

Samsung SSD PM963

An NVMe SSD engineered for data center environments



Performance that outclasses SATA solid-state drives

Experience approximately four times the performance of SATA SSDs

Data center SSDs typically tend to run around the clock every day of the year. Furthermore, they involve multiple parallel users performing various types of mixed random read and write operations on huge datasets.

Today's data centers also must serve requests from diverse applications and databases with numerous kinds of services, each requiring different performance needs.

Therefore, IT managers want SSDs that are optimized to deliver outstanding mixed workload performance, which is typical of diverse data center applications that simultaneously access the same device.

Engineered for outstanding performance under varying data center workloads, the Samsung PM963 NVMe SSD is optimized for mixed workloads, especially read-intensive environments.

It provides instant responsiveness to the host system by applying the PCIe® 3.0 interface standard, as well as the highly efficient NVMe protocol.

The PM963 also delivers superb reliability for continual operation, regardless of unanticipated power loss.

Using its proven expertise and wealth of experience in cutting-edge SSD technology and memory solutions, Samsung SSDs help data centers operate continually at the highest performance levels. Samsung has the added advantage of being the sole manufacturer of all of its SSD components, ensuring end-to-end integration, quality assurance and the utmost compatibility.

The Samsung PM963 NVMe SSD delivers:

- **High performance:** Approximately four times the performance of SATA SSDs for sequential and sustained random read, using the Samsung state-of-the-art V-NAND flash memory. A new optimized Samsung NVMe controller, with a PCIe Gen3 x 4 host interface, supplies up to 32 Gb/s of bandwidth.
- **High density:** Flash memory chips in the Samsung V-NAND architecture of the PM963 stack cells in a vertical arrangement and utilizes TLC (triple-level cell) technology to increase storage capacity, while retaining the same

footprint. The PM963 is available in 960 GB, 1.92 TB and 3.84 TB disk capacities.

- **Exceptional value:** By providing better performance-to-power efficiency, the PM963 is more affordable than SATA SSDs in terms of TCO (total cost of operation). With approximately four times the IOPS per GB, data center operators can continue to scale workload utilizations on their servers, decreasing the cost of computing.
- **Small M.2 form factor:** Almost one quarter smaller in size when compared with the HHL (half-height, half-length) card SSDs, the PM963 M.2 drive provides high capacities up to 1.92 TB thanks to Samsung V-NAND technology.

Optimized for data center environments

To meet the demand for high utilization, high-duty-cycle data centers, the PM963 uses firmware that prioritizes QoS (quality of service) for sustained random workloads. This firmware keeps all virtual machines running quickly and smoothly. The firmware is also optimized for always-on, always-busy workloads, ready to respond quickly to incoming host requests.

The Samsung SSD Toolkit, a proprietary developed tool for Samsung NVMe SSDs, enables monitoring of the essential health status of the PM963. The toolkit identifies abnormal symptoms that occur during data center operations.

Also, the PM963 leverages the same controller and NAND flash memory as high-volume laptop PCs, allowing data centers to deploy large quantities of NVMe SSDs cost effectively.

TLC V-NAND flash memory

Samsung V-NAND technology delivers reliable and consistent performance for today's demanding data-centric world. The technology features a unique design that stacks 48 layers on top of one another instead of trying to decrease the pitch size of the cells.

This revolutionary vertical design overcomes the capacity limitations of traditional NAND technology. It also applies innovative CTF (charge trap flash) technology, which prevents data corruption caused by cell-to-cell interference.

Traditional planar NAND memory requires the creation of sets of complex program algorithms to prevent data corruption caused by cell-to-cell interference.

Revolutionary Samsung V-NAND technology overcomes capacity limitations

However, Samsung V-NAND is virtually immune to cell-to-cell interference and does not need to go through a complex program algorithm to write data. By eliminating this step, the PM963 can write data up to two times faster than traditional planar NAND flash memory.

The synergy of both structural and material innovations leads to improved speed, power efficiency and endurance. Samsung has revolutionized the storage industry by shifting the planar NAND to a vertical structure.



