TOSHIBA

Lab Report

QSAN XCubeSAN “XS3324D” with Toshiba 18TB Nearline SAS HDD

Author: Rainer W. Kaese, Senior Manager
Business Development, Storage Products Division,
Toshiba Electronics Europe GmbH

One of the biggest challenges when selecting large-scale, multi-functional storage for small and medium-sized business (SMB) use is the true, attainable performance once installed. While there is a wide range of storage solutions on the market (many with similar processors, memory capacities and features), there remains one unknown when it comes to final performance metrics – the impact of the used hard disk drives (HDDs).

This lab report reviews the QSAN XCubeSAN XS3300 Series, a powerful family of hybrid (flash and HDD based) storage products targeting the needs of SMBs and enterprise users. The HDDs selected for testing were from the Toshiba Enterprise Capacity MG-Series that are well suited to enterprise storage arrays and systems.

Initial Configuration

The model provided for testing was QSAN’s XS3324D, a 4U/24-bay front loading fully-integrated storage area network (SAN) block-storage solution with dual/redundant SAS controller and power supply (as described in Figure 1). It operates as block storage for dedicated SANs. This is support over iSCSI, fibre channel (with optional FC modules – not tested), or both. Due to its dual-controller architecture, providing a dual-path from network down to HDD access, it is best suited to high-capacity, SAS Nearline HDDs. It features a number of field replaceable units (FRU) that can be hot-swapped in the unlikely event of component failure.

Figure 1: The QSAN XS3324D

Setup

To fully test its large-scale capabilities, 24 of Toshiba's top Enterprise Capacity MG-Series HDDs were installed: the 18 TB SAS 12 Gb/s model MG09SCA18TE (see Figure 2).

Each of the controllers was connected to a dedicated SAN over two out of the four 10 Gb/s SFP+ connections. In addition, one controller is connected to a 1 Gb/s LAN (yellow cable) for management access. This controller acts as master controller for the entire unit (as indicated by the green "M/S" LED).

For reference: The model version tested was XS3324D, with Firmware Version 2.2.1.

