

TOSHIBA



Enterprise Hard Drives

Designed for your business

In direct comparison:

Sixty Enterprise 3.5-inch Hard
Disks in RAID 10 versus eight
Enterprise SSDs in RAID 6

Can a 100% HDD-based solution meet the requirements of tasks such as “Storage for Virtualisation“ and “Email Serving“? Toshiba Electronics Europe (TEE) compares Enterprise Capacity HDD and SSD based solutions to build 50TB to 100TB storage.

Introduction

With the purchase of new storage systems, a question arises: whether to choose conventional technology i.e. rotating hard disks or HDDs or much faster and more robust semiconductor-based Solid State Drives (SSDs)?

Except in particularly cost-sensitive markets, SSDs are the preferred choice for workplace computers with one or two local storage components and capacity requirements in the area of 250GB to a few TB. However, the situation differs when it comes to large centralised storage servers in data centres and cloud storage. Enterprise SSDs are up to 10x more expensive than hard drives per capacity and the industry is still far away from any price parity. With large storage systems in the range of dozens of TB or more it makes sense, from a commercial point of view, to use comparatively cost-effective HDD-based storage. However, HDDs are known to be slower than SSDs: about 25% to 50% slower at sequential access and magnitudes of tens slower when it comes to random access. Thus, the demand for “agile” storage cannot be neglected when comparing HDDs against SSDs.

Approach to enterprise storage

There are two ways to look at this:

- 1) How to implement inherently fast SSD-based, large capacity storage as cost effectively as possible?
- 2) How to configure slower HDD-based storage to be as agile as possible?

The ideal solution would be a “hybrid” combination of SSD and HDD in storage systems with caching or tiering. There is a lot of news about this technology as it offers useful and attractive solutions, but the performance, however, is dependent on current implementation, data access, configurations and can rarely be predicted with certainty. To understand more it is useful to look at a direct comparison of the monolithic implementation of local RAID sets between HDDs and SSDs.

The initial questions were: Is it possible to build a storage system of 50TB to 100TB with 60 comparatively low-priced enterprise capacity (“nearline”) HDDs with 7200rpm to meet the needs of “storage for virtualisation” and “e-mail serving” or do you have to choose SSD-based solutions? Which architecture is suitable for HDDs to achieve maximum possible performance and how big is this system? How should SSDs be used as an alternative?

The test case

Sixty 2TB 7200rpm 12GB/s SAS HDDs (Toshiba MG04SCA20EE) were installed in a Chenbro 4U 60 bay “RM43260” JBOD as RAID10 on a Microsemi Adaptec ASR8885 and with all cache functions turned off for comparability. Thus, 60TB net capacity is available and the controller can distribute data to various disks concurrently thanks to the RAID0 functionality. The protection function of the RAID1 – solely mirroring – takes complex parity calculations off the controller.

Solution	HDDs – many spindles in parallel
Number of devices	60
Type	2TB 12GB/s SAS 3.5" Enterprise
Configuration	RAID10
Model	Toshiba MG04SCA20EE
Chassis	Chenbro 4U 60 Bay ”RM43260“
RAID Controller Model	Microsemi Adaptec ASR8885
Net storage capacity	60TB

Results for the application as an e-mail server

With the typical load of an e-mail server – “fio“-emulated in the lab by 64k byte blocks, randomly writing and reading from 12 jobs – the 60TB storage system achieves 8300MB/s data throughput at 1300 I/O-operations per second (IOPS).

Comparison with SSD solutions

For comparison, a storage system with enterprise SSDs was built, but kept to the commercial framework set by the 60 HDD storage. For the budget of 60 2TB HDDs and a simple 60-bay JBOD it is possible to obtain between eight to ten enterprise SSDs with 1.6TB storage capacity and a write endurance of 3 DWPD. Three DWPD should be the minimum for “agile multi-purpose” enterprise systems, as the HDDs to be compared have no write endurance limitation due to the technology. Eight SATA 6GB/s eSSDs of 1.6TB storage capacity (Toshiba THNSN81Q60) were integrated directly into the front plane of the server and configured as RAID6 (also on an Adaptec controller). In this case on a similar, second model, to avoid interference when operating both RAIDs. RAID6 was chosen because the comparatively expensive SSDs should be optimised for minimal redundancy overhead.

