

Upgrade Your Enterprise with the SSD that Redefines Performance

PM9A3 SSD

Whitepaper



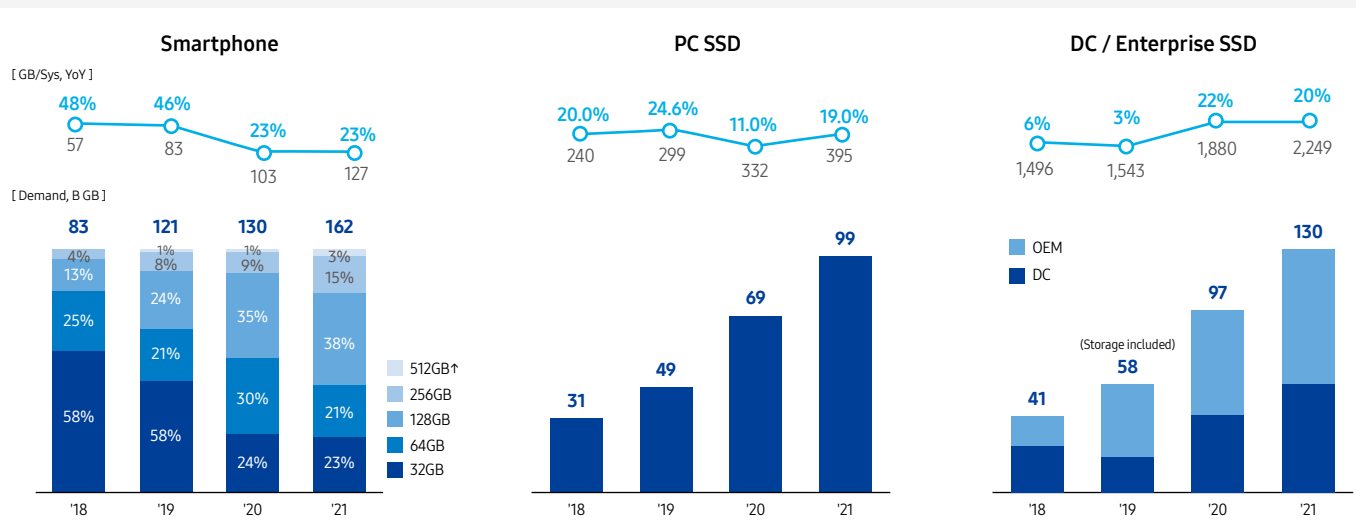
Executive Summary

Today, an explosion of data and an ever-expanding number of computing devices – spanning everything from IoT to 5G-enhanced edge computing, to exascale supercomputers – is redefining computing and storage demands.

These ongoing trends have wide-ranging ramifications for enterprise computing and web-scale data centers. And with digital services experiencing exponential growth, these trends are expected to continue in the post-pandemic era. This means that now more than ever, the world needs more powerful data centers, enterprise servers and storage solutions.

NAND Solution Market Outlook

Enterprise SSD market to see continuing growth due to data explosion



* Source : Samsung, Forward Insights

What's more, modern server workloads continue to change thanks to an influx of machine-generated unstructured data. These new workloads need to be analyzed in real time to facilitate everything from AI-assisted diagnostic imaging to high-frequency stock trading, as well as efforts like banking and credit card fraud detection, and to enable software to make real-time decisions about how best to deliver packages. Solid-state devices are better suited for meeting the storage demands of use cases like these than traditional spinning disks – hard disk drives (HDD).

Scalable, high-performance, energy-efficient storage devices must continue to keep pace with the market's evolving needs. Fortunately, industry innovations in the form of standards and technologies such as PCIe® Gen.4 and our sixth-generation V-NAND are proving equal to the task. Samsung's state-of-the-art PM9A3 NVMe™ SSD is a testament to this fact – one that is available in more form factors than any competing SSD.

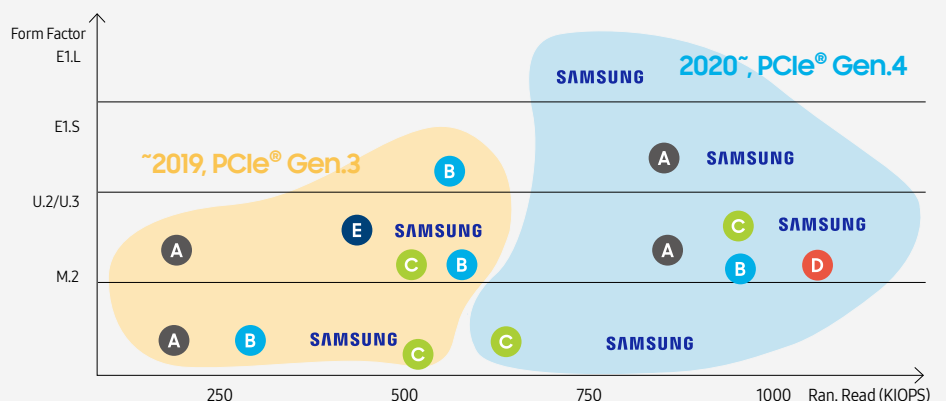
Before discussing the ins and outs of the PM9A3, we'll begin this paper with a brief history of the SSD industry, along with details on how the industry has grown and how its products are being used.

Technical trend of DC NVMe™ SSDs

From 2020, PCIe® Gen.4 based DC NVMe™ SSDs have been introduced continuously

- SSD's own performances are getting higher and various types of new form-factors are being adopted to meet host level requirements.

Samsung is the leading company to support best in class performances and all kinds of form-factor



SSD History & Market Growth

A Brief History of SSDs

SSD data storage uses integrated circuit assemblies to store data persistently, typically using flash or NAND flash memory. It has been close to 50 years since semiconductor researchers began developing the first solid-state drives, which store data on memory chips that are becoming increasingly faster, more powerful and energy-efficient, as opposed to the spinning magnetic disks used by traditional HDDs.

SSDs store data in thousands of tiny memory cells that are arranged on circuit boards. As the “solid-state” in their name suggests, SSDs have no moving parts, which allows them to be inherently faster than HDDs, as well as smaller, more durable, and more reliable.

However, when they were first introduced, SSDs were far too expensive, and initially lacked the storage capacities necessary for widespread use. For example, when Storage Technology Corp. introduced what some considered to be the first SSDs to be compatible with IBM mainframe computers in 1978, the company’s car-engine-sized machines were only capable of storing a meager 45 megabytes (MB) – roughly equivalent to the amount of data involved in 20 average-sized web pages today. On top of that, the exorbitant \$400,000 price tag per system scared away many potential buyers.

By comparison, today, SSD buyers can choose from dozens of offerings with huge, one-terabyte (TB) storage capacities for less than \$100. This is reflective of the power of Moore’s Law, which describes how each generation of technology utilizes denser circuitry, resulting in higher performance, less energy use, and lower cost per MB.

Not surprisingly, the SSD market has been growing rapidly as more SSDs replace traditional HDD devices. SSD technology, too, has rapidly evolved, with storage densities increasing and the price per MB steadily dropping. With the growth of data over the last decade, many enterprises have been accelerating their transition from HDDs to high-capacity SSDs.

SSD Market Growth

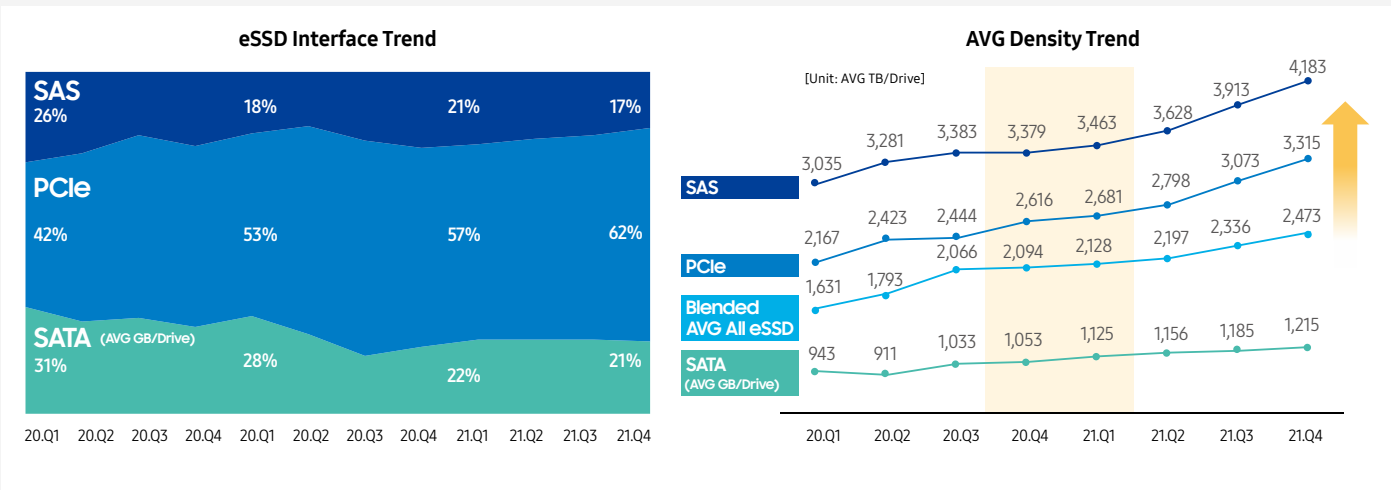
The industry’s overall rate of SSD sales growth has been driven by PCIe performance benefits, as well as increasing densities, improving performance, and lowering costs. Last year there was an impressive 22 percent growth in worldwide SSD sales, which rose from \$23 billion to \$32 billion according to the London-based market researcher firm Omdia. Moreover, in its Worldwide 2020–2024 Enterprise SSD and HDD Combined Market Overview, IDC projects that SSDs will steadily displace performance-optimized HDDs as primary storage in servers and storage systems.

