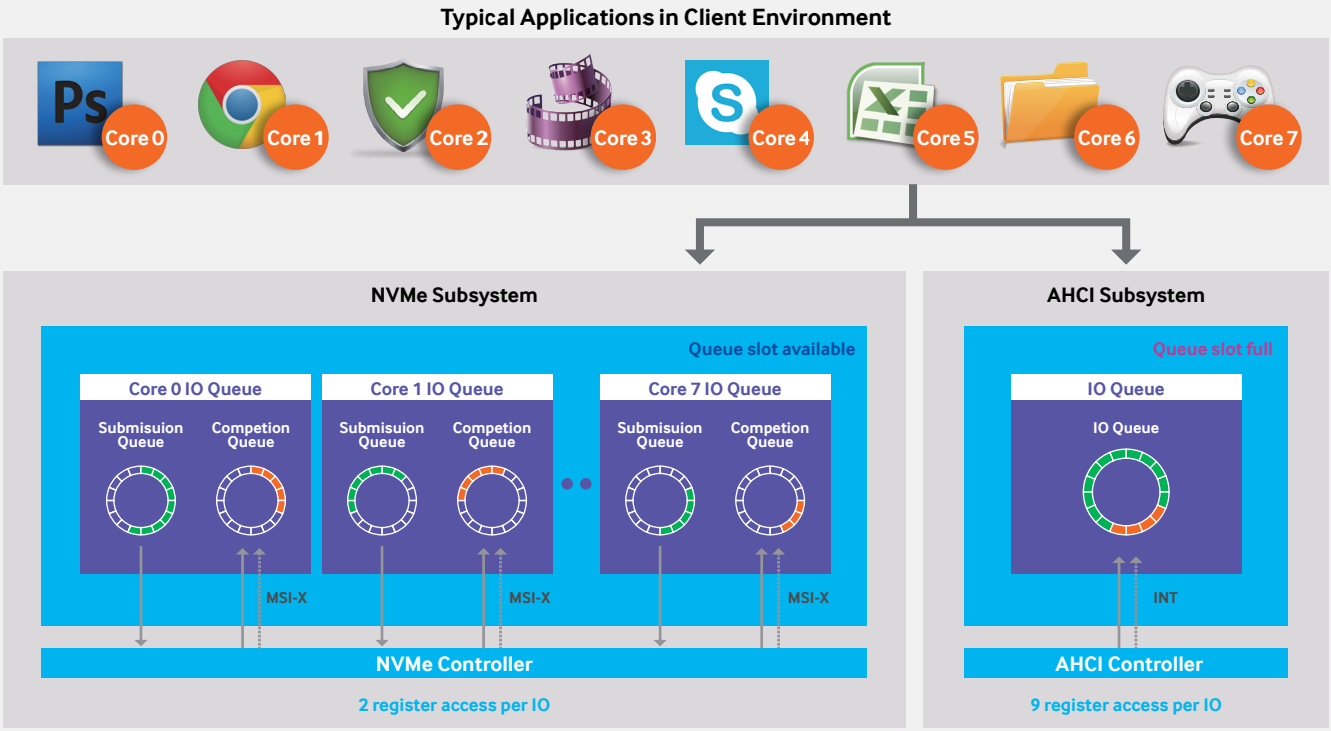

Figure 6. NVMe queues per core

1.5.1. 64K queues in NVMe versus one queue in AHCI

The advantage of having many submission and completion queues in NVMe is fairly obvious. On multiprocessor systems, the ability to have multiple queues allows the user to have one pair of submission and completion queues per processor core. Having a unique pair of queues per core permits the NVMe driver to process I/O initiations and completions without acquiring and holding any locks. Furthermore, NVMe strongly supports the use of Message Signaled Interrupt-Extended (MSI-X) interrupts, which enables the NVMe device to direct its interrupts to a particular core. With the help of MSI-X interrupts, a request that is submitted by the NVMe driver on a particular core will result in an interrupt through a completion queue unique to that core, indicating its completion.

AHCI, on the other hand, has only one queue shared with all the available cores. With the maximum number of outstanding I/Os (queue depth) being only 32, the queue gets full quickly and the applications need to wait before sending data for the queue to become empty. Moreover, the interrupts for any I/O completion are by default sent to core 0, thus causing it to perform intermittent operations and throttling the maximum throughput.

1.5.2. 64K command queue depth in NVMe versus 32 queue depth in AHCI

Today's PCs with high-end processors are capable of driving hundreds of thousands of IOPS. The load placed by such a system is too huge for the device to keep pace, resulting in the need for I/O queues with larger queue depths. If the queue depth is less, the system will spend most of its time replenishing the command queue in order to keep the device busy. As storage devices get faster consistently, the processor has to return repeatedly to replenish the command queue, resulting in waste of time. To eliminate this bottleneck, deep queues were needed, which take more time to be drained, allowing the processor to spend time on other important tasks.

1.5.3. Two register R/Ws in NVMe versus nine R/Ws in an AHCI

To submit and process a completed command, an AHCI requires nine-register R/W if the operation is a native command queue (NCQ) operation and six-register R/W if the command operation is a non-NCQ operation. In contrast, NVMe requires just two register writes for the command issue/completion cycle. In memory-mapped I/O access, read is non-posted, which means it takes longer to read the register, whereas write is a posted transaction.

- **Command submission**: Write the new value of the submission queue head pointer to the submission queue doorbell register.

In queue processing, queue depth, latency and register R/Ws, NVMe shows superior performance over the AHCI

- **Command completion:** Write the new value of the completion queue head pointer to the command completion queue doorbell register.

Higher register access overhead in an AHCI results in higher latencies, which can work well with the relatively slower command issue/completion rate of SATA HDDs. However, with the advent of PCIe SSDs that require the lowest latency possible, an interface with a good queue management protocol was required. NVMe is the perfect solution to that requirement.

1.6. Samsung 950 PRO for ultra-fast speeds on client systems

The 950 PRO is a high-performance SSD with low power consumption, and efficiently designed for use in both ultra-slim notebook PCs and workstations. It features an M.2 form factor, which is an improved revision to the legacy mSATA connector design. The M.2 has a smaller and more flexible physical specification along with some advanced features, which makes it a more suitable candidate for solid-state storage applications. Through the M.2 port, the 950 PRO is capable of delivering sequential R/W speeds up to 2,500 MB/s and 1,500 MB/s respectively, and random R/W speeds up to 300K IOPS and 110K IOPS respectively.

In general, the industry-standard M.2 slots for SSDs rely on PCIe Gen 2 x1, x2, x4 and PCIe Gen 3 x4 lanes. Samsung recommends using the 950 PRO with new M.2 ports based on PCIe Gen 3 x4 lanes to realize the full potential of the SSD. Figure 7 illustrates how one can distinguish between the latest M.2 ports with PCIe Gen 3 lanes (shown as Ultra M.2) and typical M.2 ports with PCI Gen 2 (shown as M.2 x4).

Figure 7. PCIe Gen 3 x4 versus Gen 2 x4 M.2 slots

2. Understanding client system requirements

The requirements of client-based SSDs differ greatly from those of enterprise-class and data center SSDs and hence demand discrete attention. Client SSDs typically serve a single system with a single user running various client software applications for approximately eight hours per day, seven days a week, on small datasets, mostly in a read-intensive manner. In contrast, enterprise SSDs tend to run round the clock every day of the year, with multiple parallel users performing write-intensive operations on much bigger datasets. Quite obviously, the design of an SSD and its included features differs vastly in both the segments.

Major concerns in client-grade systems are low power consumption, high performance and high endurance. This section discusses these important performance factors in a client environment and how the design of the 950 PRO addresses these concerns.

2.1. Low power consumption

Today's consumers demand power-efficient storage and speed in their devices. Consumer SSDs are used in laptops, desktops and mobile devices, where preserving power is an important criterion. For typical home laptop and desktop users, workloads may range from lightweight applications like storage of email, data, pictures, music and other digital content to heavy workloads like running development applications (Microsoft Visual Studio®, Adobe® Photoshop®, and so on), video editing, web hosting, high-resolution gaming applications, and so on.

The 950 PRO boasts low power consumption without sacrificing consistently high performance throughout its lifecycle

Because an SSD is very fast, it is either in the idle or near-idle state for the majority of these applications. Therefore, low power consumption in the idle state becomes an important factor on client systems. Keeping this in mind, the 950 PRO was designed to consume only 2.5 mW of power in the idle state and just 5.7 W even at peak workloads. As a result, it delivers a longer run time before requiring a charge, making it a perfect candidate for client applications.

2.2. High performance

Typical applications in a client environment mainly involve many small random transactions of 4 KB/8 KB read or write operations. As these transaction numbers range in the thousands, latency for each transaction becomes a very important factor in client systems. The 950 PRO has been specifically tuned to consistently provide very high transaction rates in the range of 300K IOPS for random read operations and 110K IOPS for random write operations with latency as low as 3.3 microseconds. Moreover, for small random reads (4 KB) with a low queue depth (queue depth 1), the 950 PRO delivers as many as 13.9K IOPS as compared to a mere 100 IOPS from traditional hard disk storage. For sequential workloads, it clocks speed up to 2,500 MB/s for read and 1,500 MB/s for write operations. It also features dramatic improvements in time taken for Microsoft Windows® startup, application loading, importing pictures, video editing, gaming and file copying.

2.3. High endurance

Consumer data is important and thus an SSD should be able to sustain high levels of load for a considerable duration of time. The 950 PRO has been tested for endurance on client systems with very heavy workloads to ensure that the customer receives enterprise-grade durability at the price of a consumer-grade SSD. The 512 GB 950 PRO is rated for 110 GB of writes per day for 5 years, which is approximately 200 TB of total writes. Moreover, the active use of the TRIM command allow the SSD to maintain high performance levels throughout its lifecycle.

Samsung V-NAND technology provides more capacity, better performance, outstanding endurance and superb power efficiency

3.Samsung 950 PRO characteristics

	Samsung 950 PRO
Usage applications	Client PCs
Form factor	M.2 (22 x 80)
Capacity	256/512 GB
Host interface	PCIe Gen 3 x4 Lanes (up to 32 Gb/s) NVMe 1.1b
NAND flash memory	V-NAND
L1.2 Power-saving mode	Supported
Bootable (UEFI Option ROM)	Supported
Power consumption	Active R/W: 5.7 W, Idle: 70 mA
TBW	256 GB : 200 TBW / 512 GB : 400 TBW
Sequential R/W (MB/s)	2,500/1,500
Random R/W (IOPs)	300K/110K
Physical dimensions	22 x 80 (2.3 mm T)
Weight	Under 10 g
Warranty	5 Year Limited Warranty

Table 4. Product specifications

3.1. 2Bit/Cell Samsung V-NAND flash memory

The 950 PRO is equipped with the Samsung second generation . By stacking 32 layers of 2-bit memory cells vertically in a three-dimensional structure, new potential for 3D memory capacities is created. This cutting-edge V-NAND technology eliminates performance and reliability issues resulting from capacity limitations found in 2D planar NAND technologies. The Samsung V-NAND technology delivers reliable and consistent performance at lower costs for today's demanding data-centric world. This innovative architecture improves memory in the following essential areas[6]:

- **More capacity**: Fits more memory cells in a NAND chip in less space to provide significantly more capacity
- **Better performance**: Writes data much faster by virtually eliminating cell-to-cell interference
- **Outstanding endurance**: Experiences less stress with insulators more resistant to wear to provide greater endurance
- **Superb power efficiency**: Reduces power consumption by reducing the number of programming steps

Over the past 15 years, the NAND flash memory cell structure has gone from a 120 nm scale to a 19 nm scale whereas capacity has grown by 100 times. Shrinking is a fundamental technological advancement in almost any field, but it gains significant importance in memory engineering. As memory structures shrink, so does the multitude of mobile devices used by people worldwide on a daily basis, which eventually leads to the two biggest problems inherent in shrinking technology, cell-to-cell interference and patterning.

3.1.1. Cell-to-cell interference

When an electric charge flows into one cell, some of that charge flows into a neighboring cell, too (known as the coupling effect). This extraneous charge to the neighboring cell changes the stored data, resulting in data corruption. This interference does not occur when the space between cells is greater than 30 nm, but as that space shrinks smaller than 20 nm, which is the case in most of the SSDs today, the chance for interference drastically increases, inevitably making the cell unreliable.

The 950 PRO addresses this issue by adopting Charge Trap Flash (CTF) technology, which results in a cell-to-cell interference-free structure. In 2D planar NAND memory, the cell stores the electron in the conductor. With V-NAND memory, the conductor is replaced by an insulating layer of silicon nitride (SiN), which temporarily traps electrical charges to maintain cell integrity and thereby lowers the coupling. In addition, the V-NAND memory has a larger channel area than planar NAND memory, which improves the initial electron dispersion, creating an almost coupling-free cell structure, making the 950 PRO extremely reliable. Figure 8 illustrates the

A compact form factor, new controller and innovative thermal management enhance the appeal of the 950 PRO

design of the V-NAND and the percentage decrease in natural width/cell coupling with respect to 2D planar NAND:

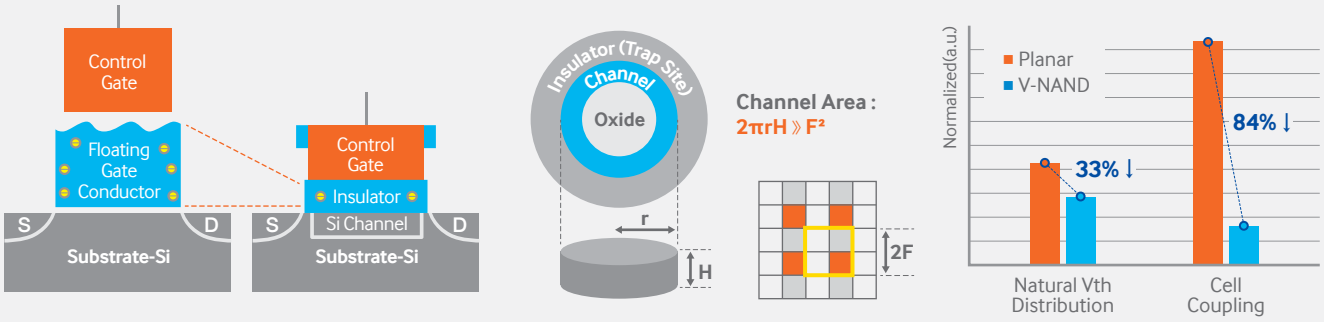

Figure 8. Cell characteristics of V-NAND versus planar NAND memory

3.1.2. Patterning

Patterning is a manufacturing technology developed for photolithography to enhance density. In an effort to shrink the cells to fit more cells in less space in 2D planar NAND flash memory, it becomes difficult for light to penetrate the mask to transfer the pattern to a photoresist. The reduction in light constricts the patterning process, thereby limiting the use of 2D planar NAND flash memory in today's demanding memory environment.

However, the 950 PRO, equipped with V-NAND design, uses the vertical stacking of cells instead of photolithography to increase capacity. This design results in a wider gap between each cell. The cell-to-cell space in a planar 1Ynm NAND typically ranges from 15 nm to 16 nm, but Samsung V-NAND flash has from 30 nm to 40 nm of space between cells, thereby eliminating patterning limitations as depicted in Figure 9.

