

Lab Report

QSAN XCubeNXT XN8000D with Toshiba 16TB Nearline SAS HDDs.

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 businesses (SMB) is the true, attainable performance once installed. While there is a wide range of storage solutions on the market, many with similar processors, memory, and features, there remains one unknown when it comes to final performance metrics – the impact of the used hard disk drives (HDD). This lab report reviews the QSAN XCubeNXT XN8000D, a powerful family of unified 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 XN8024D, a 4U/24-bay front loading fully integrated unified storage solution with dual/redundant SAS controller and power supply (Figure 1). It can operate both as network-attached storage (NAS), providing traditional shared files and folders, and simultaneously as block storage for dedicated storage area networks (SAN). This is support over iSCSI, fibre channel, 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 XN8024D



Figure 2:
Toshiba HDD
16TB SAS
MG08SCA16TE

To fully test its large-scale capabilities, 24 of Toshiba's top Enterprise Capacity MG-Series HDD were installed: the 16 TB SAS 12Gb/s model MG08SCA16TE (Figure 2). Each of the controllers was connected to a dedicated storage area network over 10GB/s SFP+ connections using the installed 4-port module. For the purposes of evaluation, only a single-path network connection was used. In a production environment, a dual-path connection could be used, providing a stable and reliable connection from the HDDs to the client/user. In addition, two further 10GbE LAN connections were made, while the 1GbE management ports were connected to allow access to the browser-based configuration interface of the QSM operating system.



Figure 3: Two 10GB/s SFP+ connections (left) together with the 10GbE LAN (yellow cables) and 1GbE management (black cables) ports (right).

Initial setup and testing: single pool of 24 HDDs

To establish performance figures for a single storage pool, the 24 HDDs (384 TB gross capacity) were configured in a RAID10 configuration to form a pool of 192 TB net capacity. Two volumes were installed on the pool, each consisting of a 48 TB shared folder and a 48 TB iSCSI target. A production system would typically include some SSD caching but the performance gain this provides is heavily dependent on the workload. With a focus on establishing baseline performance, initial testing focused on this pure HDD array.

Testing examined the read- and write-bandwidth individually for NAS (shared network folder) and SAN (iSCSI with logical

volume mapped into a server). This was followed by testing both operations in parallel. In the results, shown in Table 1, the “1 × Writing” is one write process from one server, while “2 × Writing” indicates one write process from one server to one folder with a second write process from a second server to a second folder. The same approach applies to the reading tests. Further testing examined the achievable bandwidth with both NAS and SAN operations in parallel with the results shown in Table 2.

Overall, a total bandwidth of around 800 to 900 MB/s was observed distributed across the different tasks.

Single Pool Configuration		
Test	NAS Operation (Shared Network Folder)	SAN Operation (iSCSI Dedicated Storage Network)
1 × Writing	600 MB/s	600 MB/s
2 × Writing	400 + 400 MB/s	450 + 450 MB/s
1 × Reading	800 MB/s	600 MB/s
2 × Reading	425 + 425 MB/s	450 + 450 MB/s

Table 1: Single pool configuration bandwidth test results for NAS or SAN operation.

Single Pool Configuration – NAS & SAN Parallel Operation	
Test	Bandwidth
1 × NAS Read	180 MB/s
1 × NAS Write	120 MB/s
1 × SAN Read	650 MB/s
1 × SAN Write	470 MB/s

Table 2: Single pool configuration bandwidth test results for parallel NAS and SAN operation.

iSCSI Disk Benchmarking

The raw bandwidth measurements on their own do not accurately reflect the workload of a real, production storage solution. Toshiba regularly tests individual HDDs and SSDs under conditions that more accurately reflect real storage workloads of storage for email servers, databases, and video surveillance systems. These use a combination of random reads and writes for 4 kB block sizes, sequential reads and

