

SMB3.1.1 POSIX/Linux Extensions

Steven French

Principal Software Engineer (Azure Storage Microsoft)
and member of the Samba team





SMB3.1.1

POSIX

SMB3.1.1



Outline

- What is POSIX?
- POSIX vs. Linux
- Why SMB3.1.1?
- What are the SMB3.1.1 POSIX/Linux Extensions?
- What works today? And how to try them out?
- Wireshark status
- Improving the documentation ...
- Where do we go from here?



SMB3.1.1 POSIX Extensions status

- For the Linux client POSIX extensions were available experimentally starting in the 5.8 Linux kernel (in 2020) and much improved by the 5.15 kernel (released in October 2021)
- Samba's "smbclient" tool has support for POSIX extensions
- Ksmbd server has had support for the POSIX extensions for a few years
- Samba server now has partial support for the POSIX extensions
- And there are more ...



What is POSIX?

- POSIX was created in 1988

POSIX means Portable Operating System Interface based on UNIX

- Later standardized via "The Austin Group" in the late 90s
 - Current is <https://pubs.opengroup.org/onlinepubs/9699919799/>
- SMB was created four years earlier, in 1984 by Dr. Barry Feigenbaum at IBM
- SMB3.1.1 was created in 2015 and has been extended multiple times



- POSIX is old and the file system calls haven't changed much

- But Linux keeps evolving



Linux > POSIX

- Linux API is so much bigger than POSIX!

Currently huge number of syscalls!

- Try “git grep SYSCALL_DEFINE”
 - currently 879!
 - 18 more than when we checked at SDC!
 - 500+ are even documented
“man syscalls”
 - Linux FS layer alone has 223!
- Vs. only about 100 POSIX API calls

```
/stat.c:SYSCALL_DEFINE2(fstat64, unsigned long, fd, struct stat64
/stat.c:SYSCALL_DEFINE4(fstatat64, int, dfd, const char __user *, f
/stat.c:SYSCALL_DEFINE5(statx,
/stat.c:COMPAT_SYSCALL_DEFINE2(newstat, const char __user *, filena
/stat.c:COMPAT_SYSCALL_DEFINE2(newlstat, const char __user *, filen
/stat.c:COMPAT_SYSCALL_DEFINE4(newfstatat, unsigned int, dfd,
/stat.c:COMPAT_SYSCALL_DEFINE2(newfstat, unsigned int, fd,
/statfs.c:SYSCALL_DEFINE2(statfs, const char __user *, pathname, st
/statfs.c:SYSCALL_DEFINE3(statfs64, const char __user *, pathname,
/statfs.c:SYSCALL_DEFINE2(fstatfs, unsigned int, fd, struct statfs
/statfs.c:SYSCALL_DEFINE3(fstatfs64, unsigned int, fd, size_t, sz,
/statfs.c:SYSCALL_DEFINE2(ustat, unsigned, dev, struct ustat __user
/statfs.c:COMPAT_SYSCALL_DEFINE2(statfs, const char __user *, pathr
/statfs.c:COMPAT_SYSCALL_DEFINE2(fstatfs, unsigned int, fd, struct
/statfs.c:COMPAT_SYSCALL_DEFINE3(statfs64, const char __user *, pat


/statfs.c:COMPAT_SYSCALL_DEFINE3(fstatfs64, unsigned int, fd, compa
/statfs.c:COMPAT_SYSCALL_DEFINE2(ustat, unsigned, dev, struct compa
/sync.c:SYSCALL_DEFINE0(sync)
/sync.c:SYSCALL_DEFINE1(syncfs, int, fd)
/sync.c:SYSCALL_DEFINE1(fsync, unsigned int, fd)
/sync.c:SYSCALL_DEFINE1(fdatasync, unsigned int, fd)
/sync.c:SYSCALL_DEFINE4(sync_file_range, int, fd, loff_t, offset, l
/sync.c:SYSCALL_DEFINE4(sync_file_range2, int, fd, unsigned int, fl
/timerfd.c:SYSCALL_DEFINE2(timerfd_create, int, clockid, int, flags)
/timerfd.c:SYSCALL_DEFINE4(timerfd_settime, int, ufd, int, flags,
/timerfd.c:SYSCALL_DEFINE2(timerfd_gettime, int, ufd, struct __kern
/timerfd.c:SYSCALL_DEFINE4(timerfd_settime32, int, ufd, int, flags,
/timerfd.c:SYSCALL_DEFINE2(timerfd_gettime32, int, ufd,
/userfaultfd.c:SYSCALL_DEFINE1(userfaultfd, int, flags)
/utimes.c:SYSCALL_DEFINE4(utimensat, int, dfd, const char __user *,
/utimes.c:SYSCALL_DEFINE3(futimesat, int, dfd, const char __user *,
/utimes.c:SYSCALL_DEFINE2(utimes, char __user *, filename,
/utimes.c:SYSCALL_DEFINE2(utime, char __user *, filename, struct ut
/utimes.c:SYSCALL_DEFINE2(utime32, const char __user *, filename,
/utimes.c:SYSCALL_DEFINE4(utimensat_time32, unsigned int, dfd, cons
t, flags)
/utimes.c:SYSCALL_DEFINE3(futimesat_time32, unsigned int, dfd,
/utimes.c:SYSCALL_DEFINE2(utimes_time32, const char __user *, filen
/xattr.c:SYSCALL_DEFINE5(setxattr, const char __user *, pathname,
/xattr.c:SYSCALL_DEFINE5(lsetxattr, const char __user *, pathname,
/xattr.c:SYSCALL_DEFINE5(fsetxattr, int, fd, const char __user *, r
/xattr.c:SYSCALL_DEFINE4(getxattr, const char __user *, pathname,
/xattr.c:SYSCALL_DEFINE4(lgetxattr, const char __user *, pathname,
/xattr.c:SYSCALL_DEFINE4(fgetxattr, int, fd, const char __user *, r
/xattr.c:SYSCALL_DEFINE3(listxattr, const char __user *, pathname,
/xattr.c:SYSCALL_DEFINE3(llistxattr, const char __user *, pathname,
/xattr.c:SYSCALL_DEFINE3(flistxattr, int, fd, char __user *, list,
/xattr.c:SYSCALL_DEFINE2(removexattr, const char __user *, pathname
/xattr.c:SYSCALL_DEFINE2(lremovexattr, const char __user *, pathnar
/xattr.c:SYSCALL_DEFINE2(fremovexattr, int, fd, const char __user *

french@smfrench-ThinkPad-P52:~/cifs-2.6$ git grep SYSCALL_DEFINE |
850      5070      69194
```


Linux File System API grows

Recently e.g.: very important feature: "folios" added and many changes to internal APIs as well (netfs, fscache ...) and io_uring continues to improve

| Syscall name | Kernel Version introduced |
|---------------------|----------------------------------|
| epoll_pwait2 | 5.11 |
| mount_setattr | 5.12 |
| faccessat2 | 5.8 |
| close_range | 5.9 |



Goals: Fast! Easy! Transparent!