Figure 1. Samsung V-NAND architecture overcomes the capacity limitations of traditional NAND technology

NVMe SSDs versus SATA SSDs

Today, SSDs are available with a variety of system interfaces, based on the performance requirements of the applications. Common interfaces include SATA, SAS (serial attached SCSI) and PCIe.

AHCI/SATA interface limitations

The AHCI mode set for SATA fulfills its intended architecture and design goals quite well for devices such as HDDs and optical drives. However, 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). In an SSD, the information stored on the device actually resides on NAND chips.

Earlier NAND chips had speeds of approximately 50 MB/s per chip with usually 4 to 8 chips on a typical SSD. This accounted for speeds up to 400 MB/s (8 x 50 MB/s). Since a 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 sufficient for older NAND chips.

However, the NAND chips currently used in SSDs usually run at around 200 MB/s. With the advent of the latest NAND chips, such as Toggle[®] 2.x and ONFI (open NAND flash interface) 3.x, PCIe SSDs are capable of delivering speeds up to 400 MB/s per chip, resulting in overall speeds 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.

NVMe is built for the SSD

NVMe is a scalable host controller interface designed to address the needs of enterprise, data center and client systems that utilize PCIe-based SSDs. NVMe capitalizes on the low latency and parallelism of PCIe SSDs, mirroring the parallelism of contemporary CPUs, platforms and applications.

A SATA 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, thereby eliminating any intermediate protocols and allowing faster data flow and processing.

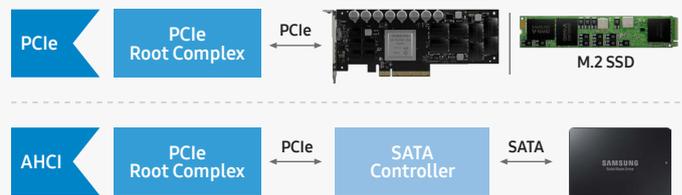


Figure 2. SSD architecture - AHCI versus PCIe interface

Superb performance over SATA SSDs

The PM963 delivers a wide bandwidth of up to 2,000/1,200 MB/s sequential R/W speeds respectively, using less than 7.5 W of power. With the help of the Toggle 2.0 NAND flash interface, the PM963 delivers random performance of up to 430K IOPS for 4 KB reads and up to 40K IOPS for 4 KB writes in the sustained state.^{1,2}

1. The PM963 performance is based on PCI Gen3 x 4, random performance measured using Fio[®] 2.1.3 in Linux[®] RHEL 6.5 with 4 KB of data transfer size in queue depth 32 by 4 workers and sequential performance with 128 KB of data transfer size in queue depth 32 by one worker.
2. SATA SSD performance measured using Fio 2.1.3 with queue depth 32, C216 Intel[®] SATA 6 G port.

The PM963 excels in performance when compared with the SATA SSD

The PM963 was compared with the Samsung SATA SSD for the performance comparison. Both SSDs used the same TLC V-NAND flash memory. The PM963 provided 3.8 times more bandwidth than the SATA SSD for sequential reads up to 2,000 MB/s and 4.3 times more IOPS than the competitor for sustained random reads up to 430K IOPS.