The migration away from HDD technology and toward SSD enables systems to boot faster and load applications quicker, and improves overall system responsiveness.

SSD technology is now primed to become the leading storage medium for the big data age. Its inherent attributes of scalability, reliability and energy efficiency make it the premier option for the server market.

Industry Interface Trend: NVMe™ Drives eSSD Growth

PCI Express™ performance benefits drive interface conversion and higher density SSD – More than 8/16TB



* Source : Samsung, Forward Insights

Next-Generation SSDs

PM9A3: Gen.4 Drives Dynamic Data Traffic

Traditionally, data centers and enterprise storage have represented two separate markets for SSDs. While the former require storage solutions that are optimized for data center architectures and offer greater efficiency, SSDs for enterprises tend to emphasize special features and robust reliability.

Samsung's PM9A3 NVMe™ SSD offers best-in-class PCIe® Gen.4 performance, multiple platform efficiencies, a highly scalable design encompassing several form factors, and greater reliability with advanced V-NAND technology. The SSD is specifically designed for the next era of servers, both in data centers and for enterprises in general, meets the needs of servers for mainstream use, and is suitable for the web, applications, and file servers.

The following table highlights the PM9A3's advancements and compares them with the previous generation's enterprise SSD leader, the PCIe® Gen.3-based PM983.

PM9A3 Overview

		PM9A3	PM983
Interface		PCIe® 4.0 x 4	PCIe® 3.0 x 4
Controller		Elpis (8ch)	Phoenix (8ch)
NAND		V6 TLC	V4 TLC
Form-Factor		U.2, U.3, M.2, E1.S, E1.L	M.2, U.2
Capacity		15.36TB/7.68TB/3.84TB/1.92TB/960GB	7.68TB/3.84TB/1.92TB/960GB
Logical Block Address		512B(Default), 4KB	512B(Default), 4KB
NVMe™ Specification	NVMe™ version	1.4(Mandatory Only)	1.3c
	Sanitize	NVMe™ 1.3 Compliance (Block Erase, Crypto Erase)	Not Support
	Device Self Test	Support	Not Support
SED		TCG/Opal	TCG/Opal
Drive Writes per Day		1@5 years	1.3@3 years

Open Compute Project®

Meets Open Compute Project® Specification

The new PM9A3 fully complies with the Open Compute Project®'s (OCP) NVMe™ Cloud SSD Specification, which proves that it satisfies the rigorous demands of enterprise workloads and hyperscalers. This includes an emphasis on efficient, scalable operations and support, along with commitments to open source and to creating a positive impact.

EFFICIENCY

- All OCP contributions must have an efficient design some examples of efficiency may be
 - power delivery and conversion
 - thermal efficiency
 - platform performance(per-W for example)
 - reduction in overall infrastructure costs
 - reduction in code weight
 - reduction in latencies

SCALABILITY

- All OCP contributions must be scalable. This means that the technology is designed with the right supporting features to allow for its maintenance in large scale deployments. This may include:
 - physical maintenance
 - remote management
 - upgradability
 - error reporting
 - maintenance service (remote or on-premise)
- Management tools should be open-sourced or made available to adopters
- Supporting documentation for each contribution should enable adopters towards
 - installation guidance
 - initialization process
 - configuration information
 - how to obtain service support

OPENNESS

- All OCP contributions must be open. This encourages as much open source contribution as possible. OCP understands that in certain cases 100% open source contribution may not be possible
- Whether fully open source or not, a contribution should strive to comply with a set of already existing open interfaces, at the very least be able to provide one
- Providing a solution compatible with already existing OCP contributions is one way to implement existing(open) interfaces

IMPACT

- All OCP contributions must create meaningful positive impact within the OCP ecosystem. This may be attained by the introduction of:
 - efficiency gains
 - utilization of new technologies
 - products that are valuable for scale out computing, creating a multiplier effect by building on top of already existing OCP solutions
 - enabling a more robust supply chain by contributing alternative compatible solutions

* Source: <https://www.opencompute.org/files/OCP-Tenets-FINAL2-1.pdf>

Design Capability: New Reference design based on OCP specs

Samsung is the leading company to provide server and datacenter system designers' capability.

- PM9A3 fully implements OCP's specification(Poseidon project)
- Offering the most 1U server-optimized form-factor, the PM9A3 will improve space utilization, add PCIe® Gen.4 speeds, enable increased capacity and more



Poseidon SW

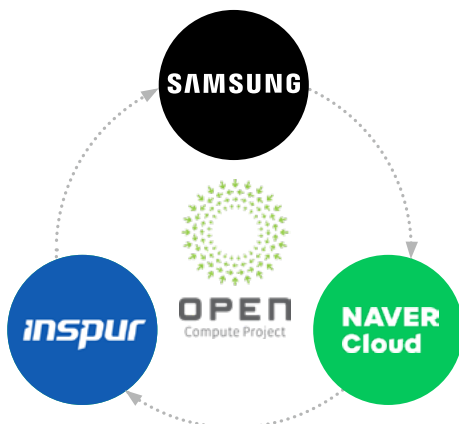
Open-source Storage OS for NVMe™F



Poseidon HW

PCIe® Gen.4 E1.S SSD Ref. System

OCP based 3-way Partnership



Tremendous Performance in Storage System

PM9A3 delivers up to 6,950 MB/s in sequential read speed with form factor flexibility

Category	Model	Value	Improvement
Sequential Read	PM983a	3,100 MB/s	▲ 3,850
	PM9A3	6,950 MB/s	
Power Efficiency	PM983a	188 MB/s Per watt	▲ 95
	PM9A3	283 MB/s Per watt	

* Source : <https://www.samsung.com/semiconductor/newsroom/tech-leadership/samsung-encourages-open-source-innovation-through-the-open-compute-project/>

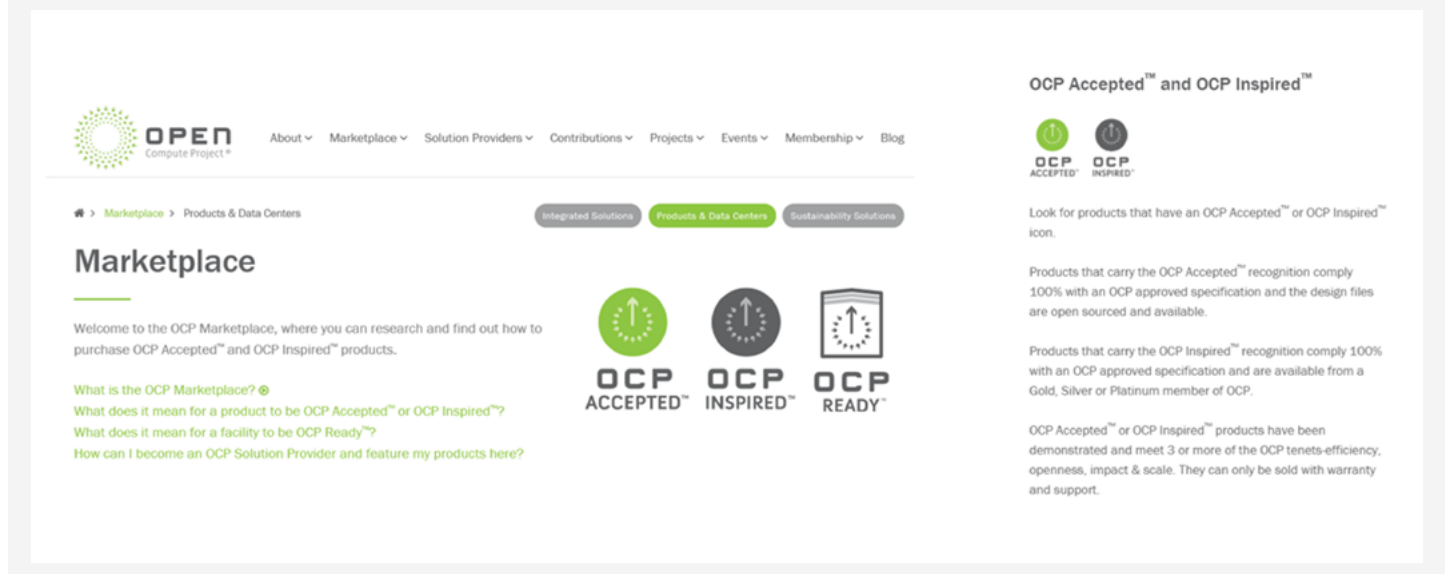
Open Compute Project®

Recently, the OCP Foundation recognized the PM9A3 as an OCP Inspired™ product. The OCP Inspired™ certification attests that the PM9A3 has been formally reviewed by the OCP Foundation and meets the approved specification. Beginning in the third quarter of 2021, the OCP Foundation will provide a marketplace that offers major server manufacturers, as well as general customers such as SMBs (small and medium-sized businesses) access to channels that will enable them to purchase Samsung SSDs.

