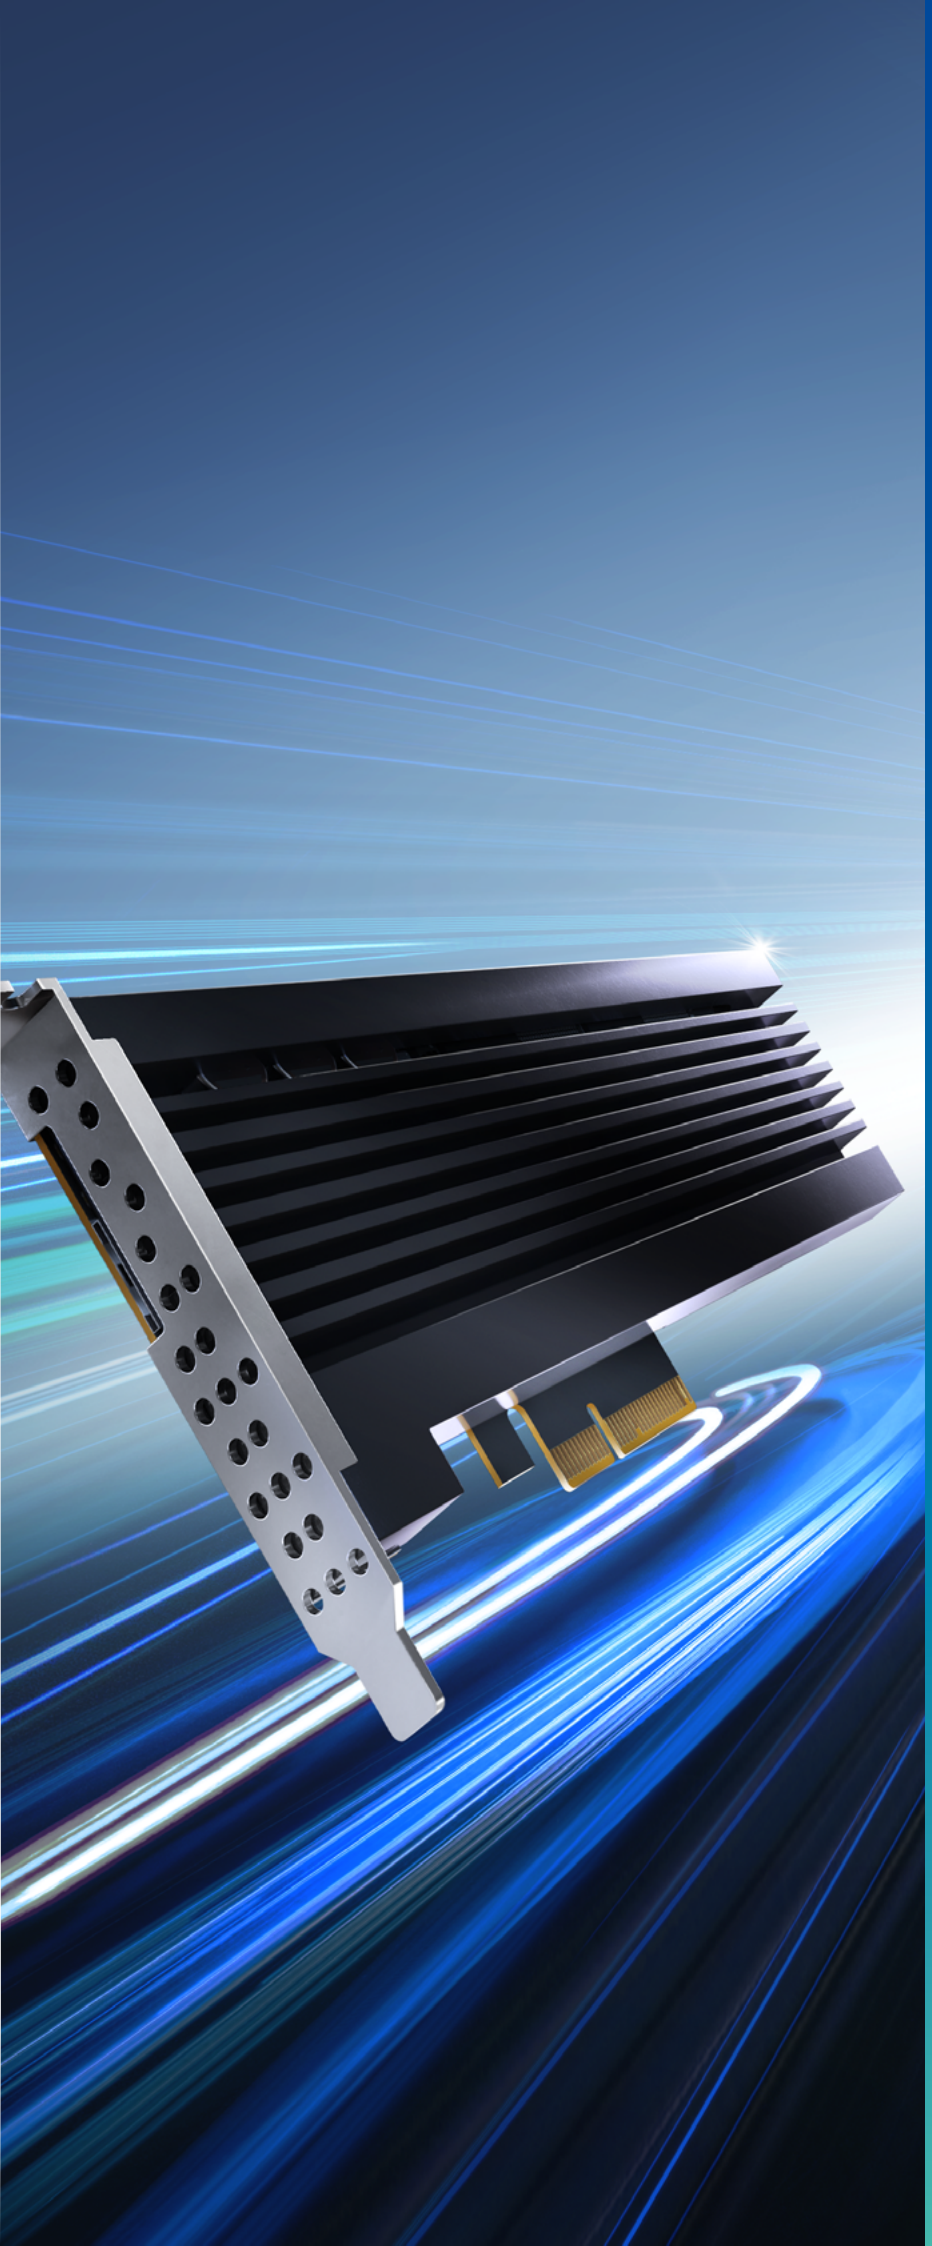

Samsung Z-SSD SZ985

Ultra-low Latency SSD for Enterprise and Data Centers

Brochure

SAMSUNG

A high-speed storage device from the SSD technology leader

Samsung Z-SSD SZ985 offers more capacity than PRAM-based data center SSDs, with comparable latency and higher endurance

We live in a world where new applications are continuously debuting, and Big Data, IoT, and other applications generate enormous amount of data every day. Meanwhile, data centers are always looking for better storage architecture to manage data generated by a variety of applications. SATA SSDs are taking the place of HDDs in datacenters and recently, high-performance NVMe SSDs are gaining momentum in premium storage services and in storage for high-end virtual machines. However, another level of requirements for high performance storage targeting cache and real-time data analytics is rising. Fulfilling those requirements using memory is too expensive to scale and the existing storage technology is far below the required performance standards.

Samsung Z-SSD SZ985, a new type of performance SSD, is specially designed for enterprise applications such as Database Management Systems (DBMS), data analysis, and cache, where high throughput and low latency are a must. It is an ultra-low latency NVMe SSD containing Samsung Z-NAND flash memory and a next-generation NVMe controller, supporting PCIe Gen3 x 4 lanes. The SZ985 provides sequential read and write speeds of up to 3,200 MB/s and 3,000 MB/s respectively, and is able to perform a 4 KB random read operation in 20 µs. Based on Samsung's innovative V-NAND technology, the SZ985 offers high capacities of up to 3.2 TB of storage [1] and high endurance of 30 Drive Writes Per Day (DWPD) for up to five years. Thanks to its ultra-low latency, high read and write speeds, high capacity, and high endurance, the SZ985 is ready to create a new segment in the memory and storage market.

