


HDD Lab Report

AIC JBOD "J4024-04-35X"

Rainer Kaese, Senior Manager Business Development Storage Products,
Toshiba Electronics Europe GmbH



Introduction

A glance inside any data center might reveal thousands of HDDs working side by side in glittering arrays, blinking as they smoothly transfer millions and millions of bits and bytes.

Compact, and easy to install and maintain, large scalable storage is the key building block for large-enterprise and scientific data centers as well as for cloud-scale computing. 4U 24x 3.5 in. frontloader "just a bunch of disks" (JBOD) enclosures are widely used for this, because they – unlike the alternative 4U toploader – will fit into any rack and allow easy access to the drives.

Here's where our lab report starts:

AIC has recently launched the fourth generation of its 4U 24-bay frontloader JBOD. It features a robust and reliable design, with no cables on the signal path. With a depth of just 38.55 cm, it's the most compact model in its class, leaving ample room for cabling and other installations in any rack. The dual I/O modules (IOMs) based on Broadcom's SAS35X-series expanders support another reduction in power consumption for the backplanes and connectivity.

Toshiba had the chance to evaluate a sample J4024-04 in the European HDD lab in Düsseldorf, Germany.



Picture 1: JBOD in the Toshiba lab

Dimensions & mechanical features

The J4024-04 JBOD has the standard 19 in. form factor dimensions in terms of width (438 mm) and height (4U = 174.4 mm). The depth was reduced from its predecessor, from 534 mm to just 385.5 mm, making it the most compact model in its class.

The J4024 has four fans and up to two I/O modules, plus two power supply modules, all hot-swappable and serviceable from the back. The I/O modules come with 4x 12 GB/s mini-SAS HD ports, and each of them can be controlled via in-band SES commands, but also via a dedicated management LAN port or serial connection.

The front panel carries an on/off switch and some signal lights.

The 24 robust metal-frame hot-swap HDD trays can be installed conveniently without tools, as a 3.5 in. HDD can

be snapped into place without any need for screws. However, three screw holes are available in each tray, and we recommend fixing the drives with at least one screw to keep them from vibrating. The trays can also be used for 2.5 in. drives or SSDs, but in this case four screws are mandatory.

Setup in the Toshiba lab

Model:	AIC J4024-04-35X
Firmware:	0c3b
Host OS:	Linux (Centos 7.9)
Host OS:	Windows (Windows Server 2019 Standard)
Host bus adapter (HBA):	Broadcom Avago HBA 9500-16e (Host IF: 8x PCIe-Gen4)
RAID controller:	Microchip Adaptec® SmartRAID Ultra 3254-16e/e (16x PCIe-Gen4)



Picture 2: Toshiba HDD MG09SCA18TE in JBOD tray.

Tests with enterprise capacity ("nearline") SAS drives

Model name: Toshiba MG09SCA18TE
Block size 512 byte emulated
Firmware 0104

Data-rate outer diameter: 282 MB/s

Power consumption


Idle_B:	3.36 W
Sequential write:	7.62 W
Sequential read:	8.71 W
Random write:	6.64 W
Random read:	9.47 W

Basic JBOD functions

Basic function:	OK
SAS IOM detected:	OK
Hot-plug/re-insert:	OK
Smart reading:	OK
Enclosure management:	OK, tested SES in-band setting

For precise measurement of the power consumption, we used a high-accuracy professional power analyser (R&S HMC8015).

JBOD on, no drives, SAS link to host on:	64 W
JBOD with drives, maximum startup power over 500 ms:	325 W
JBOD with raw drives at HBA:	280 W
Lambda (ratio of active and reactive power):	0.96
Noise at 1 m distance:	84 dB



Picture 3: Power measurement setup in the Toshiba HDD lab.

64 W for the JBOD without drives is an amazingly low value, and the maximum startup power and the idle JBOD with around 300 W are also excellent. A high lambda factor of 0.96 (this JBOD does not create significant reactive power) is typical of today's state-of-the-art power supply units, but is still worth mentioning here.

A noise level of 84 dB is rather high, but noise is not a factor in the data center environment. As AIC is generally known for good airflow and good cooling (which results in higher noise levels), we were curious to know how the HDD temperature figures would develop under heavy loads.

Performance measurements in the Toshiba lab

We connected the JBOD with 2x mini-SAS HD cables to the HBA or RAID controller. For different HDD configurations we measured sequential and random performance and power consumption under the different load scenarios, and looked at lowest and highest internal HDD temperatures.