- Repeating an older slide about goals of SMB3.1.1:
 - Fastest, most secure general purpose way to access file data, whether cloud or on premises or virtualized
 - Implement all reasonable Linux/POSIX features - so apps don't know they run on SMB3 mounts (vs. local)
 - As Linux evolves, and needs new features, quickly add to Linux kernel client and Samba and ksmbd




Why Not Other Protocols?

- SMB3.1.1 is easily extensible
- SMB3.1.1 works tightly with a set of protocols which can do more than any other file system protocol
- SMB3.1.1 has the best, most exhaustive set of testcases (not just smbtorure ...)
- SMB3.1.1 and related protocols have more documentation (and documentation that has been tested and verified)
- SMB3.1.1 is proven across multiple client types, OS, architectures (and POSIX extensions have been a moving target, done before ...) and has been extended before ...
- (And don't forget ... SAMBA rocks! And multiple open source server choices that support SMB3.1.1 (including ksmbd and Samba) And cifs.ko is one of most active FS)



Why are extensions needed?

SMB2 and later (including SMB3.1.1) default file semantics are largely based on Windows behavior which differs from POSIX (and Linux) in some ways



What about SMB1/CIFS? It had Unix Extensions...

- DO NOT USE OLD INSECURE DIALECTS (lookup Ned Pyle's presentations ...)
- SMB3.1.1 is faster, more secure, simpler, better
- And broadly implemented by modern clients and servers
- And the SMB3.1.1 POSIX extensions are simpler than the old CIFS Unix Extensions and intended to completely replace their use



What works today without POSIX Extensions

- Normal file and directory operations (open, read, write, fsync, close) to all servers, and hardlinks and even client handled symlinks (“mfsymlinks”), case preserving file name behavior, mapping almost all problematic characters in filenames (“\” is the one exception)
- To most servers:
 - Sparse file operations: setsparse, query allocated ranges, punch hole
 - copy_range and clone_range (clone range is less commonly supported)
 - Special file handling via reparse points (or xattrs ala “sfu”)
 - Xattrs
- Emulation of mode bits via various alternatives (cifsacl, modefromsid)



What can be emulated today without POSIX Extensions

- fcollapse and finsert
- Most delete and rename scenarios (some exceptions is where the rename fails with access denied with rename onto an existing file)
- Most byte range (easier with OFD rather than “posix” BRLs) and whole file lock scenarios
- Most special mode bits (e.g. sticky bit)



But POSIX apps expect

- Case sensitive file and directory names
- Primitive mode bits returned (“0777” perms) rather than a “rich ACL”
- Delete of open files (and these not seen by readdir)
- Rename of open file
- Advisory (rather than mandatory) byte range locks
- 2 additional fields returned by statfs (total/free inodes)



What if POSIX and Windows semantics collide?

If the same file is being used by Windows and Linux clients – how do we deal with semantic differences?

- Lessons from the CIFS (SMB1) Unix Extensions
- What about RichACLs and mode bits?



What about the Apple Mac extensions to SMB3.1.1?

... They only address part of what Linux needs



- Implement AAPL context
 - Improved Mac interop is another benefit
 - Samba even has a `vfs_fruit` module that adds other interesting features (spotlight integration e.g.)
- Subset of POSIX requirements can be solved
- `kAAPL_SERVER_CAPS = 0x01`,
 - `kAAPL_SUPPORTS_READ_DIR_ATTR = 0x01`,
 - `kAAPL_SUPPORTS_OSX_COPYFILE = 0x02`,
 - `kAAPL_UNIX_BASED = 0x04`
 - `kAAPL_SUPPORTS_NFS_ACE = 0x08`
- `kAAPL_VOLUME_CAPS = 0x02`,
 - `kAAPL_SUPPORT_RESOLVE_ID = 0x01`,
 - `kAAPL_CASE_SENSITIVE = 0x02`
- `kAAPL_MODEL_INFO = 0x04` (pad, length, model string)



How can I try the SMB3.1.1 POSIX extensions?

- Client implementations include: Linux kernel (cifs.ko), especially 5.8 kernel and later, and newer versions of Samba's smbclient (thanks Volker)
- Server versions include Ksmbd server (thanks Namjae) as well as newer Samba server (thanks JRA, Volker, David etc.) ...
- Third party implementations also were tested at SNIA SDC SMB3 plugfest



Demo

Some examples of where POSIX Extensions are needed ...



Case Sensitivity – without extensions can open the wrong file ...
and owner and mode bits usually set to defaults (not real owner)

```
root@smfrench-ThinkPad-P52:/home/smfrench# ls /mnt1/small-dir -l
total 12
-rwxr-xr-x 1 root root  0 Sep 13 02:42 747-file
-rwxr-xr-x 1 root root 19 Sep 13 02:44 casesensitiveexample
-rwxr-xr-x 1 root root 11 Sep 13 02:44 CaseSensitiveExample
-rwxr-xr-x 1 root root  6 Sep 13 02:44 CASESENSITIVEEXAMPLE
-rwxr-xr-x 1 root root  0 Sep 13 02:12 file1-root
-rwxr-xr-x 1 root root  0 Sep 13 02:41 file-as-smfrench
root@smfrench-ThinkPad-P52:/home/smfrench# cat /mnt1/small-dir/CaseSensitiveExample
mixed case
root@smfrench-ThinkPad-P52:/home/smfrench# cat /mnt1/small-dir/casesensitiveexample
mixed case
root@smfrench-ThinkPad-P52:/home/smfrench# cat /mnt1/small-dir/CASESENSITIVEEXAMPLE
mixed case
root@smfrench-ThinkPad-P52:/home/smfrench#
```



SMB3.1.1 POSIX Extensions are Easy to Understand

- A simple negotiate context, an open context, a new file info level and a new fsinfo level
- Everything else relies on existing SMB3.1.1 features



Demo

Some examples to Samba and KSMBD (kernel server) using the SMB3.1.1 POSIX Extensions



Setting up Samba and ksmbd shares are easy

NB: Samba requires "smb3 unix extensions = yes" in smb.conf

```
root@smfrench-ThinkPad-P52: /home/smfrench/cif...  smfrench@smfrench-ThinkPad-P52: ~/ks
See smb.conf.example for a more detailed config file
[global]
server multi channel support=yes
[global]
workgroup = SAMBA
map to guest = Bad User
passwd backend = tdbSAM
printing = cups
printcap name = cups
host msdfs = yes
server multi channel support = yes
log level = 4
smb3 unix extensions = yes
[scratch]
comment = scratch share for testing
browseable = yes
path = /scratch
guest ok = yes
read only = no
ea support = yes
create mask = 0777
local/samba/etc/smb.conf" 84L, 1690B  2,0-1  "/etc/ksmbd/smb.conf" 11L, 154B
[global]
server multi channel support=yes
[test]
path = /test
writeable=yes
read only = no
[scratch]
path = /scratch
writeable=yes
read only = no
~
~
~
~
~
~
~
~
~
~
```



Additional Samba server configuration advice

- Build Samba with `./configure.developer`
- To Support for reporting mode bits currently Samba server requires disabling `acl_xattr` (do not enable saving ACLs to xattr this in `vfs objects` on the share if you want to test returning mode bits)
- Consider relaxing the `create mask`
 - `smb.conf` parameters `create mask = 0777` and `directory mask = 0777`



