SNIA DEVELOPER CONFERENCE
SDC 25
By Developers FOR Developers

Hyatt Regency Santa Clara, CA
September 15-17, 2025

Scaling Secure File Access: What's New in Linux SMB3.1.1: Performance, Security and Beyond

Steven French
Principal Software Engineer –
Microsoft Azure Storage
Also Member of the Samba Team

www.sniadeveloper.org

- This work represents the views of the author(s) and does not necessarily reflect the views of Microsoft Corporation
- Linux is a registered trademark of Linus Torvalds.
- Other company, product, and service names may be trademarks or service marks of others.

Who am I?

- Steve French smfrench@gmail.com
- Author and maintainer of Linux cifs vfs (for accessing Samba, Azure, Windows and various SMB3/CIFS based NAS appliances)
- Co-maintainer of the kernel server (ksmbd)
- Member of the Samba team (co-creator of the “net” utility)
- coauthor of SNIA CIFS Technical Reference, former SNIA CIFS Working Group chair
- Principal Software Engineer, Azure Storage: Microsoft

Outline

- Overview of Linux FS activity
- Recent client improvements
- Coming soon … what to look forward to
- Testing improvements

Linux Kernel: A year ago and now …

- Now 6.17-rc5 (still called “Baby Opossum Posse”)
- Then 6.11-rc6
- 78454 changesets (non-merge commits) over period (up)
- 41824 files changed (up more than 10%)
- 3.3 million insertions, 1.5 million deletions
  - Up more than 10%
- Kernel is HUGE 28.8 million lines of code (measured last week)

Linux Storage, Filesystem, MM, BPF summit

- Last year was in May (at Salt Lake City), this Spring was in Montreal

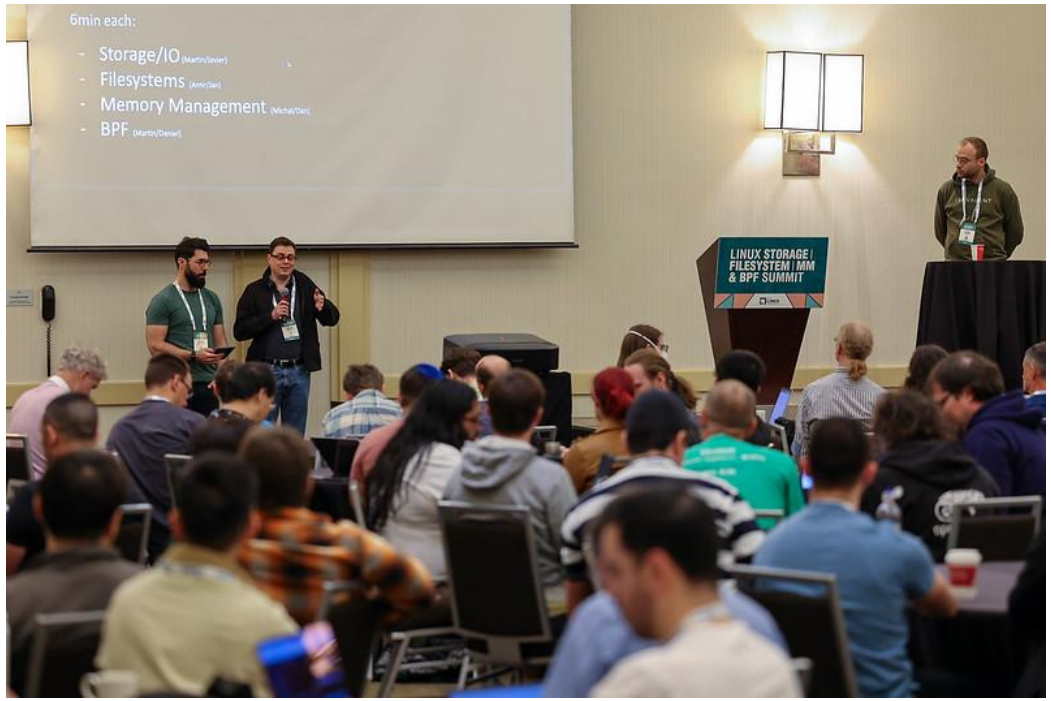

- Great group of talented Linux Linux developers working on storage and FS

Some Linux FS topics of interest discussed recently

- Testing … testing … and more automated testing … (e.g. kdevops)
- Continue work on folios, netfs mapping layer, variable size pages, fscache, improving page caching
- Leveraging eBPF better
- Improvements to statx (for example for subvol/snapshot), improvements to swap
- fine grained timestamps
- How to move Linux kernel drivers and fs to Rust eventually
- Leveraging eBPF more and more (not just dynamic tracing)
- Extending in kernel encryption: e.g. new kernel QUIC driver (helps SMB3.1.1 and other)
- Extending use of compression (e.g. to page cache)
- Shift to cloud and Better support for faster storage (NVME) and net (RDMA/smbdirect)

