

TUXERA

TUXERA

- 40 years in the business
- Finland HQ, global offices
- Mission critical embedded file systems, flash controllers, TCP stack, crypto libs
- Enterprise SMB since 2016
- Working with IBM for 4 years



IBM

TUXERA FUSION

TUXERA

Enterprise-grade data access layer for modern storage platforms

FUSION
SMB




FUSION
NFS

Fusion TUXERA

- Distro-agnostic Linux file servers
- Enterprise, commercial
- Ultra high performance & scale
- Samba/Ganesha/NFSD replacement
- Use OS-native clients, not add-ons
- **Maximize storage, productivity, value**

TUXERA FUSION



RDMA &
multichannel
performance




Scale-out
reliable
cost saving



Architecture
without
compromise



Rolling
upgrade zero-
downtime



Multi-
protocol
flexibility



RDMA & multichannel performance



Perf TUXERA

- SMB Direct + SMB Multichannel / NFS RPC over RDMA and client trunking
- Direct memory access
- CPU bypass & offload
- Lowest latency
- Up to 800Gb per-port networking
- Table stakes for ML, HPC

Single client FIO 1MB IO CX6 200GbE on latest Samba & Fusion (Direct IO/no buffer)



24.4GB/s at 8MB IO on 200Gb CX6 – testing 400Gb CX7s soon

Fusion NFS raw 200Gb performance vs Ganesha & NFSD

Demo

Ganesha

13.6 GB/s

```
demo@tuxus-02: /mnt/nfs/fin $ sudo mount -t nfs -o vers=4.1,rdaa,max_connect=16,nconnect=16 10.2.0.1:/DemoShare /mnt/nfs/
mount.nfs: Protocol not supported for 10.2.0.1:/DemoShare on /mnt/nfs
demo@tuxus-02:~$ sudo mount -t nfs -o vers=4.1,max_connect=16,nconnect=16 10.2.0.1:/DemoShare /mnt/nfs/
demo@tuxus-02:~$ cd /mnt/nfs/fin
demo@tuxus-02:/mnt/nfs/fin$ fio --size=10G --rw=read --direct=1 --bs=1M --ioengine=libaio --runtime=15 --numjobs=16 --time_based --group_reporting --name=seq_read --iodepth=16
```

NFSD

22.5 GB/s

```
demo@tuxus-02: /mnt/nfs/fin $ sudo mount -t nfs -o vers=4.1,rdaa,max_connect=16,nconnect=16 10.2.0.1:/mnt/raid0/nfs/share /mnt/nfs/
demo@tuxus-02:~$ cd /mnt/nfs/fin
demo@tuxus-02:/mnt/nfs/fin$ fio --size=10G --rw=read --direct=1 --bs=1M --ioengine=libaio --runtime=15 --numjobs=16 --time_based --group_reporting --name=seq_read --iodepth=16
```

Fusion NFS

22.7 GB/s

```
demo@tuxus-02: /mnt/nfs/fin $ sudo mount -t nfs -o vers=4.1,rdaa,max_connect=16,nconnect=16 10.2.0.1:/DemoShare /mnt/nfs/
demo@tuxus-02:~$ cd /mnt/nfs/fin
demo@tuxus-02:/mnt/nfs/fin$ fio --size=10G --rw=read --direct=1 --bs=1M --ioengine=libaio --runtime=15 --numjobs=16 --time_based --group_reporting --name=seq_read --iodepth=16
```




Recycle Bin

Search

Search

Search





*Samba cannot support AU's above 90% usage

*Ganesha does not support RDMA NFS 4.1

Supermicro SYS-121H-TNR
2x Intel Xeon® 2.80 GHz 6526Y CPUs / 256GB RAM
NVIDIA ConnectX-6 200GbE RoCEv2
8 Gen4 Samsung NVME software RAID-0 with XFS
Ubuntu 24.04 server

MLPerf llama3-8b checkpoint test: Fusion SMB vs NFSD vs Ganesha vs Samba




Benchmark	Checkpointing
Model	llama3-8b
Processes Sweep	8 16
Checkpoint size	107.6 GB
Checkpoints write	10
Checkpoints read	10

*Ganesha & Samba results are TCP, they do not support RDMA transport

	Server	Client
Hostname	tuxus-01	tuxus-02
CPU	INTEL(R) XEON(R) GOLD 6526Y 16 physical cores x 2 socket(s), 64 logical proc	INTEL(R) XEON(R) GOLD 6526Y 16 physical cores x 1 socket(s), 32 logical proc
OS	Ubuntu 24.04.4 LTS, 6.8.0-107- generic	Ubuntu 24.04.4 LTS, 6.17.13- smbdirect
Filesystem	XFS, /mnt/raid0, 8x MZQL21T9HCJR-00A07	SMB 3.1.1/NFS 4.1
NIC	Nvidia MT2910 [ConnectX-7]	Nvidia MT2910 [ConnectX-7]