Mounting from the Linux kernel client

- Remember to add "posix" on mount command
- Also consider "mfsymlinks" if want client only symlinks

```
root@smfrench-ThinkPad-P52:/home/smfrench/mulder-posix/samba# mount -t cifs //local
/mnt1 -o username=testuser,password=testpass,mfsymlinks,posix
root@smfrench-ThinkPad-P52:/home/smfrench/mulder-posix/samba# cat /proc/fs/cifs/De
Display Internal CIFS Data Structures for Debugging
-----
CIFS Version 2.38
Features: DFS,FSCACHE,STATS,DEBUG,ALLOW_INSECURE_LEGACY,CIFS_POSIX,UPCALL(SPNEGO),
WITNESS
CIFSMaxBufSize: 16384
Active VFS Requests: 0

Servers:
1) ConnectionId: 0x4 Hostname: localhost
Number of credits: 443 Dialect 0x311 posix
TCP status: 1 Instance: 1
Local Users To Server: 1 SecMode: 0x1 Req On Wire: 0
In Send: 0 In MaxReq Wait: 0

Sessions:
1) Address: 127.0.0.1 Uses: 1 Capability: 0x300047 Session Status: 1
Security type: RawNTLMSSP SessionId: 0x35ada477
User: 0 Cred User: 0

Shares:
```



Can display owner and mode bits

(bug in group owner vs. user owner – fixed now, see next slide)

```
root@smfrench-ThinkPad-P52:/home/smfrench# ls /scratch/small-dir -l ; mount | grep smb3 ; ls /mnt1/small-dir -l
total 12
-rwxr--rwx 1 smfrench smfrench  0 Sep 13 02:42 747-file
-rwxr--r-- 1 smfrench smfrench 19 Sep 13 02:44 casesensitiveexample
-rwxr-xr-x 1 smfrench smfrench 11 Sep 13 02:44 CaseSensitiveExample
-rwxrwxrwx 1 smfrench smfrench  6 Sep 13 02:44 CASESENSITIVEEXAMPLE
-rw-r--r-- 1 root     root      0 Sep 13 02:12 file1-root
-rw-rw-r-- 1 smfrench smfrench  0 Sep 13 02:41 file-as-smfrench
//localhost/scratch on /mnt1 type smb3 (rw,relatime,vers=3.1.1,cache=strict,username=testuser,uid=0,noforceuid,gid=0,noforcegid,addr=127.0.0.1,file_mode=0755,dir_mode=0755,soft,posix,posixpaths,serverino,mapposix,mfsymlinks,rsize=4194304,wsiz=4194304,bsize=1048576,echo_interval=60,actimeo=1,closetimeo=5)
total 2
-rwxr--rwx 1 root smfrench  0 Sep 13 02:42 747-file
-rwxr--r-- 1 root smfrench 4096 Sep 13 02:44 casesensitiveexample
-rwxr-xr-x 1 root smfrench 4096 Sep 13 02:44 CaseSensitiveExample
-rwxrwxrwx 1 root smfrench 4096 Sep 13 02:44 CASESENSITIVEEXAMPLE
-rw-r--r-- 1 root root      0 Sep 13 02:12 file1-root
-rw-rw-r-- 1 root smfrench  0 Sep 13 02:41 file-as-smfrench
root@smfrench-ThinkPad-P52:/home/smfrench#
```



Note exact mode bits and owner reported w/POSIX Extensions

| No. | Time | Source | Destination | Protocol | Length | Info |
|-----|-------------|-----------|-------------|----------|--------|---|
| 2 | 14.00553... | 127.0.0.1 | 127.0.0.1 | SMB2 | 454 | Create Request File: ;GetInfo Request FILE_INFO/SMB2_FILE_F |
| 3 | 14.01069... | 127.0.0.1 | 127.0.0.1 | SMB2 | 670 | Create Response File: ;GetInfo Response;Close Response |
| 5 | 14.01149... | 127.0.0.1 | 127.0.0.1 | SMB2 | 360 | Create Request File: ;Find Request SMB2_FIND_POSIX_INFO Pat |
| 6 | 14.02095... | 127.0.0.1 | 127.0.0.1 | SMB2 | 974 | Create Response File: ;Find Response |
| 7 | 14.02139... | 127.0.0.1 | 127.0.0.1 | SMB2 | 168 | Find Request File: SMB2_FIND_POSIX_INFO Pattern: * |

Blob Length: 542

Info: 800000000000000000000008d31e82a40c7d8011765b03244c7d8018d31e82a40c7d8018d31e82a...

- FilePosixInfo
- FilePosixInfo
- FilePosixInfo
- FilePosixInfo
 - Next Offset: 0
 - Create: Sep 13, 2022 02:43:30.175737000 CDT
 - Last Access: Sep 13, 2022 02:43:34.911762800 CDT
 - Last Write: Sep 13, 2022 02:43:30.175737000 CDT
 - Last Change: Sep 13, 2022 02:43:30.175737000 CDT
 - Allocation Size: 0
 - End Of File: 0
 - File Attributes: 0x00000000
 - Inode: 0x00000000601ec8af
 - File Id: 0x00000000000010302
 - Reserved: 00000000
 - Number of Links: 2
 - Reparse Tag: REPARSE_TAG_RESERVED_ZERO (0x00000000)
 - POSIX perms: 0777
 - Owner SID: S-1-22-1-0
 - Group SID: S-1-22-2-0

```
root@smfrench-ThinkPad-P52: ~  
root@smfrench-ThinkPad-P52:~# ls /mnt1 -la  
total 139  
rwxrwxrwx  3 root      root           0 Sep 13 02:12 .  
rwxr-xr-x 32 root      root          4096 Sep  3 19:15 ..  
rwxr--r--  1 testuser1 testuser1 139264 Sep  7 10:43 ls_on_scratch  
rwxrwxrwx  2 root      root           0 Sep 13 02:43 small-dir
```



Query FS Info – includes additional posix fields

```
root@smfrench-ThinkPad-P52:/home/smfrench# mount -t cifs //local
-o username=testuser,password=testpass,mfsymlinks,posix
root@smfrench-ThinkPad-P52:/home/smfrench# stat -f /mnt1
File: "/mnt1"
  ID: c7df5aa0f1e89eff Namelen: 255      Type: smb2
Block size: 4096      Fundamental block size: 4096
Blocks: Total: 139092115  Free: 48993190   Available: 48993190
Inodes: Total: 278320128  Free: 273211966
root@smfrench-ThinkPad-P52:/home/smfrench#
```