writes for 1 MB block sizes, and mixed read/writes across a range of block sizes. The measurements obtained resulted from the tool 'fio' using the scripts shown in Listing 1. The results in Table 3 are for a Windows logical volume on an iSCSI target implemented in a RAID10 of a 24 HDD pool. The sequential read of 1 MB blocks reached the theoretical 1.12 GB/s bandwidth limitation of the 10GbE network.

```

fio --filename=test --size=200G --direct=1 --rw=randread --bs=4k --iodepth=16 --time_based --runtime=200
--group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=16 --norandommap --randrepeat=0
--output=randread.txt

fio --filename=test --size=200G --direct=1 --rw=randwrite --bs=4k --iodepth=16 --time_based --runtime=200
--group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=16 --norandommap --randrepeat=0
--output=randwrite.txt

fio --filename=test --size=200G --direct=1 --rw=read --bs=1024k --iodepth=16 --time_based --runtime=200
--group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=4 --norandommap --randrepeat=0
--output=seqread.txt

fio --filename=test --size=200G --direct=1 --rw=write --bs=1024k --iodepth=16 --time_based --runtime=200
--group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=4 --norandommap --randrepeat=0
--output=seqwrite.txt

fio --filename=test --size=200G --direct=1 --rw=randrw --bssplit=4k/20:64k/50:256k/20:2M/10 --iodepth=16
--time_based --runtime=200 --group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=8
--norandommap --randrepeat=0 --output=mixed.txt
    
```

Listing 1: Parameters of iSCSI tests using 'fio'.

Single Pool Configuration – Windows logical volume	
Test	iSCSI Disk Benchmark
Random read 4 kB	2740 IOPS
Random write 4 kB	1970 IOPS
Sequential read 1 MB	1120 MB/s
Sequential write 1 MB	940 MB/s
Mix 4k/64k/256k/2M (20%/50%/20%/10%)	700 IOPS / 200 MB/s

Table 3: Single pool configuration bandwidth test results for parallel SAN and NAS operation accessing a Windows logical volume.

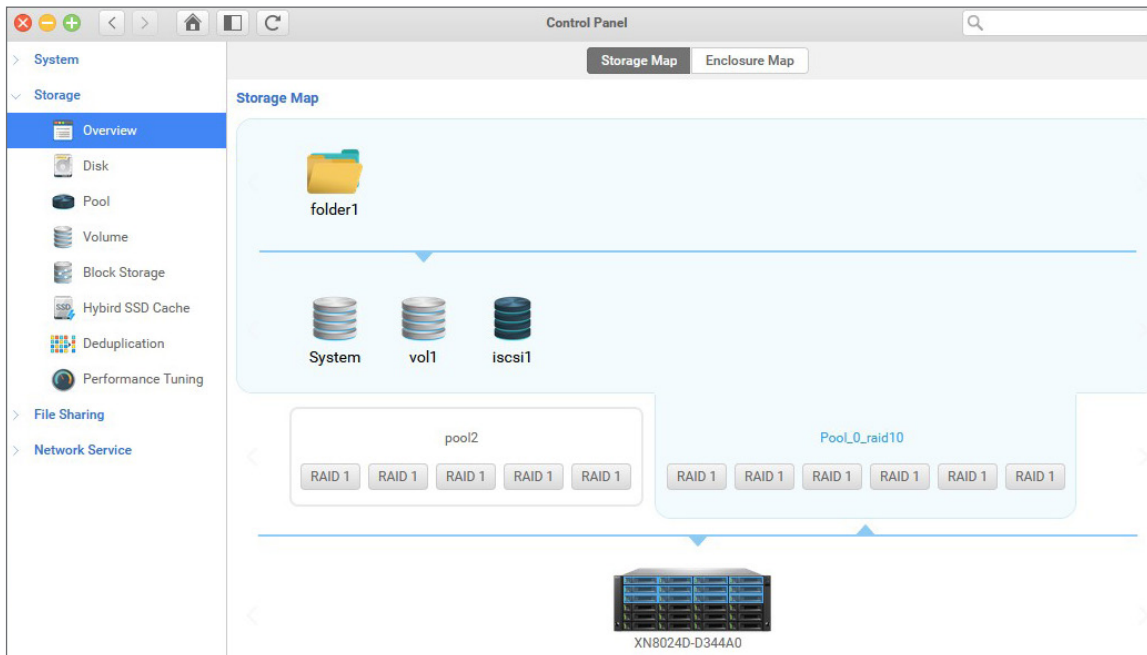


Figure 4:
The XN8024D configured as two 12 HDD pools.

Two pools of 12 HDDs

The QSAN paper “How to Adjust Performance in Windows” examines best practices using the XCubeNXT series. The document suggests configuring the unit as two pools to achieve highest performance. Bearing in mind that two controllers are available, this makes logical sense as it allows each pool to be accessed independently. However, this means that each pool is comprised of 12 HDDs, and therefore 12 spindles, compared to the 24 spindles of the single-pool configuration. Thus, it is typically expected that such a change will result in lower bandwidths and system performance. To check this, the XN8024D was configured into two pools

each of 12 HDDs in RAID10 (Figure 4). Each pool consisted of a 48 TB iSCSI target together with a 48 TB shared folder, accessed via its own controller (i.e. no dual path). As expected, the performance of a single shared folder or iSCSI target in this 12-HDD configuration was lower than that of the single pool of 24 HDDs (Table 4). However, the combined performance of the two 12-HDD pools was still better than half the performance of the single 24-HDD pool. Furthermore, when accessing both logical pools simultaneously, the results show that they shared the total performance of the 24-HDD pool.