Figure 2: Toshiba MG09SCA18TE Enterprise SAS HDD

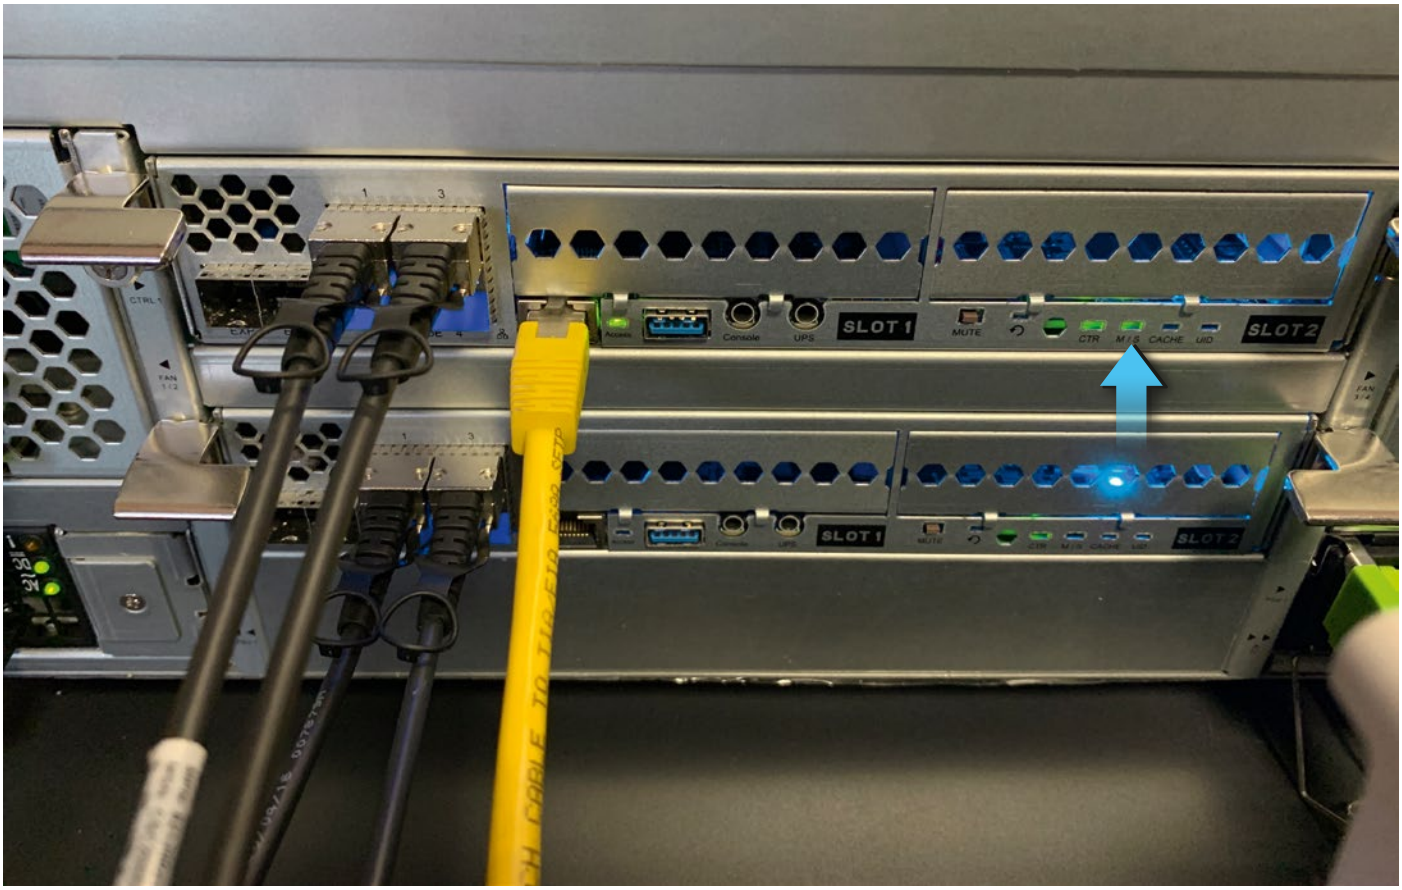

Figure 3: XS3324D Network Connections

One Pool of 24 HDD Configuration

For highest possible consistency of performance, all 24 HDDs (432 TB gross capacity) were configured into one single pool with a RAID10 configuration with 216 TB net capacity. This pool is comprised of four installed volumes of about 50 TB each.

Evaluations show the base consistent and continuous performance of the HDD array. In practical operations, some SSD caching may be implemented, but the gain over the pure base performance of the HDD array will be dependent on the actual workload, so evaluation of only the HDD base array was conducted.

Drive Benchmark with Different Workload Patterns

For iSCSI block storage, we measured not only bandwidth but also detailed drive performance for different workloads (same as we usually do in the lab to benchmark single HDDs and SSDs).

At first we exercised one volume via iSCSI 10 Gb/s connections:

Workload	IOPS	Bandwidth (MB/s)
Sequential Write 1024 kB		1111
Sequential Read 1024 kB		1059
Random Write 4 kB	6623	
Random Read 4 kB	5109	
Mixed 4 kB / 64 kB / 256 kB / 2MB	700	201

Table 1: Performance Measurements for One iSCSI Drive

The sequential speed reached the limit of the network connection (10 Gbp/s~1.14 GB/s). 5-6k IOPS is an impressive result for HDD-only based storage. It would even be suitable for certain high active storage workloads.

For two drives via iSCSI on the same server:

Workload	IOPS	Bandwidth (MB/s)
Sequential Write 1024 kB		351 + 362
Sequential Read 1024 kB		490 + 537
Random Write 4 kB	3324+3295	
Random Read 4 kB	2379+2297	
Mixed 4 kB / 64 kB / 256 kB / 2MB	276+281	81+81

Table 2: Performance Measurement for Two iSCSI Drives

Here the performance was equally split between the two active iSCSI volumes, with some (but rather low) loss for the switching/arbitration between the two volumes.

Finally, testing was conducted across four drives via iSCSI on two servers:

Workload	IOPS	Bandwidth (MB/s)
Sequential Write 1024 kB		159+199+ 221+129
Sequential Read 1024 kB		168+199+ 336+320
Random Write 4 kB	1751+1784+ 1538+1538	
Random Read 4 kB	1085+1543+ 1629+887	
Mixed 4 kB / 64 kB / 256 kB / 2MB	123+125+ 128+131	36+36+37+ 38

Table 3: Performance Measurement for Four iSCSI Drives

The sequential and random performance is still shared rather equally between all four active iSCSI volumes. With more than 1k IOPS for each iSCSI this is still suitable for active storages, especially in virtualized server environments.