Query FS Info – additional POSIX fields

Samba
server
works
now too!

```
root@smfrench-ThinkPad-P52:~/tmp# mount | grep smb3
//localhost/test on /mnt/posix type smb3 (rw,relatime,vers=3.1.1,cache=strict,username=s
=0,noforceuid,gid=0,noforcegid,addr=127.0.0.1,file_mode=0755,dir_mode=0755,soft,posix,p
rverino,mapposix,mfsymlinks,rsize=4194304,wsize=4194304,bsize=1048576,echo_interval=60,a
osetimeo=1)
//localhost/test on /mnt/non-posix type smb3 (rw,relatime,vers=3.1.1,cache=strict,userna
,uid=0,noforceuid,gid=0,noforcegid,addr=127.0.0.1,file_mode=0755,dir_mode=0755,soft,nour
,mfsymlinks,rsize=4194304,wsize=4194304,bsize=1048576,echo_interval=60,actimeo=1,closeti
root@smfrench-ThinkPad-P52:~/tmp# stat -f /test
File: "/test"
ID: 10302000000000 Namelen: 255 Type: xfs
Block size: 4096 Fundamental block size: 4096
Blocks: Total: 139092115 Free: 30106259 Available: 30106259
Inodes: Total: 247480920 Free: 240853758
root@smfrench-ThinkPad-P52:~/tmp# stat -f /mnt/posix
File: "/mnt/posix"
ID: aab31952b3127fb4 Namelen: 255 Type: smb2
Block size: 4096 Fundamental block size: 4096
Blocks: Total: 139092115 Free: 30106254 Available: 30106254
Inodes: Total: 247480880 Free: 240853716
root@smfrench-ThinkPad-P52:~/tmp# stat -f /mnt/non-posix
File: "/mnt/non-posix"
ID: aab31952b3127fb4 Namelen: 255 Type: smb2
Block size: 1024 Fundamental block size: 1024
Blocks: Total: 556368460 Free: 120425008 Available: 120425008
Inodes: Total: 0 Free: 0
```



More details: Local vs non-POSIX SMB3.1.1 (example to Samba server)

```
oot@smfrench-ThinkPad-P52:~/tmp# ls /test -l
total 256
r--r--r-- 1 root root 0 May 4 20:21 0444file
rwx---r-- 1 root root 0 May 4 20:22 0704file
rwxrw-r-x 1 root root 0 May 4 20:21 0765file
rwxrwxrwx+ 3 testuser1 testuser1 21 Apr 24 22:52 dbench
rwxrwxrwx+ 1 testuser1 testuser1 65536 Apr 24 23:39 file1
rwxrwxrwx+ 1 testuser1 testuser1 65536 Apr 24 23:39 file2
rwxrwxrwx+ 1 smfrench smfrench 0 Apr 10 16:07 file3
rwxr--rwx 2 smfrench smfrench 6 Apr 9 22:43 file-smfrench-747
rwxrwxrwx+ 2 testuser1 testuser1 19 Apr 24 22:43 fsx
rw-r--r-- 1 root root 0 Apr 9 22:42 local-file-as-root
rw-rw-r-- 1 smfrench smfrench 0 Apr 9 22:43 local-file-as-smfrench
rwxrwxrwx+ 1 smfrench smfrench 1067 May 4 20:29 mfsymlink-source
rwxrwxrwx+ 1 testuser1 testuser1 8256 Apr 24 23:55 mmap-writev-overlap
rwxrwxrwx+ 1 testuser1 testuser1 24 Apr 24 22:19 mtab_output
rwxrwxrwx+ 2 smfrench smfrench 6 Apr 9 22:38 new-dir-root
rwxrwxrwx+ 1 root root 0 Apr 9 22:38 new-file-root
rwxrwxrwx+ 1 smfrench smfrench 0 Apr 9 22:37 new-file-smfrench
rwxr-xr-x 2 root root 6 Nov 23 17:48 root-dir
rwxrwxr-x 2 smfrench smfrench 6 Nov 23 17:49 smfrench-dir
rwxrwxrwx 1 root root 6 May 4 20:27 source -> target
rwxrwxrwt 2 root root 6 May 4 20:24 stickybitdir01777
rwxr-xr-x 2 root root 22 Apr 9 22:59 target
rwxrwxrwx+ 2 testuser1 testuser1 28672 Apr 24 22:25 test
rw-r--r-- 1 root root 0 Apr 9 22:56 this-is-fifo
rwxrwxrwx+ 1 testuser1 testuser1 0 Jan 1 1960 timestamp-test.txt
rwxrwxrwx+ 1 smfrench smfrench 0 May 4 20:16 upper
rwxrwxrwx+ 1 smfrench smfrench 0 Apr 9 22:38 Upper
rwxrwxrwx 1 root root 0 May 4 20:17 UPPER
```

```
oot@smfrench-ThinkPad-P52:~/tmp# ls /mnt/non-posix -l
total 192
rwxr-xr-x 1 root root 0 May 4 20:21 0444file
rwxr-xr-x 1 root root 0 May 4 20:22 0704file
rwxr-xr-x 1 root root 0 May 4 20:21 0765file
rwxr-xr-x 2 root root 0 Apr 24 22:52 dbench
rwxr-xr-x 1 root root 65536 Apr 24 23:39 file1
rwxr-xr-x 1 root root 65536 Apr 24 23:39 file2
rwxr-xr-x 1 root root 0 Apr 10 16:07 file3
rwxr-xr-x 2 root root 0 Apr 9 22:43 file-smfrench-747
rwxr-xr-x 2 root root 0 Apr 24 22:43 fsx
rwxr-xr-x 1 root root 0 Apr 9 22:42 local-file-as-root
rwxr-xr-x 1 root root 0 Apr 9 22:43 local-file-as-smfrench
rwxrwxrwx 1 root root 6 May 4 20:29 mfsymlink-source -> target
rwxr-xr-x 1 root root 8256 Apr 24 23:55 mmap-writev-overlap
rwxr-xr-x 1 root root 24 Apr 24 22:19 mtab_output
rwxr-xr-x 2 root root 0 Apr 9 22:38 new-dir-root
rwxr-xr-x 1 root root 0 Apr 9 22:38 new-file-root
rwxr-xr-x 1 root root 0 Apr 9 22:37 new-file-smfrench
rwxr-xr-x 2 root root 0 Nov 23 17:48 root-dir
rwxr-xr-x 2 root root 0 Nov 23 17:49 smfrench-dir
rwxr-xr-x 2 root root 0 Apr 9 22:59 source
rwxr-xr-x 2 root root 0 May 4 20:24 stickybitdir01777
rwxr-xr-x 2 root root 0 Apr 9 22:59 target
rwxr-xr-x 2 root root 0 Apr 24 22:25 test
rwxr-xr-x 1 root root 0 Apr 9 22:56 this-is-fifo
rwxr-xr-x 1 root root 0 Jan 1 1960 timestamp-test.txt
rwxr-xr-x 1 root root 0 May 4 20:16 upper
rwxr-xr-x 1 root root 0 Apr 9 22:38 Upper
rwxr-xr-x 1 root root 0 May 4 20:17 UPPER
```



More details: Local vs new POSIX SMB3.1.1 (example to Samba server)