Why Samsung SZ985?

- **Ultra-low Latency:** Up until now the benefits of persistent storage have come at the cost of performance because storage devices are nowhere near as fast as modern processors and system memory. Therefore, we strive to make persistent storage media as fast as possible in order to achieve outstanding storage performance. The SZ985 is a high-speed NAND flash memory-based storage device that boasts 99.999% read QoS latency under 90 µs and 150 µs in 4 KB mixed random read (70%) and write (30%) at queue depths 1 [2] and 16, respectively.
- **Consistently high performance:** Since its first SSD release to the market, Samsung has been the leader in NAND flash memory based SSD technology. Using Samsung's proven expertise and wealth of experiences in c utting-edge SSD technology and memory solutions, Samsung SSDs have helped data centers operate continually at the highest performance levels. As the number one provider of NAND flash SSDs and the manufacturer of all of its SSD components- from the raw NAND flash to the controller and firmware- Samsung ensures steady supply of NAND flash-based high performance SZ985.
- **Outstanding Performance and Endurance:** Tests [3] show that the SZ985 achieves more than 30% higher sequential read and write throughput and 4 KB random read IOPS over PRAM-based SSDs (P-SSDs). With Samsung's V-NAND technology as its foundation, the SZ985 provides remarkable endurance of 30 DWPD for up to five years.
- **High Capacity:** Based on NAND flash memory technology, the SZ985 offers up to 3.2TB disk capacities, while P-SSDs can offer only much lower densities. This makes Samsung SZ985 better suited to enable data centers to continue to scale on their servers, decreasing the costs of computing.

1. 1.6 TB and 3.2 TB SSD availability schedule yet to be determined.
2. Read and write latency were measured by using FIO in CentOS 7.0 and 4 KB transfer size with queue depth 1 on a random and sequential workloads of sustained state, respectively.
3. Random performance was measured using FIO on CentOS 7.0 with queue depth 32 by 16 workers and sequential performance with queue depth 32 by 16 workers.

SAMSUNG

Consistent Low Latency and Incredible Performance

Samsung Z-SSD SZ985's Advanced Features

- **Z-NAND New Flash Memory for Ultra-low Latency:** Samsung developed Z-NAND, which has powerful channel speeds and read performance compared to TLC V-NAND, allowing it to achieve faster levels of data transfer. Z-NAND is based on Samsung V–NAND technology, delivering reliable and consistent performance. This new specialized NAND with a 'unique circuit design' dramatically reduces latency to levels on par with P- SSDs. Coupling an extra die with SLC NAND, which is much faster and more reliable than MLC and TLC NAND, provides tremendous performance, reliability, and SSD lifetime benefits.
- **New Generation NVMe SSD Controller:** The SZ985 is powered by Samsung's next generation NVMe controller with PCIe Gen 3 and it reaches sequential read and write performance at 3,200 MB/s and 3,000 MB/s, respectively with the SZ985. The controller provides hardware automation for read data paths which optimizes low- latency random read performance, especially with Z-NAND; automatic I/O scheduling for better QoS; higher controller operating frequency than those of previous controllers, which greatly reduces processing time; and doubled DRAM bandwidth with LPDDR4 versus the previous controller's LPDDR3 DRAM.

Ultra-low Latency for Nimble Response Times

Latency is the most powerful feature of the SZ985, which shows stable latency QoS under various types of I/O patterns.
4 KB random read and write requests are the most common I/O for SSDs. The SZ985 keeps both read and write latencies of 4 KB random I/O under 20 µs at 99.999% at queue depth 1 (QD1). Typically, SSDs reach their maximum throughput at queue depth 16 (QD16). The SZ985's 4 KB random read latency is maintained at less than 100 µs and the write latency is around 150 µs even at 99.999%, both at queue depths of 16.

SSDs have relatively poor latency in mixed workload situations compared with system memory. But the SZ985 minimizes read latency even while garbage collection is in progress. The SZ985 achieves 4 KB random read latency under 500 µs even at 99.9999% at queue depth 16 with 70% of the requests being 4 KB random reads and 30% being random writes.