FIO comparison test: Fusion SMB vs NFSD vs Ganesha vs Samba



*Ganesha & Samba results are TCP, they do not support RDMA transport

Max storage write speed on 8xRAID-0 NVME XFS is 22GB/s


Scale-out reliable cost saving



Scale-out **TUXERA**


- Clustering with all nodes active
- Aggregate performance
- Transparent failover and continuous availability
- Relies on distributed storage

Real customer configuration example – Samba (CES) & IBM ESS 3500



Single connected 200 GbE Scale 200B/sec


Real customer configuration example – Fusion SMB & IBM ESS 3500



Single connected 200 GbE Scale 200B/sec

Fusion SMB scale-out cluster transparent failover

Demo



PS C:\temp> net use
New connections will not be remembered.

Status	Local	Remote	Network
OK	Z:	\\fusion-srv.contoso.com\sh1	Microsoft Windows Network

The command completed successfully.

PS C:\temp>

sh1 (\\fusion-srv.contoso.com)

Search sh1 (\\fusion-srv.contoso.com)

New | Sort | View | Details

Name	Date modified	Type	Size
This folder is empty.			

OneDrive

Desktop | Downloads | Documents | Pictures | Music | Videos


This PC

Local Disk (C:) | DVD Drive (D:) | sh1 (\\fusion-srv.contoso.com) | DVD Drive (D:) | Network

0 items

Fusion NFS scale-out cluster transparent failover

Demo




Architecture without compromise




Architecture

TUXERA

- User-mode safety & security
- Multi-threaded, low overhead
- Physical, VM, container
- Protocol compliant SMB 2.0-3.1.1 & NFS 4.0/4.1 feature set




Open
Source




Tuxera
Fusion

16 Win11 clients running FIO 4x4K RW (64 files)



```
tsmb@gigantnode-6: ~  
0[24] 3[19] 6[23] 9[36] 12[4.] 15[5.] 18[3.] 21[26] 23[14] 26[23] 29[38] 32[4.] 35[0.] 38[5.]  
1[35] 4[17] 7[35] 10[4.] 13[5.] 16[4.] 19[4.] 21[37] 24[32] 27[18] 30[3.] 33[3.] 36[5.] 39[5.]  
2[19] 5[35] 8[22] 11[5.] 14[5.] 17[4.] 22[21] 25[40] 28[27] 31[5.] 34[6.] 37[1.]  
Mem[|||||] 14.0G/188G Tasks: 80, 86 thr, 559 kthr; 6 running  
Swp[ ] 0K/2.0G Load average: 7.07 6.35 5.26  
Uptime: 7 days, 03:02:29
```

Main	I/O	PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
		671454	root	20	0	113M	18372	15036	S	39.7	0.0	9:50.90	smbd: client [192.168.75.226]
		671435	root	20	0	113M	17656	14640	S	39.0	0.0	9:59.18	smbd: client [192.168.75.196]
		671481	root	20	0	113M	17404	14068	S	39.0	0.0	10:03.34	smbd: client [192.168.75.141]
		671393	root	20	0	113M	17736	14400	R	38.4	0.0	6:40.76	smbd: client [192.168.75.203]
		671399	root	20	0	113M	17660	14324	R	38.4	0.0	9:31.19	smbd: client [192.168.75.216]
		671462	root	20	0	113M	17652	14636	R	38.4	0.0	10:00.13	smbd: client [192.168.75.183]
		671468	root	20	0	121M	17724	14388	S	38.4	0.0	10:05.89	smbd: client [192.168.75.186]
		671475	root	20	0	113M	18036	14700	R	38.4	0.0	10:06.76	smbd: client [192.168.75.246]
		671414	root	20	0	113M	18300	14964	S	37.7	0.0	9:24.50	smbd: client [192.168.75.177]
		671488	root	20	0	113M	17724	14388	R	37.7	0.0	9:53.27	smbd: client [192.168.75.202]
		671427	root	20	0	113M	17932	14596	S	37.0	0.0	9:27.63	smbd: client [192.168.75.194]
		671448	root	20	0	113M	17060	14044	S	37.0	0.0	10:01.14	smbd: client [192.168.75.224]
		671420	root	20	0	113M	17312	13976	S	35.7	0.0	9:01.29	smbd: client [192.168.75.184]
		671494	root	20	0	113M	17336	14320	R	35.7	0.0	5:50.93	smbd: client [192.168.75.237]
		671406	root	20	0	113M	17628	14612	S	35.0	0.0	9:57.19	smbd: client [192.168.75.229]
		671442	root	20	0	113M	17308	14292	R	35.0	0.0	9:43.20	smbd: client [192.168.75.159]
		672412	tsmb	20	0	6748	5120	3520	R	2.0	0.0	0:00.52	htop --filter smbd
		670102	root	20	0	90400	24960	22080	S	0.0	0.0	0:00.83	/usr/sbin/smbd --foreground --no-process-group
		670106	root	20	0	87844	6504	3520	S	0.0	0.0	0:00.03	smbd: notifyd .
		670107	root	20	0	87860	6184	3200	S	0.0	0.0	0:00.03	smbd: cleanupd



```
tsmb@gigantnode-6: ~  
0[14] 3[7.] 6[34] 9[13] 12[7.] 15[7.] 18[7.] 20[10] 23[9.] 26[9.] 29[8.] 32[7.] 35[8.] 38[8.]  
1[8.] 4[9.] 7[16] 10[7.] 13[53] 16[6.] 19[7.] 21[12] 24[14] 27[15] 30[8.] 33[5.] 36[8.] 39[7.]  
2[8.] 5[9.] 8[9.] 11[7.] 14[9.] 17[9.] 22[10] 25[17] 28[8.] 31[7.] 34[9.] 37[6.]  
Mem[|||||] 15.3G/188G Tasks: 59, 6264 thr, 515 kthr; 2 running  
Swp[ ] 0K/2.0G Load average: 5.41 3.07 1.31  
Uptime: 6 days, 05:05:50
```

