



SMB3.1.1 POSIX Protocol Extensions: Summary and Current Implementation Status

Steve French

Azure Storage – Microsoft

Samba Team

And

Jeremy Allison

Google/Samba Team



POSIX
SMB
3.1.1

Legal Statement

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

Outline

- Linux is a lot more than POSIX ...
- Why do these extensions matter?
- Implementation Status
- What works today?
- Some details
- How to handle Linux continuing to extend APIs?
- Wireshark and Tracing

Linux > POSIX

- Currently huge number of syscalls!

(try “git grep SYSCALL_DEFINE”

well over 850 and 500+ are

even documented “man syscalls”

FS layer has 223). Verified today

VS

- Only about 100 POSIX API calls

```
/stat.c:SYSCALL_DEFINE2(fstat64, unsigned long, fd, struct stat64 __user *  
/stat.c:SYSCALL_DEFINE4(fstatat64, int, dfd, const char __user *, filename,  
/stat.c:SYSCALL_DEFINE5(statx,  
/stat.c:COMPAT_SYSCALL_DEFINE2(newstat, const char __user *, filename,  
/stat.c:COMPAT_SYSCALL_DEFINE2(newlstat, const char __user *, filename,  
/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, struct st  
/statfs.c:SYSCALL_DEFINE3(statfs64, const char __user *, pathname, size_t,  
/statfs.c:SYSCALL_DEFINE2(fstatfs, unsigned int, fd, struct statfs __user  
/statfs.c:SYSCALL_DEFINE3(fstatfs64, unsigned int, fd, size_t, sz, struct  
/statfs.c:SYSCALL_DEFINE2(ustat, unsigned, dev, struct ustat __user *, ubu  
/statfs.c:COMPAT_SYSCALL_DEFINE2(statfs, const char __user *, pathname, st  
/statfs.c:COMPAT_SYSCALL_DEFINE2(fstatfs, unsigned int, fd, struct compat_  
/statfs.c:COMPAT_SYSCALL_DEFINE3(statfs64, const char __user *, pathname,  
  
/statfs.c:COMPAT_SYSCALL_DEFINE3(fstatfs64, unsigned int, fd, compat_size_  
/statfs.c:COMPAT_SYSCALL_DEFINE2(ustat, unsigned, dev, struct compat_ustat  
/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, loff_t,  
/sync.c:SYSCALL_DEFINE4(sync_file_range2, int, fd, unsigned int, flags,  
/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 __kernel_itim  
/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 *, filename,  
/utimes.c:SYSCALL_DEFINE3(futimesat, int, dfd, const char __user *, filename,  
/utimes.c:SYSCALL_DEFINE2(utimes, char __user *, filename,  
/utimes.c:SYSCALL_DEFINE2(utime, char __user *, filename, struct utimbuf __  
/utimes.c:SYSCALL_DEFINE2(utime32, const char __user *, filename,  
/utimes.c:SYSCALL_DEFINE4(utimensat_time32, unsigned int, dfd, const char  
t, flags)  
/utimes.c:SYSCALL_DEFINE3(futimesat_time32, unsigned int, dfd,  
/utimes.c:SYSCALL_DEFINE2(utimes_time32, const char __user *, filename, st  
/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 *, name,  
/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 *, name,  
/xattr.c:SYSCALL_DEFINE3(listxattr, const char __user *, pathname, char __  
/xattr.c:SYSCALL_DEFINE3(llistxattr, const char __user *, pathname, char __  
/xattr.c:SYSCALL_DEFINE3(flistxattr, int, fd, char __user *, list, size_t,  
/xattr.c:SYSCALL_DEFINE2(removexattr, const char __user *, pathname,  
/xattr.c:SYSCALL_DEFINE2(lremovexattr, const char __user *, pathname,  
/xattr.c:SYSCALL_DEFINE2(fremovexattr, int, fd, const char __user *, name)  
french@snfrench-ThinkPad-P52:~/cifs-2.6$ git grep SYSCALL_DEFINE | wc  
850      5070      69194
```

513 syscalls with man pages!

