White Paper:

SSDs in The Data Center: The Cost Advantage

Redefining Storage Total Cost of Ownership

SAMSUNG

Cost per gigabyte: Why it isn't the whole story

Storage continues to grow in importance as nearly everything around us is being digitized. As with all investments, it's important to ensure maximum return on investment for storage purchases. This requires a solid understanding of what really matters most when it comes to storage in the data center.

Many enterprises are still relying heavily on cost per gigabyte when making storage investment decisions. While cost per gigabyte can be a useful metric as a part of the decision making, it is not accurate when used alone because solid state drives (SSDs) offer numerous advantages over hard disk drives (HDDs) that aren't measured by cost per gigabyte. A better measure is Total Cost of Ownership (TCO), which takes into account all capital and operating expenses that occur during the lifespan of an investment, making it far more reliable and accurate than any single metric.

Understanding total cost of ownership

Total Cost of Ownership, as the words themselves suggest, accounts for all costs related to a purchase. The costs can be split into capital expenditure (capex) and operating expenses (opex).

Capex in this context means the acquisition cost of servers, storage, network and software. Capex is usually easier to evaluate than opex because enterprises have certain capacity and performance targets, making it straightforward to compare different setups.

The key factor is that the more drives needed to reach capacity and performance targets, the more servers and network equipment needed, which increases total capex. While SSDs command a higher cost per gigabyte, they offer higher density and performance than HDDs, and can thus reduce capex since the cost of other equipment is lower. That's why enterprises should also consider metrics such as density per rack and performance per rack in order to calculate capex accurately.

Opex includes the cost of infrastructure, labor and support. Infrastructure costs consist of floor space, cooling and electricity, which tend to increase with the number of racks but also depend on hardware characteristics. Labor refers to the cost in man hours to setup and manage the data center, which is again proportional to the number of racks.

Support costs are more general than infrastructure and labor, as practically everything else that is required to keep the data center operating falls under support criterion. Some of the main costs include those for hardware replacements, data protection and possible expansion, as well as something like end-user IT support. Unlike other costs, technical or customer support expenses don't necessarily scale with the size of the data center, but are more dependent on the initial hardware choices and their inherent characteristics, such as reliability.

Compared to capex, opex is typically more difficult to evaluate because many of the costs are hidden and don't occur until later on. The next section discusses some of the key opex elements and provides insights on how to evaluate them properly.

Evaluating opex

Electricity

When making an investment decision, enterprises don't always fully understand the cost of electricity in the long term and how hardware choices impact that cost. Similar to capex, the cost of electricity scales with the number of racks as more electricity is needed to power every individual server.

It's important to note, also, that the components inside servers vary greatly in terms of power efficiency.

SSDs have a tremendous advantage in power efficiency compared to HDDs as SSDs typically consume less than half the power, while delivering hundreds of times higher IOPS (input/output operations per second). With orders of magnitude higher IOPS per watt, SSDs can equalize the difference in cost per gigabyte in a matter of months by reducing electricity cost, especially when used for frequently accessed data where drives are under a constant load.

	SSD[1]	15k rpm HDD[2]	10k rpm HDD[3]	7,200 rpm HDD[4]	SSD Advantage
Capacity	3.84TB	600GB	1.8TB	8TB	
4KB Random Read (IOPS)	97,000	680	437	136	14,264 - 71,323 %
4KB Random Write (IOPS)	24,000	749	574	314	3,204 - 7,643 %
Power Consumption	3.0W	7.5W	7.8W	7.4W	59 - 62 %
Watts per Terabyte	0.79	12.5	4.33	0.93	16 - 94 %
Read IOPS per Watt	32,333	91	56	18	35,431 - 179,528 %
Write IOPS per Watt	8,000	99	74	42	8,081 - 19,047 %

Reliability: SSD vs HDD

Similar to how components vary in power efficiency, they also vary in reliability. An extensive academic study[5] performed by Bianca Schroeder and Garth A. Gibson in 2007 shows that HDDs have an annual failure rate (AFR) of 3% on average. Over a five-year period, which is a typical investment period for enterprises, this means that more than 15% of the HDDs have to be replaced. Even though replacements are usually covered by the manufacturer's warranty, which is typically five years for enterprise drives, there are many additional costs including labor, risk of data loss, and the customer impact when taking servers off-line to manage drive replacement.

SSDs, in comparison, have AFR of about 0.5%[6]. In other words, the number of HDD failures in just one year is greater than the number of SSD failures in five years. Due to the higher inherent reliability, SSDs offer minimal downtime and don't need extensive amounts of redundancy like HDDs do to protect data against drive failures. Whereas HDD arrays often need RAID 6 level redundancy, RAID 5 is typically sufficient for SSD-based arrays as the failure rates are so low.