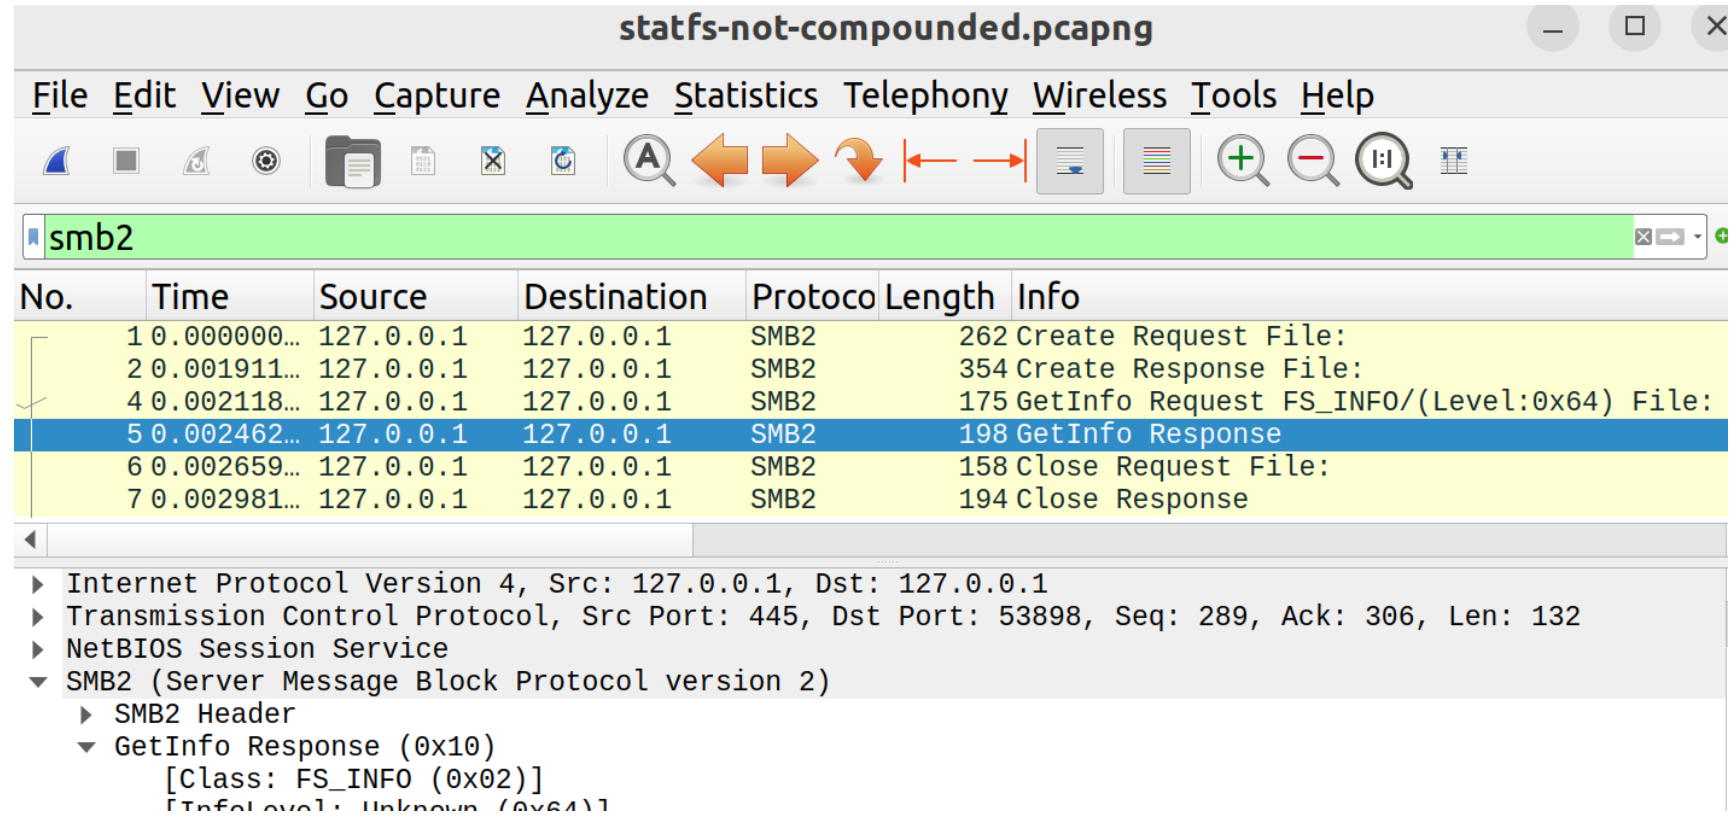
```
root@smfrench-ThinkPad-P52:~/tmp# ls /test -l
total 256
-r--r--r-- 1 root root 0 May 4 20:21 0444file
-rwx---r-- 1 root root 0 May 4 20:22 0704file
-rwxrw-r-x 1 root root 0 May 4 20:21 0765file
-rwxrwxrwx+ 3 testuser1 testuser1 21 Apr 24 22:52 dbench
-rwxrwxrwx+ 1 testuser1 testuser1 65536 Apr 24 23:39 file1
-rwxrwxrwx+ 1 testuser1 testuser1 65536 Apr 24 23:39 file2
-rwxrwxrwx+ 1 smfrench smfrench 0 Apr 10 16:07 file3
-rwxr--rwx 2 smfrench smfrench 6 Apr 9 22:43 file-smfrench-747
-rwxrwxrwx+ 2 testuser1 testuser1 19 Apr 24 22:43 fsx
-rw-r--r-- 1 root root 0 Apr 9 22:42 local-file-as-root
-rw-rw-r-- 1 smfrench smfrench 0 Apr 9 22:43 local-file-as-smfrench
-rwxrwxrwx+ 1 smfrench smfrench 1067 May 4 20:29 mfsymlink-source
-rwxrwxrwx+ 1 testuser1 testuser1 8256 Apr 24 23:55 mmap-writev-overlap
-rwxrwxrwx+ 1 testuser1 testuser1 24 Apr 24 22:19 mtab_output
-rwxrwxrwx+ 2 smfrench smfrench 6 Apr 9 22:38 new_dir_root
-rwxrwxrwx+ 1 root root 0 Apr 9 22:38 new-file-root
-rwxrwxrwx+ 1 smfrench smfrench 0 Apr 9 22:37 new-file-smfrench
-rwxr-xr-x 2 root root 6 Nov 23 17:48 root-dir
-rwxrwxr-x 2 smfrench smfrench 6 Nov 23 17:49 smfrench-dir
-rwxrwxrwx 1 root root 6 May 4 20:27 source -> target
-rwxrwxrwt 2 root root 6 May 4 20:24 stickybitdir01777
-rwxr-xr-x 2 root root 22 Apr 9 22:59 target
-rwxrwxrwx+ 2 testuser1 testuser1 28672 Apr 24 22:25 test
-rw-r--r-- 1 root root 0 Apr 9 22:56 this-is-fifo
-rwxrwxrwx+ 1 testuser1 testuser1 0 Jan 1 1960 timestamp-test.txt
-rwxrwxrwx+ 1 smfrench smfrench 0 May 4 20:16 upper
-rwxrwxrwx+ 1 smfrench smfrench 0 Apr 9 22:38 Upper
-rwxrwxrwx 1 root root 0 May 4 20:17 UPPER
```

```
root@smfrench-ThinkPad-P52:~/tmp# ls /mnt/posix -l
ls: cannot access '/mnt/posix/source': No such file or directory
total 192
-r--r--r-- 1 root root 0 May 4 20:21 0444file
-rwx---r-- 1 root root 0 May 4 20:22 0704file
-rwxrw-r-x 1 root root 0 May 4 20:21 0765file
drwxrwxrwx 3 root testuser1 0 Apr 24 22:52 dbench
-rwxrwxrwx 1 root testuser1 65536 Apr 24 23:39 file1
-rwxrwxrwx 1 root testuser1 65536 Apr 24 23:39 file2
-rwxrwxrwx 1 root smfrench 0 Apr 10 16:07 file3
drwxr--rwx 2 root smfrench 0 Apr 9 22:43 file-smfrench-747
drwxrwxrwx 2 root testuser1 0 Apr 24 22:43 fsx
-rw-r--r-- 1 root root 0 Apr 9 22:42 local-file-as-root
-rw-rw-r-- 1 root smfrench 0 Apr 9 22:43 local-file-as-smfrench
lrwxrwxrwx 1 root smfrench 6 May 4 20:29 mfsymlink-source -> target
-rwxrwxrwx 1 root testuser1 8256 Apr 24 23:55 mmap-writev-overlap
-rwxrwxrwx 1 root testuser1 24 Apr 24 22:19 mtab_output
drwxrwxrwx 2 root smfrench 0 Apr 9 22:38 new_dir_root
-rwxrwxrwx 1 root root 0 Apr 9 22:38 new-file-root
-rwxrwxrwx 1 root smfrench 0 Apr 9 22:37 new-file-smfrench
drwxr-xr-x 2 root root 0 Nov 23 17:48 root-dir
drwxrwxr-x 2 root smfrench 0 Nov 23 17:49 smfrench-dir
-????????? ? ? ? ? ? ? ? source
drwxrwxrwt 2 root root 0 May 4 20:24 stickybitdir01777
drwxr-xr-x 2 root root 0 Apr 9 22:59 target
drwxrwxrwx 2 root testuser1 0 Apr 24 22:25 test
-rw-r--r-- 1 root root 0 Apr 9 22:56 this-is-fifo
-rwxrwxrwx 1 root testuser1 0 Jan 1 1960 timestamp-test.txt
-rwxrwxrwx 1 root smfrench 0 May 4 20:16 upper
-rwxrwxrwx 1 root smfrench 0 Apr 9 22:38 Upper
-rwxrwxrwx 1 root root 0 May 4 20:17 UPPER
```