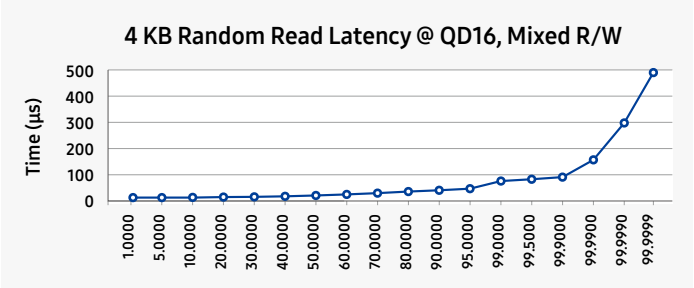

High Throughput for Maximum Performance

We compared the SZ985 with Samsung PM963 and a P-SSD. The SZ985 provides sequential read and write speeds of up to 3,200 MB/s and 3,000 MB/s, respectively, which results in 33% and 50% better sequential read and write throughput, respectively, than those of the PRAM-based SSD.

As a majority of the mission-critical workloads rely on small-sized random reads (4 KB or less), the SZ985 shows excellent performance in delivering high IOPS at low queue depth. Although the P-SSD provides higher 4KB random write IOPS, the SZ985 achieves 36% higher random read IOPS at 750K IOPS [4].

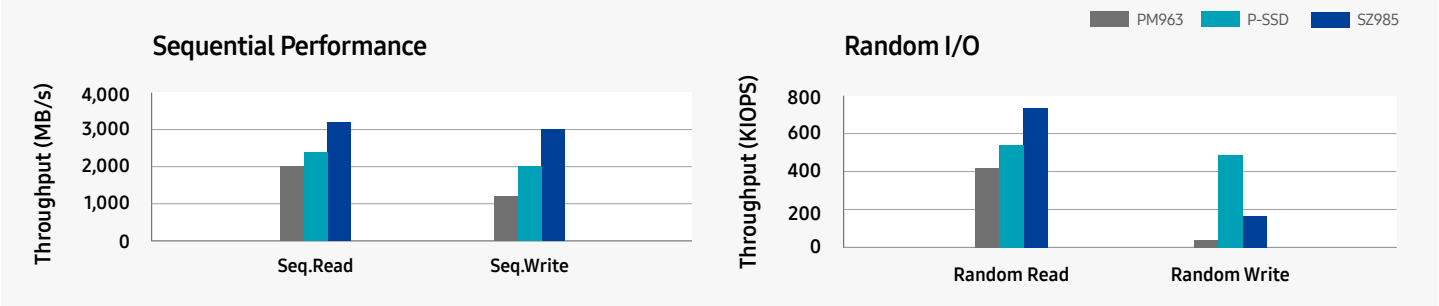

4. Sequential and random throughput comparisons are based on product specifications.

SAMSUNG

Efficient Power Consumption and Long-lasting Endurance

Each application has its own workload characteristics, but for many traditional applications like database DBMS, mixed 4 KB random I/O with 70% read operations and 30% write operations is an appropriate guideline to measure the performance of devices. The SZ985 shows that its random I/O performance in mixed situations increases greatly as the queue depth grows until queue depth 16 and reaches its maximum at queue depth 32.

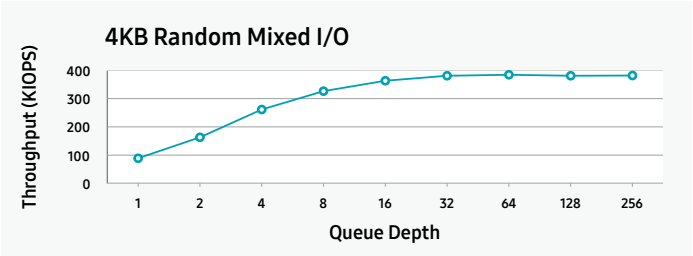

High Performance per Watt for Efficient Power Consumption

SZ985 consumes maximum power of up to about 9 W for sequential read/write, while delivering staggering performance of more than 3,000 MB/s. This equates to more than 330 MB/s per watt. In comparison, P-SSD consumed more power at almost 9 W for sequential read and 13 W for sequential write, while providing read performance of 2,700 MB/s and write performance of 2,200 MB/s. The results show that the SZ985 achieved 18% and 95% better sequential read and write performance-to-power ratios, respectively, than P-SSD.

