MacOS media & entertainment workloads with SMB

IBM Storage Scale Days 2024

March 5-7, 2024 | Stuttgart Marriott Hotel Sindelfingen

Jukka Muhonen

TUXERA Make it work.

Storage and networking technologies from micro-controllers to global public clouds

TUXERA

Fusion File Share by Tuxera

World's most advanced and scalable enterprise SMB server on Linux

Key advantages of Fusion File Share

FUSION FILE SHARE
By Tuxera

Our high-performance, highly-scalable, drop-in replacement for Samba.

- Highly threaded architecture
- High-performance – 2x to 60x faster than SAMBA
- 100% to 500% better scalability than SAMBA
- Fault tolerant with Transparent failover and Continuous Availability
- Extensive SMB-protocol support – 3.1.1
- Scale-out (active-active)
- RDMA (SMB-Direct), Multichannel, and Compression
- Low CPU and memory usage
- Low latency
- Native GPFS support

Key advantages of Fusion File Share

FUSION FILE SHARE By Tuxera

Highly threaded architecture with adjustable settings for different workloads

Each client connection is a thread, not a process:

- Data transport threads
- Meta data transport threads
- VFS data threads
- VFS meta data threads
- Minimized CPU & memory usage

Adjustable quality of service by tuning:

- Concurrent open files
- Concurrent client connections
- Concurrent open files per user-session
- Concurrent VFS threads per share

MacOS and IBM Scale specific SMB functionalities

FUSION FILE SHARE By Tuxera

MacOS & IBM Scale functionalities

IBM Scale

- Native support for IBM Scale

MacOS

- Apple extension support (AAPL)
  - Samba Fruit type of support
- Compound messaging support for performance optimization

MacOS SMB performance

TUXERA

MacOS performance test environment

FUSION FILE SHARE By Tuxera

Hardware and software environment

MacOS client

- *Processor*: Apple M2 Ultra with 24-core CPU, 60-core GPU, 32-core Neural Engine
- *Memory*: 192GB unified memory
- *Storage*: 1TB SSD storage
- *NICs*: 2x ATTO dual port 100GbE, X16 PCIE 4, LP, INTEGRATED QSFP28

Storage + SMB

- *NICs*: Mellanox 100GbE
- *VM*: Ubuntu Linux + Fusion SMB + Single IBM Scale instance (GPFS)

Network

- Mellanox 100GbE switch

MacOS performance tests

FUSION FILE SHARE By Tuxera

Metadata test

- 20000 files in a SMB share
- Use Finder to list folder content

Test results

Server	Time to completion
Samba v4.19 (with AAPL)	31.39 s
Fusion (with AAPL)	9.67 s

MacOS performance tests

FUSION FILE SHARE By Tuxera

M&E workload test

- Tframetest –w 49856 –n 2000 –t1 /SMBShare/

Test results

Server	Multichannel	Write	Read
Samba v4.19 (with AAPL)	Single interface	1,463 GB/s	1,659 GB/s
Samba v4.19 (with AAPL)	Two interfaces	1,601 GB/s	2,348 GB/s
Fusion (with AAPL)	Single interface	1,606 GB/s (+10%)	2,679 GB/s (+61%)
Fusion (with AAPL)	Two interfaces	2,013 GB/s (+25%)	4,009GB/s (+70%)

POC Environment

Test Cases

1- Single Client SMB2 Performance over TCP (IPoIB)

2- Single Client SMB3 Multi-Channel over TCP (IPoIB)

3- Single Client SMB3 Direct using RDMA (EDR IB)

4- Multi-Client SMB3 Direct using RDMA (EDR IB)

5- Single GPFS Client using RDMA (EDR IB)

POC/Benchmark Results – 4GB Filesize

FIO Write Test:

fio.exe --name=fiotest --directory=\\ESS32KSMB\ess32kshare\ --size=4G --rw=write --bs=4M --numjobs=24 --ioengine=windowsaio --iodepth=16 --group_reporting --runtime=60 --ramp_time=30 --direct=1

Test	Numjobs	xfersize	Avg MiB/s Write	Avg IOPs Write
Single Client SMB2 TCP	24	4M	2615	616
Single Client SMB3 Multi-Channel TCP	24	4M	9840	2519
Single Client SMB3 Direct RDMA	24	4M	9998	2499
Multi-Client SMB3 Direct RDMA	24	4M	TBD	TBD
Single Scale Client RDMA	24	4M	3039	685

FIO Read Test:

fio.exe --name=fiotest --directory=\\ESS32KSMB\ess32kshare\ --size=4G --rw=read --bs=4M --numjobs=24 --ioengine=windowsaio --iodepth=16 --group_reporting --runtime=60 --ramp_time=30 --direct=1

Test	Numjobs	xfersize	Avg MiB/s Read	Avg IOPs Read
Single Client SMB2 TCP	24	4M	3390	847
Single Client SMB3 Multi-Channel TCP	24	4M	10600	2718
Single Client SMB3 Direct RDMA	24	4M	11000	2816
Multi-Client SMB3 Direct RDMA	24	4M	19598	4898
Single Scale Client	24	4M	4972	1242

Name of Speaker

Mail: jukka.muhonen@tuxera.com

Mobil: +358405602722

Telefon:

LinkedIn: https://www.linkedin.com/in/jukka-muhonen-3228b535/

Thank you for using

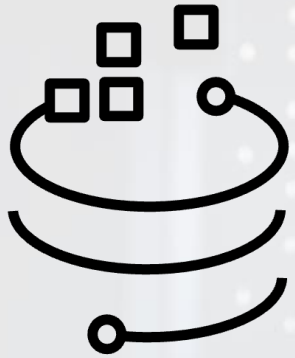

Storage Scale

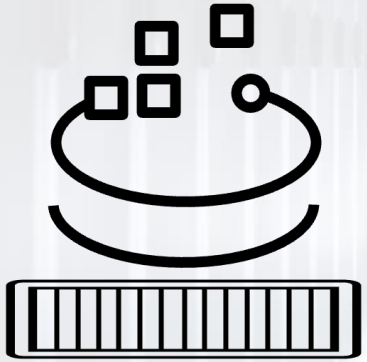

Storage Scale System