Test	NAS Operation		SAN Operation	
	1 Pool	2 Pool	1 Pool	2 Pool
1 × Writing	600 MB/s	450 MB/s	600 MB/s	350 MB/s
2 × Writing	400 + 400 MB/s	300 + 300 MB/s	450 + 450 MB/s	280 + 280 MB/s
1 × Reading	800 MB/s	600 MB/s	600 MB/s	450 MB/s
2 × Reading	425 + 425 MB/s	400 + 400 MB/s	450 + 450 MB/s	300 + 300 MB/s

Table 4: Bandwidth test results comparing 1 × pool 24 HDDs and 2 × pool 12 HDDs in NAS and SAN operation.

	1 Pool		2 Pool	
Test	1 job/1 iSCSI Target	1 job/1 iSCSI Target	2 jobs/2 Different iSCSI Targets	
Random read 4 kB	2740 IOPS	2590 IOPS	2 × 2460 IOPS	
Random write 4 kB	1970 IOPS	1790 IOPS	2 × 1750 IOPS	
Sequential read 1 MB	1120 MB/s	980 MB/s	2 × 770 MB/s	
Sequential write 1 MB	940 MB/s	630 MB/s	2 × 460 MB/s	
Mix 4k/64k/256k/2M (20%/50%/20%/10%)	700 IOPS / 200 MB/s	620 IOPS / 180 MB/s	2 × 630 IOPS / 2 × 190 MB/s	

Table 5: ‘fio’ workload tests comparing 1 × pool 24 HDDs and 2 × pool 12 HDDs.

The more realistic ‘fio’ workload tests, as show in Listing 1, were replicated on the two 12-HDD pool configuration. The results in Table 5 show that, while a single job on a single iSCSI target performed slightly less well than the 24-HDD single pool, two jobs on two different iSCSI targets, when combined, are considerably higher.

From these results it is clear that, should a single shared-folder or large block drive with highest possible performance be desired, the single pool of 24 HDDs is the optimal choice. However, the more practical case of several folders and/or blocks of individually accessed storage, is better implemented using two pools of 12 HDDs thanks to the higher-attainable workload performance.

The addition of SSD caching

As already mentioned, SSD caching is used in production systems to boost performance and primarily improves the system’s random read and write access benchmarks. The two pool configuration was retained. In order to accommodate a 1.6 TB SAS Enterprise SSD (10 DWPD write endurance), one of the 12 HDDs has to be removed to make space. Due to the RAID10 configuration, the 11th HDD is unusable but can remain, marked as a global hot spare drive (Figure 5).

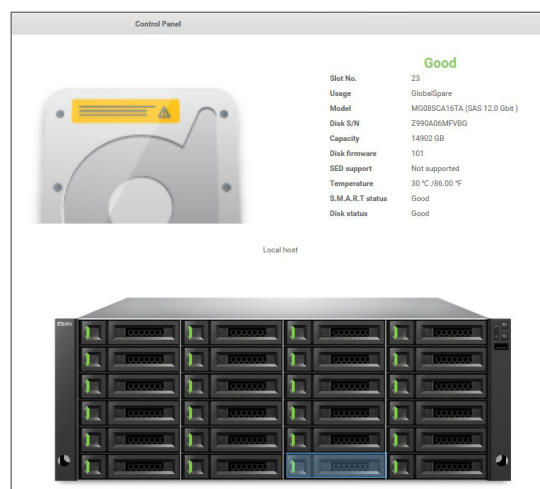
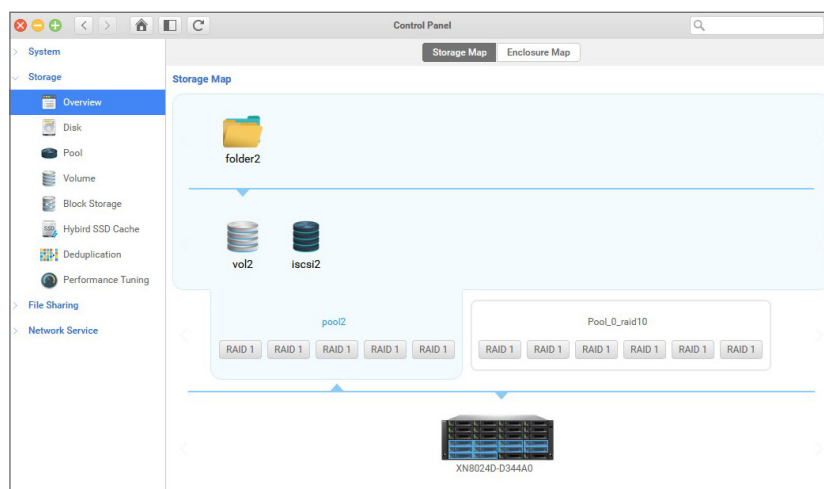