The four different configurations were:

1. All drives connected with HBA, exercised in parallel. This scenario was relevant for software-defined storage systems.
2. All drives connected to a RAID controller, configured as RAID0 = all drives striped in parallel as one large physical device. This was to evaluate the highest performance, as the management of the block device was done by the hardware RAID controller. But without any redundancy it had less practical relevance.

3. All drives connected to a RAID controller, configured as RAID10 = striped mirrors. This was a practical configuration, and it turned out to be the fastest RAID configuration with proper redundancy.
4. Same configuration as 3, but formatted as Windows logic volume.

1. All drives parallel as single physical devices (multipath):

OS:	Linux (Centos 7.9)
HBA/controller:	Broadcom HBA9500-16e
HDD:	24x Toshiba MG09SCA18TE
Configuration:	Dual IOM 2x mini-SAS HD cables (3 m long) Multipath setup for disks

Workload	Power (W)	IOPS	Bandwidth (MB/s)
Sequential write 1024K	280		4540
Sequential read 1024K	290		4610
Random write 4K	240	12100	
Random read 4K		15600	
Mixed 4K/64K/256K/2M	260	9810	1070
Temperature ambient	26°C		
Temperature HDD min.	28°C		
Temperature HDD max.	30°C		

Almost 5 GB/s in bandwidth and more than 10K IOPS showed the agility of the SAS Enterprise Capacity HDD that was used. Power consumption stayed consistently below 300 W. The

Script (all drives parallel as single physical devices):

```
fio --direct=1 --bs=1m --iodepth=16 --size=32g --ioengine=libaio --group_reporting --rw=write --output=seqwrite.log -- name=/dev/dm-{2..25}

fio --direct=1 --bs=1m --iodepth=16 --size=32g --ioengine=libaio --group_reporting --rw=read --output=seqread.log -- name=/dev/dm-{2..25}

fio --direct=1 --bs=4k --iodepth=16 --size=512m --ioengine=libaio --group_reporting --rw=randwrite --output=randwrite.log --name=/dev/dm-{2..25}

fio --direct=1 --bs=4k --iodepth=16 --size=512m --ioengine=libaio --group_reporting --rw=randread --output=randread.log --name=/dev/dm-{2..25}

fio --direct=1 --bssplit=4k/20:64k/50:256k/20:2M/10 --iodepth=16 --size=8g --ioengine=libaio --group_reporting --rw=randrw --output=mixed.log --name=/dev/dm-{2..25}
```

temperature values confirmed what we were expecting from the AIC's remarkable cooling abilities: a temperature spread of just 2°C between the coldest and the hottest drive, and just +4°C maximum above ambient temperature are amazing results.

2. All drives as RAID0, Windows physical drive:

OS: Windows Server 2019
 HBA/controller: Microchip Adaptec® SmartRAID Ultra 3254-16e/e (16xPCIe4)
 HDD: 24x Toshiba MG09SCA18TE
 Configuration: Dual IOM 2x mini-SAS HD cables (3 m long)

Workload	Power (W)	IOPS	Bandwidth (MB/s)
Sequential write 64K	290		5010
Sequential read 64K	290		5630
Random write 4K	255	12800	
Random read 4K	320	4530	
Mixed 4K/64K/256K/2M	320	2520	750
Idle (raid background)	240		
Temperature ambient	26°C		
Temperature HDD min.	28°C		
Temperature HDD max.	30°C		

Bandwidth was even higher than the first configuration, as the aggregation of the single HDD performance was done by the hardware RAID controller.

Script (all drives as RAID0, Windows physical drive):

```
fio --filename=\\.\\Physicaldrive1 --direct=1 --rw=write --bs=64k --iodepth=16 --time_based --runtime=300
--group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=1 --norandommap --randrepeat=0
--output=seqwritephysical.log

fio --filename=\\.\\Physicaldrive1 --direct=1 --rw=read --bs=64k --iodepth=16 --time_based --runtime=300
--group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=1 --norandommap --randrepeat=0
--output=seqreadphysical.log

fio --filename=\\.\\Physicaldrive1 --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=\\.\\Physicaldrive1 --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=\\.\\Physicaldrive1 --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
```

3. All drives as RAID10, Windows physical drive:

OS: Windows Server 2019
 HBA/controller: Microchip Adaptec® SmartRAID Ultra 3254-16e/e (16xPCIe4)
 HDD: 24x Toshiba MG09SCA18TE
 Configuration: Dual IOM 2x mini-SAS HD cables (3 m long)

Workload	Power (W)	IOPS	Bandwidth (MB/s)
Sequential write 64K	290		3020
Sequential read 64K	270		2950
Random write 4K	260	5500	
Random read 4K	320	3940	
Mixed 4K/64K/256K/2M	280	1860	550
Idle (raid background)	240		
Temperature ambient	26°C		
Temperature HDD min.	28°C		
Temperature HDD max.	30°C		

Here we saw a drop in performance due to the RAID10 mirroring not 24 drives, but just 12 of them were working in parallel. Still the power consumption stayed less than 300 W except for the heavy-workload worst-case scenario of random reading over full capacity of 200 TB. (For random reading, the moving of the HDD's actuator arm to find a small random piece of data consumes a lot of power.)

