White Paper:

SSDs Deliver More at the Point-of-Processing

How increased capacity and durability expand
SSD applications in the data center

SAMSUNG
BUSINESS

SSDS DELIVER MORE AT THE POINT-OF-PROCESSING

Historically, the tradeoff of hard disk drives (HDDs) versus solid-state drives (SSDs) in enterprises has revolved around three variables: capacity, endurance, and price. Improved flash memory fabrication processes and capital expenditures on fab facilities have produced denser chips with more flash less expensively. This has raised SSD capacities and lowered costs substantially, and the latest V-NAND flash technology promises even more flash at lower costs. More durable flash and improved write-management tactics have vastly improved endurance, and in most cases, enterprise SSDs meet or exceed user requirements.

Higher capacity, longer lasting SSDs have shifted the tradeoff from simply exchanging an HDD for an SSD to rethinking how enterprises deploy and use storage. Faster flash memory, better flash controllers, new form factors, enhanced high-speed interfaces, and advanced multi-threading drivers have tipped the scales of performance and flexibility in favor of SSDs for many applications.

Cutting-edge SSD technology now offers opportunities for enterprises to reduce data center footprints, or increase capability in a similar footprint. Exploring three use cases – scale-out architecture with converged modular servers, tiers servicing high-value data, and moving toward all-flash data centers – illustrates how IT architects can put the right SSD technology to work at the point-of-processing.

SCALING-OUT WITH CONVERGED MODULAR SERVERS

One strategy for servers is a scale-up approach, consolidating workloads on a single big CPU or a tight cluster of several CPUs, perhaps with multiple large cores in each. For example, the 2Q2015 variants of the Intel Xeon processor E7 v3 range from 115-165 W TDP, with 4 to 18 cores running at 1.9 to 3.2 GHz. Power and cooling permitting, a 1U rackmount package with these processors can carry quite a processing punch. Several 2.5" drives can be integrated in a 1U rackmount for direct-attached storage, and Ethernet allows network-attached storage configurations. Larger packages up to 4U height allow several processors and more drives in a single rackmount unit.

Many applications are now moving to scale-out architecture, where workloads are dispersed as threads to a group of many smaller cores. Increasing the speed of a scale-out server complex means adding more cores, each less expensive and power hungry, all interconnected with a high-speed fabric such as 10GigE. These lower TDP cores, often variants of Intel Atom processors or ARMv8-based processors, can be packaged on a module or blade, which is then hot-plugged into a converged modular server.

Density increases also come through new storage form factors. New M.2 SSD technology can cut the footprint for processor modules in converged modular servers. Instead of designing to the physical size of the much larger 2.5" drive and associated cabling, modules can now use the smaller M.2 footprint. Free of SATA limitations, M.2 SSDs with PCIe x2 or x4 interfacing maintain a low profile installed flush to a processor board, while delivering performance similar to 2.5" SSDs with PCIe.