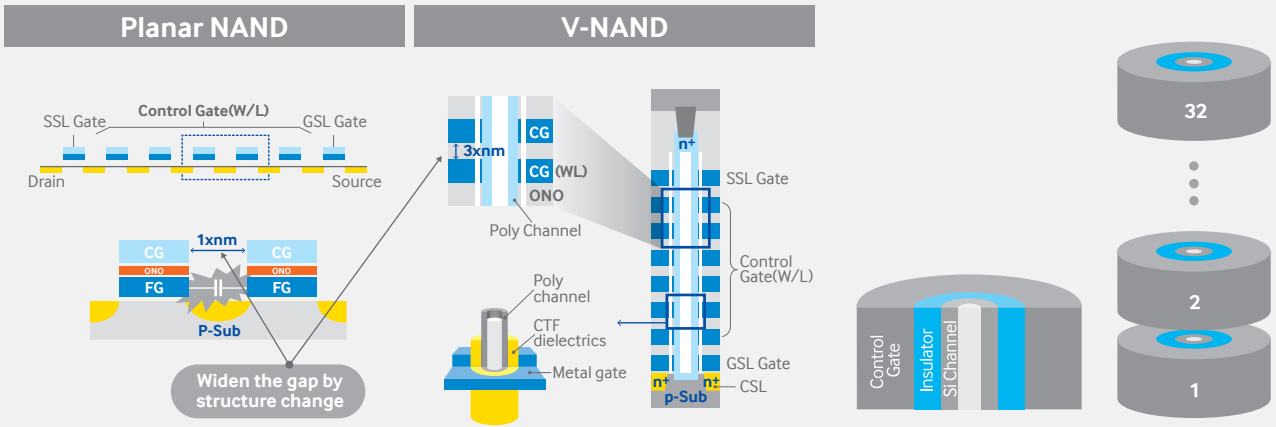

Figure 9. V-NAND versus planar NAND stacking

3.2. Small M.2 form factor

As modern-day computers, particularly laptops and notebooks, continue to get smaller, storage drives also need to get smaller, while continuing to provide the same or a higher level of performance at the same time. The 950 PRO comes in the latest M.2 form factor using PCIe lanes supported universally by almost all vendors in the market today. With just a 22 mm width and 80 mm length, the 950 PRO is almost one-fourth smaller in size when compared to the Samsung PM1725 PCIe SSD, as shown in Figure 10.

A compact form factor, new controller and innovative thermal management enhance the appeal of the 950 PRO

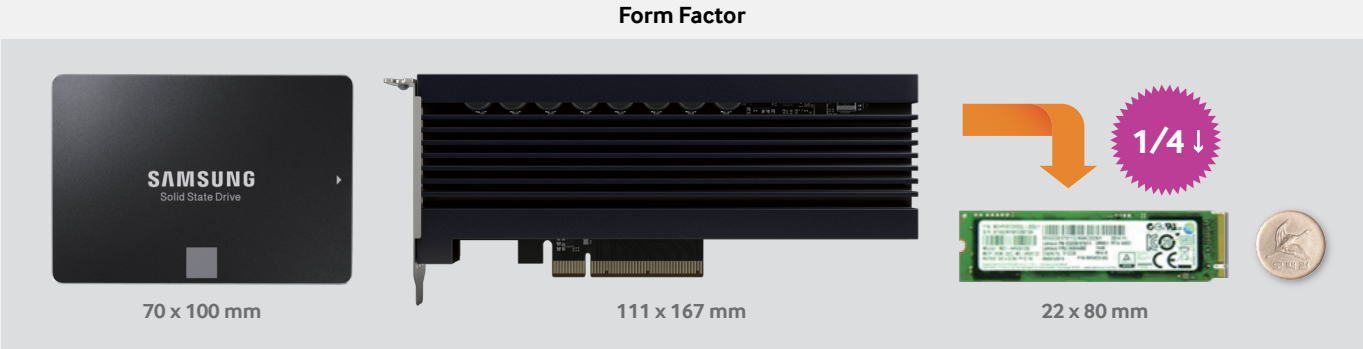

Figure 10. 950 PRO M.2 form factor

3.3. New SSD controller

The Samsung new generation NVMe SSD controller in the 950 PRO supports the following features:

- Thermal management
- Power management
- S.M.A.R.T. features

3.3.1. Thermal management

Although an SSD is considerably fast and performs tasks within a matter of seconds, sometimes when it experiences a large number of requests continuously, its internal temperature may rise beyond a certain threshold and could damage the hardware components. To address this issue, the 950 PRO is equipped with a sensor, which periodically monitors the device's temperature. It also features Dynamic Thermal Throttling (DTT) technology. DTT controls the temperature of the device to ensure it does not exceed its maximum limit by reducing its performance in a phased manner.

Figure 11 depicts DTT stepping technology in effect during sequential write operations on the 950 PRO when responding to increased temperature conditions. The red line denotes temperature levels and the green line corresponds to sequential write performance. Notice that when the 950 PRO exceeds the threshold, it quickly returns to below the threshold level and recovers to its maximum performance level.

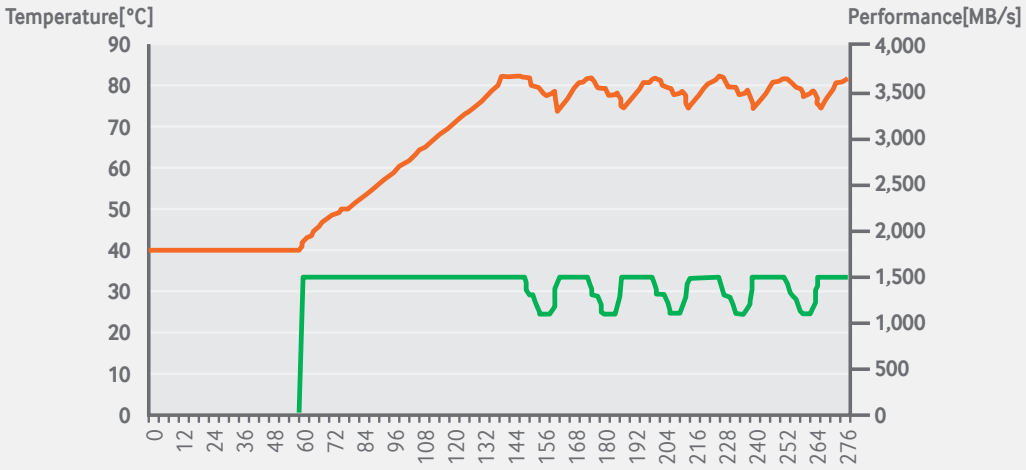

Figure 11. Performance management of the 950 PRO using DTT

Aggressive power-saving features conserve battery life and reduce power consumption in idle, sleep and even active mode

As the continuous ongoing writes increase device temperature, the 950 PRO follows a 10-step performance reduction DTT methodology. With level 0 being the maximum performance state and level 9 being the minimum performance state, it protects itself from overheating while maintaining a good level of performance. The DTT stepping algorithm is activated when the temperature exceeds 165.2°F (74°C), reducing the performance of the 950 PRO by one level, thereby allowing the device to cool down. It measures the temperature level again and, if it is greater than or equal to the previous reading, it goes one step further, reducing the performance by one more level. This step-by-step method is repeated until the device temperature comes to a low steady state, permitting a performance drop only to the extent necessary for recovering a stable temperature, as well as protecting user data and the device. Table 5 illustrates DTT performance levels.

Entry Temperature	Level	Performance
Normal (below 167 °F [75 °C])	0	100 %
DTT entry temperature (167 °F [75 °C])	1	87.5 %
Temperature 167.9 °F (75.5 °C)	2	75 %
Temperature 168.8 °F (76 °C)	3	62.5 %
Temperature 169.7 °F (76.5 °C)	4	50 %
Temperature 170.6 °F (77 °C)	5	37.5 %
Temperature 171.5 °F (77.5 °C)	6	25 %
Temperature 172.4 °F (78 °C)	7	12.5 %
Temperature 173.3 °F (78.5 °C)	8	DTT performance
Temperature 174.2 °F (79 °C)	9	Meltdown (10 MB/s)

Table5. Thermal management of the 950 PRO using DTT

3.3.2. Power management

SSDs in general provide some significant reductions in the energy consumption of the storage component and thus are able to improve battery life to some extent. As more systems are designed to go into sleep mode when closed or turned off, rather than a complete shutdown, there is a constant draw on the battery to keep some essential data active for quick recovery of the system when awakened. To minimize this power consumption, different methods are deployed in various interfaces. In SATA devices, a feature called DevSleep reduces the amount of power used by SSDs in the sleep mode by creating a low power state. In PCIe devices, the Active State Power Management (ASPM) protocol is used to manage less active PCIe serial links for lower power consumption.

ASPM in the 950 PRO follows the Intel guidance of having three operational and two idle (low power) states (see Appendix A). Furthermore, the 950 PRO supports Host-Initiated Power Management (HIPM) using ASPM, which means the host has the ability to choose the active state of PS0, PS1 or PS2 according to the power needs of the system and Autonomous Power State Transitions (APST) for automatic transitions in different power states.

The 950 PRO has adopted the new L1.2 low-power standby mode as defined by the PCIe standards, which allows all high-speed circuits to be turned off to lower the power consumption of the storage device further during sleep mode. By embracing the L1.2 level of standby operation, the power consumption of the 950 PRO has been drastically reduced to approximately the 2.5 mW low power mode. This reduction equates to approximately a 97 percent decrease from the 50 mW consumed using L1 state and 50% less than the 5 mW consumed using SATA DevSleep mode as shown in Figure 12. Moreover, the 950 PRO supports hardware automation for faster wake-up from the L1.2 sleep state, allowing it to wake up in just 700 µs. This will allow laptops and mobile devices to remain in the dormant mode with nearly no battery drain, but to wake up immediately whenever required.

Figure 12. Power management of the 950 PRO using L1.2

Aggressive power-saving features conserve battery life and reduce power consumption in idle, sleep and even active mode

3.3.3. Self-Monitoring, Analysis and Reporting Technology (S.M.A.R.T.)

The Samsung S.M.A.R.T. feature tracks the health status of the SSD by monitoring the SSD to detect and report on various reliability indicators. This technology is designed to anticipate failures and warn users of an impending drive failure, giving the user time to replace an ailing drive to avoid data loss and/or unexpected outages. All data and underlying sectors are tested during periods of inactivity on the device to confirm the health of the drive. Refer to Appendix B for the S.M.A.R.T. features supported by the 950 PRO.

When compared with SATA-based HDDs and SSDs, the 950 PRO set new benchmarks in performance

4. Samsung 950 PRO power and performance

4.1. Performance evaluation environment

For evaluating the SSD's performance, the 950 PRO was used in a non-RAID configuration with no pre-conditioning done on the device in the fresh-out-of-box (FOB) state. In order to accurately assess the device's performance, synthetic as well as real-world workloads were performed on the device. These tests were conducted using inbox NVMe drivers in the Microsoft Windows operating system. The test environment and tools used are listed in Table 6 below.

Processor	Intel Core® i7-4790k @ 4.0 GHz
Memory	Samsung DDR3 8 GB RAM
Motherboard	ASUS™ Z97 PRO
Operating system	Windows 8.1 PRO K x64 with inbox driver
Test suite and workloads	Iometer1.1.0, PCMark® Vantage, PCMark 7
HDD and SSD precondition	Fresh-out-of-box (FOB)

Table 6. Experimental setup and test parameters

4.2. Performance benchmarks

Most of the client applications, such as video streaming and file copying, use sequential I/O patterns. Figure 13 shows sequential R/W performances of the 256 GB 950 PRO compared to one popular HDD and two best-in-class SATA SSDs demonstrated by the tests.

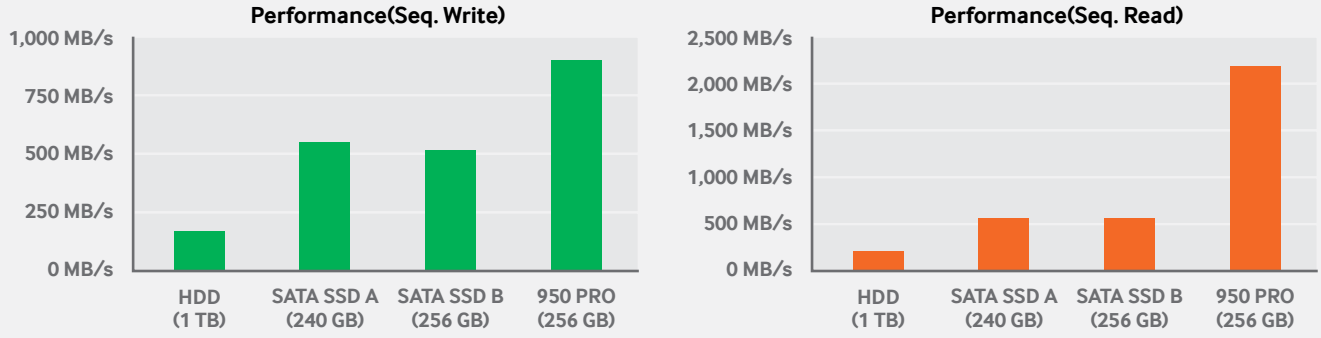

Figure 13. Sequential R/W performance of the 256 GB 950 PRO versus SATA HDD/SSDs

SAMSUNG

When compared with SATA-based HDDs and SSDs, the 950 PRO set new benchmarks in performance

As shown in Figure 13, the 256 GB 950 PRO is almost 100% faster in sequential writes and approximately 300% faster in sequential reads than competitor SATA SSDs. Figure 14 compares the 950 PRO with the density of the 512 GB to similar densities of competitor SSDs. An approximately 200% improvement in sequential writes and approximately 350% improvement in sequential read performance was observed with the 512 GB 950 PRO.