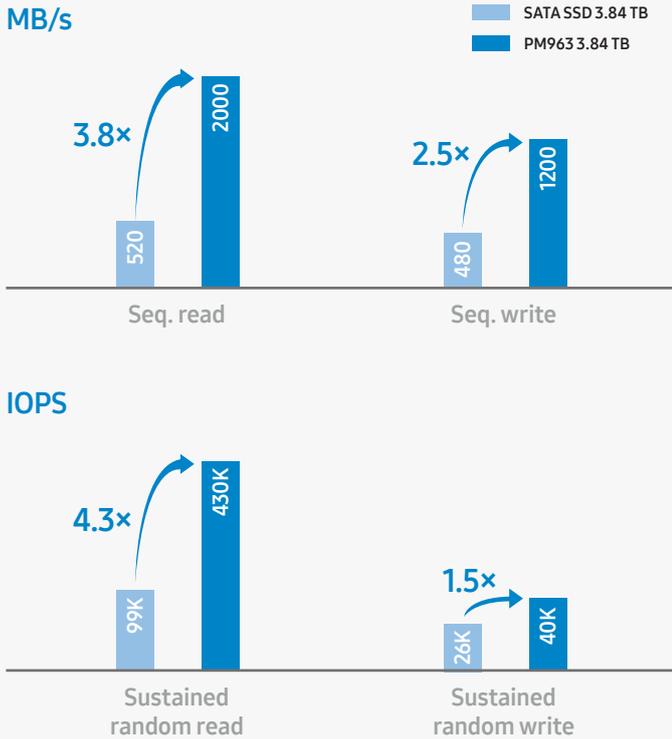


Figure 3. R/W performance comparisons between the Samsung SATA SSD and the PM963

Better performance-to-power ratio

The PM963 consumes maximum power of up to 7.5 W for sequential R/W in the active state, while delivering staggering performance improvements. In comparison, the SATA SSD consumes up to 2.8 /3.7 W for sequential R/W respectively to achieve its maximum performance. However, the PM963 achieves 47/31 percent better sequential R/W performance-to-power ratios respectively than the SATA SSD.

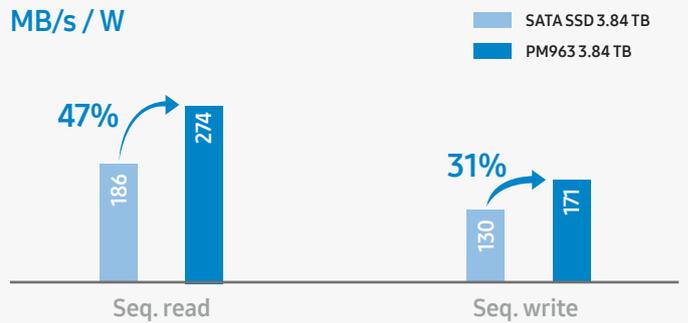


Figure 4. Performance-to-power ratios of the PM963 and the SATA SSD

Small M.2 form factor

As modern-day computers continue to get smaller, storage drives also need to get smaller, while at the same time continuing to provide the same or a higher level of performance. The PM963 comes in 2.5-inch and M.2 form factors using PCIe lanes, supported universally by almost all vendors in the market today. With just 22 mm width and 110 mm length, the PM963 M.2 SSD is almost a quarter smaller in size when compared to the HHL card SSDs, which are popular in data center and enterprise environments.



Figure 5. PM963 M.2 form factor

Enterprise-grade power-loss protection

During normal power-off periods, the host server allocates time to maintain data integrity by transmitting standby commands to each device. In the event of an unexpected power loss, though, the cached data in the internal buffers (DRAM) of the storage device can be lost. This event can occur with power outages or users unplugging devices from the system.

However, the PM963 has been designed to prevent data loss with its PLP (power-loss protection) architecture. Upon detection of an external power failure, the PLP solution helps reduce the risk of data loss. Using the stored energy from tantalum capacitors, it immediately provides enough time to transfer all user data and meta-data cached in DRAM to the flash memory.

Fail-safe technologies protect data from unexpected disruptions

Dynamic thermal throttling

Although an SSD is considerably fast and performs tasks within a matter of seconds, sometimes its internal temperature may rise beyond a certain threshold. A rise in temperature can happen when the SSD experiences a large number of requests continuously. The increase in temperature could damage the hardware components and, eventually, the data in the storage media.

To address this issue, the PM963 is equipped with DTT (dynamic thermal throttling), which controls the operation of the device based on temperature. When the device reaches a certain threshold, the DTT mode maintains the maximum level of performance that can be achieved while keeping the temperature of the device below its throttling limit.

As continuous operations increase the device's temperature, the PM963 follows a three-step performance reduction DTT methodology to protect itself from overheating. The DTT is activated when the NAND temperature exceeds the threshold and reduces the performance of the device by one level, enabling 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 and reduces the performance one more level. This step-by-step method is repeated until the temperature of the device comes to a low steady state. This three-step process helps ensure minimal performance drops, as well as data and device protection.