M.2 SSD

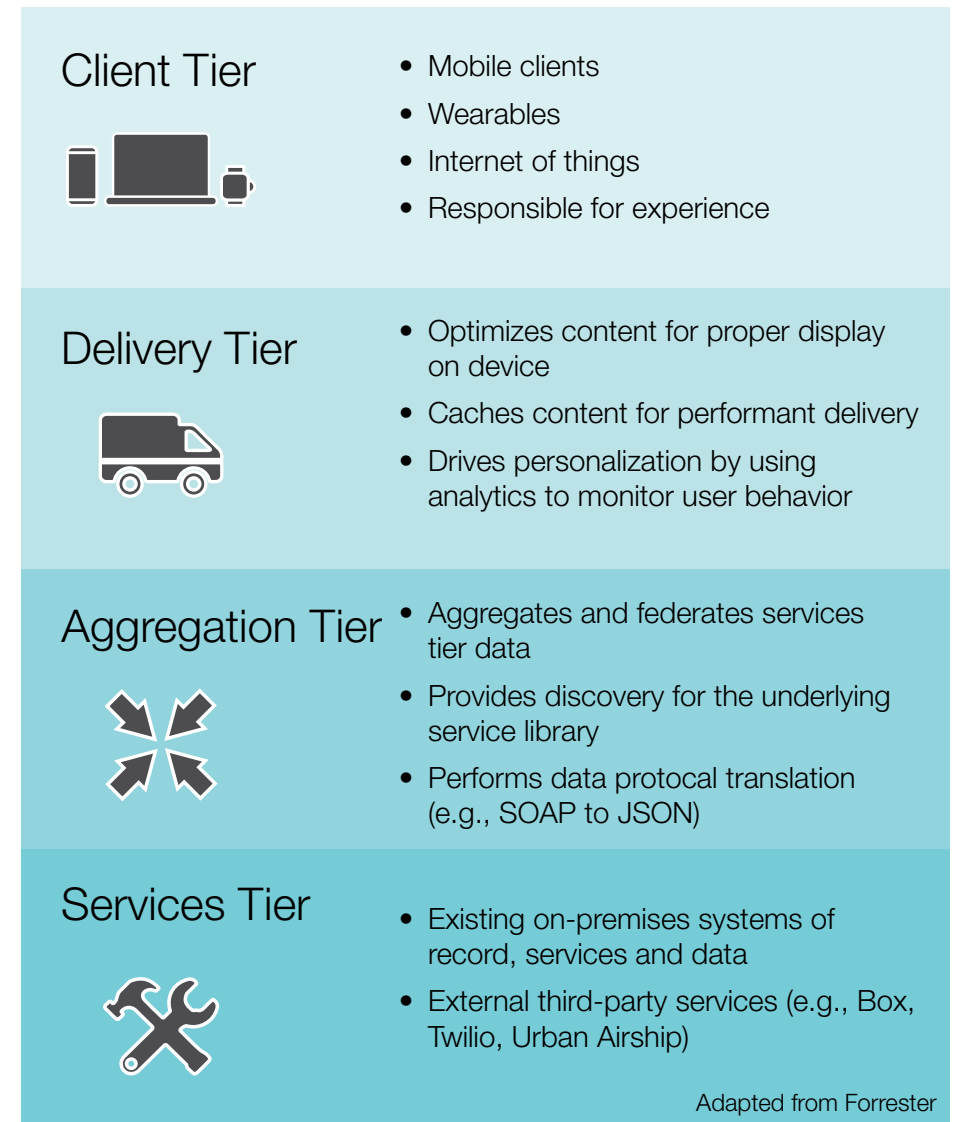

22mm

80mm

2.5" SSD

70mm

100mm
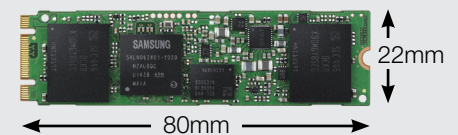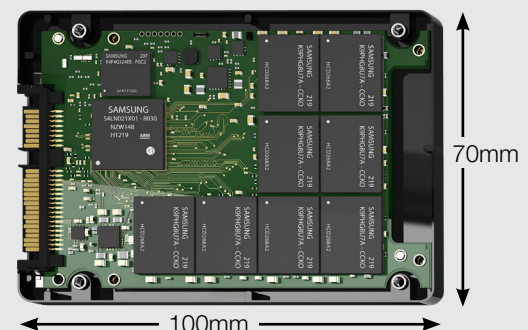
RETHINKING TIERS FOR SERVICING HIGH-VALUE DATA

Pairing storage directly with processor cores becomes even more important in the case of high-value data and tiered architecture. In a traditional IT scale-up scheme, Tier 1 is presentation and aggregation, Tier 2 is application, and Tier 3 is storage. The architecture evolved around cost-effective archival of volumes of data, infrequently accessed somewhere in a "data lake".

In newer schemes designed for the Internet of Things and rapid response for mobile and social applications, a stream of data is processed closer to its source using real-time analytics. These heavy transactional workloads move data out of the quiescent HDD lake in Tier 3 into SSDs paired with processors in Tier 1 and 2. Data center SSDs with excellent quality of service (QoS) figures offer low latency and high I/O operations per second (IOPS) under mixed read/write workloads.