← → ↻ ⓘ Not secure | man7.org/linux/man-pages/dir_section_2.html

man7.org > Linux > man-pages

Linux/UNIX system programming |

Linux man pages: section 2

```
accept(2)
accept4(2)
access(2)
acct(2)
add_key(2)
adjtimex(2)
afs_syscall(2)
alarm(2)
alloc_hugepages(2)
arch_prctl(2)
arm_fadvise(2)
arm_fadvise64_64(2)
arm_sync_file_range(2)
bdflush(2)
bind(2)
bpf(2)
break(2)
brk(2)
cacheflush(2)
capget(2)
capset(2)
chdir(2)
chmod(2)
chown(2)
chown32(2)
chroot(2)
clock_getres(2)
clock_gettime(2)
clock_nanosleep(2)
clock_settime(2)
clone(2)
clone2(2)
__clone2(2)
clone3(2)
close(2)
connect(2)
copy_file_range(2)
creat(2)
create_module(2)
delete_module(2)
dup(2)
dup2(2)
dup3(2)
epoll_create(2)
epoll_create1(2)
epoll_ctl(2)
```

```
ioctl_xfs_bulkstat(2)
ioctl_xfs_fsbulkstat(2)
ioctl_xfs_fscounts(2)
ioctl_xfs_fsgetxattr(2)
ioctl_xfs_fsgetxattr(2)
ioctl_xfs_fsnumbers(2)
ioctl_xfs_fsop_geometry(2)
ioctl_xfs_fssetxattr(2)
ioctl_xfs_getbmap(2)
ioctl_xfs_getbmapa(2)
ioctl_xfs_getbmapx(2)
ioctl_xfs_getresblks(2)
ioctl_xfs_goingdown(2)
ioctl_xfs_inumbers(2)
ioctl_xfs_scrub_metadata(2)
ioctl_xfs_setresblks(2)
io_destroy(2)
io_getevents(2)
ioperm(2)
ioprio(2)
ioprio_get(2)
ioprio_set(2)
io_setup(2)
io_submit(2)
ipc(2)
isastream(2)
kcmp(2)
kexec_file_load(2)
kexec_load(2)
keyctl(2)
kill(2)
killpg(2)
lchown(2)
lchown32(2)
lgetxattr(2)
link(2)
linkat(2)
listen(2)
lstat(2)
lstatxattr(2)
llseek(2)
__llseek(2)
lock(2)
lookup_dcookie(2)
lremovexattr(2)
lseek(2)
```

```
rt_sigqueueinfo(2)
rt_sigreturn(2)
rt_sigsuspend(2)
rt_sigtimedwait(2)
rt_tgsigqueueinfo(2)
s390_guarded_storage(2)
s390_pci_mmio_read(2)
s390_pci_mmio_write(2)
s390_runtime_instr(2)
s390_sghi(2)
sbrk(2)
sched_getaffinity(2)
sched_getattr(2)
sched_getparam(2)
sched_get_priority_max(2)
sched_get_priority_min(2)
sched_getscheduler(2)
sched_rr_get_interval(2)
sched_setaffinity(2)
sched_setattr(2)
sched_setparam(2)
sched_setscheduler(2)
sched_yield(2)
seccomp(2)
security(2)
select(2)
select_tut(2)
semctl(2)
semget(2)
semop(2)
semtimedop(2)
send(2)
sendfile(2)
sendfile64(2)
sendmsg(2)
sendmmsg(2)
sendto(2)
setcontext(2)
setdomainname(2)
setgid(2)
setuid(2)
setfsuid(2)
setfsgid(2)
setfsgid32(2)
setfsuid(2)
setfsgid32(2)
setgid(2)
```

