Samba 2025: Enterprise-Ready, Cloud-Optimized

Ralph Böhme, Samba Team, SerNet

2025-09-17

Outline

SNIA DEVELOPER CONFERENCE
SDC 25
By Developers FOR Developers

Samba in 2025

Enterprise-Ready, Cloud-Optimized

Where are we in 2025/2026?

- Features
- Performance
- Scalability

SNIA DEVELOPER CONFERENCE
SDC 25
By Developers FOR Developers

Samba in 2025

- Highly scalable: clusters up to 32 nodes
- Thousands of connections per node
- Cluster-wide coherent Windows locking semantics
  - not everyone does this, check with your vendor
- Mature and battle tested implementation
- Participating at CIFS and SMB plugfests since the dawn of time

No vendor lock in

- Mature Open Source project with a large community
- Used in many commercial products
- Enterprise support: SerNet, RedHat, SuSE, ...

Samba Cluster

SNIA DEVELOPER CONFERENCE SDC 25 By Developers FOR Developers

Resilient

- One process per client, helper threads for IO
- If a process crashes, only a single client is affected

Adaptable and customizable

- Pluggable VFS architecture
- Allows close integration with Open Source filesystems like CephFS
- IBM is putting their eggs in this basket for the future of scaleout enterprise SMB

Sovereign Tech Fund

SerNet

Souvereign Tech Fund

- A German government initiative that supports the development and maintenance of open digital infrastructure
- Some example funded projects: Gnome, FreeBSD, Log4j, PyPi, ...
- STF invests 680k € into Samba via SerNet
- Three SerNet Samba developers hacking for 18 months

SNIA DEVELOPER CONFERENCE SDC 25 By Developers FOR Developers

Features

SMB3 Enterprise Features

All SMB3 Enterprise features shipping or shipping in 2026:

- SMB3 UNIX Extensions
- SMB3 Directory Leases
- SMB over QUIC
- SMB3 Persistent Handles / Transparent Failover
- SMB-Direct
- Metrics: Prometheus Exporter

SNIA DEVELOPER CONFERENCE SDC 25 By Developers FOR Developers

SMB3 UNIX Extensions in Samba

- Now enabled by default in Samba 4.23
- Ready to compete with NFS for the Linux datacenter
  - SMB3 includes encryption and signing as core features
  - Strong authentication model (Kerberos, LDAP/AD integration).
  - Secure out-of-the-box
  - Advanced features like Persistent Handles, Compression, ...
  - Builtin scale-out file serving with `ctdb`

SMB3 UNIX Extensions Specification

- Three documents: `POSIX-SMB2`, `POSIX-FSA` and `POSIX-FSCC`
- `POSIX-FSCC` contains on-the-wire protocol
- `POSIX-SMB2` contains server and client behavior
- `POSIX-FSA` contains filesystem semantics where they differ from NTFS
- All three WIP, but work is funded to finish the specification in 2026