Linux Filesystems Activity over past year (since 6.11-rc6 kernel)

- 6969 filesystem related changesets (up significantly over the past year), 9% of total kernel changesets over this period even though only 4.1% of the lines of code! Decrease from last year,
  - Lots of developer attention on filesystems …
  - Linux kernel fs are 1.19 million lines of code total (measured last week, unchanged)
- Lots of progress over last few years!
  - Bcachefs added in 6.7 kernel now tagged “externally supported” and likely to be removed in longer term
  - Old ntfs fs removed (in favor of NTFS3), Reiserfs marked ‘obsolete’ and removed in 6.13 kernel, system V fs in 6.15, will JFS be next?

Most Active Linux Filesystems over the past year

- VFS (mapping layer) 684 changesets (up 25%)
- The top filesystems and VFS dominate the activity
- Most active local fs are bcachefs 1274 (!) now frozen and not taking merge request, BTRFS 1103 (up > 30%), XFS 669 (down), f2fs 518 (up a lot), ext4 305 (up more than 50%)
- Most active non-local fs is SMB3.1.1 (cifs.ko) 393 (up)
- Then NFSD (server) 281 (up) and NFS (client) 217
  - cifs.ko also had many more lines changed. It has been a VERY active year for cifs.ko
- Other: ksmbd server (147, up), erofs (112), OCFS2 (108), GFS2 (107), ceph (80), ntfs3 (66)

Largest File systems (by lines of code)
Key FS have code which is carefully watched

- VFS (mapping layer) 56KLOC
- XFS 150KLOC
- Btrfs 112KLOC
- SMB 91KLOC (mostly client code, expected to grow a lot in 6.18)
- Bcachefs 88KLOC (new changesets frozen, so not likely to increase)
- OCFS2 67KLOC
- NFS 65KLOC
- Ext4 47KLOC

SMB3.1.1 Activity was strong this year

- cifs.ko activity was strong, 393 changesets, activity up
  - cifs is 66KLOC kernel code (not counting user space utilities), up slightly but likely to increase a LOT in 6.18
  - Almost double more activity than NFS, and more lines changed to. Lots of great contributions from many different developers (Thank you!)
- ksmbd activity up
  - 29KLOC kernel code (not counting user space tools and libraries)

Note that Samba server (userspace) is over 4.1 million lines of code (orders of magnitude bigger than the kernel smbd server or any of the NFS servers) and is much more active, and includes key security tools and services (not just file server)

Repeating our Goals for SMB3.1.1 on Linux

- Be the fastest, most secure general-purpose way to access file data, whether in the cloud or on premises or virtualized
  - Improve directory lease support
  - Keep improving compounding, multichannel
- Support more Linux/POSIX features – so apps don't know they run on SMB3 mounts (vs. local)
  - SMB3.1.1 POSIX extensions, new FSCTLs
  - Use xfstests to locate new features to emulate
- As Linux evolves, quickly add features to Linux kernel client and Samba and ksmbd
  - More test automation and keep adding more tests

Linux File API still growing (7 recently)

e.g.: memory management continue to improve, and lots of changes to internal APIs (netfs e.g.). 238 fs related syscalls now, much more than POSIX had

Syscall name	Kernel Version introduced
setxattrat, getxattrat, listxattrat, removexattrat	6.13
open_tree_attr	6.15
file_getattr, file_setattr	6.17
And new fs ioctls:	
FS_IOC_GETUUID	6.9
FS_IOC_GETFSSYSFSPATH	6.9
F_DUPFD_QUERY fsctl	6.10
F_CREATED_QUERY fsctl	6.12

One of the strengths of SMB3.1.1 is broad interop testing

- In-person plugfests are back!
- SMB3.1.1 plugfest colocated with SDC
  - And at SambaXP now as well!
- Many exciting things being tested

Progress and Status update for Linux Kernel Server (ksmbd)

See Namjae Jeon's (linkinjeon@kernel.org) talk, from last year's SambaXP, and my update at SDC2024 (Elevating Linux File Access: Recent Enhancements to the SMB 3.1.1 Client | SDC 2025)

Recent improvements in the kernel client

(cifs.ko)

Recent Debugging Improvements

- More dynamic trace points added
- Many more debugging tools now available (see Meetakshi's presentation for more details)