These changes are leading some to the conclusion that the traditional tiering approach is outdated, especially for mobile-enabled applications. Analysts at Forrester have suggested a four-tier "engagement" platform – clients (what others call the "edge"), delivery, aggregation, and services. Their observation is important: data is no longer a commodity stored post-processing, it is a dynamic component woven throughout an architecture, and undergoes transformation at different points in the scheme.[1]

Client Tier

- Mobile clients
- Wearables
- Internet of things
- Responsible for experience

Delivery Tier

- Optimizes content for proper display on device
- Caches content for performant delivery
- Drives personalization by using analytics to monitor user behavior

Aggregation Tier

- Aggregates and federates services tier data
- Provides discovery for the underlying service library
- Performs data protocol translation (e.g., SOAP to JSON)

Services Tier

- Existing on-premises systems of record, services and data
- External third-party services (e.g., Box, Twilio, Urban Airship)

Adapted from Forrester

Further enhancement of high-value data architectures comes through adoption of the NVMe logical interface. Under NVMe, a multi-threading operating system can manage storage more effectively, leveraging more requests with 64K queues each carrying up to 64K commands. From a system performance perspective, the latest SSDs supporting NVMe block far less often waiting for requests to complete compared to previous versions of SSDs with only AHCI.

The sum of these effects implies that fewer processor/SSD pairs located close to points of high-value data interaction can service a greater number of transactions. Infrastructure costs are rightsized accordingly. The most important aspect may be transformative – high-value data, analyzed and actioned quickly, can lead to improved decision making in the enterprise, and enhanced satisfaction and loyalty among consumers choosing to use services.

MOVING TOWARD ALL-FLASH DATA CENTERS

Enterprise teams are also thinking about the all-flash data center. Where large amounts of infrequently accessed data exist, HDDs are still a lower-cost option, but the gap is closing. Where high-value data requires greater responsiveness, SSDs are clearly better. In between lies the discussion of hybrid or all-flash architectures. As with any newer technology, there are lingering myths from early SSD issues now solved, and there are new ideas from innovative applications in early stages of adoption.

Increases in total byte writes (TBW) and device writes per day (DWPD) driven by V-NAND flash technology have removed most concerns over SSD endurance. Some objections have surfaced over data retention in power-off situations. The JC-64.8 subcommittee of JEDEC has defined a data retention target for enterprise SSDs specified at the end of the TBW limit of three months in power down.

HDFS Data Distribution

Input file

	Node A	Node B	Node C	Node D	Node E
1					
2	2	1	1	2	1
3	4	2	3	3	3
4	5	5	4	5	4
5					

Adapted from Cloudera

Data Retention Target for Enterprise SSDs

Application Class	Client	Enterprise
Active Use (power on)	40ºC 8 hours a day	55ºC 24 hours a day
Retention Use (power off)	30ºC 1 year	40ºC 3 months
Functional Failure Rqmt (FFR)	≤3%	≤3%
UBER	≤10⁻¹⁵	≤10⁻¹⁶

Adapted from JEDEC
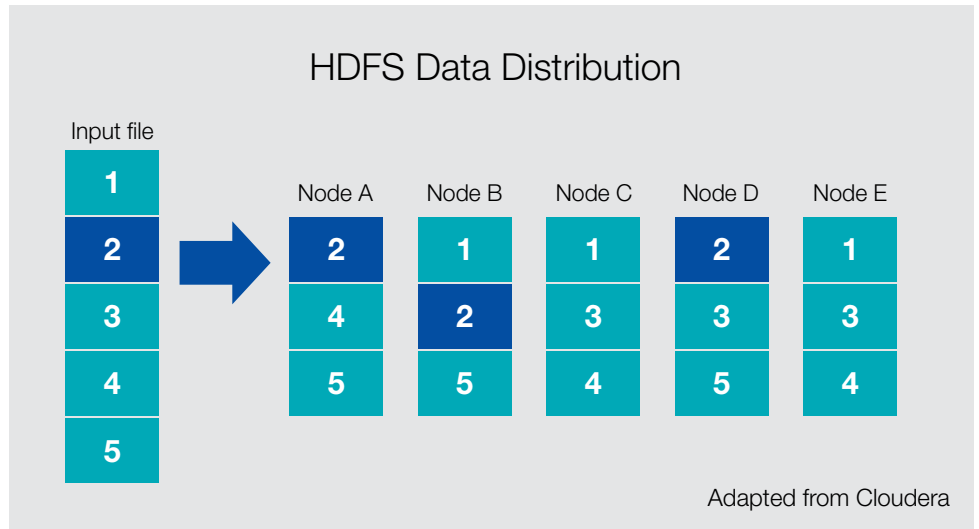
This removes any concern over abrupt data loss, providing plenty of time for maintenance measures. With data de-duplication and compression techniques, SSD efficiency and lifetime can be increased.[2]