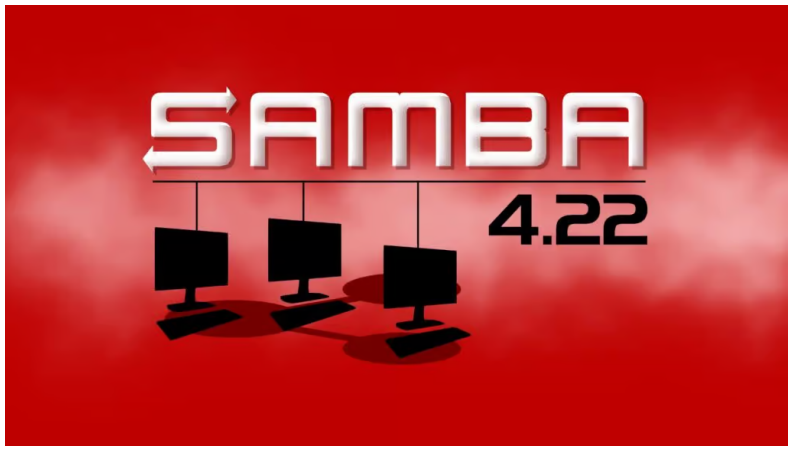

SMB3 Directory Leases: Directory Caching done Right

- Reliable and efficient directory cache
- Shipping with Samba 4.22:
  - `smb3 directory leases = yes | no | auto`
  - enabled by default, but currently still disabled on clusters
  - reason: renaming a directory now correctly triggers scanning for H-leases which is an expensive operation on cluster
  - waiting for real world experience from users that use it on clusters. . .
- Significant frame reduction for metadata heavy workloads
- Linux kernel client still doesn't fully leverage it

SNIA DEVELOPER CONFERENCE SDC 25 By Developers FOR Developers

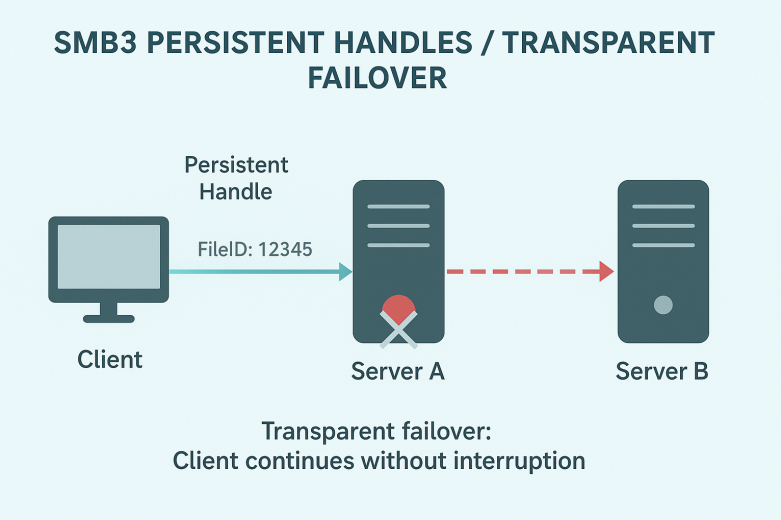

SMB3 Persistent Handles

- Mount your Hyper-V .vhdx over SMB3
- VM-level continuous availability
- Will be shipping in 2026 with Samba 4.24
- Currently discussing interesting behaviour on Windows Server
  - when combining Persistent Handles with oplocks and byterange locks

SMB over QUIC

- Samba 4.23 ships it, waiting for kernels to ship `quic.ko`
- `quic.ko`: a new Linux kernel module implementing QUIC in the kernel
- New socket type in the UNIX socket API: `IPPROTO_QUIC`
- Samba continues to use the socket API with minimal changes
- New transport option: `server smb transports = +quic`
  - currently disabled by default in the server
- See metze's talk for details
  (you need to travel back in time if you missed it)

SNIA DEVELOPER CONFERENCE SDC 25 By Developers FOR Developers

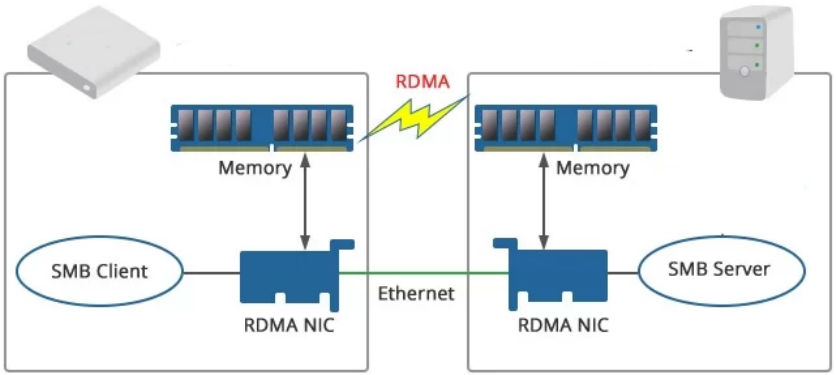

SMB Direct

- Scheduled for 2026
- Consolidate SMB direct support in the Linux kernel
  (both `cifs.ko` and `ksmbd` ship their own code)
- Expose it to userspace via another UNIX socket type `IPPROTO_SMBDIRECT`
- Finally use it in Samba with minimal changes
- See metze's talk for details
  (you need to travel back in time if you missed it)