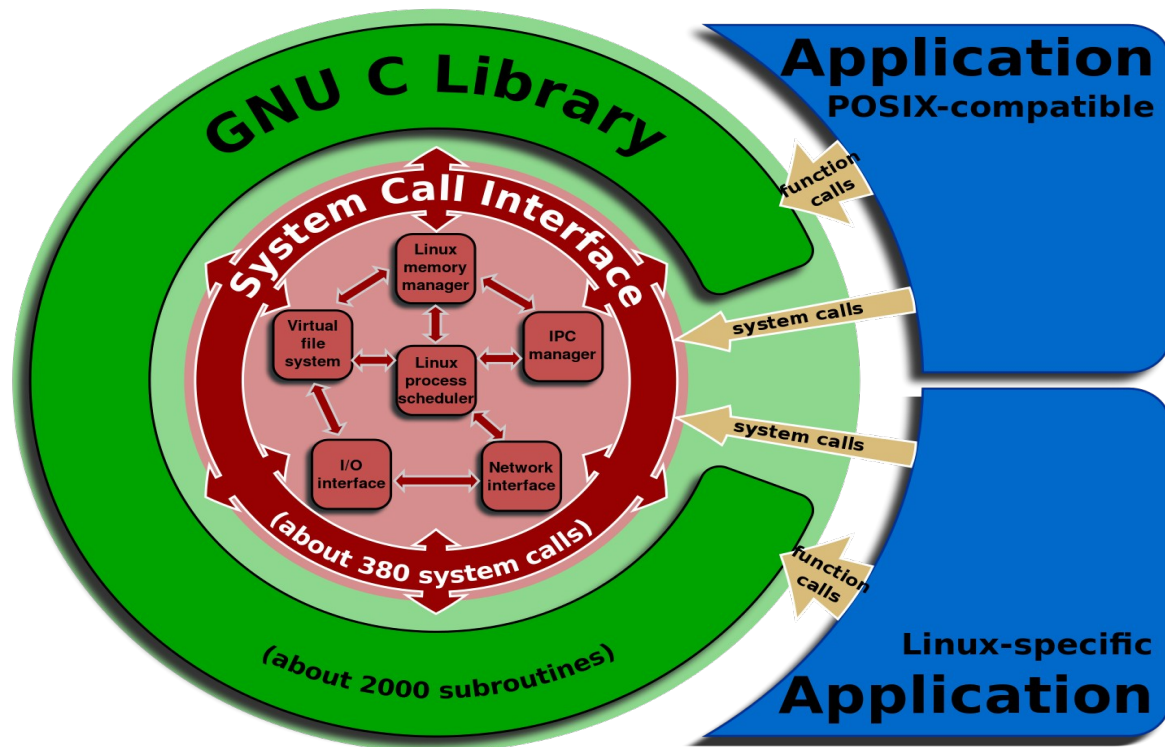
+12 just since last year's SDC!

Some examples of new fs ones from past 9 months ...

Syscall name	Kernel Version introduced
io_uring_enter	5.1
io_uring_register	5.1
io_uring_setup	5.1
move_mount	5.2
open_tree	5.2
fsconfig	5.2
fsmount	5.2
fsopen	5.2
fspick	5.2


Repeating an old slide ...

- Remember LINUX > POSIX




And not just new syscalls ... new flags ...

- ❑ 2 examples of richer Linux vs. simpler POSIX
- ❑ `fallocate` has 7 flags
 - Insert range
 - Unshare range
 - Zero range
 - Keep size
 - But POSIX `fallocate` has no flags
- ❑ `Rename` (`renameat2`) has 3 flags
 - `noreplace`, `whiteout` and `exchange`
 - POSIX `rename` has none



- **Network File systems matter**
 - these extensions to most popular network fs protocol (SMB3) are important
 - block devices struggle to do file system tasks: locking, security, leases, consistent metadata
- Linux Apps need to work over network mounts and continue to work as Linux evolves
- Improve common situations where customers have Linux and Windows and Mac clients
- Make sure extensions work with most secure, most optimal SMB3.1.1 dialect (don't encourage less secure network file systems, or even SMB1/CIFS)



Quick Overview of Status


- Linux kernel client:
 - 5.1 kernel or later can be used. Enable with mount option “posix”
 - Readdir, create, mkdir, statfs complete
 - POSIX queryinfo and support for new reparse tags for special files (in progress)
- Samba (experimental tree available, enable with smb.conf parm)
 - Server
 - All major features work. Merge delayed due to time consuming conflicts with other large changes. Special file handling (Sockets, FIFOs, char device handling) needs to be updated
 - Client tools (smbclient)
 - Major features work. Additional options could be added to cmd set
- SMB3 Kernel server (cifs's ksmbd.ko)
 - Partially implemented: it supports the POSIX negotiate context and partially parses POSIX open context
- 3rd party prototypes
- Wireshark patches available (network analysis)

Why Samba?