Planned and in-progress Improvements (To-dos)

(cifs.ko)

Coming soon (hopefully some tested at collocated test event)

- Improving directory entry caching (better leveraging directory leases)
- Improved use of compounding (e.g. for “ls” could save a roundtrip)
- Support for SMB3.1.1 over QUIC
- New common RDMA/smbdirect code, helper module shared by ksmbd, cifs.ko and userspace such as Samba server and tools (Thank you Metze!)
- SMB3.1.1 Compression support (thank you Enzo!)
- Support for O_TMPFILE (creating temp files – thank you Jun Ma))
- Improved MacOS SMB3.1.1 interop
- Improve content crypto (great idea David Howells has been investigating)
- Rewrite of the way cifs.ko calls TCP (improves perf, thank you David Howells)
- Support for additional auth mechanisms (to make it easier to move away from NTLMv2 when not domain joined and able to use sec=krb5)

Recent improvements in the user space tools

(cifs-utils)

Testing Improvements

Test … test … test …

Thank you for your time

- Future is very bright!

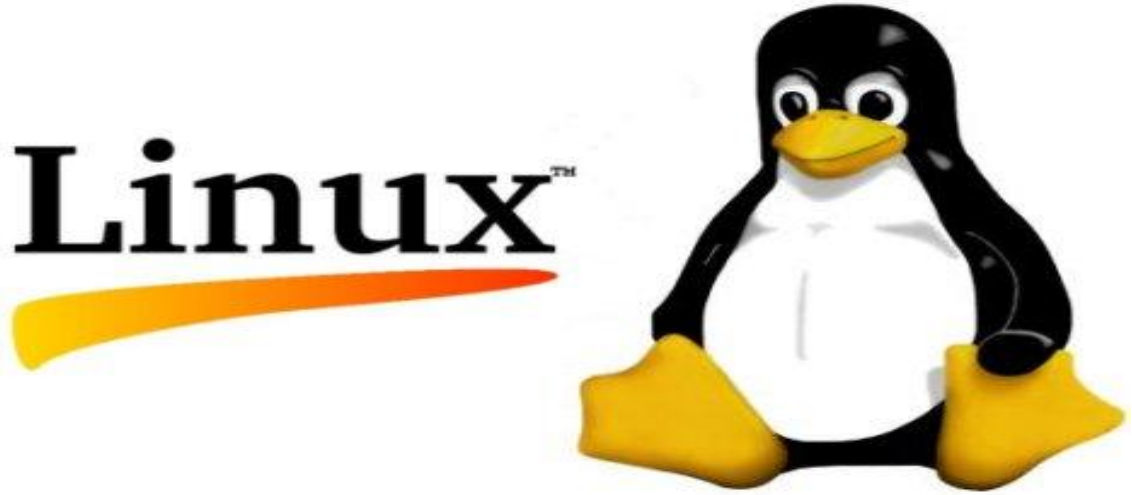

+ ***SMB3***

Additional Resources to Explore for SMB3 and Linux

- https://msdn.microsoft.com/en-us/library/gg685446.aspx
  - In particular MS-SMB2.pdf at https://msdn.microsoft.com/en-us/library/cc246482.aspx
- https://wiki.samba.org/index.php/Xfstesting-cifs and test results
  - http://smb311-linux-testing.southcentralus.cloudapp.azure.com/#/
- Linux CIFS client https://wiki.samba.org/index.php/LinuxCIFS
- Samba-technical mailing list
- And various presentations at http://www.sambaxp.org and Microsoft Learn (learn.microsoft.com) and of course SambaXP archives and SNIA … http://www.snia.org/events/storage-developer
- And the code:
  - https://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/tree/fs/smb
  - For pending changes, soon to go into upstream kernel see:
    - https://git.samba.org/?p=sfrench/cifs-2.6.git;a=shortlog;h=refs/heads/for-next
  - Userspace client tools code: https://git.samba.org/?p=cifs-utils.git (master branch)
  - Kernel server code: https://git.samba.org/ksmbd.git/?p=ksmbd.git (ksmbd-for-next branch)