SNIA DEVELOPER CONFERENCE SDC 25 By Developers FOR Developers

Metrics: Prometheus Exporter

`smb_prometheus_endpoint`

- Exports Samba server metrics in Prometheus-compatible format
- Seamless integration with Prometheus & Grafana monitoring setups
- New service `smb_prometheus_endpoint`
- See the manpage for usage and configuration

SNIA DEVELOPER CONFERENCE SDC 25
By Developers FOR Developers

Performance

Performance with io_uring

Huge performance boost coming in 2026:

- It's not a matter of architecture: processes vs threads
- It's a matter of using modern Linux kernel APIs to avoid copying data

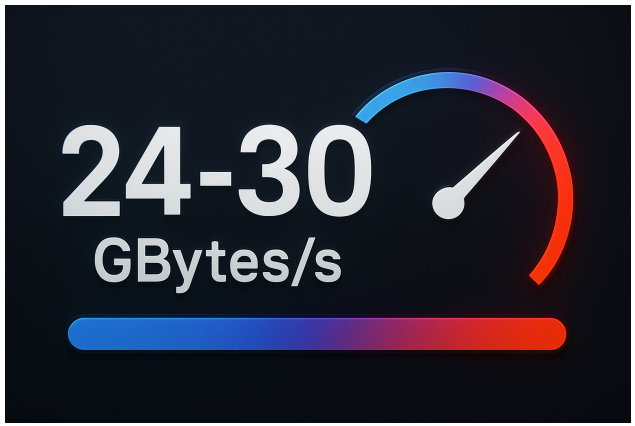

4 loopback connections, 320% CPU (smbd: 20%, kernel io-wrk threads: 300%)

SNIA DEVELOPER CONFERENCE SDC 25 By Developers FOR Developers

Performance with `io_uring`, cont.

Linux 5.1 introduced a new scalable AIO infrastructure: `io_uring`

- Designed to avoid syscalls as much as possible
- Kernel and userspace share mmap'ed rings:
  - Submission queue (SQ) ring buffer
  - Completion queue (CQ) ring buffer
- Delegates work to kernel threads
  - `IORING_OP_SENDMSG`
  - `IORING_OP_RECVMSG`
- Optimisations to avoid copying data in the IO path:
  - `IORING_OP_SENDMSG_ZC`
  - `IORING_OP_SPLICE`

Scalability

Scalability

Cloud scale and Resiliency: 32 nodes and beyond!

- `ctdb` database implementations restricts scalability
- Samba has a pluggable database API: `dbwrap`
- Allows replacing `ctdb` with arbitrary modern scaleout Key-Value stores
- Candidates: Ceph Rados, FoundationDB, ...
  - or write our own based on RAFT
- Perfect building blocks for a Cloud-SMB server
- See my presentation from 2024@SambaXP

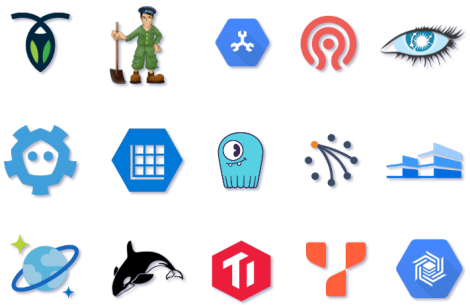

Q&A

SerNet

Thank you!
Questions?

Ralph Böhme
slow@samba.org
rb@sernet.de

SNIA DEVELOPER CONFERENCE
SDC 25
By Developers FOR Developers