Hot swap

Hot swap involves replacing an SSD, HDD, CD-ROM drive, power supply or other device while the computer system is running, requiring no shutdown and rebooting. A device is replaced when it fails or when the data needs to be replaced with other data. Hot plug is a term similar to hot swap, which means that the operating system normally recognizes the device even when the device is removed (hot removal) or inserted (hot add) while the system is running.

For the PM963, the 2.5-inch form factor SSD supports the hot plug and hot swap functions. And multiple tantalum capacitors ensure stable data integrity—even when the system is in a SPOR (sudden power off recovery) state.

Remote health monitoring

The PM963 provides a remote health-monitoring feature through the Samsung SSD Toolkit. The Samsung SSD Toolkit is proprietary software designed to help users with easy-to-use SSD management and diagnostic features for server and data center use.

The CLI (command line interface) tool currently supports Samsung data center and enterprise NVMe SSDs, and supports Linux. The Samsung SSD Toolkit is available on the Samsung SSD web site.³ The following table highlights the health-monitoring information, which can be customized.

3. http://www.samsung.com/global/business/semiconductor/minisite/SSD/global/html/support/server_downloads.html

	Contents
Media or NAND status	Uncorrectable ECC status
	RTBB (run-time bad block) information
Interface status	Transfer, link, buffer, phy status
Performance status	Link speed down, DTT monitor, performance indicator, performance parameter
Tantalum status	Tantalum capacitor status
Thermal status	DDT status
Remote FA (failure analysis)	The reason why issue happens

Table 1. List of customizable health-monitoring information

Remote health monitoring supports:

- **Health monitoring:** Provides vital drive status information and helps users update firmware, measure drive performance, initialize drives, calculate drive lifetime and more.
- **Remote health monitoring for FA:** Provides a smart way to resolve field issues using the remote health-monitoring feature. Handling issues this way reduces turnaround time compared to the traditional way of handling issues, such as dispatching an engineer, and provides quicker resolutions. The customer can receive the initial resolution in two steps: first by running the SSD Toolkit and second by sending the debug information back to Samsung and effectively getting the resolution.

Specifications

	Samsung PM963		
Form factor	2.5 inch, M.2		2.5 inch
Capacity	960 GB	1.92 TB	3.84 TB
Host interface	PCIe Gen3 x 4 Lanes		
Spec compliance	NVMe spec rev. 1.2 (partial), PCI Express CEM spec rev. 3.0, PCI Express base specification rev. 3.0		
NAND flash memory	Samsung V-NAND		
Power consumption	Active read/write: Up to 7.5/7.5 W, Idle: 2.5 W (typical)		
TBW (@4 KB random write)	1,366 TB	2,733 TB	5,466 TB
UBER	1 sector per 10^{17} bits read		
MTBF	2,000,000 hours		
Endurance	1.3 DWPD for 3 years		
Sequential read	Up to 2,000 MB/s		
Sequential write	Up to 1,200 MB/s		
Random read	Up to 430K IOPS		
Random write	Up to 40K IOPS		

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. The company is redefining the worlds of TVs, smartphones, wearable devices, tablets, cameras, digital appliances, printers, medical equipment, network systems, and semiconductor and LED solutions. For the latest news, please visit the Samsung Newsroom at news.samsung.com.

For more information

For more information about the Samsung Semiconductor, visit samsung.com/semiconductor.



Copyright © 2016 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.

Fio is a registered trademark of Fio Corporation.

Intel is a trademark of Intel Corporation in the U.S. and/or other countries.

Linux is a registered trademark of Linus Torvalds.

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

Toggle is a registered trademark of Toggle, Inc.

Samsung Electronics Co., Ltd. (Maetan-dong) 129 Samsung-ro, Yeongtong-gu, Suwon-si, Gyeonggi-do 16677, Korea www.samsung.com 2016-08