Design Capability: New Reference design based on OCP specs

The OCP foundation recognizes PM9A3 as OCP Inspired™ product

- Customers easily can reach out Samsung SSD PM9A3 through OCP Marketplace from 3Q



* Source : <https://www.opencompute.org/products>

PCIe® Gen.4 Performance

Industry-Leading PCIe® Gen.4 Performance

The bandwidth of PCIe® Gen.4 is twice that of the prior Gen.3 interface, which means that the new PCIe® Gen.4-based PM9A3 offers greater throughput and higher performance than competing PCIe® data center SSDs.

The PM9A3 also leverages Samsung’s technological know-how to deliver enormous performance gains across sequential and random operations. With random write performance being one of the most important factors when designing servers, the PM9A3 enables 200,000 IOPS random writes, and features hardware automation technology to further enhance random write performance.

PCIe® Gen.4 Performance

PM9A3 provides best-in class PCIe® Gen.4 performance for all kinds of form-factor

- Random write performance is the key performance of PM9A3.

Even though high performance, PM9A3 consumes small amount of power for datacenters’ power efficiency

Form Factor		U.2 (2.5”mm 7mmT)				M.2 (22x110mm)		
Capacity		7.68TB	3.84TB	1.92TB	960GB	3.84TB	1.92TB	960GB
Sequential (128KB, MB/s)	Read	6,700	6,900	6,800	6,500	5,500	5,500	5,000
	Write	4,000	4,100	2,700	1,500	2,000	2,000	1,400
Random (4KB, KIOPS)	Read	1,100	1,000	740	580	800	800	550
	Write	200	180	130	70	85	85	60
Power (Watt)	Active (Read/Write)	11/13.5	11/13.5	10/12.5	9.5/8	8.2/8.2	8/8.2	7.5/6.5
	Idle	3.5	3.5	3.5	3.5	2.5	2.5	2.5

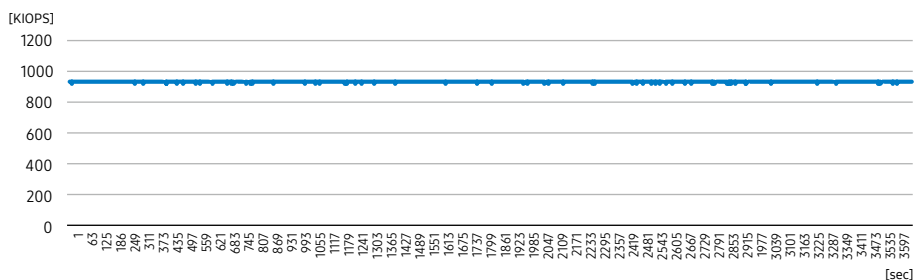
* Performances of E1.S, E1.L and U.3 form-factor are similar to U.2 based products and are providing to contracted and selected OEM companies For processing this project, Samsung will not provide detail numbers.

Moreover, in addition to offering industry-leading performance, the PM9A3 offers better IOPS consistency than market alternatives. An SSD’s performance is generally governed by controller channels – the physical paths through which data is bidirectionally transferred between the controller and the NAND.

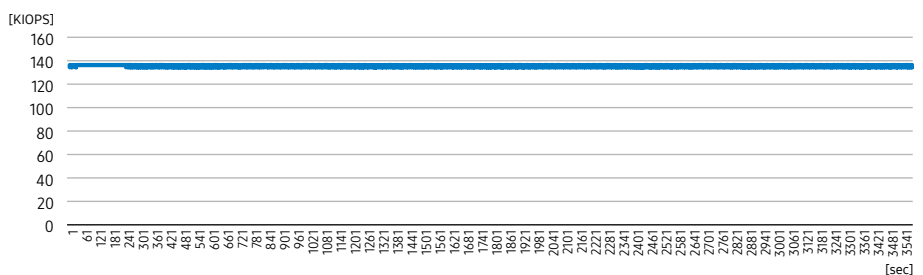
One might expect that SSDs with eight channels tend to be more power efficient and cheaper than those with 16 channels or more. However, thanks to Samsung’s know-how and hardware automation, the eight-channel PM9A3 has been optimized for system-level integration, and manages to provide consistent or better performance than competing 16-channel SSDs.

PM9A3 shows better stable IOPS consistency with higher performance

4KB Ran.Read



4KB Ran.Write

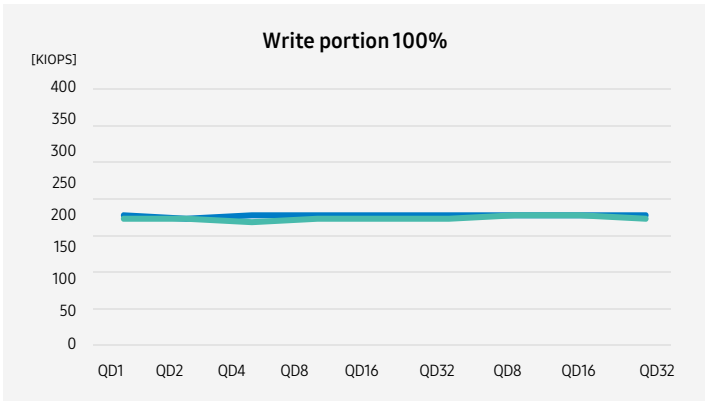
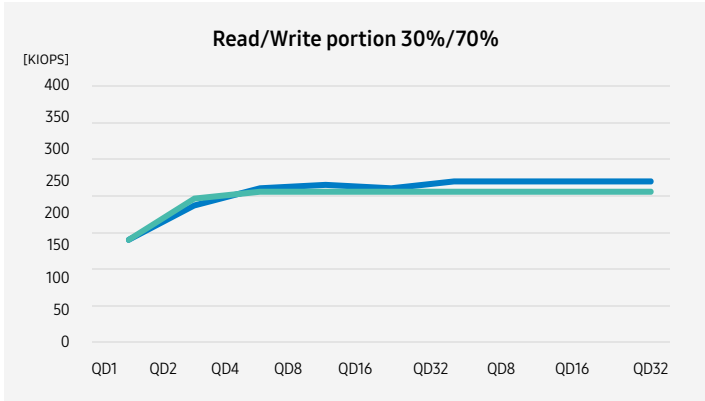
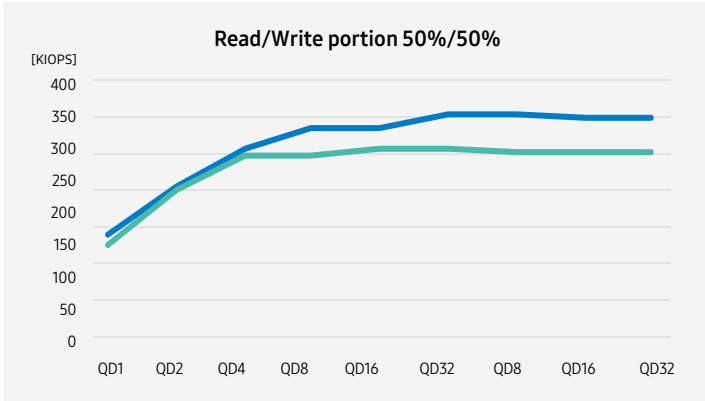
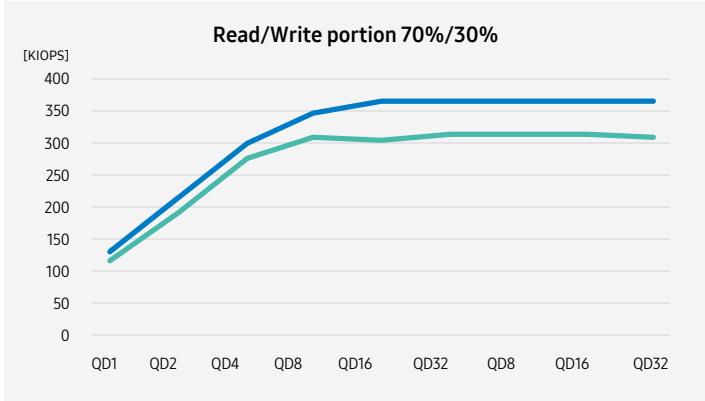
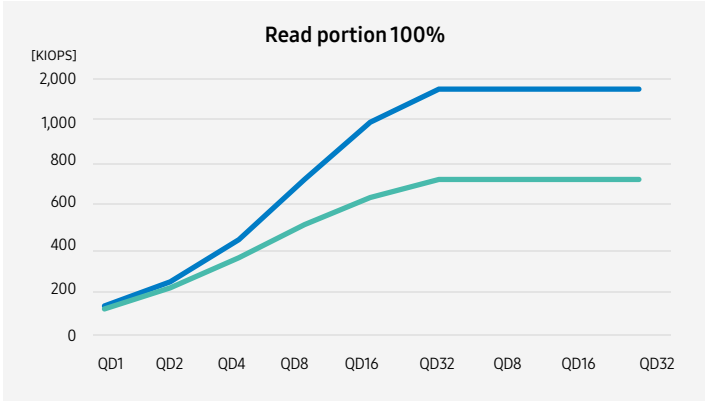