The IOPS advantage of SSDs applies to more than user-facing applications. An emerging approach to managing storage more effectively is logical snapshots, which can be used in several ways. Snapshots can serve as an incremental backup and rollback mechanism. Combined with enhanced metadata, they can create and manage copies of data used by different subsystems – for example, an IoT system that performs real-time analytics followed by updates to operational systems.

Big data applications also benefit. A typical Hadoop system not only distributes data storage, but also creates multiple copies of data based on the HDFS replication factor – the default is 3, but some are setting the factor as high as 8 or 10 in larger clusters. Again, the objective is to reduce applications all hammering on a single database table location, and spread access requests across nodes. Data can also be partitioned into subsets for different uses. For example, a real-time analytics reporting subsystem may have its own copy of data. Live production data can be cloned into a development sandbox safely.[3]

By minimizing the number of applications pursing data in a given physical location, accessing copies made quickly in SSDs with little overhead instead, bottlenecks are reduced and performance increases. Rather than simply trading HDDs for SSDs, the performance benefits of SSDs can be leveraged to create gains difficult to achieve with HDDs and implement a more efficient system architecture.

DELIVERING RESULTS WITH ENTERPRISE SSDS

The goal for all these ideas is to deliver a suite of services in the minimum profile of real estate, power, and cooling in a data center, with the right levels of performance, management overhead, and longevity. Enterprise SSD technology has evolved from an alternative for bulk data storage to delivering point-of-processing support for innovative new data center architectures.

M.2 SSDs take high-performance storage into smaller places, and change how mobile and infrastructure platforms are designed. 2.5" SSDs offer similar levels of IOPS performance with larger storage capacity and legacy packaging compatibility.

Both now deliver the benefits of PCIe physical and NVMe logical interfacing, and achieve endurance goals with V-NAND flash. Enterprise variants have the needed QoS and reliability features to assure performance under all conditions.

Enterprise IT architects can now evaluate the right technology for each point-of-processing, implementing solutions more efficiently with better scalability. The result in some cases will be a smaller data center footprint, while in other cases it will be greatly enhanced capability in a similar footprint. Either way, enterprise SSDs deliver benefits today and will grow and evolve for future needs.

SAMSUNG ENTERPRISE SSD PORTFOLIO

PM863 Series Data Center SSDs
- Designed for read-intensive applications
- Built in Power Loss Protection, low latency, sustained performance
- SATA 6Gb/s Interface
- Form-factors: 2.5"

SM863 Series Data Center SSDs
- Designed for write-intensive applications
- Built in Power Loss Protection, low latency, sustained performance
- SATA 6 Gb/s Interface
- Form-factors: 2.5"

950 Pro Series Client PC SSDs
- Designed for high-end PCs and workstations
- PCIe Interface
- NVMe protocol
- Form-factors: M.2

850 Pro Series Enterprise SSDs
- Designed for light enterprise and Client PCs
- SATA 6Gb/s Interface
- Form-factors: 2.5"

[1]"Mobile Needs a Four-Tier Engagement Platform", Ted Schadler, Forrester, November 20, 2013, http://blogs.forrester.com/ted_schadler/13-11-20-mobile_needs_a_four_tier_engagement_platform

[2]"The Truth About SSD Data Retention", Kristian Vatto, AnandTech, May 13, 2015, http://www.anandtech.com/show/9248/the-truth-about-ssd-data-retention

[3]"Evolution of All-Flash Array Architectures", David Floyer, Wikibon, December 20, 2014, http://wikibon.com/evolution-of-all-flash-array-architectures/

Learn more: samsung.com/enterprisessd | 1-800-SAMSUNG | 1-866-SAM4BIZ

Follow us: ▶ youtube.com/samsungbizusa | 🐦 @SamsungBizUSA

SAMSUNG BUSINESS

WHP-SSD-CAPACITY-SEP15J