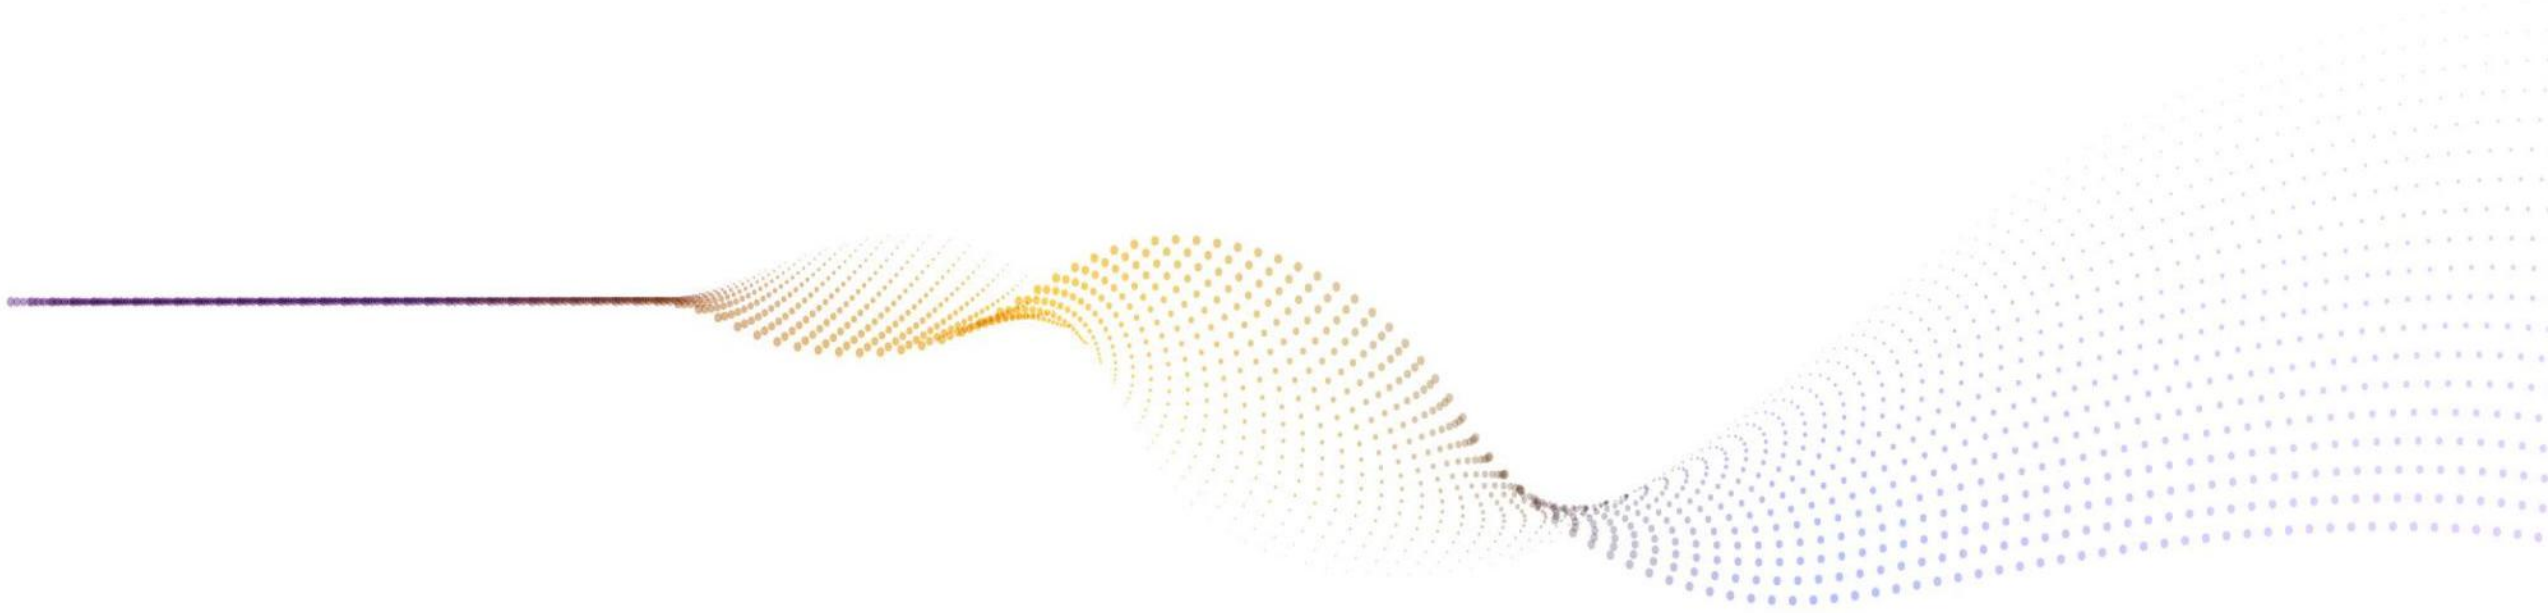

Thank you for attending!

Please remember to rate this session. You get access the presentations at
http://sniadeveloper.org/conference