PCIe® Gen.4 Performance

4KB Ran. Performance of 3.84TB varying workload patten compared to one of competing 16-channel based SSD.

PM9A3 provides robust 4KB Ran. performance in various types of mixed workload pattern
Datacenter customers are usually focusing on 70% Read case(mixed workload 7:3), PM9A3 shows up to 20% higher IOPS.

PM9A3 A



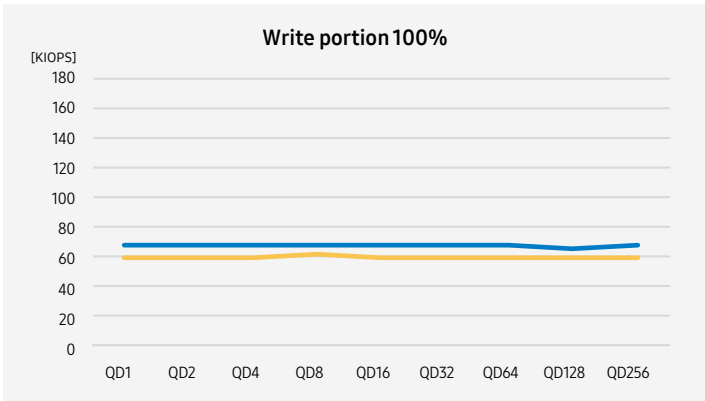
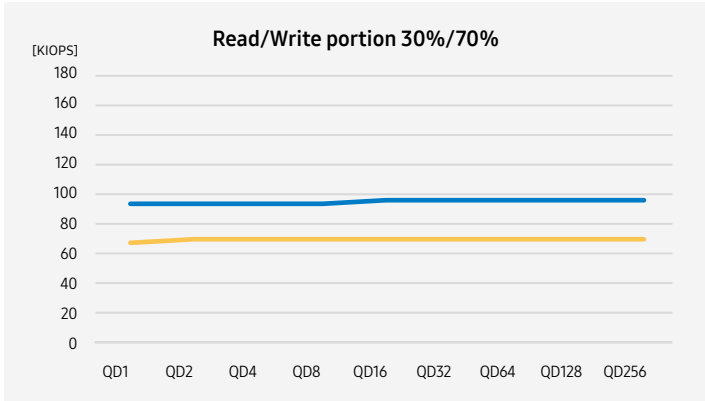
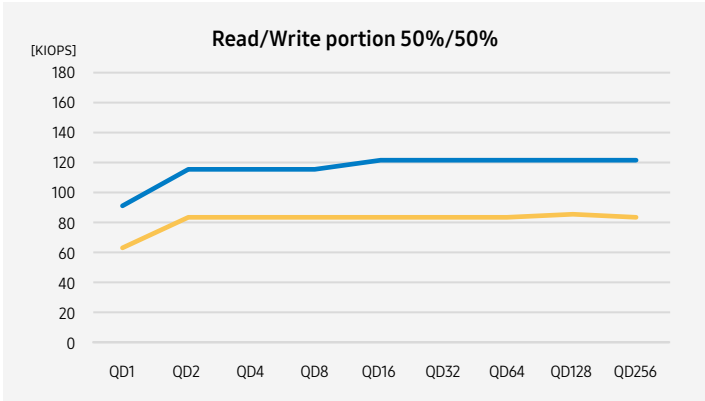
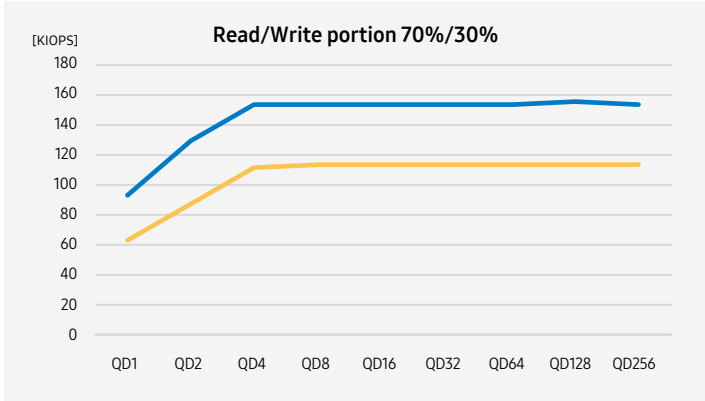
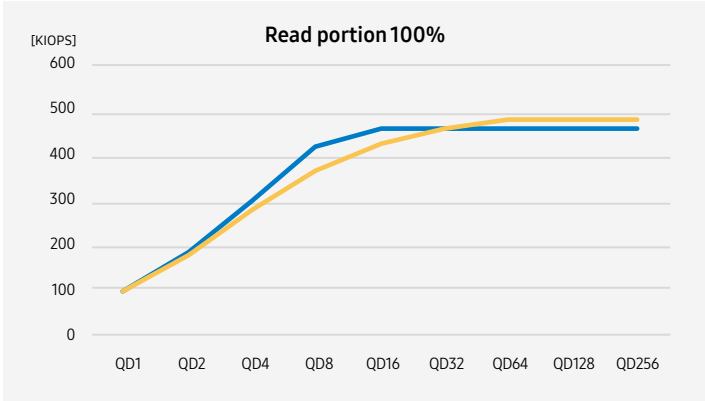
PCIe® Gen.4 Performance

In addition, when managing mixed workloads, the PM9A3 offers significantly better random read/write performance than conventional eight-channel SSDs.

PM9A3 provides robust 4KB Ran. performances in various types of mixed workload pattern

Datacenter customers are usually focusing on 70% Read case(mixed workload 7:3), PM9A3 shows up to 27% higher IOPS