Better performance (POSIX QFS Info now compounded)

Before:



statfs-not-compounded.pcapng

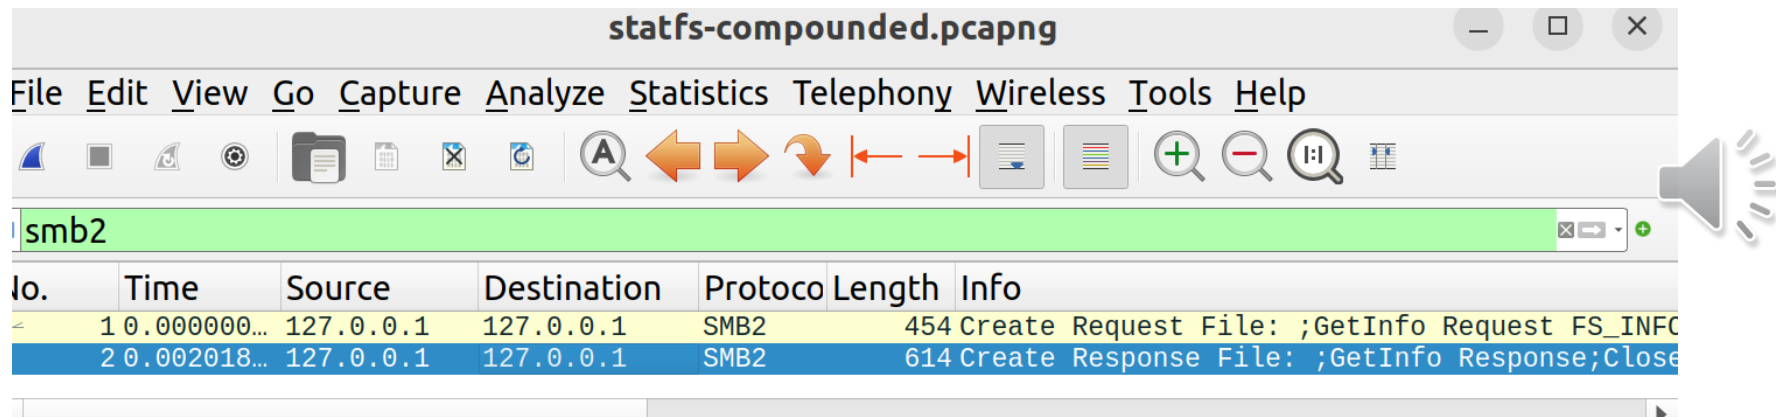
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help

smb2

| No. | Time | Source | Destination | Protocol | Length | Info |
|-----|-------------|-----------|-------------|----------|--------|--|
| 1 | 0.000000... | 127.0.0.1 | 127.0.0.1 | SMB2 | 262 | Create Request File: |
| 2 | 0.001911... | 127.0.0.1 | 127.0.0.1 | SMB2 | 354 | Create Response File: |
| 4 | 0.002118... | 127.0.0.1 | 127.0.0.1 | SMB2 | 175 | GetInfo Request FS_INFO/(Level:0x64) File: |
| 5 | 0.002462... | 127.0.0.1 | 127.0.0.1 | SMB2 | 198 | GetInfo Response |
| 6 | 0.002659... | 127.0.0.1 | 127.0.0.1 | SMB2 | 158 | Close Request File: |
| 7 | 0.002981... | 127.0.0.1 | 127.0.0.1 | SMB2 | 194 | Close Response |

Internet Protocol Version 4, Src: 127.0.0.1, Dst: 127.0.0.1
Transmission Control Protocol, Src Port: 445, Dst Port: 53898, Seq: 289, Ack: 306, Len: 132
NetBIOS Session Service
SMB2 (Server Message Block Protocol version 2)
SMB2 Header
GetInfo Response (0x10)
[Class: FS_INFO (0x02)]
[InfoLevel: Unknown (0x64)]

Now:



statfs-compounded.pcapng


File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help

smb2

| No. | Time | Source | Destination | Protocol | Length | Info |
|-----|-------------|-----------|-------------|----------|--------|---|
| 1 | 0.000000... | 127.0.0.1 | 127.0.0.1 | SMB2 | 454 | Create Request File: ;GetInfo Request FS_INF |
| 2 | 0.002018... | 127.0.0.1 | 127.0.0.1 | SMB2 | 614 | Create Response File: ;GetInfo Response;Close |

SMB3.1.1 POSIX Extensions details





smb2 Expression...