Since 1992 ...

Top Server on Linux

- Proven
- Broadly Implemented
- Extensible
- Secure
- Well Tested
- Implements richest File Protocol
- Enormous Client Base
(Mac, Windows, now better on Linux! And more)
- 3.5 Million LOC!








smb2

Source	Destination	Protocol	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 01 00 26 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 Mzu . % . . . w
00e0 95 f7 40 b3 00 00 02 00 06 00 00 00 00 02 00 . @ . . . .
00f0 02 00 01 00 00 00 05 00 12 00 00 00 00 06 00 . . . . .
0100 6f 00 63 00 61 00 6c 00 68 00 6f 00 73 00 74 00 o - c a l h o s
0110 00 00 00 00 00 00 00 01 10 00 00 00 00 93 ad . . . . #
0120 25 50 9c b4 11 e7 b4 23 83 de 96 8b cd 7c %P . . . .
```

- > 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

Source	Destination	Protocol	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  ....@.s...
0010 00 e8 c8 34 40 00 40 06 73 d9 7f 00 00 01 7f 00  ...4@.s...
0020 00 01 a0 cc 01 bd b5 6f c3 fc 26 6c fe 77 80 18  ....o..&l

```

```
>-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
```


Smbclient now has experimental support for SMB3.1.1 POSIX Extensions


```
smfrench@smfrench-ThinkPad-P52: ~  
smfrench@smfrench-ThinkPad-P52:~$ /usr/local/samba/bin/smbclient //localhost/test -U testuser  
Unable to initialize messaging context  
Enter SAMBA\testuser's password:  
Try "help" to get a list of possible commands.  
smb: \> posix  
SMB2 unix extensions supported  
smb: \> posix_mkdir newdir 0777  
posix_mkdir created directory \newdir  
smb: \> help  
?  
blocksize      cancel          case_sensitive  cd              chmod  
chown          close           del             deltree        dir  
du             echo            exit            get             getfacl  
geteas         hardlink        help            history         iosize  
lcd            link            lock            lowercase      ls  
l              mask            md              mget           mkdir  
more           mput            newer           notify          open  
posix          posix_encrypt   posix_open      posix_mkdir     posix_rmdir  
posix_unlink   posix_whoami    print           prompt          put  
pwd            q               queue           quit            readlink  
rd             recurse        reget           rename          reput  
rm             rmdir          showacls        setea           setmode  
scopy          stat            symlink         tar             tarmode  
timeout        translate      unlock          volume          vuid  
wdel           logon           listconnect     showconnect     tcon  
tdis           tid             utimes          logoff          ..  
!  
smb: \> posix_rmdir 0777  
Failed to unlink directory \0777. NT_STATUS_OBJECT_NAME_NOT_FOUND  
smb: \> posix_rmdir newdir  
posix_rmdir deleted directory \newdir  
smb: \> █
```

A year ago ... and now ...
kernel SMB3 client cifs.ko rapidly improving

- A year ago Linux 5.0-rc6 “Shy Crocodile”



Last month: 5.5 “Kleptomaniac Octopus”



2019 Linux FS/MM summit (in Puerto Rico in May)

- Great group of talented developers



Most Active Linux Filesystems this year

- 5338 kernel filesystem changesets last year (since 5.0-rc4 kernel) (flat)
 - FS activity: 6.2% of overall kernel changes (which are dominated by drivers) up slightly as % of activity
 - Kernel is huge (> 19 million lines of code, measured Saturday)
- There are many Linux file systems (>60), but six (and the VFS layer itself) drive 69% of activity (btrfs, xfs and cifs are the three most active)
 - File systems represent 4.8% of kernel source code (940KLOC) but among the most carefully watched areas
- cifs.ko (cifs/smb3 client) activity is strong
 - #3 most active of all fs with 375 changesets
 - 54.9KLOC (not counting user space cifs-utils which are 11.6KLOC and samba tools which are larger)

Linux File System Change Detail for past year (5.0-rc4 to 5.6)