PM9A3 1.92TB B 1.92TB

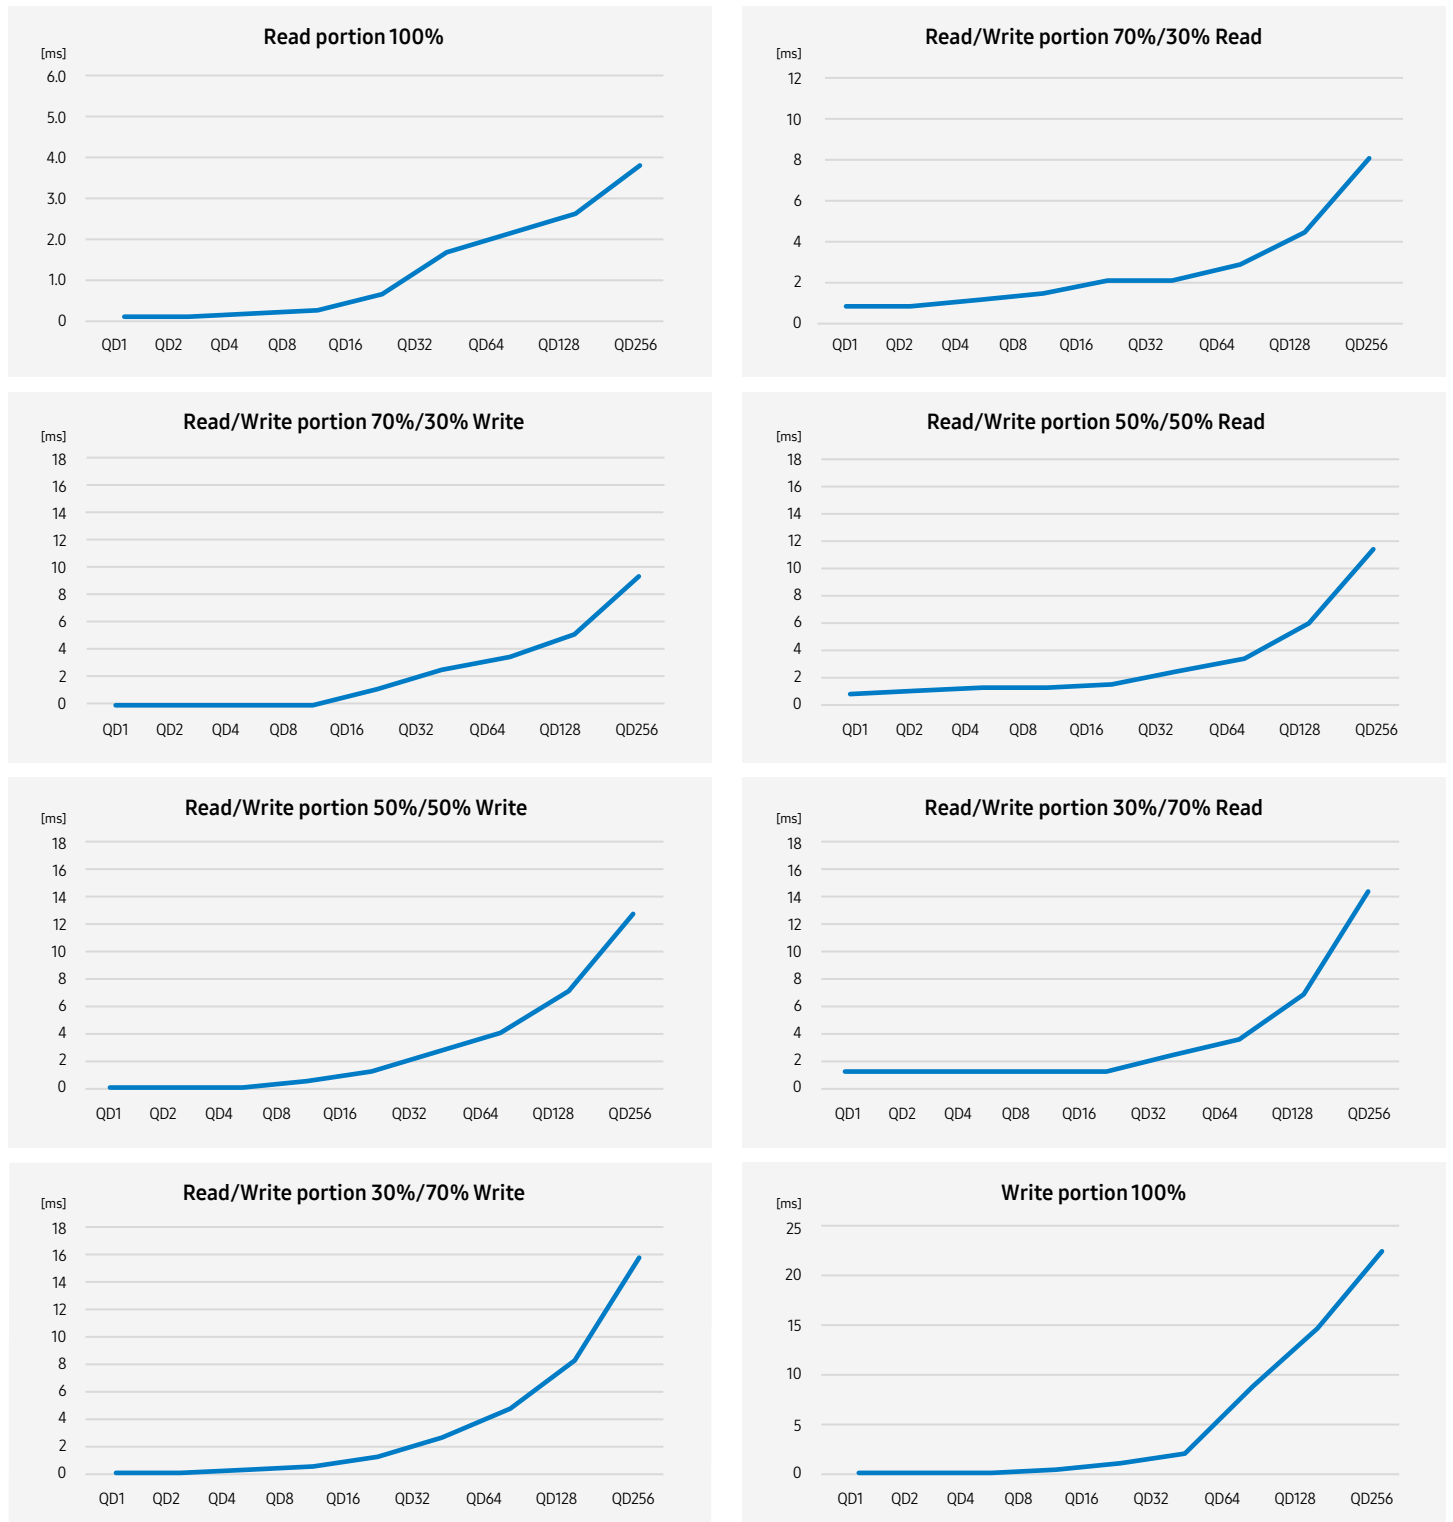


PCIe® Gen.4 Performance

The PM9A3's random average latency scores, too, are consistent and competitive when managing various workloads, while its random write latency scores routinely exceed those of conventional PCIe® Gen.4 SSDs. In quality of service (QoS), an important measurement of SSD performance, the PM9A3 offers an average latency that's consistently more stable across various QoS workloads. Providing significantly lower QoS at 99.99 percent, the PM9A3 is capable of enhancing host systems with robust stability.

4KB QoS 99.99%(ms) of 3.84TB varying workload patten

PM9A3 provides much lower QoS in 99.99%, it can provide robust stability for host system



Multiple Form Factor

A Form Factor for Every Need

Samsung's PCIe® Gen.4 NVMe™ SSD portfolio is robust in the market, and includes five different form factors – U.2, U.3, M.2, E1.S, E1.L – of various widths and lengths. This adds flexibility to installations and enables the SSDs to meet virtually any server requirements. Capacities range from 960GB at the lower end to 15.36TB for U.2 – more than double the capacities of earlier technologies. The M.2 version offers up to 3.84TB of storage capacity, while the E1.S model provides up to 7.68TB of storage space.



Design Capability: Multiple Form Factor

Samsung can meet various form-factor needs from server manufacturers and DC architecture designers – M.2, U.2, E1.S, E1.L and U.3

SSD Form factors

"Form factor" refers to the size and physical configuration of a device.

- At first, server manufacturers needed to standardize storages' dimensions for their systems' design
- As much higher capacities are needed, F/F has become as constraints for scalability, so many new F/Fs have been introduced to meet new demands of server market, for example, high capacity servers and new interface for better performance

	U.2	E1.S	E1.L	E3.S/E3.L	M.2
					
Dimension (WxL, mm)	69.85 x 100.2	33.6 x 118.75 (Others)	38.4 x 318.75	76 x 113 (Short) 76 x 142 (Long)	22 x 110
Connector	SFF-8639				M.2 Connector
Max Power	25W	12/16/20/25W	25/40W	20/35/40/70W	8.25W
Hot-Swap	0	0	0	0	X
Dual-Port	0	0	0	0	X
Max. Capacity	32TB	16TB	32TB	32TB	4TB
Number of PCIe® Lane	4	4	4/8	4/8/16	4

The New Height of Dynamic Speed

Uncompromising Performance

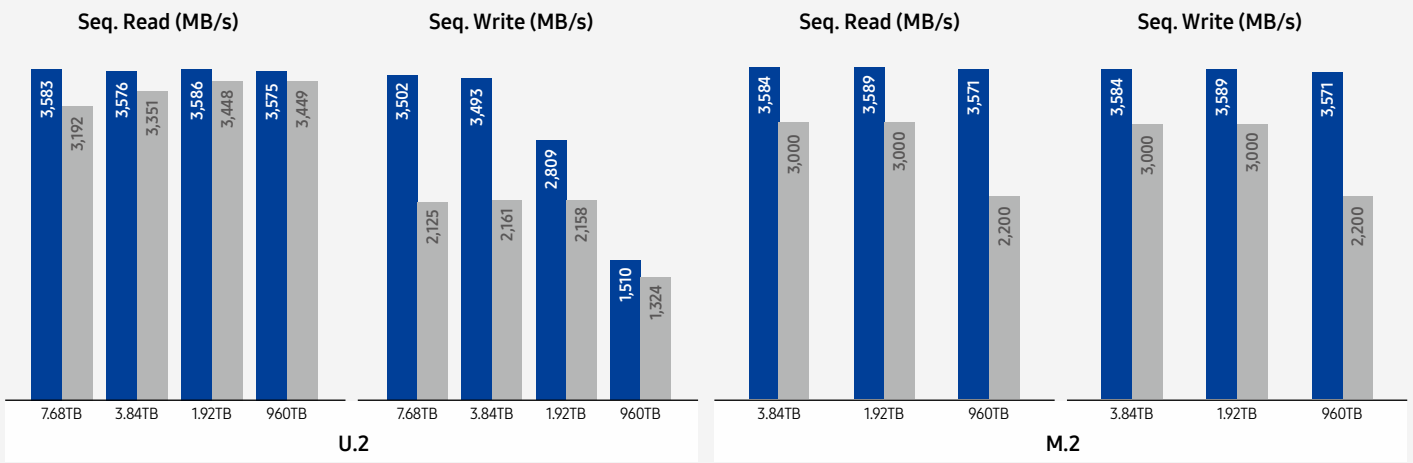
Unlike other PCIe® Gen. 4 SSDs, which can face compatibility issues when utilized with previous-generation servers, the PM9A3 offers superior write speeds even when paired with PCIe® Gen. 3 servers. This increase is especially impressive when it comes to random writes, and can be attributed to PCIe® Gen. 4's enhanced performance. The results are a sequential write speed of 3,500MB/s, a random read speed of 800,000 IOPS, and a random write speed of 200,000 IOPS.

