


TOSHIBA



Enterprise Hard Drives

Designed for your business

Comparing cost and
performance of HDD- and
SSD-based Storage Systems

Comparing cost and performance of
HDD- and SSD-based Storage Systems

Hard disc drives (HDD) are the obvious choice for storing significant amounts of data, especially where cost is a primary consideration. However, in applications where performance is the main consideration, then solid-state drives (SSD) are preferred.

By their very nature, SSDs do not need to seek data, making them several times faster in terms of IO operations per second (IOPS) in applications such as databases where random performance is the main consideration. In sequential performance applications such as video streaming, SSDs offer at least double the performance of HDDs, measures in Megabytes per second (MB/s).

However, enterprise-class SSDs remain significantly higher in terms of cost per unit of storage capacity:

	Sequential Performance	Random Performance	Cost per TB
HDD 10.5krpm	~250 MB/s	~300 IOPS	1x
Enterprise SATA SSD 3DWPD	500 MB/s	>30000 IOPS	>5x

Where the application only requires a relatively small amount of storage and performance is the primary consideration then SSD is the obvious choice. This includes uses such as boot devices, tier-0 storage for working data, databases and other similar applications. Here, SSDs are typically implemented as single devices, or maybe a few in a set.

For applications that require larger amounts of data storage, multiple HDD or SSD must be combined into a single storage solution, using RAID or software defined storage management. In this way, many inexpensive but relatively slow devices can be combined to deliver performance that may be able to approximate the performance of a few high-speed components. As a side benefit, the available and usable capacity will be significantly larger for the same cost.

Evaluating trade-offs

Toshiba's customers have often discussed the trade-offs between HDD and SSD technologies and the impact on architectures and cost. Toshiba crystallised the comments into this question: "Can an array of many 10krpm HDDs better solve requirements for cost-per-capacity AND performance than a solution based upon fewer Enterprise SSDs? Where is the crossing point between the two approaches, and which architecture should be used for optimal performance?" Toshiba then set out to find a definitive answer.

Test scenarios

To begin to answer the question, Toshiba selected a typical enterprise SSD storage configuration as a reference system. This included a total of eight 1.6TB / 3DWPD SATA enterprise SSDs in a RAID6 (double parity) configuration for protection. (3DWPD means a write endurance of three overwrites per day – the minimum recommendation for active enterprise storage applications.) These SSDs were installed in the front plane of the reference server and driven by an internally connected RAID controller.

To address the question, an HDD-based comparison was needed. Twenty-four HDDs with capacities up to 2.4TB have approximately the same cost as the eight SSDs, so these were used. They were installed in a 2U, 24bay JBOD connected to a second RAID controller in the evaluation server using an external dual path SAS connection. The HDDs were configured as RAID10, which combines RAID0 striping that uses drives in parallel to aggregate the performance with RAID1 mirroring for data protection. RAID10 is the fastest configuration with protection, albeit with a 50% redundancy overhead due to the mirroring. As HDDs are relatively inexpensive when compared to SSDs, this was considered to be equitable and a valid test.



Storage	8x eSSD SATA 1.6TB 3DWPD	24x HDD SAS 2.4TB 10krpm
Configuration	RAID 6	RAID 10
Controller	Microsemi Adaptec SmartRAID 3154-8i	Microsemi Adaptec SmartRAID 3154-8e
Chassis	Supermicro CSE-826BAC4 (in Server)	AIC XJ1-20242-05 2U 24 x 2.5" hot swap bays, hot swap JBOD with dual SAS 12G expander controller, Tool-less HDD tray, 549W 80+ Platinum RPS (external)
Net Capacity (TB)	9,6	28,8

The first point noted from the testing was that the HDDs offered three times the net capacity for approximately the same cost.

Performance evaluation

The performance of both solutions was evaluated using “fiio” to generate a synthetic random workload that included 16 concurrent reading and writing tasks.