TOSHIBA

Scripts Used for Test Sequences

Each test sequence executes:

- sequential write 100% blocksize= 1024 kB
- sequential read 100% blocksize= 1024 kB
- random write 100% blocksize= 4 kB
- random read 100% blocksize= 4 kB
- random read 50% / write 50% blocksize= 4 k / 20%, 64 k / 50%, 256 k / 20% , 2 M / 10%

```
fio --filename=\\.\Physicaldrive4 --direct=1 --rw=write --bs=1m --iodepth=16 --time_based --
runtime=300 --group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=1 --
norandommap --randrepeat=0 --output=seqwritephysical.log

fio --filename=\\.\Physicaldrive4 --direct=1 --rw=read --bs=1m --iodepth=16 --time_based --
runtime=300 --group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=1 --
norandommap --randrepeat=0 --output=seqreadphysical.log

fio --filename=\\.\Physicaldrive4 --direct=1 --rw=randwrite --bs=4k --iodepth=16 --time_based --
runtime=300 --group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=64 --
norandommap --randrepeat=0 --output=randwritephysical.log

fio --filename=\\.\Physicaldrive4 --direct=1 --rw=randread --bs=4k --iodepth=16 --time_based --
runtime=300 --group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=64 --
norandommap --randrepeat=0 --output=randreadphysical.log

fio --filename=\\.\Physicaldrive4 --direct=1 --rw=randrw --bssplit=4k/20:64k/50:256k/20:2M/10 --
iodepth=16 --time_based --runtime=300 --group_reporting --name=job1 --ioengine=windowsaio --
thread --numjobs=64 --norandommap --randrepeat=0 --output=mixedphysical.log
```

Power Measurements

Power was measured for the different workloads (with 4 iSCSI drives active):

Workload	Power
Sequential Write 1024 kB	380 W
Sequential Read 1024 kB	370 W
Random Write 4 kB	350 W
Random Read 4 kB	410 W
Mixed 4 kB/64 kB/256 kB/2 MB	400 W
Maximum Power at Startup	450 W
Power When Unit is Idle (no read/write)	320 W
Standby Power (unit off)	9 W

Table 4: Power Measurements

The baseload of the XS3324D system running 24 disks is about 300W. Additional power depends on activity/workload. Power at random reading is highest, since the HDDs have to constantly seek for random data, so the voice-coil-motor driving the actuator arm draws the highest power levels.

A Rohde&Schwarz HMC8015 high precision power analyser was used for carrying out the testing. This was connected to both power mains of the XS3324D for accurate power measurements.

Figure 4: Power Measurement Setup

Temperature and Noise

To round out the results of this lab report, we measured the HDD temperature and noise level figures:

Noise 1 m distance rear: ..... 74 dB
Noise 1 m distance front: ..... 60 dB

Ambient Lab Temperature: ..... 23 degC
HDD Temperature (lowest at idle): ..... 29 degC
HDD Temperature (highest at idle): ..... 36 degC
HDD Temperature (lowest at active): ..... 33 degC
HDD Temperature (highest at active): ..... 38 degC

Conclusion

The tested QSAN XD3324D SAN storage system provides a large capacity, high performance, highly available and reliable enterprise data storage solution.

With 24 Toshiba 18 TB Enterprise SAS HDDs, the raw capacity is 432 TB, and depending on configurations, from 216 TB (1 Pool of 24 in RAID10) up to 360 TB (2 pools of 12 each in RAID6) net capacity.

Up to 1000 MB/s sequential performance and 5000 to 6000 random IOPS (even without SSD cache) are excellent performance benchmarks.

Maximum power consumption under full load does not exceed 400 W (even less than 2W per TB net capacity, including expander, controller, 10 Gb/s SFP+ and RJ network interfaces). Power efficiency is therefore always maintained at elevated levels.

The noise levels exhibited are acceptable for rackmount units.

The temperature of the HDDs remains low – as appropriate and efficient cooling is implemented. This supports long lifetime and low failure rates for the spinning disks.

TOSHIBA

Toshiba Electronics Europe GmbH

Hansaallee 181
40549 Düsseldorf
Germany

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

Issued 01/2023