Script (all drives as RAID10, Windows physical drive):

```
fio --filename=\\.\\Physicaldrive1 --direct=1 --rw=write --bs=64k --iodepth=16 --time_based --runtime=300
--group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=1 --norandommap --randrepeat=0
--output=seqwritephysical.log

fio --filename=\\.\\Physicaldrive1 --direct=1 --rw=read --bs=64k --iodepth=16 --time_based --runtime=300
--group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=1 --norandommap --randrepeat=0
--output=seqreadphysical.log

fio --filename=\\.\\Physicaldrive1 --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=\\.\\Physicaldrive1 --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=\\.\\Physicaldrive1 --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
```

4. All drives as RAID10, Windows logic volume:

OS:	Windows Server 2019
HBA/controller:	Microchip Adaptec® SmartRAID Ultra 3254-16e/e (16xPCIe4)
HDD:	24x Toshiba MG09SCA18TE
Configuration:	Dual IOM 2x mini-SAS HD cables (3 m long)

Workload	Power (W)	IOPS	Bandwidth (MB/s)
Sequential write 64K	280		2800
Sequential read 64K	270		2980
Random write 4K	240	5230	
Random read 4K	260	13400	
Mixed 4K/64K/256K/2M	260	2700	800
Windows copy	240		500
Idle (raid background)	240		
Temperature ambient	25°C		
Temperature HDD min.	26°C		
Temperature HDD max.	28°C		

Script (all drives as RAID10, Windows logic volume):

```
fio --filename=test --size=1T --direct=1 --rw=write --bs=64k --iodepth=16 --time_based --runtime=300
--group_reporting --name=job1 --ioengine=windowsaio --thread --numjobs=1 --norandommap --randrepeat=0
--output=seqwritelogical.log

fio --filename=test --size=1T --direct=1 --rw=read --bs=64k --iodepth=16 --time_based --runtime=300 --group_
reporting --name=job1 --ioengine=windowsaio --thread --numjobs=1 --norandommap --randrepeat=0 --out-
put=seqreadlogical.log

fio --filename=test --size=1T --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=randwritelogical.log

fio --filename=test --size=1T --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=randreadlogical.log

fio --filename=test --size=1T --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=mixedlogical.log
```

This practical use-case performance matched the previous results. Only the random read performance was higher, as the random seek was not done over the complete 200 TB of total disk capacity, but on a smaller test area of 1 TB (a more realistic scenario).

Summary

The test results for the AIC JBOD J4024-04 in the Toshiba lab were very convincing. The SAS Expander and IOM configuration enabled the full performance potential of the installed HDDs in the system. More than 5 GB/s of throughput and more than 10K IOPS with Toshiba's 18 TB Enterprise SAS HDDs was achieved. At the same time, the power consumption stayed below 300 W and the HDD temperature below +4°C over ambient temperature. It is especially this result, with the well-established cooling technology of AIC, that helps minimize HDD failure rates. This would result in longer system lifetime and would contribute to external cooling power savings by allowing operation with higher ambient temperature in the context of the data center.

With a depth of less than 40 cm, this JBOD is by far the shortest 4U enclosure. Even in short legacy racks it leaves plenty of space for cabling or even installation of auxiliary components such as switches in the rear part of the rack.

Note of thanks to our partners

This lab report is the result of dedicated and passionate collaboration. "I would like to thank all our partners for the support on this project. AIC provided the JBOD J4024-04-35X to us, Microchip supported with the raid controller Adaptec SmartRAID Ultra 3254-16e /e and finally Broadcom contributed the Host-Bus-Adapter HBA 9500-16e. Together with our Toshiba Hard Disk Drives, I was able to test AIC's JBOD under realistic data center settings in our laboratory and show its impressive results."

Rainer Kaese, Senior Manager Business Development,
Storage Products Division, Toshiba Electronics Europe GmbH

TOSHIBA

Toshiba Electronics Europe GmbH

Hansaallee 181
40549 Düsseldorf
Germany

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

Copyright © 2022 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.