Main	I/O	PID	USER	PRI	NI	VIRT	RES	SHR	S	CPU%	MEM%	TIME+	Command
		378991	root	20	0	68.3G	3240M	8636	R	56.2	1.7	3:05.67	tsmb-server -B
		392292	tsmb	20	0	13260	11520	3520	R	23.8	0.0	0:52.52	htop --filter tsmb-server
		379568	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:06.59	tsmb-server -B
		379579	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:06.09	tsmb-server -B
		379587	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:05.44	tsmb-server -B
		379620	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:03.27	tsmb-server -B
		379631	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:03.15	tsmb-server -B
		379634	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:03.06	tsmb-server -B
		379639	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:03.16	tsmb-server -B
		379644	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:02.97	tsmb-server -B
		379648	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:02.70	tsmb-server -B
		379650	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:02.71	tsmb-server -B
		379658	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:02.31	tsmb-server -B
		379664	root	20	0	68.3G	3240M	8636	R	2.7	1.7	0:02.16	tsmb-server -B
		379670	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:01.93	tsmb-server -B
		379671	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:01.94	tsmb-server -B
		379683	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:00.45	tsmb-server -B
		380042	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:08.86	tsmb-server -B
		380055	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:06.91	tsmb-server -B
		380075	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:04.41	tsmb-server -B
		380081	root	20	0	68.3G	3240M	8636	S	2.7	1.7	0:04.17	tsmb-server -B

CPU & memory using FIO from 2,500 SMB simultaneous active clients

Test	CPU max	CPU avg	Mem max	Mem avg
1M IO block, read	19.14%	8.72%	3.5G	2.4G
1M IO block, write	30.20%	13.43%	3.2G	2.1G
4K IO block, read	58.41%	49.99%	175M	174M
4K IO block, write	59.04%	49.36%	228M	219M
8K IO block, read	57.81%	51.22%	177M	175M
8K IO block, write	61.55%	50.24%	229M	221M

- AWS c5.2xlarge VMs (8 vCPU Intel Xeon 3.00GHz, 16GB mem)
- 1 VM for server, 8 VM clients, Debian 12 ("Bookworm")
- TCP, not RDMA (would be lower CPU)
- Linux CIFS mounts, each VM is many SMB connections
- FIO is the test tool
- 2,500 simultaneous active SMB client connections (unrealistically busy but more interesting)

Ganesha: 1,000 exports (supported)

Ganesha: 4000 clients

Fusion NFS: 10,000 exports (tested so far)

Fusion NFS: 10,000 clients (tested so far)

Why Fusion SMB instead of Samba?

Fusion SMB designed for modern enterprise systems

	Samba	Fusion SMB
Memory and CPU usage	High	Low
Concurrent active workload connections	100's	10000's
Small file creates per second	10s	1000's
Small file writes per second	100's	10000's
Max read throughput per 100GbE port (Windows/Linux)	2.8GB/s	11.2GB/s
Max read throughput per 100GbE port (macOS)	2.8GB/s	8.6GB/s
SMB cluster	4-16 nodes*	32 nodes
Cluster rolling upgrade with zero downtime	○	●
Direct IO data protection & integrity	○	●
SMB Direct (SMB over RDMA)	○	●
SMB over QUIC	○	●
SMB compression	○	●
SMB continuous availability	○	●
SMB persistent handles	○	●
SMB transparent failover	○	●
macOS rehydration blocker	○	●
Custom VFS development	○	●
Custom feature development	○	●
Microsoft SMB patent protection	○	●
Microsoft protocol compatibility	Trails years	Fast match
Support	Depends	Commercial
License	GPLv3	Commercial