- BTRFS 959 changesets (down slightly)
- VFS (overall fs mapping layer and common functions) 958 (up significantly)
- XFS 581 (up)
- CIFS/SMB2/SMB3 client 375 (since 4.18 kernel activity has gone way up)
- F2FS 271 (down)
- NFS client 283 (down)
- Others: EXT4 243(up), Ceph 164, AFS 135, GFS2 118, OCFS2 108 ...
- NFS server 137 (flat). Linux NFS server **MUCH** smaller than CIFS or Samba
- NB: Samba is as active as all Linux file systems put together (> 5000 changesets last year, over 1200 in the file server component alone!) - broader in scope (by a lot) and also is user space not kernel. Just counting the Samba testcases they are larger than many file servers. **Samba now 3.4 million lines of code**

For cifs.ko (kernel client) to experiment with POSIX
Extensions use 5.1 kernel or backport patch below

- commit 0d481325a9e5e3a31bf83bfcd3690a7a7152ece1

Author: Steve French <stfrench@microsoft.com>

Date: Sun Feb 24 17:56:33 2019 -0600

smb3: Update POSIX negotiate context with POSIX ctxt GUID

Recent example (cifs.ko over SMB3.1.1 POSIX) readdir (ls) now works

```
root@smfrench-ThinkPad-P52: ~  
root@smfrench-ThinkPad-P52:~# mount -t cifs //localhost/test /mnt1 -o username=smfrench,posix  
Password for smfrench@//localhost/test: *****  
root@smfrench-ThinkPad-P52:~# ls /mnt1  
0444-remote 0741 0750dir 0754 0755 0765dir 1GB 1GBa subdir test test0444file test0740file test0754file  
root@smfrench-ThinkPad-P52:~# ls /mnt1 -la  
total 2048006  
drwxr-xr-x  2 root root      0 Feb 11 13:35 .  
drwxr-xr-x 25 root root    4096 Feb 13 23:50 ..  
-rwxrwxrwx  1 root root 1048576 Feb 10 10:20 0444-remote  
-rwxr---x  1 root root      0 Feb  6 07:39 0741  
drwxr-x---  2 root root      0 Feb  4 03:29 0750dir  
-rwxr-xr--  1 root root      0 Feb  6 07:31 0754  
-rwxr-xr-x  1 root root      0 Feb  6 07:32 0755  
drwxrwr-r-x  2 root root      0 Feb  3 17:06 0765dir  
-rwxrwxrwx  1 root root 1049624576 Feb 11 13:48 1GB  
-rw-r--r--  1 root root 1048576000 Feb  5 06:42 1GBa  
drwxrwxrwx  2 root root      0 Feb  7 13:33 subdir  
drwxrwxrwx  2 root root      0 Feb 10 10:16 test  
-r--r--r--  1 root root 1048576 Feb 10 10:16 test0444file  
-rwxr-----  1 root root      0 Feb 10 10:17 test0740file  
-rwxr-xr--  1 root root 1048576 Feb 10 10:15 test0754file  
root@smfrench-ThinkPad-P52:~#
```

Another example cifs.ko over SMB3.1.1 POSIX:
mode bits preserved as expected on mkdir

```
root@fedora29:/mnt
[root@fedora29 mnt]# mkdir -m 0764 0764dir
[root@fedora29 mnt]# uname -a
Linux fedora29 5.5.0+ #1 SMP Tue Jan 28 16:35:20 CST 2020 x86_64 x86_64 x86_64 GNU/Linux
[root@fedora29 mnt]#

root@jra-posix:~
rmdir: failed to remove '0764dir': No such file or directory
[root@jra-posix ~]# rmdir /data/test/0764dir
[root@jra-posix ~]# ls /data/test -la | grep 0764dir
drwxrw-r-- 2 nobody nobody 6 Feb 4 06:11 0764dir
[root@jra-posix ~]#
```


Quality Much Improved – Top Priority

- ❑ More xfstests pass (> 150 and growing) even without POSIX extensions, vast majority of the rest are skipped due to missing features or being inappropriate for network file systems
- ❑ Many potential issues pointed out by static analysis addressed
- ❑ Starting 15 months ago **The “Buildbot”** ... reducing regressions. **VERY** exciting addition for CIT (thanks Ronnie, Aurelien and Paulo)
- ❑ POSIX Extensions (jra's tree) now a buildbot target for automated regression tests. Will be expanding test list run vs. it soon ...