Figure 5: One XN8024D 12 HDD pool reconfigured with an SSD cache (left); 11th HDD remains as global hot spare drive (right).

Again, the benchmarking scripts in Listing 1 were used to assess the new configuration. As is to be expected, sequence read and write performance benefited minimally from the new cache capability. However, the random read/write performance showed a significant factor of 4 to 6 improvement (Table 6).

Power, temperature and noise

To round out the results of this lab report, a range of power consumption and noise level measurements were made and recorded. Furthermore, the system's temperature at different stages of testing was also measured (Table 7).

Single Pool Configuration		
Test	12 × HDDs RAID10	10 × HDDs RAID10 + 1 × SSD Cache (1.6 TB)
Random read 4 kB	2590 IOPS	16880 IOPS
Random write 4 kB	1790 IOPS	6500 IOPS
Sequential read 1 MB	980 MB/s	990 MB/s
Sequential write 1 MB	630 MB/s	670 MB/s
Mix 4k/64k/256k/2M (20%/50%/20%/10%)	620 IOPS / 180 MB/s	780 IOPS / 220 MB/s

Table 6: 12 × HDD pool vs. 10 × HDD + cache pool benchmarking results.

Measurement	Result
Maximum power at start-up	450 W
Power under full read/write workload	400 W
Power when unit is idle (no read/write)	340 W
Standby Power (unit off)	9 W
Noise 1 m distance rear	78 dB
Noise 1 m distance front	58 dB
Ambient Lab Temperature	25 °C
Minimum HDD Temperature (smart-value)	28 °C
Maximum HDD Temperature (smart-value)	32 °C

Table 7: Power, noise and temperature measurements for the QSAN XN8024D during testing.

Conclusion

The QSAN XN8024D offers enterprise and SMB system administrators and storage consultants a large capacity, high-performance, and an highly available and reliable data storage solution. The system is efficiently cooled as noted by the low-spread HDD temperatures measured, helping to maintain the expected lifetime and low failure rate of spinning-platter drives. In addition, the noise generated is acceptable for this type of rackmount unit.

Combined with 24 Toshiba 16 TB Enterprise SAS hard disk drives, a raw capacity of 384 TB can be obtained. Depending on the configuration, this translates to a net capacity of between 192 TB (1 pool of 24 HDDs in RAID10) and 320 TB (2 pools of 12 HDDs in RAID6). The sequential read/write performance of 1,000 MB/s and 2,000 – 3,000 IOPS (without

SSD cache) is a very respectable result. For a single-block or file storage, a single pool of 24 HDDs is the preferred approach. However, when offering multiple blocks or concurrent access to shared folders in NAS operation, two pools of 12 HDDs deliver higher performance overall.

The XN8024D's dual controllers provide high availability for those that demand this. During testing it was demonstrated that this feature can also be used to eliminate the single-path bottleneck in both block storage and NAS operation. With total power consumption lying at under 400 W, equivalent to 2 W per TB net capacity, and considering that this includes the expander, controllers, 10GN/s SFP+ and RJ network interfaces, this unified storage solution also delivers excellent value in terms of power efficiency.

Ref. TSH047_D2_EN

¹ Model used: XN8024D-D344A0 with Firmware Version 3.3.2

² [https://www.qsan.com/data/dl_files/QSAN_Best%20Practice_XCubeNXT_How%20to%20Adjust%20Performance%20in%20Windows_2007_\(en\).pdf](https://www.qsan.com/data/dl_files/QSAN_Best%20Practice_XCubeNXT_How%20to%20Adjust%20Performance%20in%20Windows_2007_(en).pdf)

TOSHIBA

Toshiba Electronics Europe GmbH

Hansaallee 181
40549 Düsseldorf
Germany

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

Copyright © 2021 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.
Issued 01/2021