Another reliability benefit SSDs have is that their lifespan can be easily and accurately estimated. Because SSDs are semiconductor based and don't have any moving parts, the only components that wear out over time are the NAND chips themselves, and the wearout phenomenon and characteristics are well known as all chips have been validated to withstand a certain number of program/erase (P/E) cycles. HDDs, on the other hand, are subject to mechanical wearout of moving parts, which is practically impossible to reliably estimate. A reliable lifespan estimate makes the TCO calculation much more accurate because the drive replacement costs are not subject to high variability.

In addition, the lifespan of an SSD can be tailored to fit the company's IO workload and targeted lifespan by configuring the drives' over-provisioning level.

Reliability: Enterprise SSD vs client SSD

Some enterprises use client-grade SSDs in enterprise applications due to the lower cost per gigabyte, resulting in lower capex. However, client SSDs have numerous TCO downsides, as the opex can be substantially higher. First, manufacturers have not designed and validated client drives to be used in the enterprise, and thus don't grant any warranty if client drives are used in an enterprise application. That means the company has to pay for all replacement drives, even if the failure is premature, whereas an enterprise SSD might still be covered by the manufacturer's warranty.

Because client workloads are much lighter, the NAND chips inside client SSDs are not validated to withstand as many P/E cycles as the chips inside enterprise SSDs are. Due to lower endurance, client SSDs have a substantially lower lifespan in an enterprise environment, meaning that the failure rates are higher and drives need to be replaced more often. As in the case of HDDs, a higher failure rate means increased labor costs and higher risk of data loss and downtime.

Client SSDs also lack proper enterprise-grade data protection against unexpected power losses. Enterprise SSDs have built-in backup power circuitry in the form of capacitors that provide enough power for the SSD to complete any write operations that are in progress. Because client SSDs do not have such capacitors, the data of incomplete write operations may be lost in the case of an unexpected power loss.

	SM863a	PM863a	850 PRO	850 EVO
Target Market	Enterprise	Enterprise	Client	Client
Capacity	1.92TB	1.92TB	2TB	2TB
Endurance (TBW)	12,320TB	2,733TB	300TB	150TB
Price	$1,360	$1,160	$950	$700
TBW per Dollar	9.0	2.4	0.3	0.2

Pricing as of March 2017 and may be subject to change

Engineered for vastly different performance demands

Client SSDs are designed primarily as replacements for hard disk drives (HDDs) in PCs and are intended for more sporadic use. On the other hand, data center-class SSDs are built to sustain rapid, 24/7 responsiveness.

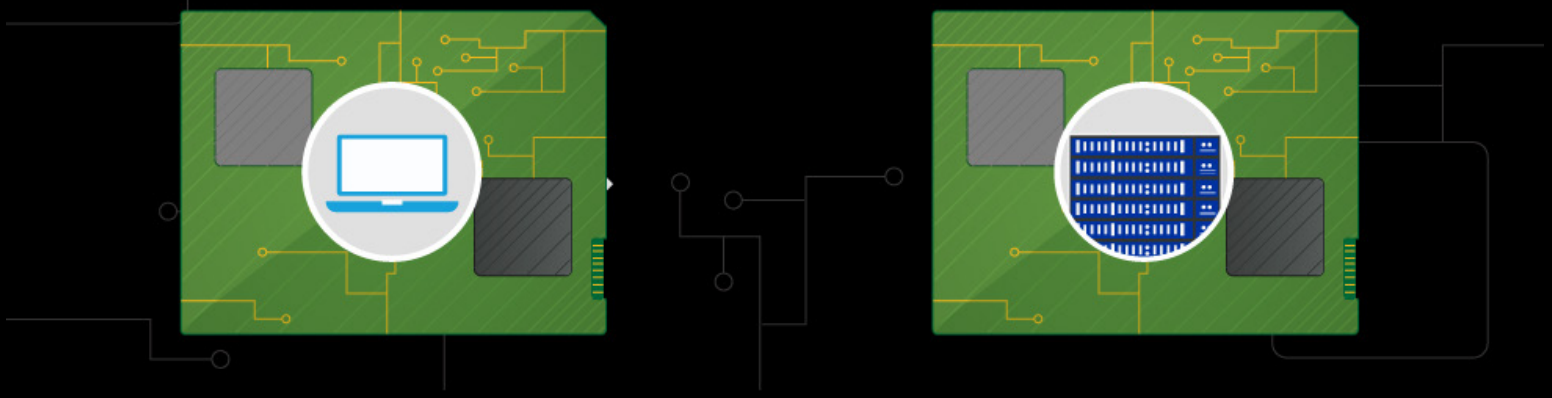

The write amplification factor

The three most important factors that affect the lifespan of an SSD are: NAND flash endurance (maximum P/E cycle), Write Amplification Factor (WAF) and the physical capacity of the SSD.