PCIe® Gen.3 Sequential Performance

PM9A3 shows up to 1.3x times faster Seq. Write than its previous generation PM983 with PCIe® Gen.3 Interface.

- Improvements of read performances around 10% are not outstanding than write due to PCIe® Gen.3 interface limit

■ PM9A3 ■ PM983

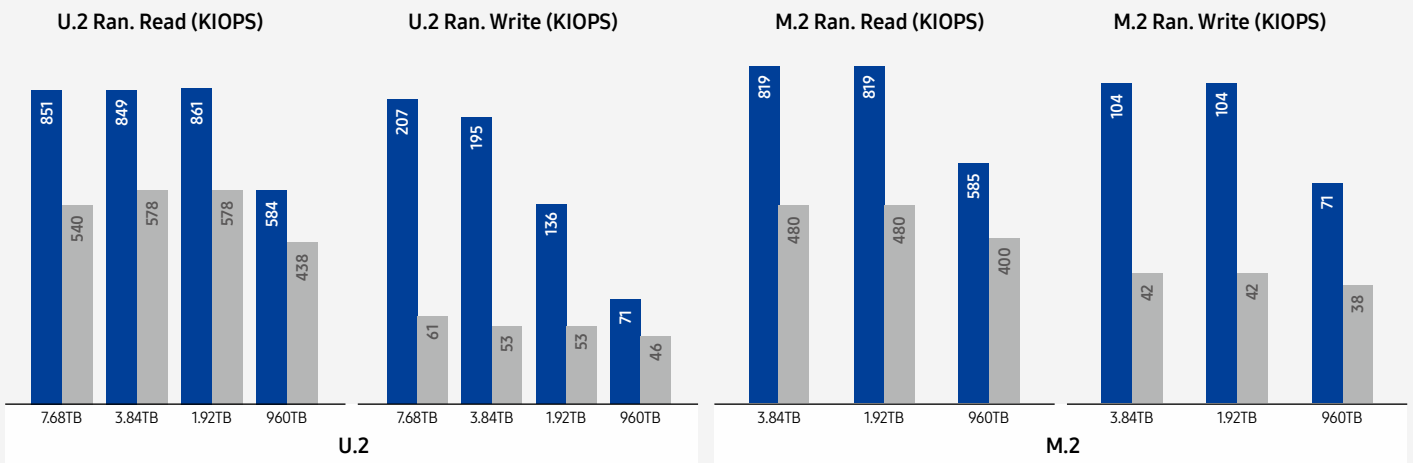


PCIe® Gen.3 Random Performance

PM9A3 shows up to 3.6x times better Ran. Write than its previous generation PM983 with PCIe® Gen.3 Interface due to improved V6 NAND.

- For the case of Ran. Read, PM983 provides up to 48% better than PM983
- NAND characteristics affects random performances directly compared to seq. performances
 - V6 NAND speed is 33% faster than V4, PM983 is based on V4 NAND

■ PM9A3 ■ PM983



Energy-Efficient Operation

Raising performance expectations with efficient power management

The PM9A3 features Samsung’s sixth-generation, 128-layer V-NAND, which offers up to 10 percent lower latency and consumes 15 percent less power than the company’s previous-generation V-NAND(96 Layer, V5). The SSD also boasts a power efficiency of 238MB per second per watt for sequential writes – a 50 percent improvement over its predecessor – which improves the power efficiency of data center operations.

Power Efficiency

Regardless PCIe® version, PM9A3 provides low power consumption with fast performances.

Lower power consumption is the key factor of TCO for datacenters.

- Lower power means lower heat generation and it directly affects cooling cost of systems and datacenters

12V supply condition		U.2 (2.5", 7mmT)				M.2 (22x110mm)		
Capacity		7.68TB	3.84TB	1.92TB	960GB	3.84TB	1.92TB	960GB
Active(Watt)	Read	11	11	11	9.5	8.2	8	7.5
	Write	13.5	13.5	13.5	8	8.2	8.2	6.5
Idle (Watt)		3.5	3.5	3.5	3.5	2.5	2.5	2.5
Form Factor Max Power		25W (Watt)				8.25W (Watt)		

1) Power consumption was measured in the 12V power pins of the connector plug in SSD. The active and idle power is defined as the highest averaged power value, which is the maximum RMS average value over 100 ms duration.

2) The measurement condition for active power is assumed for Maximum power between sequential or random performance in PCIe® Gen.4

3) The idle state is defined as the state that the host system can issue any commands into SSD at any time

Efficient Power Management

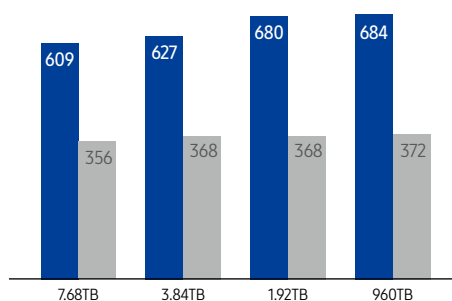
Seq. Read

PM9A3 provides up to 1.8x times better power efficiency than previous generation, PM983

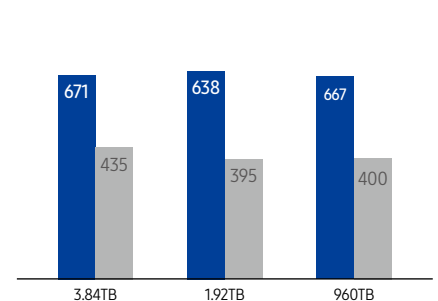
- Even though PM9A3 is PCIe® Gen.4 based SSD, less power consumption with higher performance

■ PM9A3 ■ PM983

U.2 Seq. Read (MB/s per Watt)



M.2 Seq. Read (MB/s per Watt)



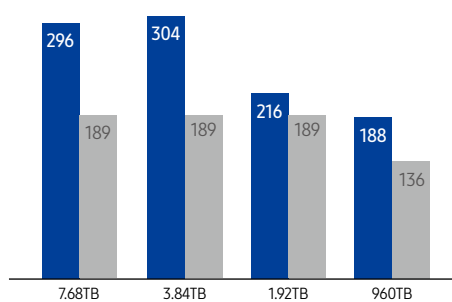
Efficient Power Management

Seq. Write

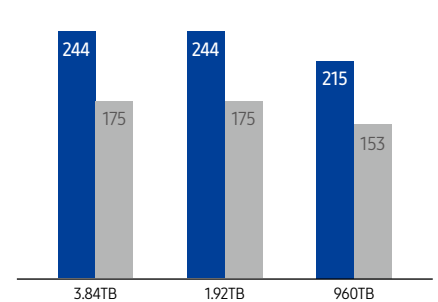
PM9A3 provides up to 1.6x times better power efficiency than previous generation

■ PM9A3 ■ PM983

U.2 Seq. Write (MB/s per Watt)



M.2 Seq. Write (MB/s per Watt)