POSIX: What could you try today?

- For obvious reasons these experimental changes not enabled by default:
 - With current mainline Linux (5.1 or later) specify new mount option “posix” and turn off remapping of reserved characters (ie append “nomapposix” to mount options)
 - Special file handling in stat output can return incorrect data (in progress)
- JRA has a tree on [samba.org](https://git.samba.org/jra/samba/.git) (git.samba.org/jra/samba/.git in branch “master-smb2”) with prototype server code
- Other vendors testing experimental distinct implementations of POSIX extensions as well. And cifs kernel server also has partial support for these

Why Isn't This Shipped Already ?

- ❑ Details are still being explored
 - Recently discovered that WSL (Windows Subsystem for Linux) defines methods of exporting special file types in reparse point tags that may be faster (and useful to be compatible with too)
- ❑ Samba VFS needs updating to match modern Linux/OpenGroup/POSIX APIs.
 - Samba server must change first to use Xxxat() (openat, mkdirat, linkat..) system calls before we can integrate this into Samba master.

What has changed ?

- In order to optimize readdir better (to avoid extra roundtrips for each file):

Symlinks and Special Files (char, block, fifo, socket) are now distinct reparse point tags (similar to what WSL does) rather than one tag for all (which requires an extra query for each special file)

New Proposed Tags for Special Files

IO_REPARSE_TAG_LX_SYMLINK

Symbolic link. Tag value is 0xA000001D.

IO_REPARSE_TAG_AF_UNIX

UNIX domain socket. Tag value is 0x80000023.

IO_REPARSE_TAG_LX_FIFO

FIFO. Tag value is 0x80000024.

IO_REPARSE_TAG_LX_CHR

Character special file. Tag value is 0x80000025.

IO_REPARSE_TAG_LX_BLK

Block special file. Tag value is 0x80000026.

Example: using it today

- On the client:
 - “mount -t smb3 //<address>/<share> /mnt -o
username=<user>,password=<pass>,
vers=3.1.1,posix,mfsymlinks,nomapposix,noperm
- On the server add to smb.conf:
 - smb2 unix extensions = yes
 - “mangled names = no”
 - “directory mask = 07777”
 - “create mask = 07777”
 - Consider removing “obey pam restrictions”

Note that directory enumeration returns a reparse point tag (for file type).

- The proposed change to add tags this would reduce roundtrips

```
1 0.0000... 127.0.0.1 127.0.0... SMB2 470 Create Request File: filetypes;GetInfo Request FILE
2 0.0015... 127.0.0.1 127.0.0... SMB2 678 Create Response File: filetypes;GetInfo Response;C
4 0.0018... 127.0.0.1 127.0.0... SMB2 278 Create Request File: filetypes
5 0.0032... 127.0.0.1 127.0.0... SMB2 354 Create Response File: filetypes
6 0.0033... 127.0.0.1 127.0.0... SMB2 168 Find Request File: filetypes SMB2_FIND_ID_FULL_DIRE
7 0.0038... 127.0.0.1 127.0.0... SMB2 1002 Find Response
8 0.0040... 127.0.0.1 127.0.0... SMB2 486 Create Request File: filetypes/blockfile;GetInfo Re
9 0.0052... 127.0.0.1 127.0.0... SMB2 702 Create Response File: filetypes/blockfile;GetInfo F
```

```
› FileIdBothDirectoryInfo: blockfile
› FileIdBothDirectoryInfo: fifo
› FileIdBothDirectoryInfo: file
› FileIdBothDirectoryInfo: dir0777
› FileIdBothDirectoryInfo: charfile
  Next Offset: 96
  File Index: 0x00000000
  Create: Sep 23, 2019 15:48:24.451803100 CDT
  Last Access: Sep 23, 2019 15:48:24.451803100 CDT
  Last Write: Sep 23, 2019 15:48:24.451803100 CDT
  Last Change: Sep 23, 2019 15:48:24.451803100 CDT
  End Of File: 0
  Allocation Size: 0
› File Attributes: 0x00000400
  Filename Length: 16
  EA Size: 2147483658
```

```
0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 04 .....
0 00 00 10 00 00 00 0a 00 00 80 00 00 00 00 c1 0f .....
0 68 00 00 00 00 00 62 00 68 00 61 00 70 00 66 00 .....
                                b o x e f
```

```
root@smf-Thinkpad-P51: ~/posix/bin
root@smf-Thinkpad-P51:~/posix/bin# ls /mnt/filetypes
blockfile charfile dir0777 fifo file somesocket symlink-to-file
```