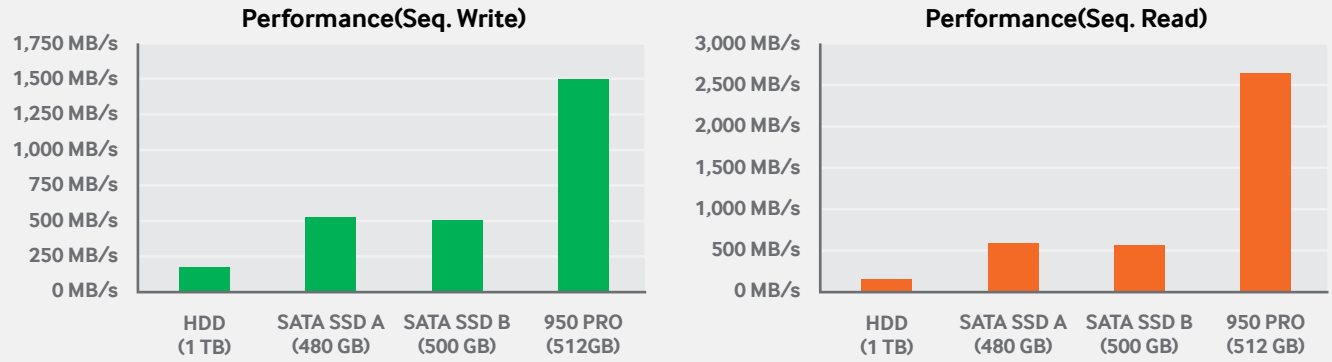

Figure 14. Sequential R/W performance of the 512 GB 950 PRO versus SATA HDD/SSDs

In addition, the 950 PRO consumes maximum power as low as 5.5 W in active state while delivering these staggering performance improvements, which directly contributes to increased battery life in client laptops and notebooks. Figure 15 compares the power consumed during sequential read and write operations by the 256 GB 950 PRO and other SATA HDD/SSDs

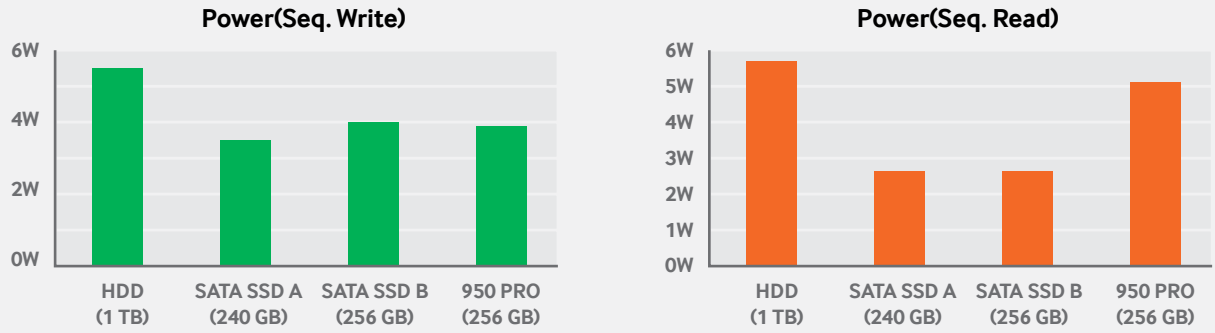

Figure 15. Power consumed during sequential R/W operations of the 256 GB 950 PRO versus SATA HDD/SSDs

Figure 16 shows a similar observation for the 512 GB 950 PRO when compared to storage devices of similar capacities.

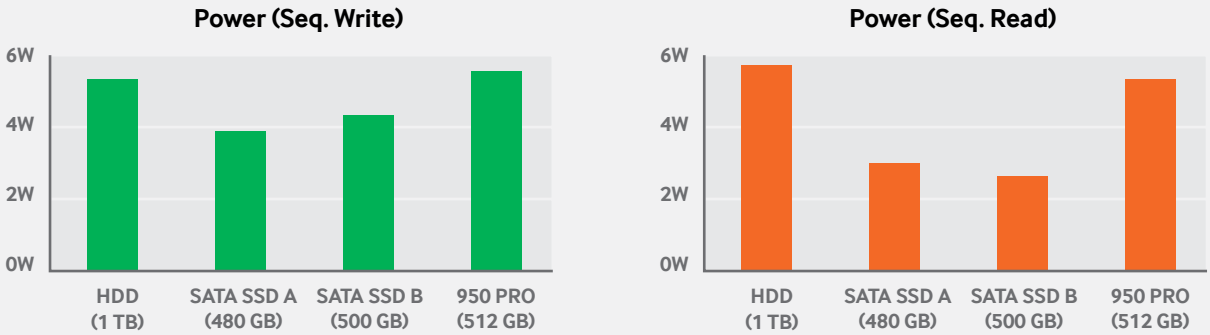

Figure 16. Power consumed during sequential R/W operations of the 512 GB 950 PRO versus SATA HDD/SSDs

The 950 PRO excels in power consumed, performance-to-power ratios and random read performance against SATA-based drives

Figure 15 and Figure 16 clearly delineate that the 950 PRO delivers higher performance than SATA drives while consuming almost the same or less power. This comparison concludes that the 950 PRO provides a superior performance-to-power ratio drive as shown in Figure 17 and Figure 18.

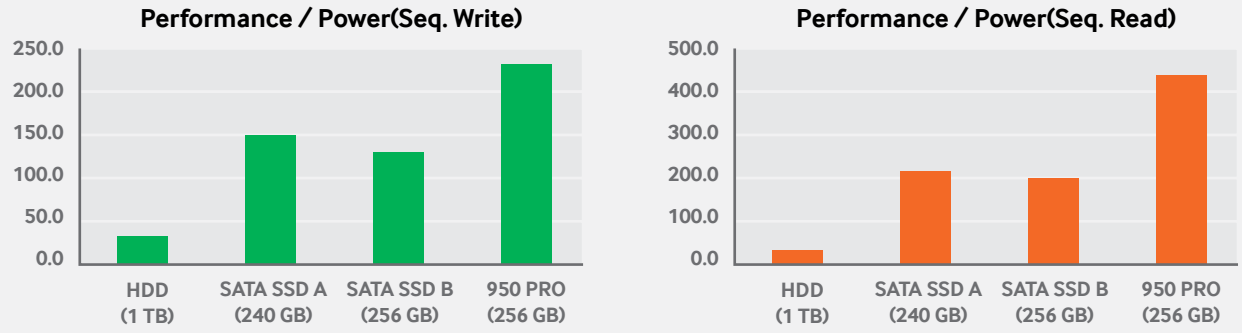

Figure 17. Performance-to-power ratio of the 256 GB 950 PRO versus SATA HDD/SSDs

Figure 18. Performance-to-power ratio of the 512 GB 950 PRO versus SATA HDD/SSDs

Apart from providing staggering sequential R/W speeds, the 950 PRO also provides consistently high random performance. As the majority of the client workload relies on small random reads (4 KB) with a low queue depth, the 950 PRO has been specifically tuned to deliver high IOPS in low queue depth. Figure 19 compares the random read performance of the 256 GB 950 PRO with respect to an HDD and SATA-based SSDs of similar capacities.

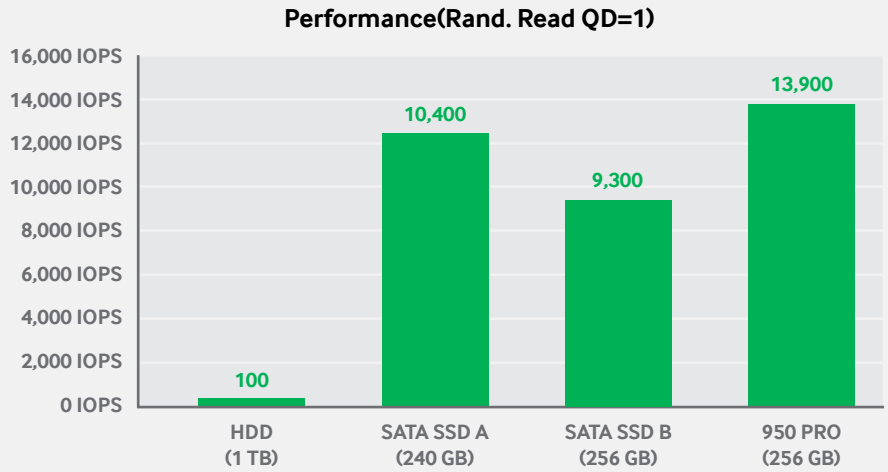

Figure 19. Random read (4 KB, QD1) performance of the 256 GB 950 PRO versus SATA HDD/SSDs

In the user experience environment, the 950 PRO outperformed the competition by a considerable margin

4.3 User experience performance

To evaluate the performance of the 950 PRO on applications commonly used in a client environment, the popular benchmark tool PCMark was used. PCMark runs traces of popular applications on a client system to report an overall system score. This score can be compared to other SSD benchmark scores run on a similar test set to evaluate an SSD's performance. The 950 PRO was compared to a traditional HDD and two SATA-based SSDs. The results are shown in Figure 20 and Figure 21. It can be clearly seen that the 950 PRO outperforms other comparable devices by a considerable margin.

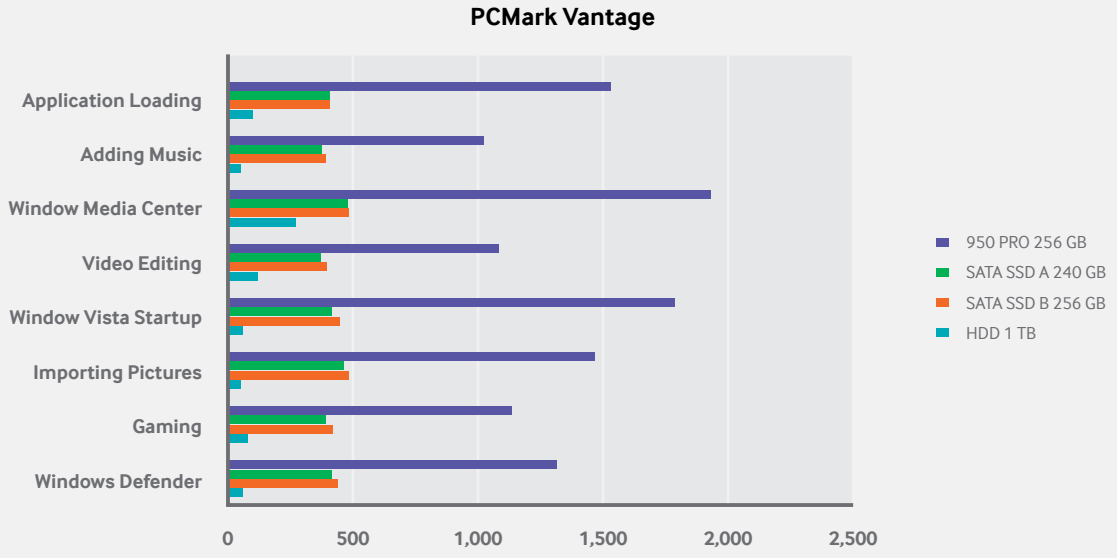

Figure 20. PCMark Vantage score comparison (950 PRO versus SATA HDD/SSDs)

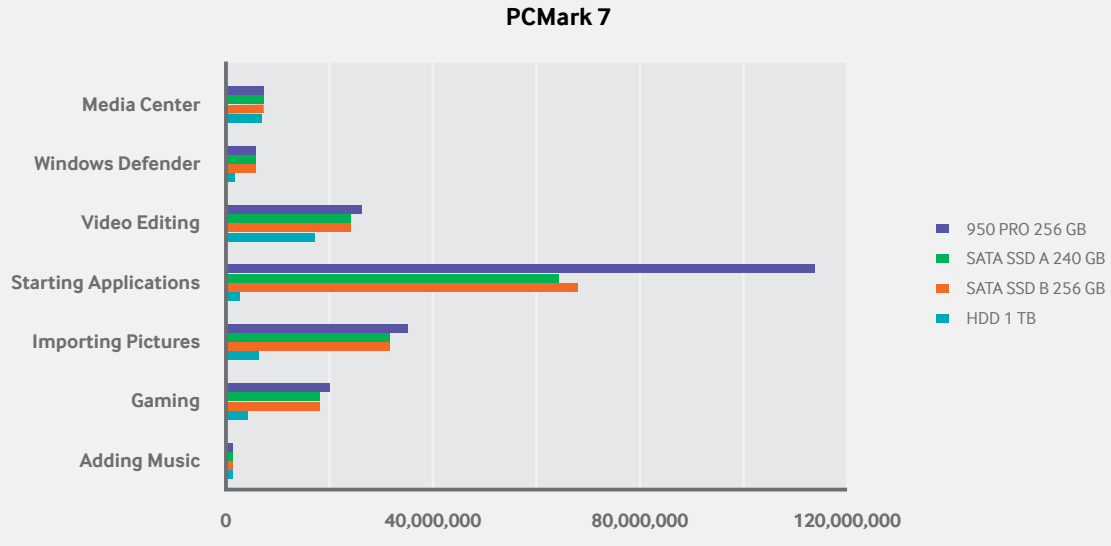

Figure 21. PCMark 7 score comparison (950 PRO versus SATA HDD/SSDs)

Samsung 950 PRO Option ROM software facilitates and simplifies the booting process

5. Samsung 950 PRO Option ROM

Option ROM is software that executes in a pre-OS environment and helps in booting an operating system[7] from a storage device, such as an SSD. A client PC can be booted either in a Basic Input/Output System (BIOS) mode or in a Unified Extensible Firmware Interface (UEFI), based on the support provided by the motherboard manufacturer. The UEFI is the successor to the legacy BIOS firmware interface, addressing its limitations, such as support for only 16-bit interfaces, limited ROM execution space, and so on.

The UEFI is backward compatible, enabling storage devices with OS installed in Legacy BIOS mode, or one that doesn't support UEFI, to boot in the UEFI mode. Booting is accomplished through a feature called the Compatibility Support Module (CSM) that emulates a BIOS environment in the UEFI mode. The Samsung 950 PRO Option ROM resides inside the SSD, enabling the SSD to be the primary boot device and supports the UEFI. Figure 22 shows the need for Option ROM in the SSD.

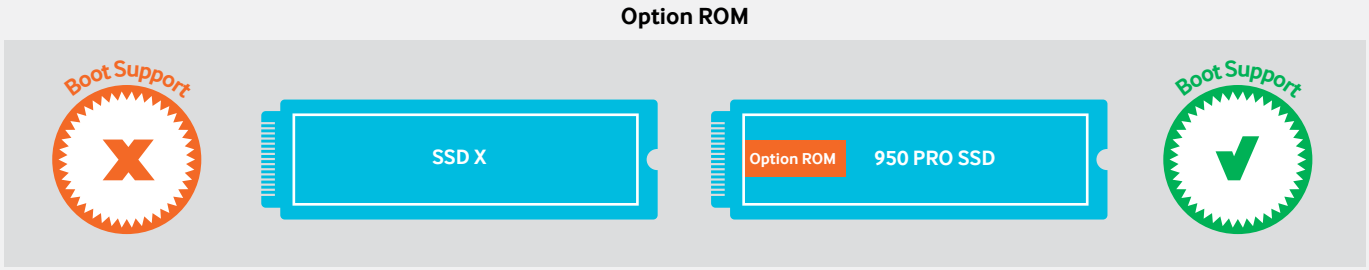

Figure 22. Importance of Option ROM

The UEFI communicates with the platform firmware and the operating system during the boot process and it is recommended to disable CSM to boot in the UEFI mode[8]. The UEFI boots by loading boot-loader programs (EFI program files with the .efi filename extension) from a partition on the disk, which is known as the EFI System Partition (ESP). During this process, Option ROM acts as a driver to interact with the SSD since the disk is a PCIe expansion card SSD. The boot manager accesses the OS loader application from the EFI, where it is executed to load the operating system and returns an error to the boot manager in case the OS loading fails. Upon loading the OS successfully, the OS loader terminates the booting services and the platform or OS takes control of the PC. Figure 23 shows the UEFI booting sequence[9].