* IBM CES supports up to 16, Samba support is 4

Fusion SMB capabilities

Anything a Microsoft Windows file server can do - running on Linux

Capability
SMB 1, 2.02, 2.1, 3.0.0, 3.1.0, 3.1.1
SMB Multichannel
SMB over QUIC
SMB Compression
Intel/AMD/ARM chipset
Signing (HMAC-SHA256, AES-128-GMAC, AES-128-CMAC)
Encryption (AES-128-CCM, AES-128-GCM, AES-256-CCM, AES-256-GCM)
Access-based enumeration (ABE)
Active Directory / OpenLDAP / Apple Open Directory integration
Kerberos
NTLM
Windows ACLs
Active-passive clustering
Volume Shadow Copy Service (VSS) support
DFSN link target support
ext4, XFS, ZFS
GPFS, Lustre, CephFS, GlusterFS, WekaFS, GFS2, OCFS2
Log/audit support
Quota support
Performance counters
Windows MMC management support
Windows-compatible event log auditing

SMB development timeline

Microsoft SMB

2012	2013-2015	2016-2018	2019-2020	2021-2023	2024-2026
SMB 3.0.0: SMB multichannel SMB Direct (RDMA) AES-CCM encryption AES-CMAC signing Scale-out file server Continuous availability	SMB 3.0.2 Tuning, fixes SMB1 deprecation	SMB 3.1.1 Pre-auth integrity AES-128-GCM encrypt	SMB 1 first removed	SMB compression SMB over QUIC (Azure) AES GMAC signing AES-256 encryption	SMB over QUIC (everywhere) SMB over QUIC client access control SMB compression with LZ4

Fusion SMB

2012	2013-2015	2016-2018	2019-2020	2021-2023	2024-2026
	In development with Microsoft	SMB 3.x SMB multichannel SMB Direct (RDMA)** AES-128 encryption AES-128 signing Pre-auth integrity Failover cluster support Microsoft partner license** Container support*	32-node cluster support*/** Scale-out file server Continuous availability	SMB compression** AES GMAC signing AES-256 encryption	SMB over QUIC Local KDC (<i>coming</i>) IAKerb (<i>coming</i>)

Samba

2012	2013-2015	2016-2018	2019-2020	2021-2023	2024-2026
	SMB 3.0.0, 3.0.2 AES encryption AES signing 4-node cluster support Container support GPLv3 license	SMB 3.1.1 Pre-auth integrity AES-128-GCM encrypt Continuous availability (limited)		SMB multichannel AES-256 encryption	

Rolling upgrade zero-downtime




Rolling Upgrade

TUXERA


- Cluster updating simplicity
- Zero downtime for new features, fixes, users, apps, & services
- Easiest change control, lowest OPEX

Cluster rolling upgrade of Fusion




- Supports $n-2$
- No node eviction needed

Cluster rolling upgrade of Fusion



- Stop Fusion service on node
- Clients move to new node
- Update Fusion service on node
- Fusion remains in “old” mode

Cluster rolling upgrade of Fusion



- Upon last Fusion upgrade, all nodes enter new functionality mode automatically
- No need to enable/configure
- Clients get new features dynamically or on next connection

Multi-protocol flexibility





Fusion NFS & SMB Multi- protocol

TUXERA

- Lock and handle awareness
- Windows ACL & NFS ACL central management
- Share data safely & transparently

Multiprotocol Fusion vs Samba/Ganesha/NFSD



FUSION SMB



RDMA



macOS 10.5GB/sec*
Win 11.5 GB/sec*



SMB
compression



SMB
over QUIC



Scale-out
high availability
failover



75% more
efficient
than Samba

*Single 100GbE port



Multi protocol support

SMB & NFS
ACL's and locks

FUSION NFS



RDMA



22.7 GB/sec*
2x Ganesha throughput
Faster than NFSD



User
mode



More scalable
than Ganesha



Scale-out
high availability
failover



More flexible
than NFSD

*Single 200GbE port



Clusters



Containers



Cloud



Virtual
Machines



Single
server




Multi-threaded



Modern
architecture

with
rolling
upgrades

Fusion vs open source throughput in GB/Sec





- Recommended by IBM
- Licensed per instance
- No client license
- No capacity license
- No client software needed
- Compliments native GPFS client
- SMB and NFS single share
- SMB Compression and SMB over QUIC
- Rolling upgrades with persistent handles
- RDMA, SMB Multichannel, NFS Multipathing
- Replaces Samba and Ganesha
- 32 Node Cluster
- Direct IO
- macOS rehydration blocker

TUXERA
FUSION
fusion@tuxera.com



Evaluation

Q & A