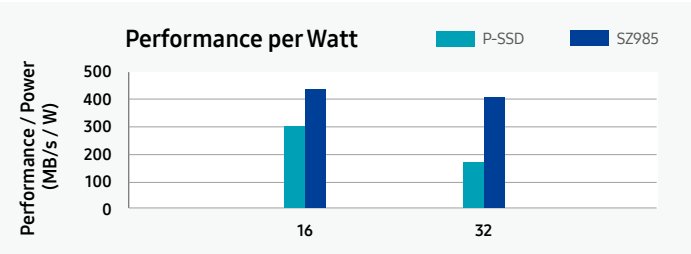

More Endurance with 30 DWPD for Five Years

Utilizing Samsung Z-NAND technology, the SZ985 provides remarkable endurance of 30 DWPD for up to five years, which equates to 43.8 Petabytes Written (PBW) for an 800 GB drive. Compared with the SZ985, the PM963 provides 1.3 DWPD for three years, which equates to only 1.37 PBW. The PRAM-based P-SSD also provides 30 DWPD, but only 12.3 PBW for the 375 GB drive, which is currently its only existing drive capacity. Assuming we are given PM963, P-SSD, and SZ985 drives of the same capacity, then the SZ985 provides over 3.5 times better endurance than the P-SSD and 32 times more than the PM963. This suggests that the SZ985 is better suited to enterprise environments, especially when longer endurance is required.

	PM963 (960 GB)	P-SSD (375 GB)	SZ985 (800 GB)
DWPD	1.3	30	30
Warranty(Years)	3	3	5
Petabytes Written	1.37	12.3	43.8

SAMSUNG

Empowering enterprise applications

Many enterprise applications are eager for better storage performance, especially with lower latency. We applied the SZ985 to various applications and the results were remarkable.

Database Management Systems

Relational database management systems (RDBMS) are one set of 'mission-critical' enterprise applications and storage latencies are vital for their performance. The low latency of the SZ985 is able to dramatically enhance the performance of Line of Business applications that use traditional DBMS. TPC-C [5] is an Online Transaction Processing (OLTP) industry standard benchmark and simulates a wholesale supplier. We measured the SZ985's TPC-C benchmark performance for DBMS workloads on PostgreSQL, one of the most popular open source RDBMS.

Samsung PM963 and a PRAM-based P-SSD were compared with the SZ985. As the number of clients increased from 16 to 32, all of the drives showed increased transaction processing power, tpmC. However, when the number of clients reached 64, all of the drives showed decreased tpmC. In any case, the SZ985 shows much higher throughput, even without any change in the DBMS engine, than the PM963: 70% higher with 16 clients and almost double with 32 and 64 clients. The SZ985 was also comparable to the P-SSD in terms of transaction power and even showed slightly better numbers with 16 clients.

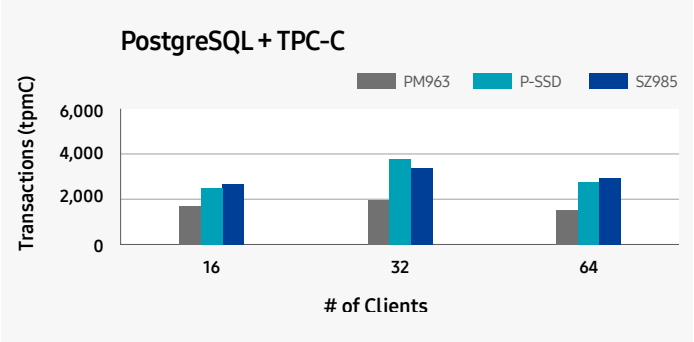

5. TPC-C is an on-line transaction processing benchmark from TPC (http://www.tpc.org/tpcc/)

Graph Analysis

As Social Networking Services spread throughout the world, another layer is added to the relationship between objects and documents via links that map them together for rapid association- for example, mapping connections on social media or finding patterns in crime data for predictive policing.

FlashX performs data analytics in the form of graphs and matrices and utilizes SSDs to scale to large datasets. It is designed for a large parallel machine with fast SSDs, but it also runs efficiently in cloud environments. FlashX has its own benchmark tool to evaluate hardware. To evaluate the SZ985's performance with the FlashX framework, we used FlashX, its Twitter graph set (42 million vertices and 1.5 billion edges), and common social analysis algorithms including PageRank.