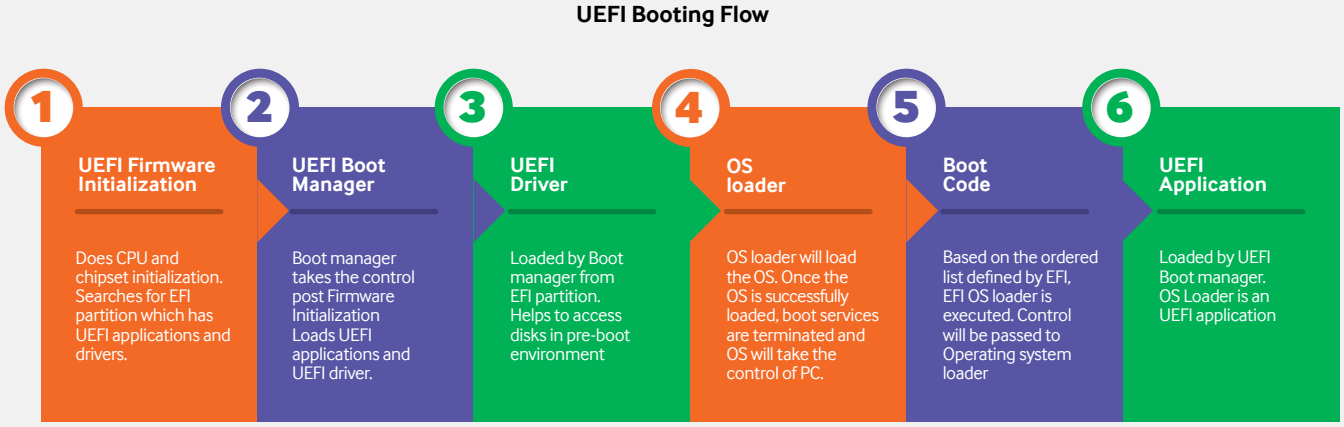

Figure 23. Booting sequence

The BIOS is the pre-existing firmware interface that was used before the UEFI evolved. When the PC is powered on, the BIOS initializes various motherboard components and makes sure that they are functioning properly with a series of diagnostic tests called a Power On Self-Test (POST)[10]. Once the POST is complete, the BIOS will load the boot sector, the Master Boot Record (MBR), from the SSD, which describes the layout of the disk partitions and boot loader. The boot loader finds the active partition and executes the Volume Boot Record (VBR), the first sector of the partition, to load the operating system.

UEFI and BIOS configuration instructions for OS installation

5.1. UEFI and BIOS configuration for OS installation

Most of the latest PCs have built-in support for the UEFI mode, except IBM® legacy-compatible systems that support booting only in the BIOS mode[10]. Windows 8.1 supports the UEFI 2.0 version or later on 32-bit (x86), 64-bit (x64) and ARM®-based PCs. A 64-bit UEFI supports only 64-bit versions and a 32-bit UEFI supports only 32-bit versions of Windows. The UEFI is supported for Windows 7 with SP1 only on a 64-bit OS by default. To install Windows 7 with SP1 32-bit editions, CSM mode has to be enabled. The following subsections illustrate steps to follow while using the 950 PRO as a bootable device in the UEFI mode.

NOTE 1: Refer to Appendix C for motherboard compatibility information for booting with the 950 PRO.
NOTE 2:To install the OS on the 950 PRO, the OS must have NVMe driver support. Windows 8.1 and 10 come with an inbox driver for NVMe, but Windows 7 does not have an inbox driver for NVMe devices. Refer to Appendix D for details on how to install the Windows 7 OS in the 950 PRO using the Samsung NVMe driver.

5.1.1. UEFI boot configuration settings

The following guidelines illustrate the UEFI settings for making the 950 PRO a bootable device in the UEFI mode on an ASRock Z97 Extreme6 PC. The same steps need to be followed for any UEFI-enabled PC.

1. Power on the PC and enter the BIOS settings by pressing F2 or the Del key.
2. Go to the Exit tab and select **Load UEFI Defaults**.

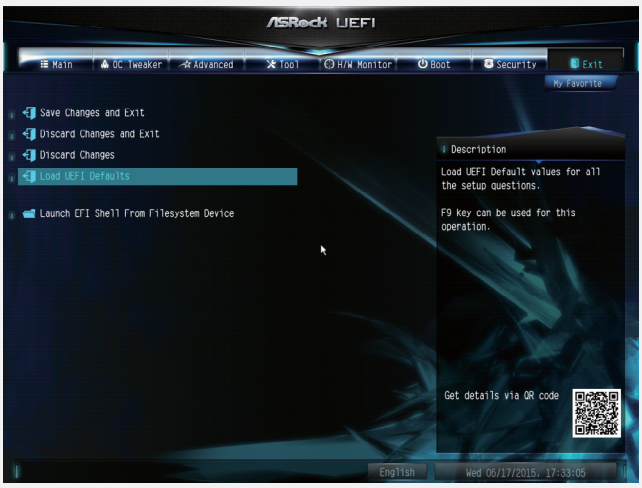

3. By default the CSM mode is enabled under the Boot tab. Disabling the **CSM** mode is recommended for the UEFI mode.

SAMSUNG

UEFI and BIOS configuration instructions for OS installation

4. Go to the **Exit** tab and select **Save Changes and Exit** to save the settings and proceed with the OS installation.
5. Once the OS is installed, reboot the PC, enter the BIOS and go to the **Boot** tab and select the Samsung 950 PRO as **Boot Option #1**. This step is needed only if there are multiple bootable devices connected to the PC.

SAMSUNG

Samsung NVMe driver features and support information

6. Samsung NVMe driver

Starting from Windows 8.1, Microsoft has included native drivers for supporting NVMe devices. Driver support is available in Linux® systems, too, from 3.3 or later kernel versions. To extend the NVMe support to more client users, Samsung provides its own proprietary NVMe driver for the following operating systems.

Operating System	32 Bit	64 Bit
Windows 7	Supported	Supported
Windows 8.1	Supported	Supported
Windows 10	Supported	Supported

Samsung NVMe drivers are Microsoft Windows Hardware Quality Labs (WHQL) tested and in compliance with all the mandatory features and commands of NVMev1.1 specification and NVMe: SCSI Translation Reference v1.4. In addition, the Samsung NVMe driver supports optional commands, which support the following features:

- Upgrade the device firmware using Firmware Activate and Firmware Image Download commands
- Change the logical block addressing (LBA) data size and/or metadata size using the Format NVM command
- Erase all user content present in the NVM subsystem so that the data is not recoverable using the Format NVM command
- Erase all user contents present in the NVM subsystem by deleting the encryption key with which the user data was previously encrypted using the Format NVM command
- Mark a logical block as invalid using the Write Uncorrectable command
- Set a range of logical blocks to zero using the Write Zeros command
- Read the command from the medium and compare the data read to a comparison data buffer transferred as part of the command using the Compare command.
- Exercise the SSD TRIM feature using Dataset Management - De-allocate command. This is equivalent to the ATA TRIM command and the SCSI UNMAP command.

Besides the above specification-compliant features, the Samsung NVMe driver also provides support for disk end of life, wherein the user data in the NVM disk can still be read even when the disk exhausts the P/E cycles and becomes read only. Samsung recommends using the Samsung NVMe driver to get optimal sustainable performance and feature-rich capabilities. Figure 25 provides the results of the PCMark 8 benchmark results of the 256 GB 950 PRO in Windows 8.1 for both the Samsung NVMe driver and the Microsoft inbox driver, clearly depicting that the Samsung driver outruns the inbox driver in almost all workloads.

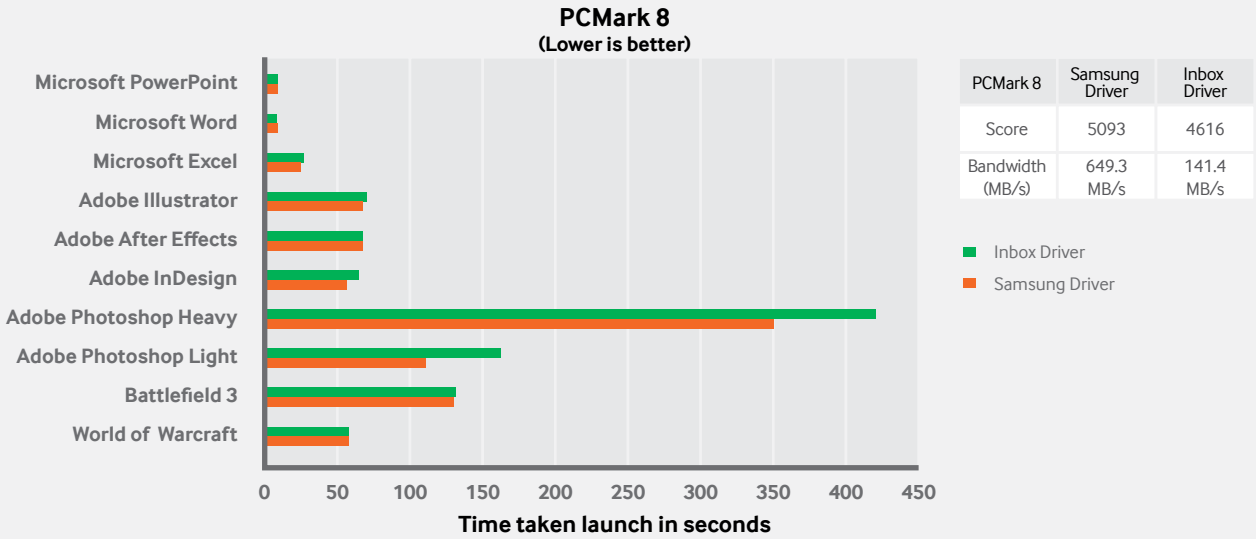

Figure 25. PCMark 8 score comparison for the 256 GB 950 PRO (inbox vs. Samsung driver)

Samsung NVMe driver features and support information

NOTE : Refer to Appendix E for the detailed procedure on how to install and uninstall the Samsung NVMe driver by using the driver installer package in Windows 8.1 and 10.

7. Conclusion

The Samsung 950 PRO proves to be the most cost-effective and lowest power-consuming SSD for client applications, delivering high performance consistently. In client environments, where power consumption is of utmost importance, the 950 PRO delivers the superb performance-to-power consumption ratio of approximately 60K random read IOPS and approximately 20K random write IOPS per watt of power consumed. Furthermore, its small form factor makes it a perfect candidate for deployment in ultra-thin notebooks as well as in desktop PCs. With enterprise-grade features, such as V-NAND flash, L1.2 power management, and an advance ECC engine, the 950 PRO exceeds all expectations for client systems, making it the superior PCIe SSD currently in the industry.

References

8. References

1. http://en.wikipedia.org/wiki/Serial_ATA
2. http://en.wikipedia.org/wiki/PCI_Express
3. http://en.wikipedia.org/wiki/PCI_Express#Lane
4. http://www.samsung.com/global/business/semiconductor/minisite/SSD/downloads/document/03_NAND_Basics.pdf
5. http://www.nvmexpress.org/
6. http://www.samsung.com/us/business/oem-solutions/pdfs/V-NAND_technology_WP.pdf
7. https://msdn.microsoft.com/en-us/library/windows/hardware/dn747882.aspx
8. https://en.wikipedia.org/wiki/Unified_Extensible_Firmware_Interface
9. http://www.uefi.org/sites/default/files/resources/UEFI_Spec_2_3_D.pdf
10. https://en.wikipedia.org/wiki/BIOS

SAMSUNG

Glossary

9. Glossary

TERM	ABBREVIATION DESCRIPTION	DEFINITION
AHCI	Advanced Host Controller Interface	A technical standard that allows software to communicate with SATA devices with support for advanced features, such as TRIM, NCQ and hot-plug.
APST	Autonomous Power State Transitions	A protocol for automatic transition of PCIe devices to different power states for low power consumption.
ASPM	Active State Power Management	A protocol that is used to manage less active PCIe serial links for low power consumption.
BIOS	Basic Input/Output System	A program that instructs the computer on how to perform a number of basic functions, such as booting, display screen, configure disk drives and keyboard control. It initializes POST. (See definition of POST below.)
CSM	Compatibility Support Module	A feature that enables storage devices with OS installed in the Legacy BIOS mode to boot in the UEFI mode by emulating a BIOS-like environment in the UEFI mode.
CTF	Charge Trap Flash	A semiconductor memory technology used in creating non-volatile NOR and NAND flash memory. It uses a silicon nitride film to store electrons, which lowers the cell coupling and results in a cell-to-cell interference-free structure.
DRAM	Dynamic Random Access Memory	A type of random-access memory, which allows processors to access any part of the memory directly. It stores each bit of data in a storage cell consisting of a capacitor and a transistor.
DTT	Dynamic Thermal Throttling	A technology that controls the temperature of the device to not exceed its maximum limit by reducing its performance in a phased manner.
Endurance	See Definition	It is the number of program/erase (P/E) cycles that can be applied to a block of flash memory before the storage media becomes unreliable.
GT/s	Giga Transfers per Second	Raw data rate - the number of operations transferring data that occur in each second. In order to calculate the data transmission rate, the transfer rate needs to be multiplied by the information channel width.
IOPS	Input/Output Operations per Second	A common measurement for computer storage performance, in particular, random access to mass storage devices and systems.
M.2 Port	See Definition	A new interface for PCIe-based devices with a low form factor and it utilizes multiple PCIe lanes (up to x4) to boost the speed of PCIe storage devices up to 4 GB/s.
MBR	Master Boot Record	A boot sector located at the very beginning of partitioned storage devices that describes the layout of disk partitions and the boot loader to start the process of loading the OS.
MSI-X	Message Signal Interrupts - Extended	Interrupts that are triggered via writing a value to a particular memory address. It features a dynamically programmable hardware table that contains entries for each of the interrupt sources in the device and can be independently masked.
NAND	Not AND	A logic method that acts as a flip-flop. The NAND flip-flop holds one bit of information and enables sequential access to memory cells.
NCQ	Native Command Queuing	A technique to increase performance of mass storage devices by allowing the individual device to internally optimize the order in which received read and write commands are executed.

Glossary continued

NVMe	Non Volatile Memory Express	A scalable host controller interface that utilizes PCIe-based solid-state drives.
OPTION ROM	See Definition	Software that executes in pre-OS environment and helps in booting an operating system from a storage device, such as an SSD.
PCIe	Peripheral Component Interconnect Express	PCI Express is a modern version of PCI. It is a high-speed serial computer expansion bus standard with numerous improvements over PCI.
POST	Power On Self-Test	A test that helps verify the computer meets the minimum requirements to boot properly.
RAID	Redundant Array of Independent Disks	A technology to improve both the data integrity and data throughput performance by using two or more independent mass storage devices.
Random Read	See Definition	A benchmark method to measure the performance of a mass storage device by doing read data access of a certain amount of data scattered in different blocks.
Random Write	See Definition	A benchmark method to measure the performance of a mass storage device by doing write data access of a certain amount of data scattered in different blocks.
SATA	Serial Advanced Technology Attachment	A serial implementation of the Parallel ATA attachment standard.
Sequential Read	See Definition	A benchmark method to measure the performance of a mass storage device by doing read data access of a large amount of contiguous blocks of data.
Sequential Write	See Definition	A benchmark method to measure the performance of a mass storage device by doing write data access of a large amount of contiguous blocks of data.
S.M.A.R.T.	Self-Monitoring, Analysis and Reporting Technology	A monitoring system to detect and report on various indicators of reliability and help in anticipating failures of an SSD device.
SSD	Solid State Drive	A data storage device that uses non-volatile solid-state memory and has no moving parts.
TRIM	See Definition	A computer command used by an SSD to inform it of data blocks in the SSD memory space which are no longer in use. This command allows the SSD controller to reorganize the memory space it is managing to allow any unused space to be optimally relocated.
UEFI	Unified Extensible Firmware Interface	A successor to the Legacy BIOS firmware interface, addressing its limitations, such as support for only 16-bit interfaces, limited ROM execution space, and so on.
V-NAND	Vertical NAND	A type of flash memory, which stacks flash memory cells vertically in three-dimensional fashion, thereby offering higher capacity and better performance.

