

STORAGE DEVELOPER CONFERENCE



Fremont, CA  
September 12-15, 2022

*BY Developers FOR Developers*

A **SNIA** Event

# smbd, quo vadis?

Ralph Böhme, Samba Team, SerNet

2022-09-14

Recap

Samba 4.17

Coming soon

Prototypes

Links & Q&A

STORAGE DEVELOPER CONFERENCE



*BY Developers FOR Developers*

## Recap

---

## 4.15, released September 2021

- VFS hardening against symlinks races
  - OEMs watch out: **lots** of changes to the VFS!
  - Watch out 2: **performance regression!**  
40% regression for meta-data heavy workloads  
The good news: Samba 4.17 addresses this.
- SMB3 Multichannel no longer experimental
  - Enterprise SMB3 Features slowly emerging, more to come. . .
- Trusted domains are no longer scanned in winbindd
  - Fixes longstanding subtle issues with Samba in complex AD scenarios
- Support for Offline Domain Join
  - Especially helpful for provisioning containers

## 4.16, released March 2022

- RPC server refactoring:
  - New RPC server process `samba-dcerpcd`
  - All RPC services now run in their own process(es) spawned by `samba-dcerpcd`
- Updated internal Heimdal copy, adds support for Kerberos FAST
- Certificate Auto Enrollment

STORAGE DEVELOPER CONFERENCE



*BY Developers FOR Developers*

## Samba 4.17

---

## File-server performance improvements:

- mainly undoing the damage of the symlinks security hardening
- completely rewritten symlink safe pathname processing
- improved internal database locking

STORAGE DEVELOPER CONFERENCE



*BY Developers FOR Developers*

Version	Time
4.11	3807 ops/s
4.12	9805 ops/s
4.13	21 ops/s
4.16	18 ops/s
4.17	8930 ops/s
4.18pre	18034 ops/s

**Table 1:** `smbtorture smb2.bench.path-contention-shared nprocs=256 path=1/2/3/4/5/6`

Version	Time
4.11	0m8.705s
4.12	0m8.859s
4.13	0m10.305s
4.16	0m12.386s
4.17	0m9.093s
4.18pre	0m8.476s

**Table 2:** `time rsync -a samba-4.17.0rc5 /mnt/dir/`



## JSON support for smbstatus

```
# smbstatus --json | jq
{
  "timestamp": "2022-09-08T07:42:17.407523+0200",
  "version": "4.18.0pre1-DEVELOPERBUILD",
  "smb_conf": "/opt/samba/etc/smb.conf",
  "sessions": {
    "1058442614": {
      "session_id": "1058442614",
      "username": "slow",
      ...
    },
    "tcons": {
      "4191952301": {
        "service": "test",
        ...
      }
    },
    "open_files": {
      "/srv/samba/test/WHATSNEW.txt": {
        "service_path": "/srv/samba/test",
        "filename": "WHATSNEW.txt",
        ...
      },
      ...
    }
  }
}
```

STORAGE DEVELOPER CONFERENCE



*BY Developers FOR Developers*

## SMB1 support can be disabled at compile time:


- New configure option `--without-smb1-server`
- Default is still to build with SMB1 support

## Python API for smbconf

- allow Python programs to read, and optionally write, Samba configuration natively

Coming soon

---



## SMB3 vs NFS4:

- SMB3: great feature set, wide deployment
- screaming performance with SMB-Direct or iouring
- NFS4: catching up feature-wise
  - defacto standard for Linux-to-Linux mounting
- everything's there to kick NFS, except UNIX Extensions

STORAGE DEVELOPER CONFERENCE



*BY Developers FOR Developers*

## Status

- slow progress in recent years
- pathname processing rewrite in 4.17 laid the foundation
  - detecting a symlink and preparing extended `STATUS_STOPPED_ON_SYMLINK` error return
- untangling DFS and symlink support in client libraries is a tough nut
  - needed for tests. . .
- project is now funded by [SerNet](#) and planned for 4.18

STORAGE DEVELOPER CONFERENCE

 SDC 27 | Fremont, CA  
September 12-15, 2022

*BY Developers FOR Developers*



**iouring**: a (not so) new async IO API on Linux

- Samba already has a VFS module that uses iouring:
  - only replaces the async dispatcher from a threadpool to iouring
  - still uses traditional socket fd to file fd buffer copying
  - no real performance gain
- iouring can be used with `splice` and `sendmsg()`
  - completely avoids copying data from network to disk memory buffers
- 35 Gbyte/s with prototype over loopback with 5 connections
- User-space vs kernel-space becomes less relevant



## Status

- Low hanging fruit, remaining work roughly a few weeks.
- Funded by SerNet and planned for 4.18


STORAGE DEVELOPER CONFERENCE



Fremont, CA  
September 12-15, 2022

*BY Developers FOR Developers*

Frame Reduction for HomeFolder Command Workload  
SDC 2010, "Evaluating SMB2 Performance for Home Directory Workloads"  
by Kruse and Lovinger



ER CONFERENCE

Fremont, CA  
September 12-15, 2022

OR Developers



## Status:

- Functional prototype with some gaps
- Passes Windows Protocol Test-suite
- Needs more CI test coverage
- directory rename lease break semantics are a challenge
- will be marked "experimental" initially (similar to early Multi-Channel support)
- Funded by **SerNet** and planned for 4.18

# Prototypes

---

Advanced prototype:

- implements SMB3 protocol
- implements complex changes at Samba's database layer
- minimal CI test coverage
- largish patch-set: some 140 patches
- needs rebasing on current master (expect serious conflicts...)
- still a lot to do...
- See my talk from SambaXP 2018
- Needs funding

## Status:

- Stefan Metzmacher started working on this as a side project
- No user/vendor/distro funding
- Large prototype consisting of several parts:
  1. New smbdirect.ko Linux kernel module
  2. User-space API, smbd and smbclient are the consumers
  3. Changes to the smbd file-server process using the API from 2
  4. Changes to the core SMB client library using the API from 2

## Todo:

- Coordinate with Linux kernel developers on the right way to expose the UAPI
- Code is still immature, several weeks needed to get it to production quality
- ftrace tracepoints, standalone test-suite, Multichannel integration
- Integration into Samba's CI for automated tests
- For details see Stefan's SDC talk from 2018
- Needs funding

## Status:

- Not needed for an SMB3 NAS box with Active Directory domain users
- Needed for impersonation of SSH public-key, sudo or nfs3 access authentication
- Requires cooperation with Heimdal and MIT which are averse of non-RFC features
- See Stefan's talk from SambaXP 2020 for details
- Several months of complex coordination and development needed

## Witness?

- New RPC service to "witness" availability of other services, in particular SMB3
- Prompt and explicit notifications about failures in highly available systems
- Controlled way of dealing with reconnects instead of detecting failures due to timeouts
- Part of the larger **Continuous Availability** feature
  - but not required for eg Persistent Handles

## Status:

- Prototype from Samuel Cabrero based on work from Günther Deschner
- Witness requires support for async RPC which we just acquired recently
- Development likely to be picked up once Persistent Handles are done

## Links & Q&A

---

- Most code available at [git.samba.org](https://git.samba.org)
- SDC 2018, Persistent Handles
- SDC 2017, SMB Direct Support within Samba and Linux
- Proposals
- Samba Roadmap



Thank you!  
Questions?

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

STORAGE DEVELOPER CONFERENCE



*BY Developers FOR Developers*