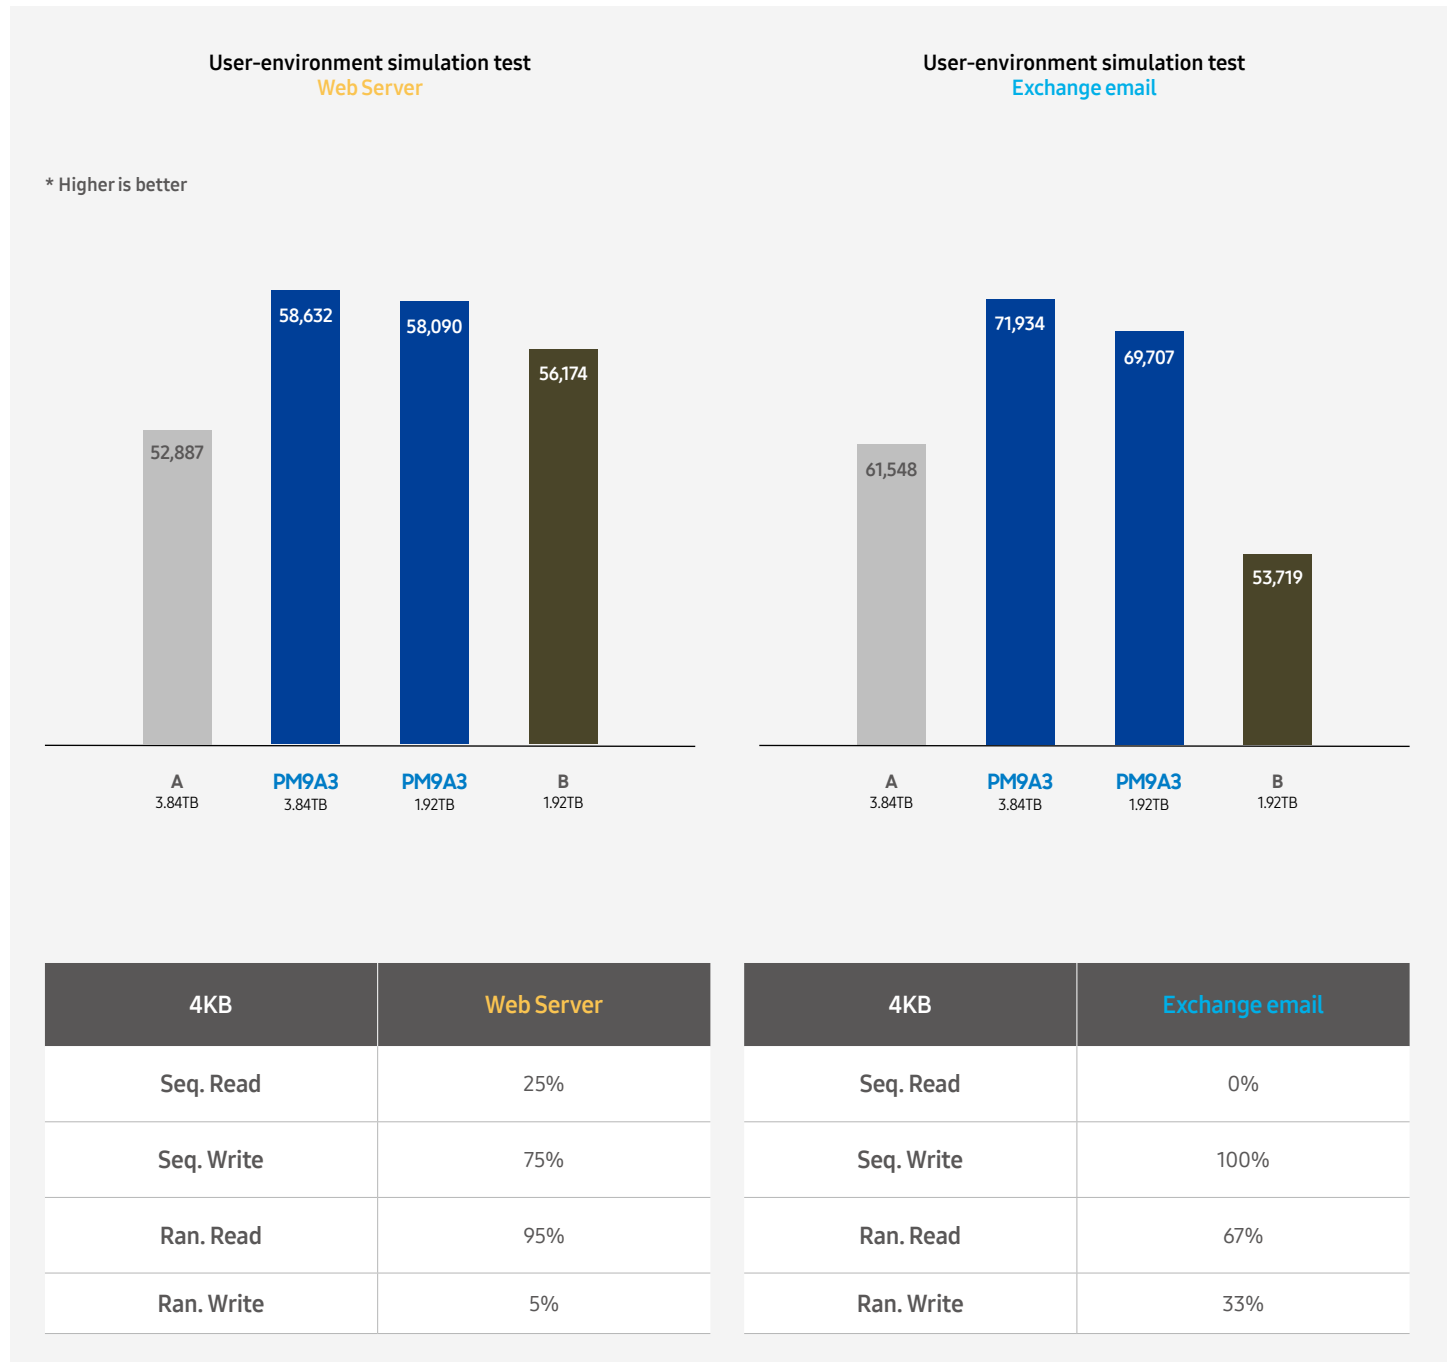
Of course, a key part of maximizing the efficiency of enterprise and data center servers involves correcting errors quickly when issues arise. The PM9A3’s Samsung Debug technology enables IT administrators to do just that. Based on NVMe™ telemetry, this convenient feature offers a basic diagnosis of a system’s status, and allows log data from failed drives to be sent to Samsung remotely to produce in-depth failure analysis (FA) reports.

Real-World Benefits

Optimal for Enterprise Workloads

Utilizing simulated tests, Samsung's PM9A3 outperformed market alternatives across multiple workloads, including a web server and an exchange email sever.¹

1) Workloads in test scripts featured 4KB transfer size with 100% access, 8KB-aligned on I/Os.



Real-World Benefits

Using a use case with a MongoDB application measured through the Yahoo Cloud Service Benchmark (YCSB), the PM9A3 completed the test in up to 32 percent less time and delivered nearly 25 percent greater throughput when compared to an alternative PCIe® Gen.4 SSD.

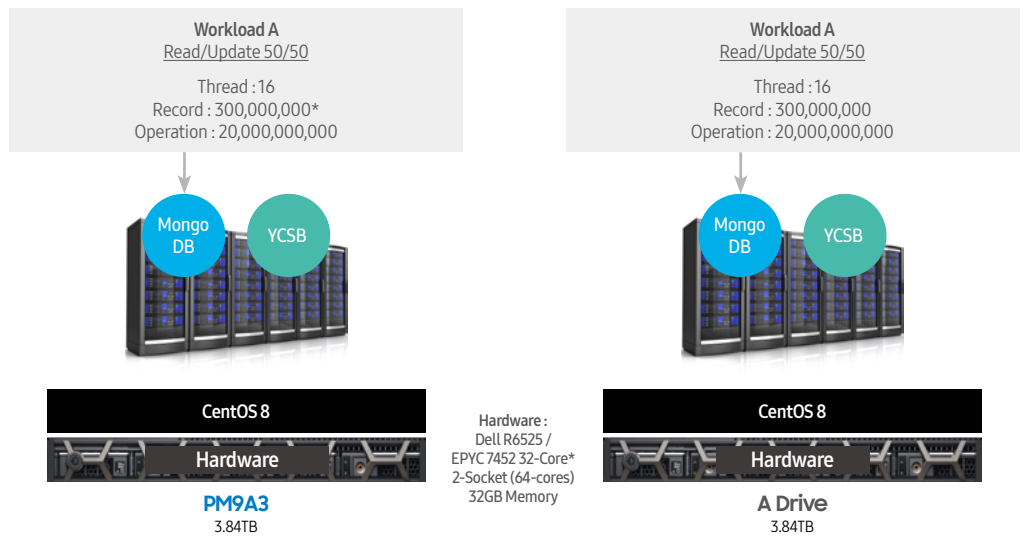
Real world example: MongoDB + YCSB (Yahoo Cloud Service Benchmark)

Samsung ran a real-world application to demonstrate the excellence of PM9A3 when running mixed workloads

- MongoDB is one of the most popular database management systems
- YCSB is a database benchmark widely used to access cloud service quality, and can be applied to various databases such as AWS S3, Aerospike, Cassandra, and so on