The SZ985 showed the best performance among SSDs for all graph algorithms of the FlashX benchmark and was faster than P-SSD by 8% to 29%. This suggests that the SZ985 would be a superb choice to execute social analytics, especially in the case of serving users with real-time results.

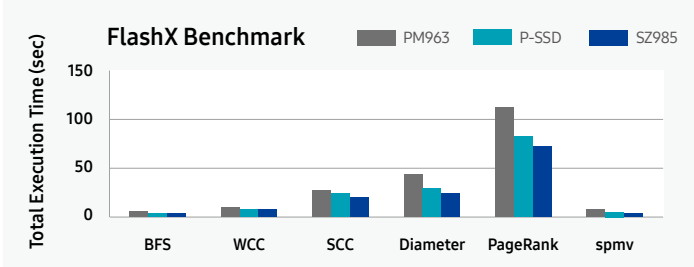

SAMSUNG

Samsung SZ985 Technical Specifications

Cache

Memcached is a distributed memory caching system. It is often used to speed up dynamic database-driven websites by caching data and objects in memory to reduce the number of times an external data source (such as a database or API) must be read. However, DRAM is volatile and provides challenges like power consumption and operational complexity in scaling architecture.

Fatcache [6] is Memcached on SSDs. It makes memory 'fat' by incorporating SSD-backed storage to overcome these limitations. Twemperf [7] is a tool for measuring Memcached server performance by generating connections and requests at a high rate. The experiment shows that the SZ985 served 66% more requests than the PM963 and 7% more than P-SSD, which confirms that Samsung SZ985 is well-suited for caching purposes.

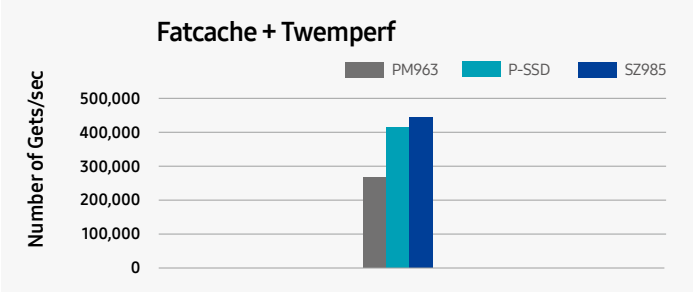

6. Fatcache is memcache on SSDs (https://github.com/twitter/fatcache)
7. Twemperf is an open source project on OpenHub (https://www.openhub.net/p/twemperf)

Form Factor	Half-height Half-length
Capacity[8]	800 GB, 1.6 TB (TBD), 3.2 TB (TBD)[1]
Host Interface	PCI Express Gen3 x4
NAND flash memory	Samsung Z-NAND Flash Memory
Sequential Read/Write	Up to 3,200 MB/s / 3,000 MB/s
Random Read/Write	Up to 750K IOPS / 170K IOPS
Reliability (MTBF)[9]	2 Million Hours
Reliability (UBER)[10]	1 sector per 1017 bits read
Endurance (DWPD)[11]	30 drive writes per day within 5 years
Latency - Read/Write (typical)	20/16 µs
QoS - Read/Write (99%)	20/20 µs

8. 1TB = 1,000,000,000,000 Bytes, 1GB = 1,000,000,000 Bytes, Unformatted Capacity. User accessible capacity may vary depending on operating environment and formatting.
9. MTBF is Mean Time Between Failure, and is the predicted elapsed time between inherent failures of a system during operation.
10. Uncorrectable Bit Error Rate (UBER) is a metric for the rate of occurrence of data errors, equal to the number of data errors per bits read as specified in the JESD218 document of JEDEC standard.
11. The endurance of SSDs in enterprise applications is defined as the maximum number of drive writes per day that can meet the requirements specified in the JESD218 document of JEDEC standard.

For more information

For more information about Samsung Z-SSD, visit **www.samsung.com/semiconductor**.

About Samsung Electronics Co., Ltd.

Samsung Electronics 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.**

Samsung Electronics Co., Ltd.
129 Samsung-ro, Yeongtong-gu, Suwon-si, Gyeonggi-do 16677, Korea **www.samsung.com** 2018-04

SAMSUNG