Appendix A

10. Appendix A - Samsung 950 PRO power states and feature list

Power States

Power State	Internal State	Average Power (RMS)	Entry Latency	PCIe Access Latency (950 PRO)
PS 0	Active State	5.5 W	N/A	N/A
PS 1	Active State	2.5 W	N/A	N/A
PS 2	Active State	1.9 W	N/A	N/A
PS 3	Low Power State	70 mW	60 ms	75 µs
PS 4	Deep Power Down	2.5 mW	10 sec	700 µs

NVMe Feature List

Admin Command Sets	O/M	Supported	
Create I/O Completion Q	M	O	Max 8ea
Create I/O Submission Q	M	O	Max 8ea
Delete I/O Completion Q	M	O	Max 8ea
Delete I/O Submission Q	M	O	Max 8ea
Get Log Page	M	O	
Identify	M	O	
Abort	M	O	
Set Features	M	O	
Get Features	M	O	
Asynchronous Event Req.	M	O	
FW Activate	O	O	
FW Image Download	O	O	3 Slot
Format NVM	O	O	

NVM Command Sets	O/M	Supported	
Flush	M	O	
Read	M	O	
Write	M	O	
Write Uncorrectable	O	O	
Write Zeros	O	O	
Compare	O	O	
Dataset Management	O	O	Deallocate Integral Dataset
Fused Operation	O	X	Compare and Write
E2E Protection	O	X	DIX/DIF

Appendix A continued

Identify Controller Data Structure

Bytes	O/M	Default Value	Description
1:00	M	144Dh	PCI vendor ID
3:02	M	144Dh	PCI subsystem vendor ID
23:04	M		Serial number
63:24	M		Model number
71:64	M		Firmware revision
72	M	2h	Recommended arbitration burst
75:73	M	0x002538	IEEE® OUI
76	O	0h	Multi-interface capabilities and namespace sharing capability
			Bit 2: 1h - Controller is associated with an SR-IOV virtual function 0h - Controller is associated with a PCI function.
			Bit 1: 1h - Device has two or more controllers 0h - Device has one controller
			Bit 0: 1h - Supported 0h – Not support
77	M	5h	Maximum data transfer size (MPS size * this field) = max transfer size
			0h: No restrictions
79:78	M	1h	CTNLID
255:80		0h	Reserved
257:256	M	7h	Optional admin command support
			Bits 15:4 - Reserved
			Bit 3: 1h - Namespace management attachment supported
			Bit 2: 1h – Firmware activate/download supported
			Bit 1: 1h Format NVM supported
			Bit 0: 0 Security send and security receive not supported
258	M	7h	Abort command limit
259	M	3h	Asynchronous event request limit
260	M	6h	Firmware updates
			Bits 7:4 – Reserved
			Bits 3:1 – Number of firmware slots
			Bit 0 – 1h slot 1 is read only
261	M	1h	Log page attributes
			Bits 7:1 – Reserved
			Bit 0: 0h S.M.A.R.T. data is global for all namespaces

Appendix A continued

262	M	3Fh	Error log page entries
263	M	4h	Number of power states support
264	M	1h	Admin vendor-specific command configuration
			Bits 7:1 – Reserved
			Bit 0 – Indicates admin vendor-specific commands use the format defined in NVM Express 1.0c Figure 8
265	O	1h	Autonomous Power State Transition Attributes
511:265		0h	Reserved
512	M	66h	Submission queue entry size
			Bits 7:4 – 6h Max SQES
			Bits 3:0 – 6h Required SQES
513	M	44h	Completion queue entry size
			Bits 7:4 – 4h Max SQES
			Bits 3:0 – 4h Required SQES
515:514		0h	Reserved
519:516	M	1h	Number of namespaces
521:520	M	1Fh	Optional NVM command support
			Bits 15:6 – Reserved
			Bit 5 - Reservations
			Bit 4 - SAVE field in Set Feature & Select field in Get Feature
			Bit 3 - 1h Write zeros supported
			Bit 2 – 1h Dataset management supported
			Bit 1 – 1h Write uncorrectable supported
			Bit 0 – 1h Compare supported
523:522	M	0h	Fused Operation support
			Bits 15:1 – Reserved
			Bit 0 – 0h Compare/write fused operation not supported
524	M	0h	Format NVM attributes
			Bits 7:3 – Reserved
			Bit 2 – 1h Cryptographic erase
			Bit 1 – 1h Secure erase per namespace
			Bit 0 – 0h Format per namespace
525	M	1h	Volatile write cache
			0h – No VWC present

Appendix A continued

527:526	M	FFh	Atomic write unit normal
			All commands atomic
529:528	M	0h	Atomic write unit power fail
530	M	1h	NVM vendor-specific command configuration
			Bits 7:1 – Reserved
			Bit 0 – Indicates NVM vendor-specific commands use the format defined in NVM Express 1.1a
531	M	0h	Reserved
533:532	O	0h	ACWU
535:534	M	0h	Reserved
539:536	O	0h	No SGL support
703:540		0h	Reserved
I/O Command Set Attributes			
2047:704		0h	Reserved
Power State Descriptors			
2079:2048	M	Refer to "Identify Power State Descriptor Data Structure"	Power state 0 descriptor
2111:2080	O		Power state 1 descriptor
2143:2112	O		Power state 2 descriptor
2175:2144	O		Power state 3 descriptor
2207:2176	O		Power state 4 descriptor
...		0h	
3071:3040	O	0h	Power state 31 descriptor
Vendor Specific			
3079:3072			Reserved
3087:3080			Reserved
3091:3088			Reserved
3092	O	1h	OEM-specific S.M.A.R.T. feature set
			Bit 2 – 1h OEM-specific S.M.A.R.T. Extended self-test supported
			Bit 1 – 1h OEM-specific S.M.A.R.T. self-test supported
			Bit 0 – 1h OEM-specific S.M.A.R.T. log supported
3278:3093		0h	Reserved

Appendix A continued

3279	0	Not implemented	Security feature set
			Bit 2 – 1h TCG supported
			Bit 1 – 1h SED supported
			Bit 0 – 1h ATA security pass-through supported
4095:3280		0h	Reserved

Identify Power State Descriptor Data Structure

Bytes	Description	Power State 0 Descriptor	Power State 1 Descriptor	Power State 2 Descriptor	Power State 3 Descriptor	Power State 4 Descriptor
255:125	Reserved					
124:120	Relative Write Latency	0	1	2	3	4
119:117	Reserved					
116:112	Relative Write Throughput	0	1	2	3	4
111:109	Reserved					
108:104	Relative Read Latency	0	1	2	3	4
103:101	Reserved					
100:96	Relative Read Throughput	0	1	2	3	4
95:64	Exit Latency	5 µs	30 µs	30 µs	300 µs	10,000 µs
63:32	Entry Latency	5 µs	30 µs	30 µs	10,000 µs	50,000 µs
31:26	Reserved					
25	Non-Operational State	0	0	0	1	1
24	Max Power Scale	0.01 W	0.01 W	0.01 W	0.0001 W	0.0001 W
23:16	Reserved					
15:00	Maximum Power	9 W	4.6 W	3.8 W	0.07 W	0.005 W

Appendix B

11. Appendix B - S.M.A.R.T. features in NVMe spec

S.M.A.R.T. features in NVMe spec		
Bytes	Description	Supported
0	**Critical Warning**:This field indicates critical warnings for the state of the controller. Each bit corresponds to a critical warning type; multiple bits may be set. If a bit is cleared to 0, then that critical warning does not apply. Critical warnings may result in an asynchronous event notification to the host.	O
2:1	**Temperature**: Contains the temperature of the overall device (controller and NVM included) in units of Kelvin. If the temperature exceeds the temperature threshold, as referred to in Section 5.12.1.4, then an asynchronous event completion may occur.	O
3	**Available Spare**: Contains a normalized percentage (0 to 100%) of the remaining spare capacity available.	O
4	**Available Spare Threshold**: When the available spare falls below the threshold indicated in this field, an asynchronous event completion may occur. The value is indicated as a normalized percentage (0 to 100%).	O
5	**Percentage Used**: Contains a vendor-specific estimate of the percentage of the device life used based on the actual device usage and the manufacturer's prediction of the device life. A value of 100 indicates that the estimated endurance of the device has been consumed, but may not indicate a device failure. This value is allowed to exceed 100. Percentage values greater than 254 shall be represented as 255. This value shall be updated once per power-on hour (when the controller is not in a sleep state). Refer to the JEDEC JESD218 standard for SSD device life and endurance measurement techniques.	O
31:6	Reserved	
47:32	**Data Units Read**: Contains the number of 512-byte data units the host has read from the controller; this value does not include metadata. This value is reported in thousands and is rounded up. When the LBA size is a value other than 512 bytes, the controller shall convert the amount of data read to 512-byte units. For the NVM command set, logical blocks read as part of Compare and Read operations shall be included in this value.	O
3:48	**Data Units Written**: CContains the number of 512-byte data units the host has written to the controller; this value does not include metadata. This value is reported in thousands and is rounded up. When the LBA size is a value other than 512 bytes, the controller shall convert the amount of data written to 512-byte units. For the NVM command set, logical blocks written as part of write operations shall be included in this value. Write Uncorrectable commands shall not impact this value.	O
79:64	**Host Read Commands**: Contains the number of read commands completed by the controller. For the NVM command set, this is the number of Compare and Read commands.	O
95:80	**Host Write Commands**: Contains the number of Write commands completed by the controller. For the NVM command set, this is the number of Write commands.	O
111:96	**Controller Busy Time**: Contains the amount of time the controller is busy with I/O commands. The controller is busy when there is an outstanding command in an I/O queue. This value is reported in minutes.	TBD
127:112	**Power Cycles**: Contains the number of power cycles.	O
143:128	**Power-on Hours:** Contains the number of power-on hours. This does not include the time that the controller was powered and in a low power state condition.	O
159:144	**Unsafe Shutdowns**: Contains the number of unsafe shutdowns. This count is incremented when a shutdown notification is not received prior to loss of power.	O

Appendix B continued

175:160	**Media Errors:** Contains the number of occurrences where the controller detected an unrecovered data integrity error. Errors such as uncorrectable ECC, CRC checksum failure or LBA tag mismatch are included in this field.	0
191:176	**Number of Error Information Log Entries**: Contains the number of error information log entries over the life of the controller.	0
511:192	**Reserved**	

Byte 0 : Critical Warning	
Bit	Definition
00	If set to 1, then the available spare space has fallen below the threshold.
01	If set to 1, then the temperature has exceeded a critical threshold.
02	If set to 1, then the device reliability has been degraded due to significant media-related errors or any internal error that degrades device reliability.
03	If set to 1, then the media has been placed in read-only mode.
04	If set to 1, then the volatile memory backup device has failed. This field is only valid if the controller has a volatile memory backup solution.
07:05	Reserved

Appendix C

12. Appendix C - Compatibility of Samsung 950 PRO Option ROM

OEM	Haswell Refresh
ASRock®	Z97-Extreme9
	Z97-Extreme6
	H97M-Pro4
	Z97-Extreme4
	Z97 KILLER
	X99 Extreme6
	Z97 PROFESSIONAL
MSI®	Z97-XPOWER
	Z97-MPOWER
	H97-GAMING
	X99 SLI KRAIT
	X99A MPOWER
	Z97 GAMING9 AC
	H97M G43
Gigabyte®	Z97X-UD3H
	H97-D3H-CF
	G1-Z97
	Z97X-UD5H
ASUS™	H97M-E
	H97-PLUS
	SABERTOOTH Z97 MARKS
	Z97-PRO

Appendix D

13. Appendix D - OS installation guide for Windows 7 with Samsung NVM Express driver

Windows 7 does not provide inbox NVMe driver support. As a result, The Samsung 950 PRO will not be listed in "selection of disk to install Windows on" during the OS installation. The following steps explain the Samsung NVMe driver installation, which is the continuation of the Windows 7 installation procedure, after Step 4 from Section 5.1.1.

1. Download the Samsung NVM Express driver package for Windows from the following URL: www.samsung.com/samsungssd.

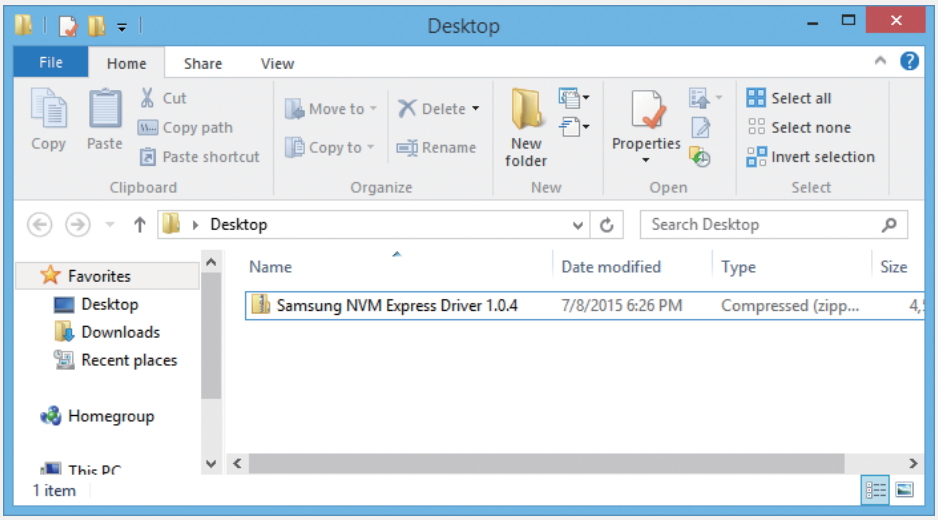

2. Extract the zip file and navigate to the **Samsung NVM Express Driver 1.0.4** folder. This folder will have Binaries and Installer folders as shown below.

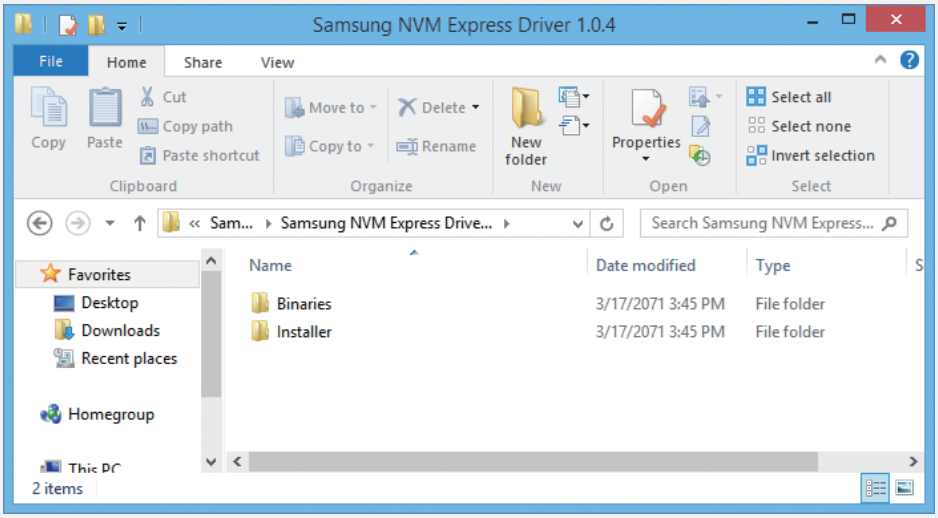

3. Copy the Binaries folder to a removable disk USB/DVD drive. Samsung NVM Express drivers are available for Windows 7, 8.1 and 10 for 32-bit and 64-bit OS as shown below.