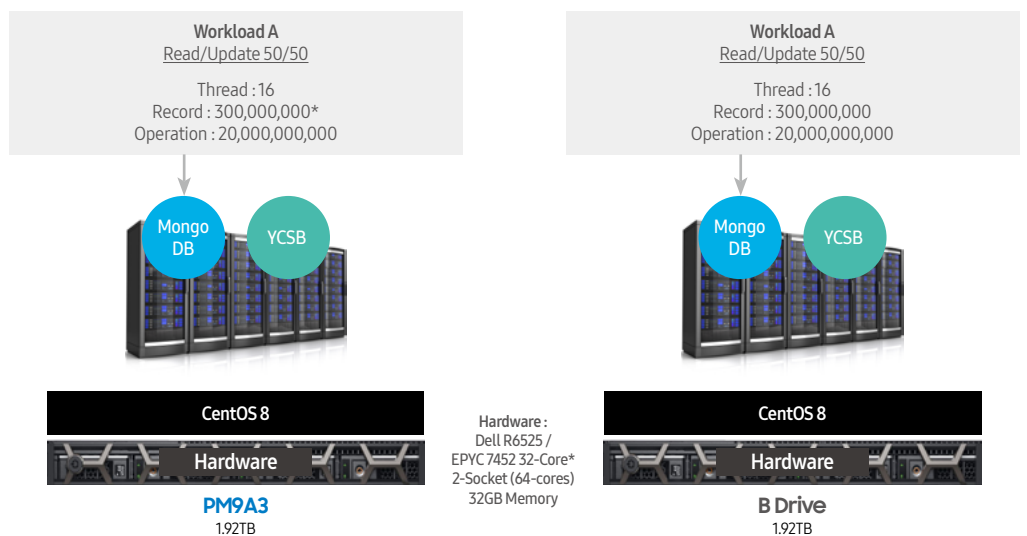
Test environments and scenario

3.84TB



Test environments and scenario

1.92TB



* 300M record counts equal to 390GB amounts of data, operation counts 2 Billion consist of Read operation 1B and update 1B.

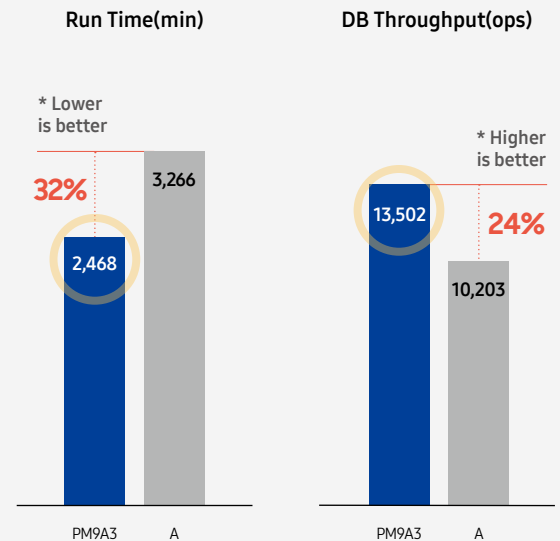
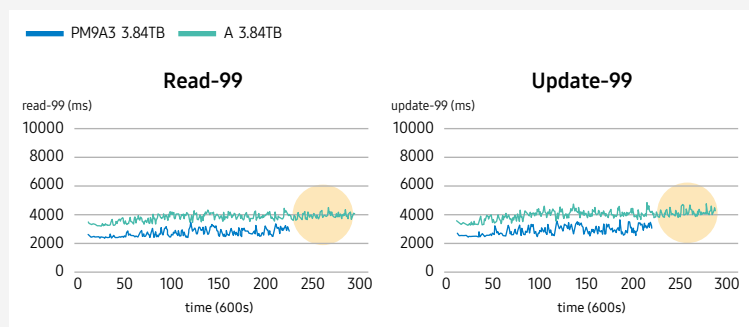
Real-World Benefits

Real world example: MongoDB + YCSB (Yahoo Cloud Service Benchmark)

PM9A3 shows 32% time saving and 25% more data operation in 3.84TB case

- Even though 8x based PM9A3 (lower datasheet performance), it provides much less runtime and more DB throughput

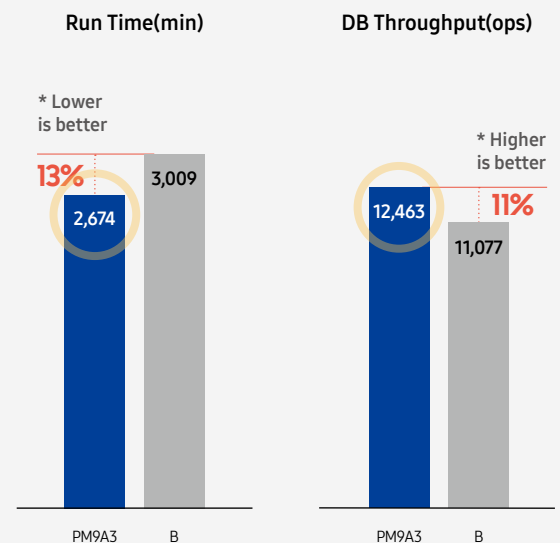
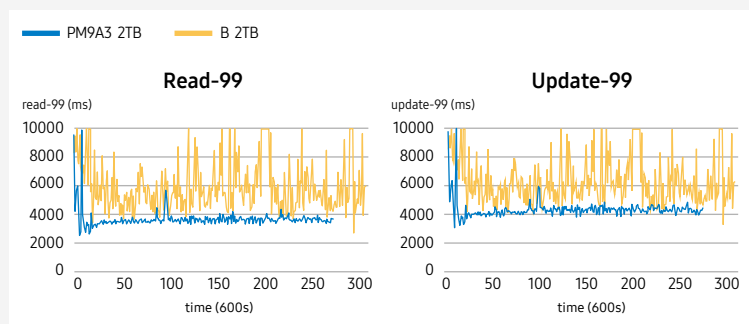
		PM9A3	A Drive
Runtime (min)		2,468 (41hr 08min)	3,266 (54hr 43min)
DB Throughput (ops)		13,502	10,203
Latency (ms)	Read Avg.	1,154	1,535
	Read 99 th	2,491	3,381
	Update Avg.	1,211	1,597
	Update 99 th	2,587	3,511



PM9A3 shows around 10% better performances in 1.92TB

- Most of all, aside from time and throughput, PM9A3 provides much more stable latency during operations

		PM9A3	B Drive
Runtime (min)		2,674 (44hr 34min)	3,009 (50hr 09min)
DB Throughput (ops)		12,463	11,077
Latency (ms)	Read Avg.	1,248	1,408
	Read 99 th	2,791	3,797
	Update Avg.	1,315	1,475
	Update 99 th	2,945	4,163



* Run Time is much shorter than compared drive under same test condition.

Conclusion

With an optimal mix of performance, reliability, design capability – thanks to a variety of available form factors – and energy-efficiency, the Samsung PM9A3 offers the industry's broadest and most capable lineup of enterprise and data center SSDs.

Combining PCIe® Gen.4 with optimal performance for mainstream (computing and storage) NVMe™ servers, the new SSDs are highly reliable, offering best-in-class QoS for various workloads. They offer improved energy efficiency, as well as enhanced design scalability under host level constraints, including 1U servers, providing multiple form factors to meet various design needs.

The SSDs also adhere to the OCP's NVMe™ Cloud SSD Specification, which includes a new OCP reference design, and are recipients of the OCP Inspired™ certification. The advantages of the devices have been demonstrated through user environment simulations across a variety of use cases, as well as through a real-world test that utilized a MongoDB application and was measured by the Yahoo Cloud Service Benchmark.

Specifically designed for the next era of servers, the PM9A3 combines best-in-class SSD attributes to enable enterprises to become more agile, scalable and performant, and significantly more efficient.

※ Evaluation results may vary depending on the server environment and settings.
Performance measurements are based on fio 3.3. Performance may vary based on SSD's firmware version, system hardware & configuration.
Test system configuration: AMD_X570, DDR4 16GB, and CentOS 7.6 with Kernel 4.15 / Physical.

※ In this paper, 1TB means 10¹² Bytes by IDEMA.

Legal Disclaimer

SAMSUNG ELECTRONICS RESERVES THE RIGHT TO CHANGE PRODUCTS, INFORMATION AND SPECIFICATIONS WITHOUT NOTICE.

Products and specifications discussed herein are for reference purposes only. All information discussed herein is provided on an "AS IS" basis, without warranties of any kind. This document and all information discussed herein remain the sole and exclusive property of Samsung Electronics. No license of any patent, copyright, mask work, trademark or any other intellectual property right is granted by one party to the other party under this document, by implication, estoppel or otherwise. Samsung products are not intended for use in life support, critical care, medical, safety equipment, or similar applications where product failure could result in loss of life or personal or physical harm, or any military or defense application, or any governmental procurement to which special terms or provisions may apply.

For updates or additional information about Samsung products, contact your nearest Samsung office.

About Samsung Electronics Co., Ltd.

Samsung inspires the world and shapes the future with transformative ideas and technologies. The company is redefining the worlds of TVs, smartphones, wearable devices, tablets, digital appliances, network systems, and memory, system LSI, foundry, and LED solutions. For the latest news, please visit the Samsung Newsroom at news.samsung.com.

Copyright © 2021 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.