| Source | Destination | Protoc | Length | Info |
|-----------|-------------|--------|--------|---|
| 127.0.0.1 | 127.0.0.1 | SMB2 | 302 | Negotiate Protocol Request |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 366 | Negotiate Protocol Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 190 | Session Setup Request, NTLMSSP_NEGOTIATE |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 400 | Session Setup Response, Error: STATUS_MORE_PROCESSING_REQUIRED, NTLMSSP_C |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 460 | Session Setup Request, NTLMSSP_AUTH, User: \smfrench |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 142 | Session Setup Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 176 | Tree Connect Request Tree: \\localhost\IPC\$ |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 150 | Tree Connect Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 176 | Tree Connect Request Tree: \\localhost\test |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 150 | Tree Connect Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 446 | Create Request File: ;GetInfo Request FILE_INFO/SMB2_FILE_ALL_INFO |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 534 | Create Response File: ;GetInfo Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 191 | Ioctl Request FSCTL_QUERY_NETWORK_INTERFACE_INFO |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 143 | Ioctl Response, Error: STATUS_INVALID_DEVICE_REQUEST |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 175 | GetInfo Request FS_INFO/FileFsAttributeInformation File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 162 | GetInfo Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 175 | GetInfo Request FS_INFO/FileFsDeviceInformation File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 150 | GetInfo Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 175 | GetInfo Request FS_INFO/FileFsVolumeInformation File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 168 | GetInfo Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 175 | GetInfo Request FS_INFO/FileFsSectorSizeInformation File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 170 | GetInfo Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 158 | Close Request File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 194 | Close Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 224 | Ioctl Request FSCTL_DFS_GET_REFERRALS, File: \localhost\test |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 143 | Ioctl Response, Error: STATUS_NOT_FOUND |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 262 | Create Request File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 354 | Create Response File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 158 | Close Request File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 194 | Close Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 262 | Create Request File: |

```

- NetBIOS Session Service
- SMB2 (Server Message Block Protocol version 2)
  - SMB2 Header
    - Negotiate Protocol Request (0x00)
      - StructureSize: 0x0024
      - Dialect count: 4
      - Security mode: 0x01, Signing enabled
      - Reserved: 0000
      - Capabilities: 0x00000077, DFS, LEASING, LARGE MTU, PERSISTENT HA
      - Client Guid: 032f6ffc-4993-c44d-8b01-425c86949469
      - NegotiateContextOffset: 0x0070
      - NegotiateContextCount: 4
      - Reserved: 0000
      - Dialect: 0x0210
      - Dialect: 0x0300
      - Dialect: 0x0302
      - Dialect: 0x0311
      - Negotiate Context: SMB2_PREAUTH_INTEGRITY_CAPABILITIES
      - Negotiate Context: SMB2_ENCRYPTION_CAPABILITIES
      - Negotiate Context: Unknown Type: (0x5)
      - Negotiate Context: SMB2_POSIX_EXTENSIONS_CAPABILITIES
        - Type: SMB2_POSIX_EXTENSIONS_CAPABILITIES (0x0100)
        - DataLength: 16
        - Reserved: 00000000
        - POSIX Reserved: 0x5025ad93
    
```

| | | | |
|------|-------------------------|-------------------------|------------------|
| 00b0 | 11 03 00 00 00 00 01 00 | 26 00 00 00 00 00 01 00 |&... |
| 00c0 | 20 00 01 00 b4 02 72 b9 | 3c ed 58 84 13 7e b1 3f |r.<X.. |
| 00d0 | 4d 7a 75 80 ee 25 a5 e0 | 81 20 cc 10 57 7e 31 65 |u.%...w |
| 00e0 | 95 f7 40 b3 00 00 02 00 | 06 00 00 00 00 00 02 00 |@..... |
| 00f0 | 02 00 01 00 00 00 05 00 | 12 00 00 00 00 00 6c 00 |c.a.l.h.o.s |
| 0100 | 6f 00 63 00 61 00 6c 00 | 68 00 6f 00 73 00 74 00 |%P.....# |
| 0110 | 00 00 00 00 00 00 00 01 | 10 00 00 00 00 00 93 ad | |
| 0120 | 25 50 9c b4 11 e7 b4 23 | 83 de 96 8b cd 7c | |

Text item (text), 24 bytes


Additional information can be found at:

- For the SMB3.1.1 POSIX Extensions see the Samba wiki at:
[SMB3-Linux – SambaWiki](https://wiki.samba.org/index.php/SMB3-Linux)
(<https://wiki.samba.org/index.php/SMB3-Linux>).
contributions to improve this page are welcome
- Information on the older SMB1 extensions can be found at
[UNIX Extensions – SambaWiki](#) e.g. and in the olderSNIA
doc




```
>- NetBIOS Session Service
  >- SMB2 (Server Message Block Protocol version 2)
    >- SMB2 Header
      >- Negotiate Protocol Request (0x00)
        >- StructureSize: 0x0024
          >- Dialect count: 4
            >- Security mode: 0x01, Signing enabled
              >- Reserved: 0000
                >- Capabilities: 0x00000077, DFS, LEASING, LARGE MTU, PERSISTENT HA
                  >- Client Guid: 032f6ffc-4993-c44d-8b01-425c86949469
                    >- NegotiateContextOffset: 0x0070
                      >- NegotiateContextCount: 4
                        >- Reserved: 0000
                          >- Dialect: 0x0210
                            >- Dialect: 0x0300
                              >- Dialect: 0x0302
                                >- Dialect: 0x0311
                                  >- Negotiate Context: SMB2_PREAUTH_INTEGRITY_CAPABILITIES
                                    >- Negotiate Context: SMB2_ENCRYPTION_CAPABILITIES
                                      >- Negotiate Context: Unknown Type: (0x5)
                                        >- Negotiate Context: SMB2_POSIX_EXTENSIONS_CAPABILITIES
                                          >- Type: SMB2_POSIX_EXTENSIONS_CAPABILITIES (0x0100)
                                            >- DataLength: 16
                                              >- Reserved: 00000000
                                                >- POSIX Reserved: 0x5025ad93
```





smb2 Expression...

