

## SAS disks for a napp-it ZFS Storageserver

High performance, high capacity,  
high number of disks, hotplug

### Hardware:

SuperMicro H12SSL-C, Epyc 7302  
128 GB RAM, 2 x SAS 3008  
passthrough (mpio)

### ESXi 7.0U3

Storage VM OmniOS 151038  
8 vCores, 80G RAM

### Multipath SAS disks

1. Samsung PM 1643a - 960GB
2. Seagate Nitro 3732 -800GB
3. WD SS530-800 10 DPWD

## Preamble

High performance, high capacity, high number of disks and hotplug capability are quite common demands for storage. For high performance the first thought is NVMe. But when it comes to number of disks say dozens or hundreds or hotplug capability you may consider 2x12G Multipath SAS or 24G SAS in future.

I have tested some SAS disks in a multipath 2 x 12G configuration. Not only from specs but also from benchmarks, the best SAS disks like Seagate Nytro or WD SS530 are near to NVMe performance with a much easier handling and expandability especially in combination with 2 x 12G SAS Multipath that offers over 2000 MB/s throughput. This is 4 times the performance of 6G Sata.

## Specs

	Samsung 1643	Seagate 3732	WD SS530	compare: Intel Optane 900
sequ read	2100 MB/s	1100 MB/s	2150 MB/s	2600 MB/s
sequ write	2000 MB/s	1000 MB/s	2120 MB/s	2200 MB/s
random read	400k iops	180k iops	440k iops	575k iops
random write	70k iops	200k iops	320k iops	550k iops
DWPD	1	10	10	10
MTBF in h	2 Mio	2,5 Mio	2,5 Mio	1,6 Mio
total write	-	14,6 PB	15,22 PB	-

From results:

- Dualpath SAS improves performance up to 100% compared to singlepath
- The best SAS SSD (WD SS530) is near to the NVMe
- The Seagate Nytro is quite as good as the WD
- Both are much faster than the cheaper Samsung 1643

SAS multipath is worth the effort.

# 1. Samsung PM 1643a (960GB) Singlepath and Multipath

Benchmark: Write: filebench\_sequential, Read: filebench, date: 10.07.2021

pool: tank

NAME	STATE	READ	WRITE	CKSUM
tank	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
c3t5002538B007B77B2d0	ONLINE	0	0	0
c4t5002538B007B77E2d0	ONLINE	0	0	0

  

host	napp-it-038
pool	tank (reccsize=128K, ssb=-, compr=off, readcache=all)
slog	-
encryption	-
remark	

  

Fb3	sync=always	sync=disabled
Fb4 fivestreamwrite.f	sync=always	sync=disabled
	2240 ops	5605 ops
	447.275 ops/s	1051.827 ops/s
	13539us cpu/op	9437us cpu/op
	11.0ms latency	4.2ms latency
	446.3 MB/s	1050.9 MB/s

---

pri/sec cache=all	randomread.f	randomrw.f	fivestreamrea
	179.9 MB/s	94.2 MB/s	1.6 GB/s

---

If something happens, reboot and select a former bootenvironment.  
Or reinstall napp-it via wget command. This preserves all setting.

same pool with mpio (2 HBA)

Fb3	sync=always	sync=disabled
Fb4 fivestreamwrite.f	sync=always	sync=disabled
	4640 ops	5673 ops
	927.913 ops/s	1134.497 ops/s
	24536us cpu/op	25328us cpu/op
	5.4ms latency	4.4ms latency
	926.9 MB/s	1133.5 MB/s

---

pri/sec cache=all	randomread.f	randomrw.f	fivestreamrea
	281.2 MB/s	489.2 MB/s	16.0 GB/s

---

## 2. Seagate Nytro 3732 – 800 mpio

Benchmark: Write: filebench\_sequential, Read: filebench, date: 10.07.2021

pool: tank

NAME	STATE	READ	WRITE	CKSUM
tank	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
c0t5000C500A19F6007d0	ONLINE	0	0	0
c0t5000C500A19F6027d0	ONLINE	0	0	0

host napp-it-038  
 pool tank (reclsize=128K, ssb=-, compr=off, readcache=all)  
 slog -  
 encryption -  
 remark mpio

Fb3 sync=always sync=disabled

Fb4 fivestreamwrite.f sync=always sync=disabled  
 5536 ops 8440 ops  
 1107.100 ops/s 1687.618 ops/s  
 24028us cpu/op 33203us cpu/op  
 4.5ms latency 2.9ms latency  
 1106.1 MB/s 1686.6 MB/s

---

pri/sec cache=all	randomread.f 209.2 MB/s	randomrw.f 464.3 MB/s	fivestreamrea 13.6 GB/s
-------------------	----------------------------	--------------------------	----------------------------

---

### 3. WD SS530 – 800 10DWPD mpio

Benchmark: Write: filebench\_sequential, Read: filebench, date: 10.07.2021

pool: tank

NAME	STATE	READ	WRITE	CKSUM
tank	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
c0t5000CCA0AB743E88d0	ONLINE	0	0	0
c0t5000CCA0AB753CE0d0	ONLINE	0	0	0

host napp-it-038  
 pool tank (reclsize=128K, ssb=-, compr=off, readcache=all)  
 slog -  
 encryption -  
 remark mpio

Fb3 sync=always sync=disabled

Fb4 fivestreamwrite.f sync=always sync=disabled  
 6316 ops 8509 ops  
 1263.157 ops/s 1701.694 ops/s  
 24045us cpu/op 32470us cpu/op  
 3.9ms latency 2.9ms latency  
 1262.2 MB/s 1700.7 MB/s

---

pri/sec cache=all	randomread.f 200.2 MB/s	randomrw.f 472.0 MB/s	fivestreamrea 13.4 GB/s
-------------------	----------------------------	--------------------------	----------------------------

---