The result of this test was as expected; for small block sizes with a significant requirement for seeking of data, the resulting performance is dominated by the IOPS spec meaning that the SSD solution offered vastly better performance.

Blocksize	IOPS 8x SSD	MB/s 8x SSD	IOPS 24x HDD 10.5k	MB/s 24x HDD 10.5k
4k Byte	80000	320	11200	45
8k Byte	52000	416	10800	86
16k Byte	30000	480	10000	160
32k Byte	15000	480	9200	294

As block size rises, so the sequential performance specification increasingly dominates the result. Although SSDs are twice as fast as HDDs, when 24 HDDs are used in parallel, the HDD solution exhibits higher performance for block sizes above 64kB.

Blocksize	IOPS 8x SSD	MB/s 8x SSD	IOPS 24x HDD 10.5k	MB/s 24x HDD 10.5k
4k Byte	80000	320	11200	45
8k Byte	52000	416	10800	86
16k Byte	30000	480	10000	160
32k Byte	15000	480	9200	294
64k Byte	7100	454	8100	518
128k Byte	3750	480	6900	883
256k Byte	2300	589	5400	1382
512k Byte	1500	768	2800	1434
1M Byte	1050	1075	1425	1459
2M Byte	590	1208	715	1464
4M Byte	330	1352	410	1679
8M Byte	170	1393	235	1925
16M Byte	90	1475	145	2376
32M Byte	45	1475	85	2785
64M Byte	23	1507	46	3015

The graphical summary below highlights the crossing point:



In real world applications, the actual workload usually comprises a mixture of smaller and larger block and file sizes. This was approximated using a generic workload comprised of a defined mix of block sizes. The results were as follows:

Block size	IOPS SSD	MB/s SSD	IOPS HDD	MB/s HDD
4k 20%, 64k 50%, 256k 20%, 2M: 10%	2400	800	3300	980

Clearly, the HDD solution offers the greatest performance as multiple spindles are working in parallel.

Operational cost

Clearly, 24 HDDs require more space and consume more power than 8 SSDs. In terms of space, the HDDs will require 2U of additional external rackspace while the SSDs can be implemented directly in the server, requiring no additional footprint in the datacenter.

However, the comparison of the power requirements is not as clearly defined so Toshiba set out to quantify this.

The power consumption of both configurations was measured and found to be 90 W for the SSDs (drives and backplane) and 230 W for the HDDs (drives and JBOD) when running the mixed workload as defined above.

Storage	8x eSSD SATA 1.6TB 3DWPD	24x HDD SAS 2.4TB 10.5krpm
Configuration	RAID 6	RAID 10
Controller	Microsemi Adaptec SmartRAID 3154-8i	Microsemi Adaptec SmartRAID 3154-8e
Chassis	Supermicro CSE-826BAC4 (in Server)	AIC XJ1-20242-05 2U 24 x 2.5" hot swap bays, hot swap JBOD with dual SAS 12G expander controller, Tool-less HDD tray, 549W 80+ Platinum RPS (external)
Net Capacity (TB)	9,6	28,8
Power for SAS BP + Drives (W)	90	230
Power/Year (kWh)	788	2015
OPEX/Year at €0.07/kWh, PUE 1.3	72	183

Assuming that datacenters pay around €0.07/kWh with a Power Usage Effectiveness (PUE) of 1.3, the HDD solution has an operational expenditure (OPEX) that is around €110 higher per year than the SSD alternative. This equates to less than €10 per month of additional cost - for a solution that offers three times the storage capacity and higher performance for most real-world workloads.

Toshiba Enterprise Performance HDD lineup

Toshiba develops and manufactures 10.5krpm Enterprise Performance HDDs within the AL14SE and AL15SE series, and is committed to continue the manufacturing and support of these models for many more years. AL14SE series drives offer capacities up to 1.8TB, while drives in the AL15SE series extend to 2.4TB.

