

Samsung's PM9D3a Solid State Drive



Today's datacenters are constantly faced with challenges of balancing power, speed and cooling efficiency with limited space in server racks. With the exponential growth of data and its many applications, the requirements of customers operating servers and datacenters are expanding. The server storage market demands constant power, space, and performance improvements through new technological innovations. In addition to developing higher performing technologies and solutions to meet greater demands, the innovations must address a number of environmental challenges that exist within storage racks located in today's data centers.

What Is PM9D3a? The PM9D3a is a highly advanced PCIe Gen5.0 Non-Volatile Memory Express (NVMe) server SSD, offering exceptional performance, power-efficiency, and reliability for enterprise-server and multicore processing data system environments. These improvements not only address the demands of today's data storage acceleration, but also lower the cost of maintaining the data center environment.

PM9D3a Features and Benefits

Best in Class Performance

The PM9D3a features a Storage Networking Industry Association (SNIA) based 2.5" form factor with full PCIe Gen5 support. This device is able to harness the production efficiencies of Samsung's eighth-generation three-bit V-NAND technology to boost random read speeds up to 2.4 times. Sequential read and write speeds can reach 12,000MB/s and 6,800MB/s respectively. With random reads up to 2,000K IOPS, random writes up to 400K IOPS, and a 15.36TB capacity, the PM9D3a is one of the highest-performing SSDs in the industry. These maximum speeds are based on the U.2/E3.S form factors.

Form Factors to Meet Any Data Center Requirement

The PM9D3a comes in a variety of form factors and capacity ranges. Included are E1.S (25/15/9.5mmT), E1.L (9.5mmT), E3.S (1T/2T), U.2/U.3 2.5" (15mmT), and M.2. This wide selection will provide the flexibility to meet storage specific needs in almost any data center without increasing the size of the data center. Within the M.2 form factor, capacities of 240GB, 480GB, 960GB, and 1.92TB are offered. For other form factors, capacities of 960GB, 1.92TB, 3.84TB, 7.68TB, 15.36TB, and 30.72TB are offered.

Enhanced Power and Cooling Efficiency

The PM9D3a offers a 45% more power-efficient solution and improved cooling efficiency by applying 2.5" 15mmT fins for better thermal characteristics, compared to the previous generation. Additionally, E1 drives come with a built-in heatsink and embedded thermal interface material. These improvements help data centers decrease management costs, maximize profits and mitigate environmental risks.

To further enhance power efficiency and reduce cooling requirements, Flexible Data Placement (FDP) support is available as an option for systems capable of utilizing its benefits. Moreover, solutions compliant with OCP v2.5 are available by request.

Low Failure Rates with High Reliability

The device includes error-correcting code (ECC) to provide the low failure rate and high reliability in the industry. It offers a mean time between failures (MTBF) of 2.5 million hours.

Built-In Cyber Security

Samsung PM9D3a SSD is validated against the TCG Storage Opal Family Test Cases. In addition, this Samsung SSD is undergoing FIPS 140-3 validation. Common Criteria certification is not commonly required of devices like SSDs at this time, but Samsung has extensive experience with CC and mobile devices; CC certification of SSDs is something that Samsung is monitoring, including The EU Cybersecurity Act (Article 49 Regulation EU 2019/881).

Telemetry Features for Data Center Management

Over a hundred data set analyses are provided in advance for data centers run by our clients to monitor and detect on-site risks and issues and handle them efficiently. This proactive approach to Telemetry allows measures to be implemented that ensure efficient server operation while optimizing operational expenses and TCO.

Summary

Samsung's PM9D3a SSD is opening up new frontiers in Artificial Intelligence, Machine Learning, and Large Language Model processing. It is ideally equipped to deal with dynamic workloads and allows data centers to process, store, and deliver data more effectively.