| Source | Destination | Protoc | Length | Info |
|-----------|-------------|--------|--------|---|
| 127.0.0.1 | 127.0.0.1 | SMB2 | 170 | GetInfo Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 158 | Close Request File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 194 | Close Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 224 | Ioctl Request FSCTL_DFS_GET_REFERRALS, File: \\localhost\test |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 143 | Ioctl Response, Error: STATUS_NOT_FOUND |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 262 | Create Request File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 354 | Create Response File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 158 | Close Request File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 194 | Close Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 262 | Create Request File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 354 | Create Response File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 158 | Close Request File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 194 | Close Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 446 | Create Request File: ;GetInfo Request FILE_INFO/SMB2_FILE_ALL_INFO |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 534 | Create Response File: ;GetInfo Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 174 | GetInfo Request FILE_INFO/SMB2_FILE_ALL_INFO File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 244 | GetInfo Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 158 | Close Request File: |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 194 | Close Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 462 | Create Request File: .Trash;GetInfo Request FILE_INFO/SMB2_FILE_ALL_INFO; |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 310 | Create Response, Error: STATUS_OBJECT_NAME_NOT_FOUND;GetInfo Response, Er |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 470 | Create Request File: .Trash-1000;GetInfo Request FILE_INFO/SMB2_FILE_ALL |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 310 | Create Response, Error: STATUS_OBJECT_NAME_NOT_FOUND;GetInfo Response, Er |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 462 | Create Request File: 0760;GetInfo Request FILE_INFO/SMB2_FILE_ALL_INFO;C |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 310 | Create Response, Error: STATUS_OBJECT_NAME_NOT_FOUND;GetInfo Response, Er |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 246 | Create Request File: 0760 |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 310 | Create Response File: 0760 |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 158 | Close Request File: 0760 |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 194 | Close Response |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 462 | Create Request File: 0760;GetInfo Request FILE_INFO/SMB2_FILE_ALL_INFO;C |
| 127.0.0.1 | 127.0.0.1 | SMB2 | 678 | Create Response File: 0760;GetInfo Response;Close Response |

```

>-NetBIOS Session Service
  >-SMB2 (Server Message Block Protocol version 2)
    >-SMB2 Header
      >-Create Request (0x05)
        >-StructureSize: 0x0039
          -Oplock: No oplock (0x00)
          -Impersonation level: Impersonation (2)
          -Create Flags: 0x0000000000000000
          -Reserved: 0000000000000000
        >-Access Mask: 0x00000100
        >-File Attributes: 0x00000000
        >-Share Access: 0x00000007, Read, Write, Delete
        -Disposition: Create (if file exists fail, else create it) (2)
        >-Create Options: 0x00000001
      >-Filename: 0760
        -Blob Offset: 0x00000078
        -Blob Length: 8
        -Blob Offset: 0x00000088
        -Blob Length: 40
      >-ExtraInfo SMB2_POSIX_CREATE_CONTEXT
        >-Chain Element: SMB2_POSIX_CREATE_CONTEXT "5025ad93-b49c-e711-1
          -Chain Offset: 0x00000000
          >-Tag: 5025ad93-b49c-e711-b423-83de968bcd7c
            -Blob Offset: 0x00000010
            -Blob Length: 16
            -Blob Offset: 0x00000020
            -Blob Length: 4
          >-Data: POSIX Create Context request
            -POSIX perms: 0740
    
```

```

0000 00 00 00 00 00 00 00 00 00 00 00 00 08 00 45 00
0010 00 e8 c8 34 40 00 40 06 73 d9 7f 00 00 01 7f 00
0020 00 01 a0 cc 01 bd b5 6f c3 fc 26 6c fe 77 80 18
    
```

```
>- NetBIOS Session Service
  >- SMB2 (Server Message Block Protocol version 2)
    >- SMB2 Header
      >- Create Request (0x05)
        >- StructureSize: 0x0039
          | - Oplock: No oplock (0x00)
          | - Impersonation level: Impersonation (2)
          | - Create Flags: 0x0000000000000000
          | - Reserved: 0000000000000000
        >- Access Mask: 0x00000100
        >- File Attributes: 0x00000000
        >- Share Access: 0x00000007, Read, Write, Delete
          | - Disposition: Create (if file exists fail, else create it) (2)
        >- Create Options: 0x00000001
        >- Filename: 0760
          | - Blob Offset: 0x00000078
          | - Blob Length: 8
          | - Blob Offset: 0x00000088
          | - Blob Length: 40
        >- ExtraInfo SMB2_POSIX_CREATE_CONTEXT
          >- Chain Element: SMB2_POSIX_CREATE_CONTEXT "5025ad93-b49c-e711-l
            | - Chain Offset: 0x00000000
            >- Tag: 5025ad93-b49c-e711-b423-83de968bcd7c
              | - Blob Offset: 0x00000010
              | - Blob Length: 16
              | - Blob Offset: 0x00000020
              | - Blob Length: 4
            >- Data: POSIX Create Context request
              | - POSIX perms: 0740
```



What Next?
Missing Features and
More Linux optimizations ...



What Next?

- Examine the xfstest skips (and failures) in much detail and add small incremental changes
 - “xfstests” is the standard Linux fs functional test suite and no one file system can pass all tests due to various fs optional features.
 - Some can be emulated some need new flags
- Where that is not possible, consider adding new POSIX extensions version (simply adding additional uuid to the POSIX negotiate context)



What Next?

- What about minor extensions to reduce roundtrips and provide better/safer emulation?
 - Fcollapse and finsert are two examples
 - NTFS fsctls like FSCTL REARRANGE_FILE and **SHUFFLE_FILE** could help if available over SMB3
 - What about exposing Windows's FILE_FLAG_POSIX_SEMANTICS
 - More compounding can help too
 - Add fsctl for rename swap (rename exchange)?




Improving server side symlink handling

- When encountering client only symlinks (mfsymlinks) this is not an issue
- But for server side symlinks rely on server supporting reparse points to report symlinks
- This can be improved by simply improving the handling of “stopped on symlink” error (which includes sufficient information in most cases to avoid needing to query reparse points)



Examples from xfstest investigations

- Add support for renameat2 and rename exchange
- POSIX ACLs (can be emulated and there is pushback on implementing primitive POSIX ACLs)
- Support for additional chattr flags (“immutable” and “noatime” updates e.g.)
- fallocate –collapse-range
- Dedupe support
- Defragmentation support (may require VFS changes) 

Examples from xfstest investigations

- Richacl support (tests 362 through 370) ??
- O_TMPFILE support (emulatable, but VFS changes would help)
- FITRIM support (may be emulatable)
- Quota support (may be emulatable already)
- Support for NFS export (nfs server on smb3 mounts)
- Case sensitive xattrs (EAs)
- SELinux support



Examples from xfstest investigations

- Support for online ‘label manipulation’ (see e.g. [xfstest generic/492](#))
- Support for casefolding (“chattr +F”)
- Would native (rather than emulated) BSD flock (whole file lock) support help?



More details (with example xfstest #)

- atime options irrelevant (test 003)
- O_TMPFILE (generic/004)
- Defragmentation (018)
- Renameat2 (025)
- POSIX ACLs (026)
- FITRIM (038)
- Metadata journaling (049)
- Freezing fsctl (068) - <https://lwn.net/Articles/287435/>



More details (continued)

- Chattr +ia (079) (“immutable”, “append only”)
- Chattr +A (277) (“no atime updates”)
- Linux disk quotas (**082**)
- Security (093) and trusted (097) xattr namespaces
- preallocated extent not marked with
FIEMAP_EXTENT_UNWRITTEN (094)
- Dedupe (**121**)
- Advisory locks (**131**)



More details (continued)

- suid/sgid bits are cleared after direct write (test **355**)
- Richacl (362)
- Encryption support (395)
- Timestamp bounds unknown (402)
- chattr +d (“nodump”) (424)
- Information about fiemap of attribute fork (425)
- NFS export (open by inode #) (426)
- Backslash in name (“Key urk`?``?`moo does not exist for FAKESLASH test`??`” in test 453)



More details (continued)

- Conflicting xattrs (test **454**)
- XATTR_REPLACE (test **486**)
- xfs_io label (492)
- Lsattr -d (**508**)
- Xattrs with slashes in name (**523**)
- Casefolding support (**556**)
- Dupremove utility (**559**), acton utility (**596**)
- Fsverity (**571**)



Thank you for your time

- A very exciting time for ...

Linux[™]



+

**S
M
B
3
.
1
.
1**