WAF refers to the multiplying effect of data being written to the drive due to the internal operation, such as garbage collection (GC) and the wear-leveling process, and is highly dependent on the data workload pattern and the over-provisioning (OP) ratio.

Considering all three factors, the lifetime of SSD can be calculated on a TBW basis.

TBW = (SSD physical capacity X NAND endurance)/WAF

Unlike Client PC, Enterprise Drives traditionally come with higher OP ratios, and a firmware that manages the internal operations reducing WAF impact on the drive.

Expansion: Maximizing data center floor space

Outgrowing a data center is a very difficult and expensive problem. Increasing floor space by expanding the facility is both expensive and risky because of glitches that can occur during construction, such as power outages. In addition, having external parties working next to the servers that hold company and customer private data is not ideal. Another facility in a different location, on the other hand, is even more expensive and time consuming, and may present additional complexity due to networking requirements.

As a result, enterprises do everything they can to avoid physical expansion by focusing on capacity and performance enhancements per rack, as a way to maximize the available floor space. As shown earlier, SSDs offer tremendously higher performance and, thanks to V-NAND, SSD capacities are growing at a faster rate. SSDs surpassed 2.5” high-performance enterprise HDDs years ago, making SSDs the best option to expand an existing data center and the go-to storage for anything performance-oriented.

SSDs have in recent years overtaken 3.5” HDDs in terms of capacity, which opens up another door for expansion and lowering the TCO of SSDs in high-capacity deployments like cloud storage.

Conclusion

It's easy to see why cost per gigabyte is a popular storage evaluation. It's very simple and easy to understand and calculate, whereas TCO is much more complex with its multiple variables. While cost per gigabyte has its role in TCO, it's only a very small part of a much bigger picture, and it's clear that no decision should be made solely on cost per gigabyte. As explained, there are many other metrics – such as IOPS per watt and capacity and performance per rack – that do a better job of evaluating TCO. The truth is that TCO can't be measured with any single metric; it's a mix of several metrics, with each having unique weight in every enterprise as ultimately every workload is different.

Do you want to find out how SSDs can lower your enterprise's TCO? Check out Samsung's Storage TCO Calculator and get TCO calculations tailored for your workload.

About the Author

Kristian Vättö is responsible for SSD marketing at Samsung Semiconductor Europe. Prior to joining Samsung, he was a Senior Editor at AnandTech.com in charge of producing highly-detailed and technical SSD reviews. Kristian holds a BSc in Economics from University of Tampere in Finland.

SAMSUNG ENTERPRISE SSD PORTFOLIO

PM863a Series Data Center SSDs

- 3 bit MLC NAND
- Designed for read-intensive applications
- SATA 6Gb/s Interface
- Form-factors: 2.5"

SM863a Series Data Center SSDs

- 2 bit MLC NAND
- Designed for write-intensive applications
- SATA 6 Gb/s Interface
- Form-factors: 2.5"

Additional Resources

- RAID Strategies in SSD Deployments
- Using SMART Attributes to Estimate Drive Lifespan
- Power Loss Protection in SSDs

[1]Samsung PM863a 3.84TB. Specifications from http://www.samsung.com/us/business/computing/solid-state-drives/MZ-7LM3T8NE
[2]HGST Ultrastar C15K600 600GB. Specifications from http://www.tomsitpro.com/articles/hgst-ultrastar-c15k600-hdd,2-906-4.html
[3]HGST Ultrastar C10K1800 1.8TB. Specifications from http://www.tomsitpro.com/articles/hgst-10k-ultrastar-c10k1800-hdd,2-875-3.html
[4]HGST Ultrastar He8 8TB. Specifications from http://www.storagereview.com/hgst_ultrastar_helium_he8_8tb_enterprise_hard_drive_review
[5]https://www.usenix.org/legacy/events/fast07/tech/schroeder/schroeder.pdf
[6]http://www.hardware.fr/

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

Follow us: youtube.com/samsungbizusa | @SamsungBizUSA

SAMSUNG

WHP-SSD-SSDANDTCO-APR17J