HDDs in both series are available with a 4kB block size, or alternatively, with an external 512B emulated on internal 4kB blocks for systems that require 512B formats. The exact size is 512B, 520B or 528B to support the SAS protection overhead, so the terminology used is '5xx'. To support legacy systems and replacements within older installations, a native 512B block size model is also available for the smaller capacity models.

Model Number	Specification						Performance		Reliability		
	Form-factor	Capacity	Block Size	Spindle Speed	Buffer Size	Interface	Data Rate (sustained)	Operating Power	MTTF	Unrecov.	Duty
		[GB]	[Byte]	[rpm]	[MiByte]		[MB/s]	[W] typ	[hrs]	Error Rate	
AL15SEB030N	2.5"	300	5xxn	10500	128	SAS 12Gbit/s	236	6.5	2.0M	1 in 10 ¹⁶	24/7
AL15SEB060N		600					245	7.0			
AL15SEB090N		900						7.0			
AL15SEB120N		1200						8.1			
AL15SEB06EQ		600	5xxe				273	6.5			
AL15SEB09EQ		900						7.0			
AL15SEB12EQ		1200						7.0			
AL15SEB18EQ		1800						8.1			
AL15SEB24EQ		2400	8.7								
AL15SEB06EP		600	4kn				273	6.5			
AL15SEB09EP		900						7.0			
AL15SEB12EP		1200						7.0			
AL15SEB18EP		1800						8.1			
AL15SEB24EP		2400						8.7			

Summary and conclusion

A solution built around 24 10.5krpm Enterprise Performance HDDs from Toshiba's AL14SE or AL15SE series offer up to three times the net capacity when compared to a solution consisting of eight Enterprise SSDs. However, despite the difference in capacity, the cost of both approaches is about the same.

For workloads based upon smaller block sizes (between 4kB and 32kB), the SSD solution is faster. This workload profile typically occurs in applications such as databases, high performance analytics and financial transactions. However, the required storage capacity in these applications is generally quite low (no more than a few Terabytes), so the cost of SSD technology is justified by the performance gain delivered by SSD based solutions.

For mixed workloads or workloads dominated by block sizes of 64kB or more, the multi-spindle HDD solution with its higher capacity (at approximately the same cost) is actually faster than the SSD approach. Most of workloads such as web hosting, email serving, generic storage, virtual storage, cloud storage, document management, shared drives, surveillance storage, streaming, content delivery, backup and archiving fall into this workload category.

Further considerations

Having confirmed that multiple inexpensive, relatively slow storage devices running in parallel can outperform theoretically faster solutions, other topics come to the fore. For example, it would be interesting to evaluate if greater quantities of Enterprise Capacity / Nearline HDDs, which are inexpensive but even slower, can achieve similar performance.

With capacities up to 14TB per drive and very competitive pricing, such a solution could support storage requirements up to Petabytes as is required for surveillance storage, document management, streaming, content delivery, backup and archiving.

One area which would be worthy of further testing and evaluation is whether a large array of inexpensive Enterprise Capacity / Nearline HDDs can also achieve the performance required for more agile storage workloads such as web hosting, email serving, generic storage, virtual storage, and shared drives.

Ever curious, and wanting to provide customers with the best application knowledge, Toshiba engineers have already set up a 4U, 60 bay JBOD with 3.5" Nearline Drives, and intend to evaluate system performance and power consumption to see what is possible. Perhaps this will be addressed in a forthcoming white paper, sometime soon? In any case, this evaluation has proved that HDD still has a lot of life left in it.

Contact us for more information:

www.toshiba-storage.com/contact/





TOSHIBA

Toshiba Electronics Europe GmbH

Hansaallee 181
40549 Düsseldorf
Germany

info@toshiba-storage.com
toshiba-storage.com

Copyright © 2019 Toshiba Electronics Europe GmbH. All rights reserved. Product specifications, configurations, prices and component / options availability are all subject to change without notice. Product design, specifications and colours are subject to change without notice and may vary from those shown. Errors and omissions excepted.

10/2018