Appendix D continued

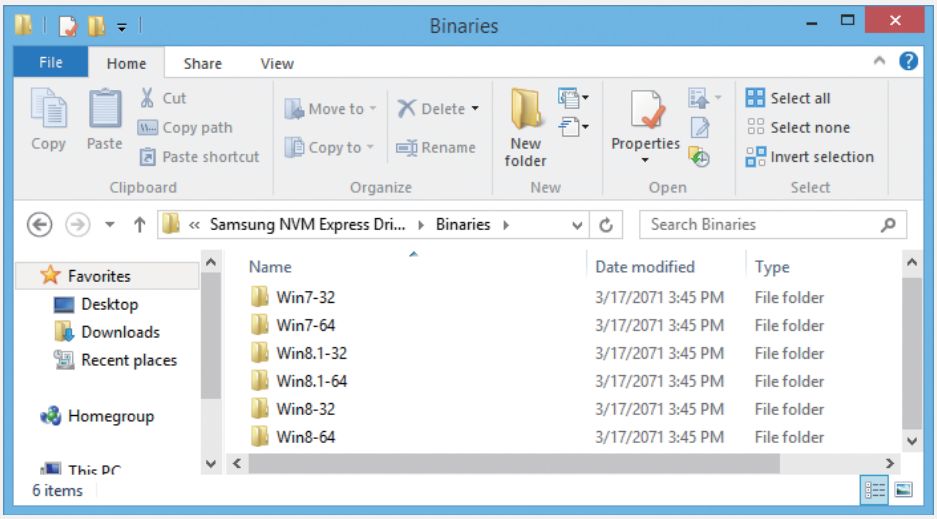

4. Insert the removable disk/DVD drive into the PC where the OS installation of the Samsung 950 PRO will be done.
5. The Windows 7 installation starts as shown in the screen shot below. Click **Next** to proceed with the installation.

6. After the following screen appears, click the **Install now** option.

Appendix D continued

7. Check the license terms and conditions box and click **Next**

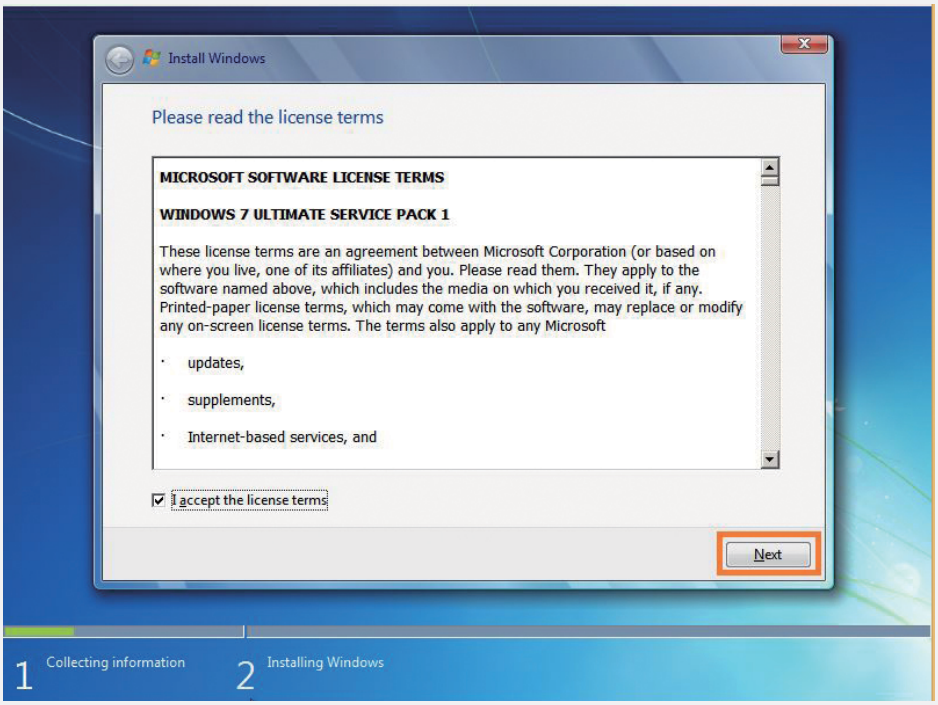

8. Since the NVMe driver is not available, the Samsung 950 PRO will not be displayed. Click the **Load Driver** option to install the driver as shown in the screen below.

Appendix D continued

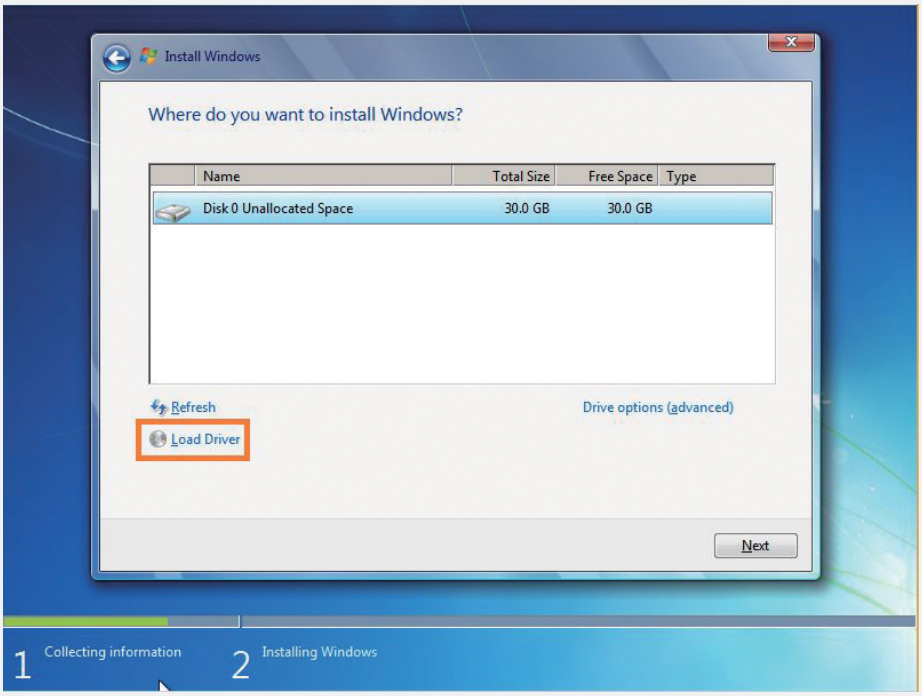

9. The Load Driver option will be displayed as shown in the screen below. Click **Browse.**

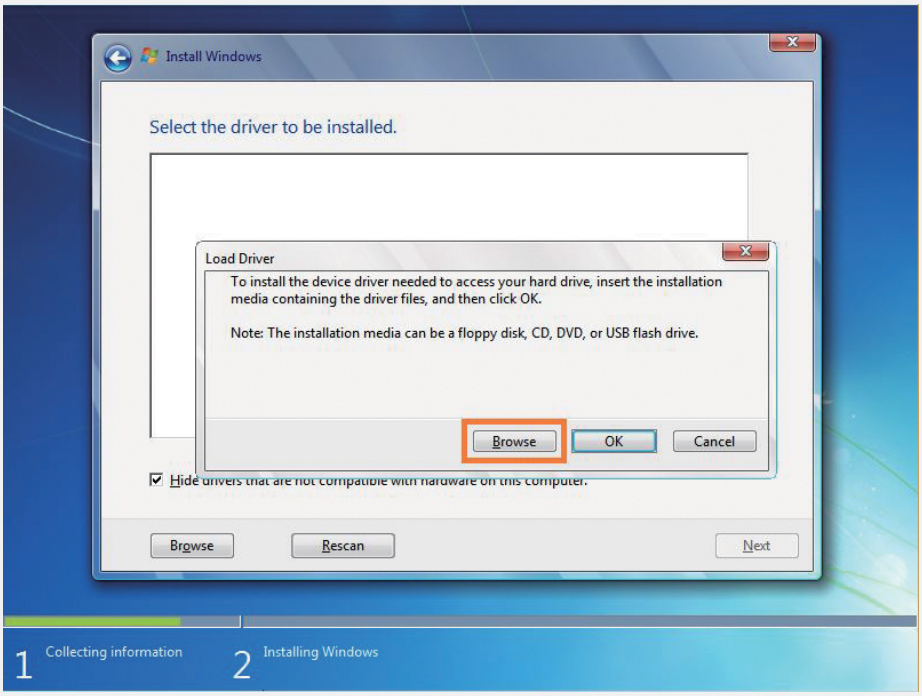

10. Choose the removable disk/DVD drive that contains the previously copied Samsung NVM Express driver in Step 3 and proceed with the driver installation.

Appendix D continued

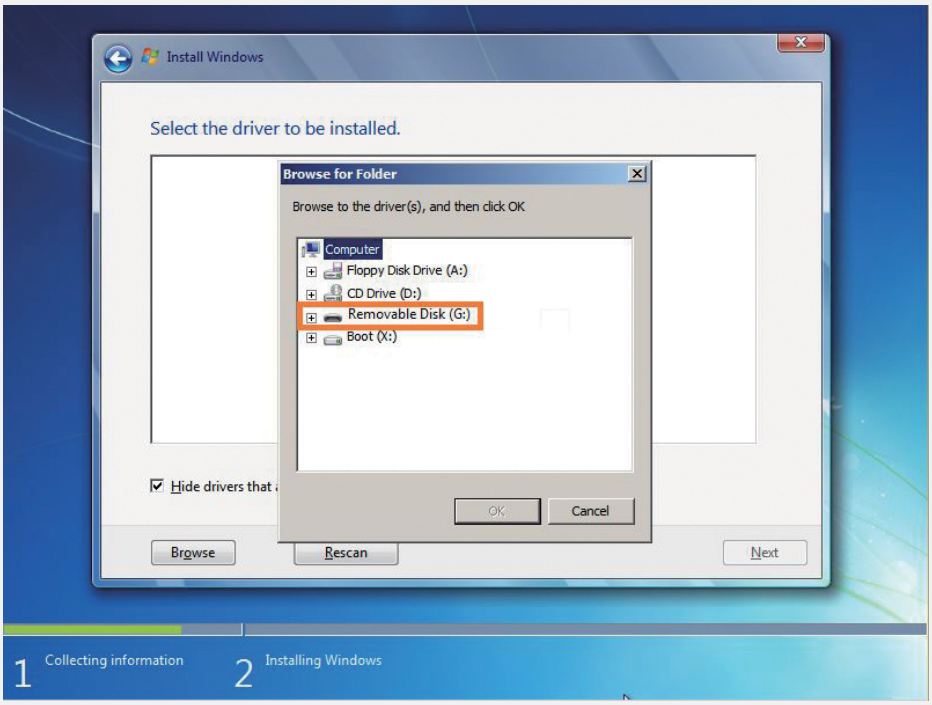

11. After Samsung NVM Express driver installation, the Samsung 950 PRO will be displayed and the user can continue with the Windows 7 installation.

Appendix E

14. Appendix E - Samsung NVM Express Driver installer package usage

This section covers the Samsung NVM Express driver installation and uninstallation procedure for Windows 7, 8.1 and 10.

NOTE: Samsung strictly recommends using the Samsung NVM Express installer package for installation and uninstallation of the Samsung NVM Express driver. The user shall not install or uninstall the Samsung NVM Express driver using the device manager or any third-party utilities and, if tried, the system stability could be compromised.

1) Installation Procedure

The following steps explain the procedure for the Samsung NVMe Installer Package installation in Windows 7, 8.1 and 10. Before proceeding with the steps below, make sure the Samsung 950 PRO is connected to the PC and the user has **Administrator Privileges**.

1. Download the Samsung NVMe driver package for Windows from the following URL: www.samsung.com/samsungssd.

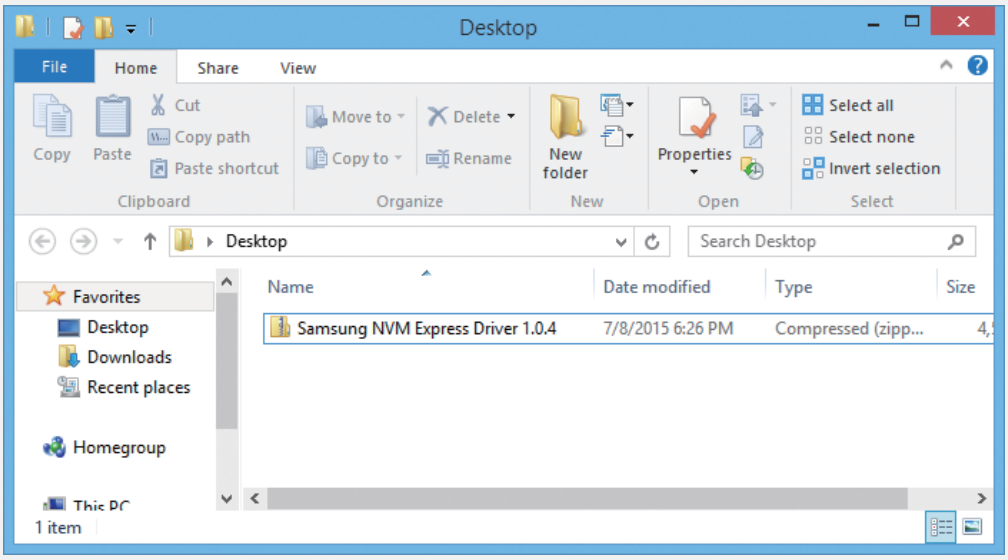

2. Extract the zip file and navigate to the **Samsung NVM Express Driver 1.0.4** folder. This folder will have **Binaries** and **Installer** folders as shown below.