E3.S 1T drives

Table 2. Technical specifications

Feature	1.92 TB drive	3.84 TB drive	7.68 TB drive	15.36 TB drive
Interface	PCIe 5.0 x4	PCIe 5.0 x4	PCIe 5.0 x4	PCIe 5.0 x4
Capacity	1.92 TB	3.84 TB	7.68 TB	15.36 TB
SED encryption	TCG Opal 2.02	TCG Opal 2.02	TCG Opal 2.02	TCG Opal 2.02
Endurance (drivewrites per day for 5 years)	1 DWPD	1 DWPD	1 DWPD	1 DWPD
Endurance (total bytes written)	3,504 TB	7,008 TB	14,016 TB	28,032 TB
Data reliability (UBER)	< 1 in 10 bits read17			
MTBF	2,500,000 hours	2,500,000 hours	2,500,000 hours	2,500,000 hours
Performance & Power - PCIe 5.0 host interface				
IOPS reads (4 KB blocks)	1,700,000	1,700,000	2,000,000	2,000,000
IOPS writes (4 KB blocks)	150,000	250,000	300,000	350,000
Sequential read rate (128 KB blocks)	12,000 MBps	12,000 MBps	12,000 MBps	12,000 MBps
Sequential write rate (128 KB blocks)	3500 MBps	6800 MBps	6800 MBps	6800 MBps
Latency (random R/W)	65 µs / 9 µs			
Latency (sequential R/W)	9 µs / 9 µs			
Typical power (R/W)	11.0 W / 10.8 W	12.5 W / 15.3 W	14.0 W / 16.0 W	14.7 W / 16.0 W
Performance & Power - PCIe 4.0 host interface				
IOPS reads (4 KB blocks)	1,700,000	1,700,000	1,700,000	1,700,000
IOPS writes (4 KB blocks)	150,000	250,000	300,000	350,000
Sequential read rate (128 KB blocks)	7000 MBps	7000 MBps	7000 MBps	7000 MBps
Sequential write rate (128 KB blocks)	3500 MBps	6200 MBps	6200 MBps	6200 MBps
Latency (random R/W)	65 µs / 9 µs			
Latency (sequential R/W)	9 µs / 9 µs			
Typical power (R/W)	9.8 W / 10.2 W	10.4 W / 14.1 W	11.5 W / 14.8 W	12.3 W / 14.8 W

2.5-inch drives

Table 3. Technical specifications

Feature	960 GBdrive	1.92 TBdrive	3.84 TBdrive	7.68 TBdrive	15.36 TBdrive	30.72 TBdrive
Interface	PCIe 5.0 x4					
Capacity	960 GB	1.92 TB	3.84 TB	7.68 TB	15.36 TB	30.72 TB
SED encryption	TCG Opal2.02	TCG Opal2.02	TCG Opal2.02	TCG Opal2.02	TCG Opal2.02	TCG Opal2.02
Endurance (drive writes per day for 5 years)	1 DWPD					
Endurance (total bytes written)	1,752 TB	3,504 TB	7,008 TB	14,016 TB	28,032 TB	50,604 TB
Data reliability (UBER)	< 1 in 10 bits read					
MTBF	2,500,000 hours					
Performance & Power - PCIe 5.0 host interface						
IOPS reads (4 KB blocks)	1,050,000	1,700,000	2,000,000	2,000,000	2,000,000	1,550,000
IOPS writes (4 KB blocks)	75,000	150,000	250,000	400,000	400,000	300,000
Sequential read rate (128 KB blocks)	12,000MBps	12,000MBps	12,000MBps	12,000MBps	12,000MBps	9,000 MBps
Sequential write rate (128 KB blocks)	1700 MBps	3500 MBps	6800 MBps	6800 MBps	6800 MBps	6450 MBps
Latency (random R/W)	65 µs / 12 µs	65 µs / 9 µs	65 µs / 9 µs	65 µs / 9 µs	65 µs / 9 µs	70 µs / 9 µs
Latency (sequential R/W)	9 µs / 9 µs					
Typical power (R/W)	10.8 W / 7.8 W	11.0 W / 10.8 W	12.5 W / 15.3 W	14.0 W / 16.0 W	14.7 W / 16.0 W	14.1 W / 16.6 W
Performance & Power - PCIe 4.0 host interface						
IOPS reads (4 KB blocks)	1,050,000	1,700,000	1,700,000	1,700,000	1,700,000	1,550,000
IOPS writes (4 KB blocks)	75,000	150,000	250,000	400,000	400,000	300,000
Sequential read rate (128 KB blocks)	7000 MBps					
Sequential write rate (128 KB blocks)	1700 MBps	3500 MBps	6350 MBps	6350 MBps	6350 MBps	6350 MBps
Latency (random R/W)	65 µs / 12 µs	65 µs / 9 µs	65 µs / 9 µs	65 µs / 9 µs	65 µs / 9 µs	70 µs / 9 µs
Latency (sequential R/W)	9 µs / 9 µs					
Typical power (R/W)	7.8 W / 7.2 W	9.8 W / 10.2 W	10.4 W / 14.1 W	11.5 W / 14.8 W	12.3 W / 14.8 W	13.4 W / 15.9