Solution	HDDs	All Flash SSD
Number of devices	60	8
Type	2TB 12GB/s SAS 3.5" Enterprise	1.6TB SATA Enterprise SSD 3DWPD
Configuration	RAID10	RAID6
Model	Toshiba MG04SCA20EE	Toshiba THNSN81Q60
Chassis	Chenbro 4U 60 Bay "RM43260"	Supermicro 826BAC4
RAID Controller Model	Microsemi Adaptec ASR8885	Microsemi Adaptec ASR8885
Net storage capacity	60TB	9.6TB
Overall budget	Equal	Equal

With the same synthetic "e-mail server" load, the SSD array offers 600MB/s at 9400 IOPS, which is slightly less than the HDD solution. Although a single SSD is much faster than the HDD, the HDD solution scores due to the large number of thirty parallel disks. In the SSD model there are eight disks working in parallel and the controller is still engaged in calculating the parity information of the RAID6. The main contributor in this application, however, is the available net storage capacity: the HDD array offers 60TB, the SSD array at the same cost and with the same performance only 10TB.

Configuration		60x HDD in RAID10	8x SSD in RAID6
Net storage capacity		60TB	9.6TB
Application: e-mail server	IOPS	13000	9400
(64k Byte Random r/w 12 jobs)	MB/s	830	600

Storage systems for virtualisation

Let's look at other applications for the same configurations: As an approximate synthetic workload of a storage server for virtualized environments, we assume 16k byte blocks, random read/write through 12 jobs.

The HDD RAID10 offers 230MB/s with 14400 IOPS, therefore, due to the smaller block size, it takes longer finding the data than the actual read/write. Owing to its technology, the SSD array does not have to search for data on rotating disks but delivers 36200 IOPS and 560MB/s – which is two times faster. However, there is still six times less storage capacity. The decision should be: "It depends" – Do I want to optimise for processing speed or high storage capacity?

Configuration		60x HDD in RAID10	8x SSD in RAID6
Net storage capacity		60TB	9.6TB
Application: virtualization	IOPS	14400	36200
(16k byte Random r/w 12 jobs)	MB/s	230	560

With the typical load of a highly active business database – 4k byte blocks random write/read over 12 jobs – the SSD array is clearly ahead: 55000 IOPS versus 15000 IOPS with HDDs. The MB/s value is not really relevant here. As databases in most cases do not extend to dozens of TB, the argument for large capacity here is usually not crucial.

Configuration		60x HDD in RAID10	8x SSD in RAID6
Net storage capacity		60TB	9.6TB
Application: virtualization (4k Byte Random r/w 12 jobs)	IOPS MB/s	15000 60	55000 220

Storage systems for archiving and (video) streaming

Finally, the engineering team also looked at the SSD/HDD comparison for “classic HDD” loads: archives and video streaming:

For data archiving purposes only, large blocks are written sequentially, so high-speed writing into large low-cost storage capacity is essential. With 1Mbyte data blocks written sequentially, the HDD array achieves 2900MB/s, while the SSD array can be written at 2500MB/s. Again, we see the effect of the “many spindles” reinforced by the slowing down of the SSD RAID6 by the need to perform a complex parity data calculation when writing all data.

Video streaming/playout reads large blocks in multiple parallel streams. With a workload of 1Mbyte blocks being read sequentially by four jobs, the HDD array will achieve 6400MB/s and the SSD array only 4200MB/s: at the same cost, but only with 1/6 of the net storage capacity. Thus, also for archiving and streaming tasks (amongst others) using HDDs will still be justified in the long term future.

Configuration		60x HDD in RAID10	8x SSD in RAID6
Net storage capacity		60TB	9.6TB
Application: archiving (1Mbyte sequential write, 1 job)	MB/s	2900	2500
Application: video streaming (1Mbyte sequential read, 4 jobs)	MB/s	6400	4200

Further considerations

It should be noted that 60 spinning HDDs in a separate 4U 60 bay toploader JBOD naturally consume more power and space than eight SSDs in the front panel of the server, so SSDs are cheaper in terms of operating costs. Furthermore, statistically you can expect more disk failures over the lifetime of 60 HDDs than with eight SSDs, which is another SSD advantage in terms of operating costs. This should be the decisive factor for SSD-based solutions in borderline cases where high-performance requirements are needed in combination with small to medium capacity. Alternatively, it is possible to also use classical database storage systems, for example 24 enterprise 2.5” performance HDDs with 10.500rpm as RAID10 in the front panel of the server. In this case, the performance is higher than 7200rpm 3.5” nearline HDDs but the cost per TB is also higher. However, this is still more cost effective than a mere SSD solution.

Conclusion

It has been shown that with all storage systems requiring “medium-high” speed and high storage capacity at low cost per capacity there is still an argument for relying on HDD systems with many spindles. In particular, RAID10 configurations with many spindles deliver the highest capacity at low purchase prices while offering competitive performance compared to SSD solutions.

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.

02/2018