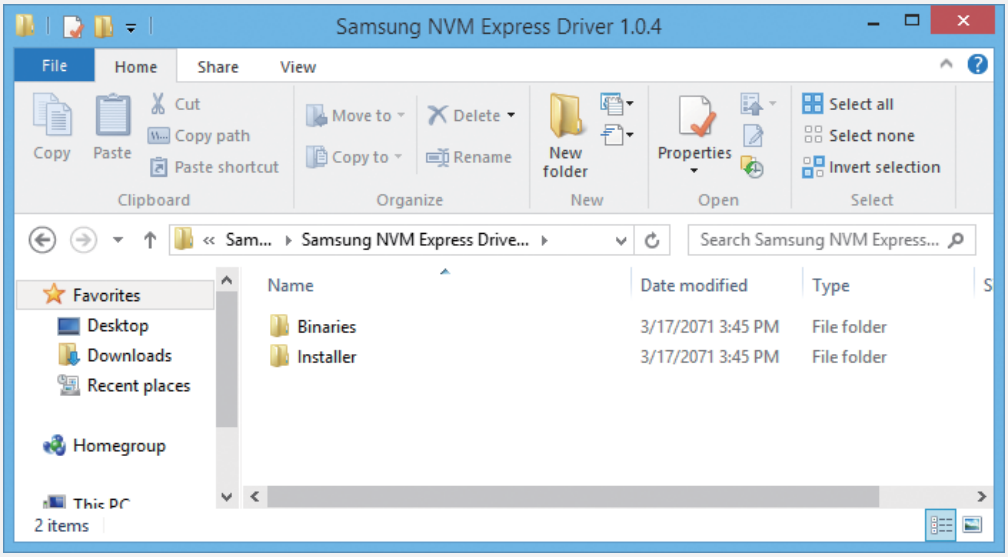

Appendix E continued

3. The **Installer** folder will have the Samsung NVM Express Driver installer executable as shown below.

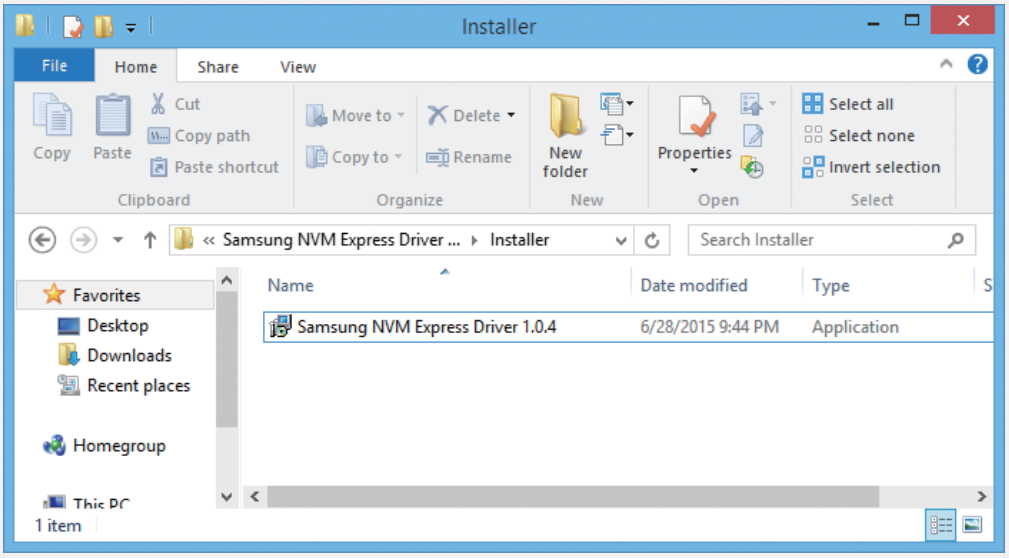

4. In Windows 8.1, the inbox driver instance for the Samsung 950 PRO in the device manager can be seen as shown below.

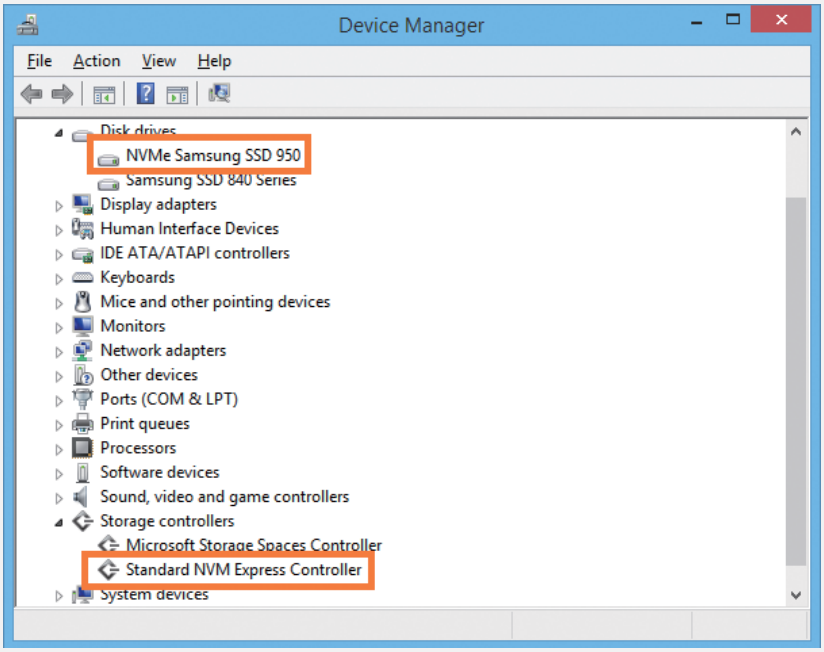

Appendix E continued

5. Run the installer file to invoke the wizard for installing the Samsung NVM Express Driver. Click the **Next** button.

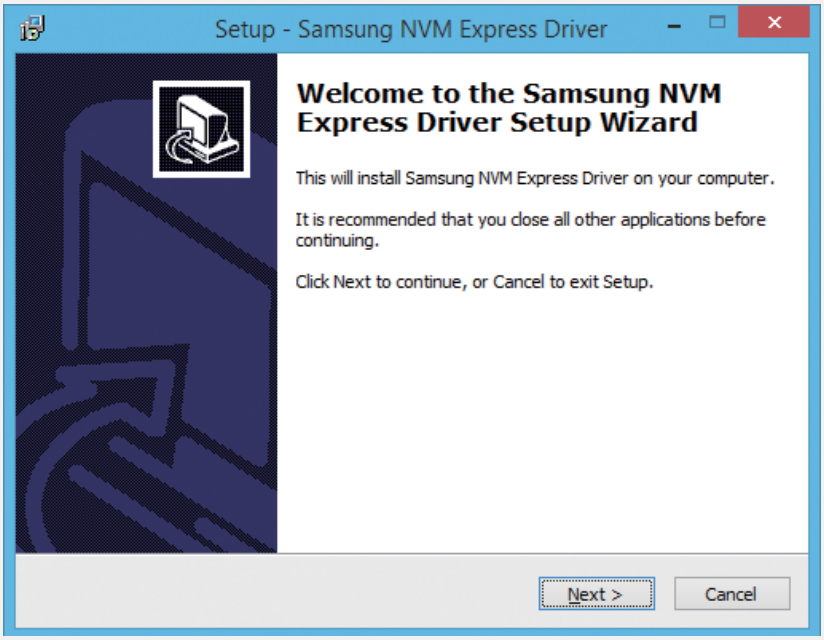

6. The license agreement will be displayed as shown below. Read the license terms carefully. If you agree, choose the **I accept the agreement** option and click the **Next** button to proceed with the installationlow.

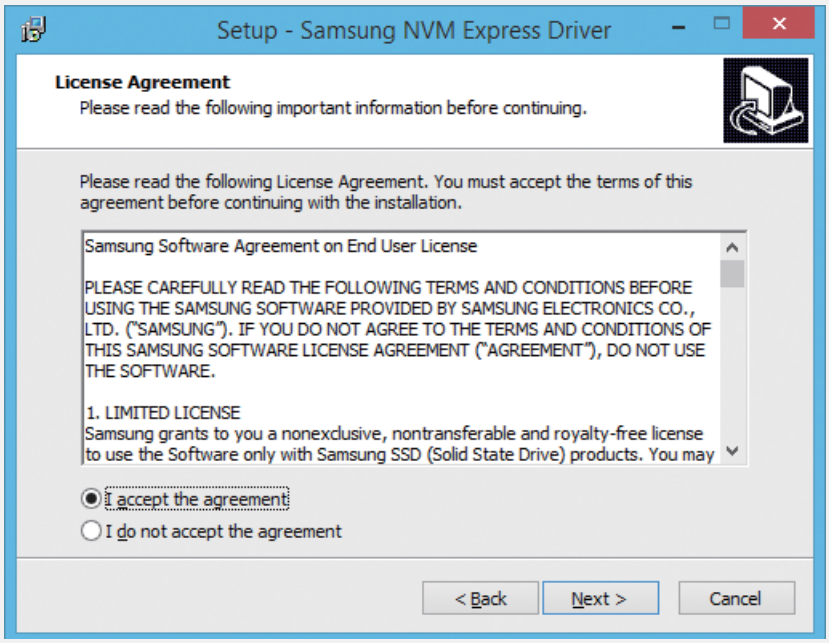

Appendix E continued

7. The driver installation path will be displayed as shown below. Click the **Next** button to continue.

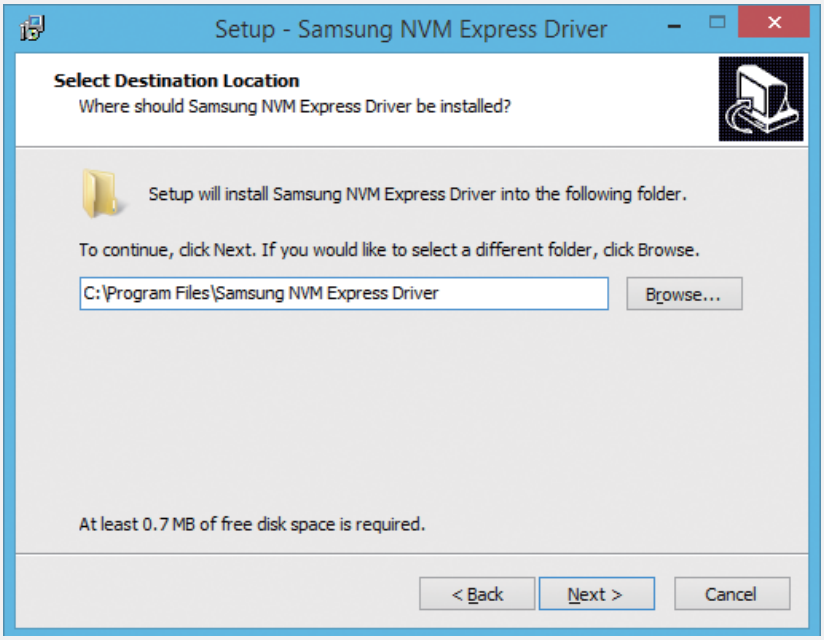

8. Click the **Install** button to **install** the Samsung NVM Express driver

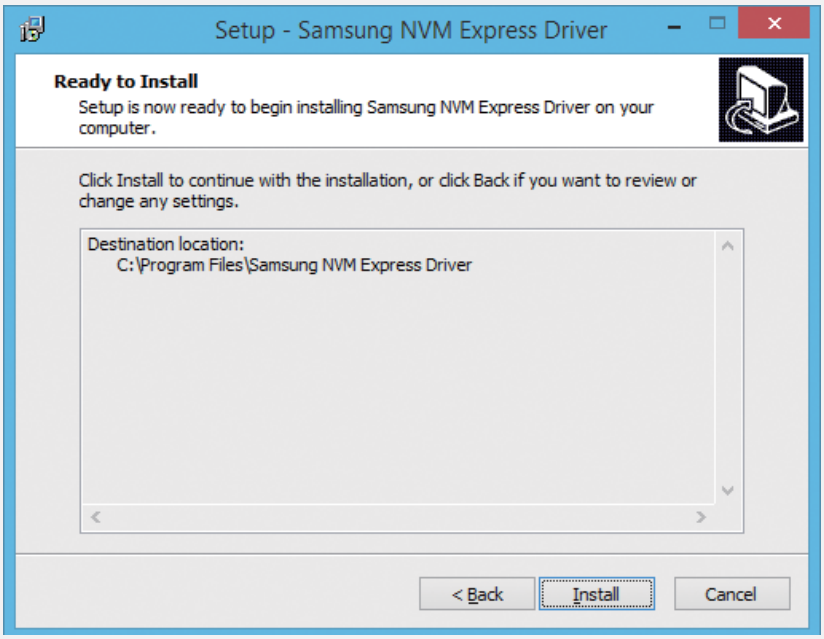

Appendix E continued

9. The driver installation progress will be displayed as shown below.

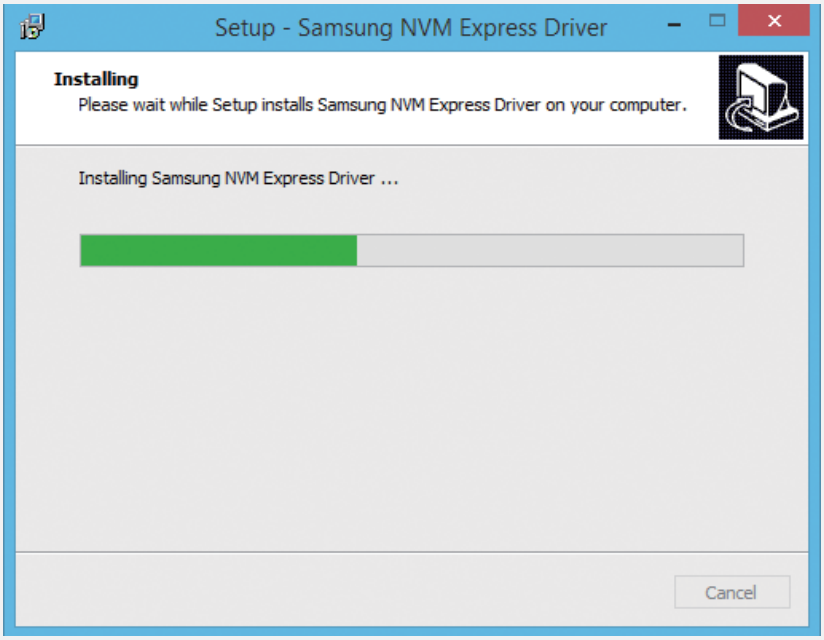

10. Once the installation is complete, choose one of the options below and click the Finish button. Samsung strictly recommends choosing the **Yes, restart the computer now** option to restart the PC. If the restart is not done, the system stability could be compromised

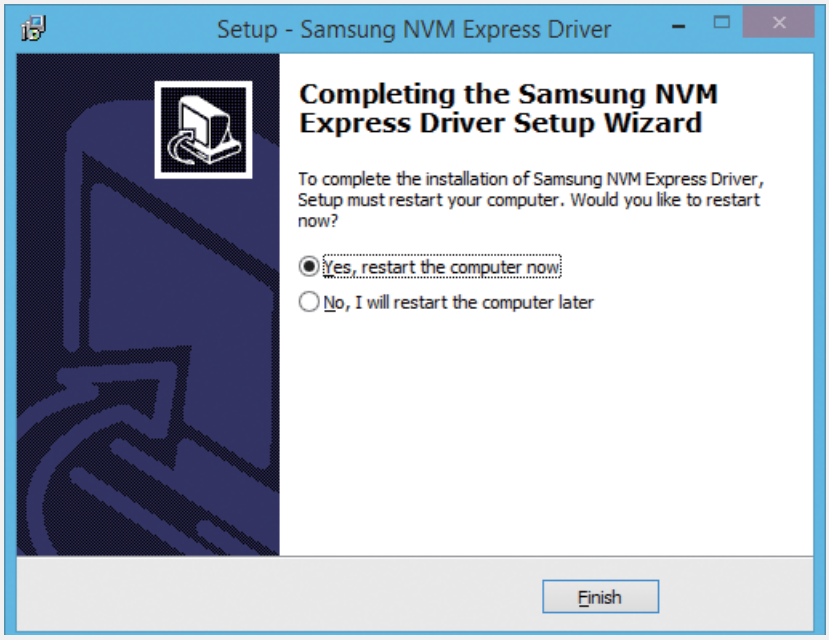

Appendix E continued

11. The Samsung NVM Express driver instance can be seen in the device manager as shown below.

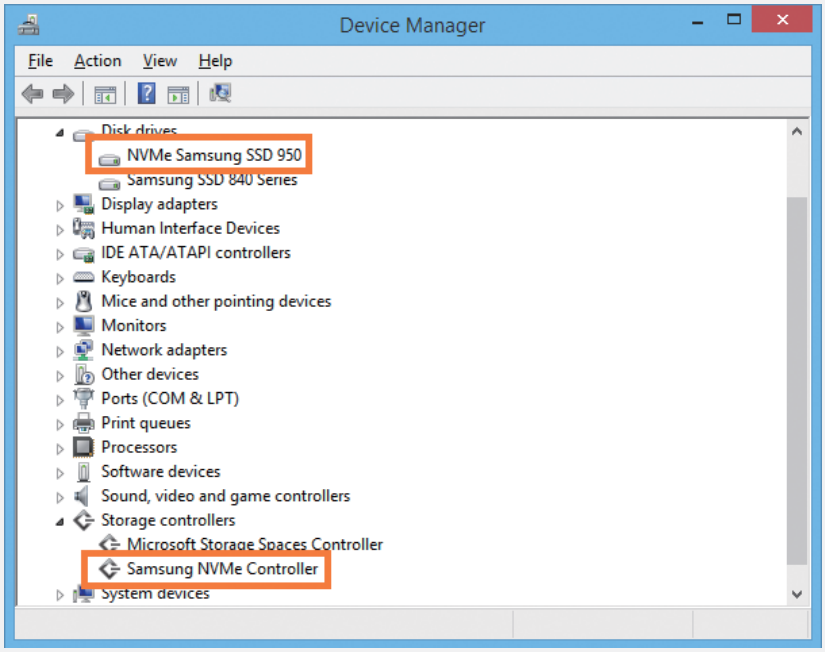

2) Uninstallation Procedure

The following steps explain the procedure for uninstalling the Samsung NVMe Installer Package in Windows 7, 8.1 and 10. Before proceeding with the following steps, make sure the user has **Administrator Privileges.**

1. Open the Windows control panel and choose the **Uninstall a program** option as shown below.

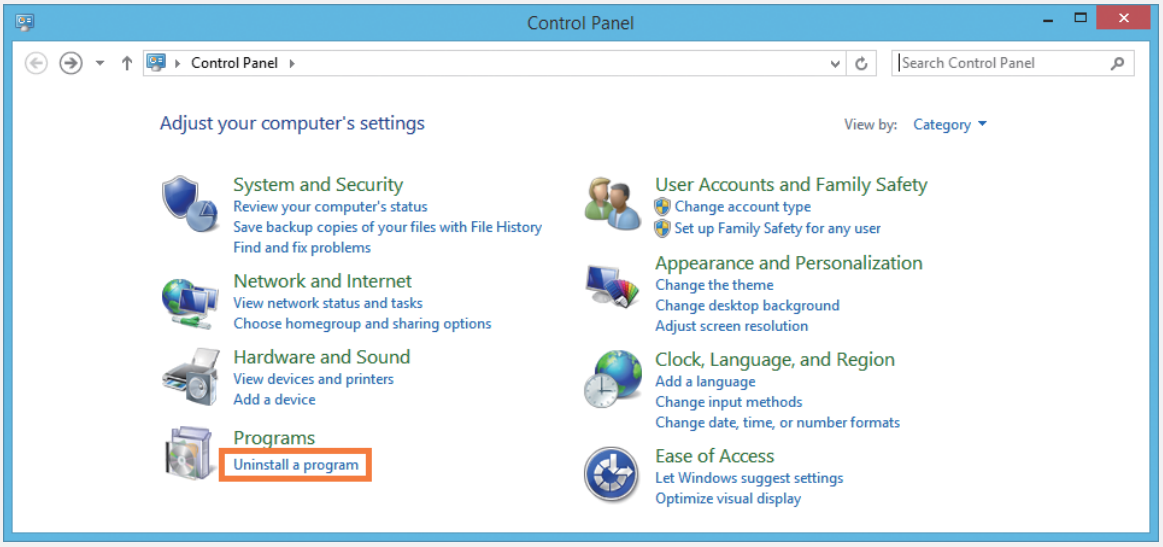

Appendix E continued

2. Select **Samsung NVM Express Driver** from the list of installed items and click **Uninstall.**

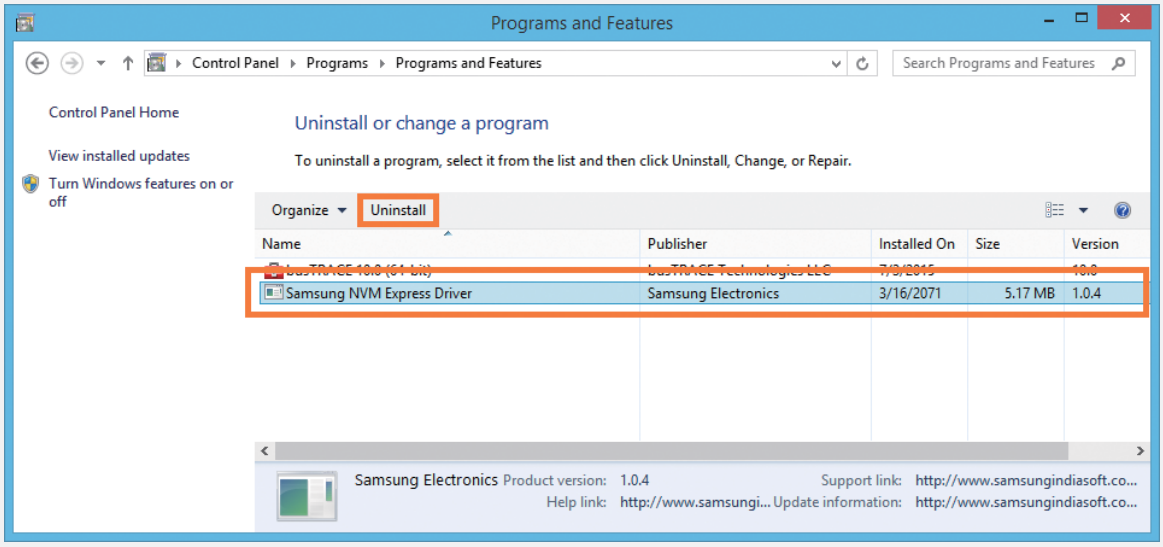

3. Confirm the driver uninstallation by clicking the **Yes** option as shown below.

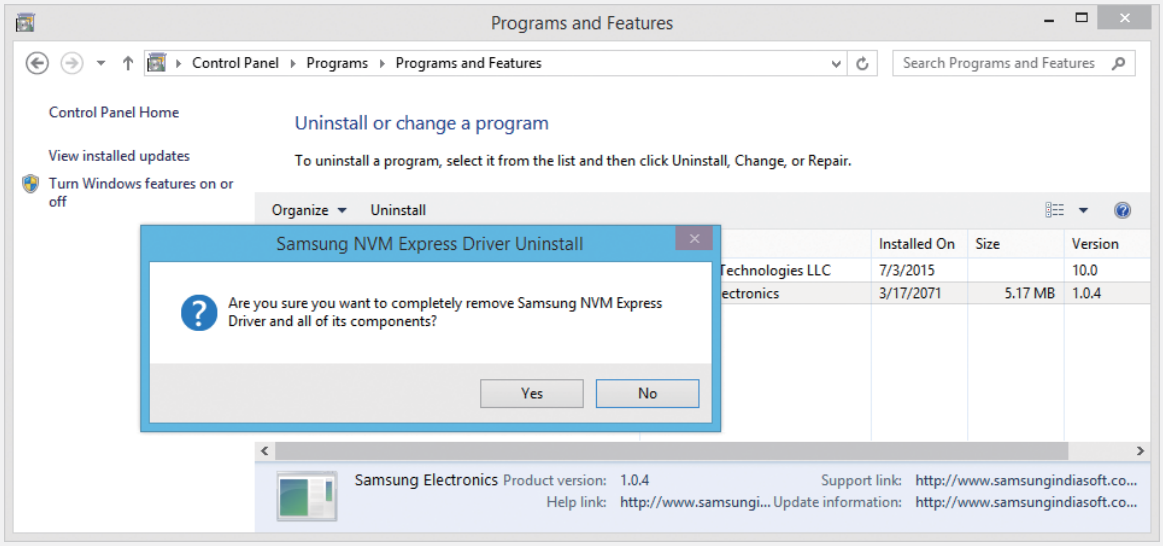

Appendix E continued

4. The uninstallation progress will be shown as indicated below.

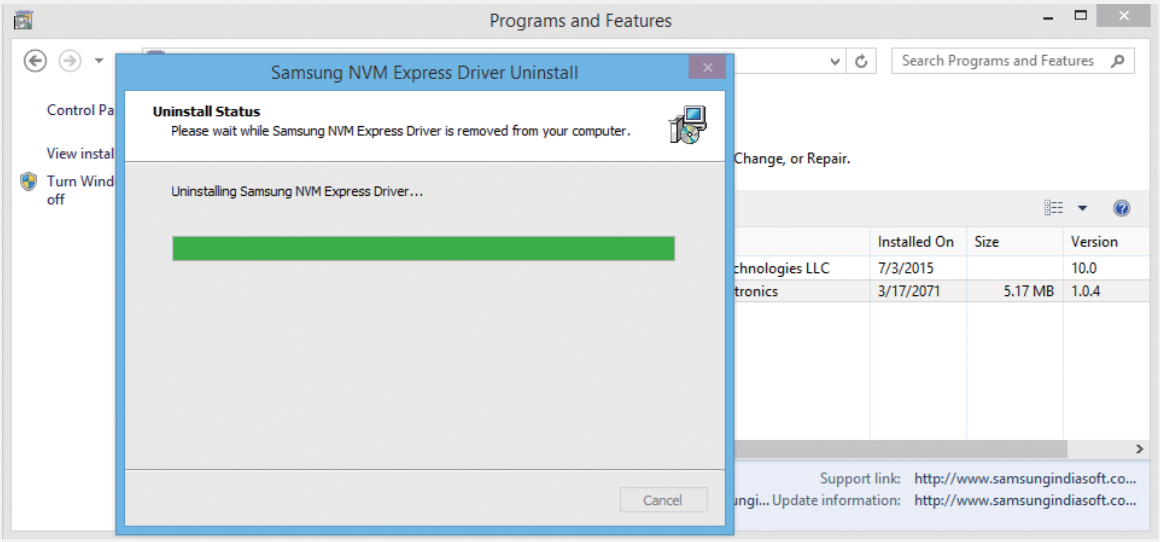

5. Once the uninstallation is completed, choose **Yes** to restart the PC. Samsung strictly recommends restarting the PC after the driver uninstallation. If the restart is not done, the system stability could be compromised.

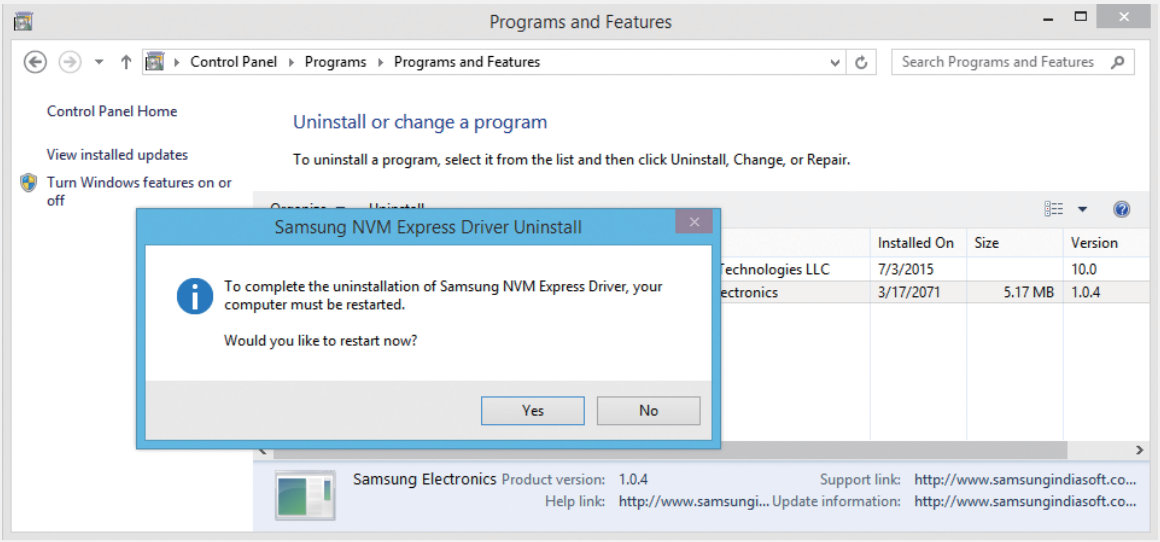

Legal and additional information

About Samsung Electronics Co., Ltd.

Samsung Electronics Co., Ltd. inspires the world and shapes the future with transformative ideas and technologies, redefining the worlds of TVs, smartphones, wearable devices, tablets, cameras, digital appliances, printers, medical equipment, network systems, and semiconductor and LED solutions. We are also leading in the Internet of Things space through, among others, our Smart Home and Digital Health initiatives. We employ 319,000 people across 84 countries with annual sales of US$196 billion. To discover more, please visit our official website at www.samsung.com and our official blog at global.samsungtomorrow.com

For more information

For more information about the Samsung SSD 950 PRO, visit www.samsung.com/semiconductor.

Copyright © 2015 Samsung Electronics Co., Ltd. All rights reserved. Samsung is a registered trademark of Samsung Electronics Co., Ltd. Specifications and designs are subject to change without notice. Nonmetric weights and measurements are approximate. All data were deemed correct at time of creation. Samsung is not liable for errors or omissions. All brand, product, service names and logos are trademarks and/or registered trademarks of their respective owners and are hereby recognized and acknowledged.

Adobe and Photoshop are registered trademarks of Adobe Systems Incorporated in the United States and/or other countries.

ARM is a registered trademark of ARM Ltd.

ASRock is the official trademark for ASRock Rack, Inc.

ASUS is a trademark of Asustek Computer Inc.

AT&T and AT&T Toggle are registered trademarks of AT&T Intellectual Property.

Gigabyte is a registered trademark of Giga-Byte Technology Co., Ltd.

IBM is a registered trademark of International Business Machines Corporation.

IEEE is a registered trademark of the Institute of Electrical and Electronic Engineers, Inc.

Intel and Intel Core are trademarks of Intel Corporation in the U.S. and/or other countries.

Linux is a registered trademark of Linus Torvalds.

Microsoft, Visual Studio and Windows are registered trademarks of Microsoft Corporation in the United States and/or other countries.

MSI is a registered trademark of Micro-Star International Co., Ltd.

PCI Express and PCIe are registered trademarks of PCI-SIG.

PCMark is a registered trademark of Futuremark Corporation.

Samsung provides this white paper for information purposes only. All information included herein is subject to change without notice. Samsung Electronics is not responsible for any direct or indirect damages, arising from or related to use of this white paper.

Samsung Electronics Co., Ltd.

129 Samsung-ro, Yeongtong-gu Suwon-si, Gyeonggi-do 16677, Korea

www.samsung.com

2015-11

SAMSUNG