SMB3.1.1 POSIX Query Directory

smb2						
No.	Time	Source	Destination	Protocol	Length	Info
1	0.0000...	127.0.0.1	127.0.0...	SMB2	470	Create Request File: filetypes;GetInfo Request FILE_INFO/SMB2_
3	0.0015...	127.0.0.1	127.0.0...	SMB2	678	Create Response File: filetypes;GetInfo Response;Close Respons
5	0.0017...	127.0.0.1	127.0.0...	SMB2	278	Create Request File: filetypes
7	0.0031...	127.0.0.1	127.0.0...	SMB2	354	Create Response File: filetypes
9	0.0033...	127.0.0.1	127.0.0...	SMB2	168	Find Request File: filetypes (Level:0x64) Pattern: *
11	0.0058...	127.0.0.1	127.0.0...	SMB2	1406	Find Response
13	0.0059...	127.0.0.1	127.0.0...	SMB2	168	Find Request File: filetypes (Level:0x64) Pattern: *
14	0.0061...	127.0.0.1	127.0.0...	SMB2	143	Find Response, Error: STATUS_NO_MORE_FILES
16	0.0063...	127.0.0.1	127.0.0...	SMB2	158	Close Request File: filetypes

Frame 11: 1406 bytes on wire (11248 bits), 1406 bytes captured (11248 bits) on interface 0

Ethernet II, Src: 00:00:00_00:00:00 (00:00:00:00:00:00), Dst: 00:00:00_00:00:00 (00:00:00:00:00:00)

Internet Protocol Version 4, Src: 127.0.0.1, Dst: 127.0.0.1

Transmission Control Protocol, Src Port: 445, Dst Port: 47916, Seq: 901, Ack: 719, Len: 1340

NetBIOS Session Service

SMB2 (Server Message Block Protocol version 2)

▸ SMB2 Header

▾ Find Response (0x0e)

[Info Level: Unknown (100)]

▸ StructureSize: 0x0009

Blob Offset: 0x00000048

Blob Length: 1264

▾ Info: 80000000000000000056f4d36f5072d501a6c8b5705072d501...

Unknown: 80000000000000000056f4d36f5072d501a6c8b5705072d501...

0000	00 00 00 00 00 00 00 00	00 00 00 00 08 00	45 00E.
0010	05 70 ad 00 40 00 40 06	8a 85 7f 00 00 01 7f 00		·p·@·@·.....
0020	00 01 01 bd bb 2c d2 61	74 d6 bd 6a 42 d5 80 18	,a t·jB·
0030	02 00 03 65 00 00 01 01	08 0a d2 89 be d0 d2 89		...e·.....
0040	be ce 00 00 05 38 fe 53	4d 42 40 00 01 00 00 00	8·S MB@.....
0050	00 00 0e 00 0a 00 01 00	00 00 00 00 00 00 00 00	1a 00
0060	00 00 00 00 00 00 5f 7b	00 00 6b 0c fb c1 63 46	_{ ··k··cF

Summary of What works

- Without Extensions:
mode bits emulated, client followed
symlinks, hardlinks
- With Extensions
allows posix unlink and rename semantics, case
sensitivity

Example (w/o POSIX extensions) (thank you to Aurelien at SuSE!)

- ❑ “modefromsid” and “idsfromsid” mount options
- ❑ Mode bits (from special ACE
 - Mode bits on left
 - File name on right
- ❑ NB mkfifo not finished
- ❑ Ownership works too

444 mnt/smb/0444

540 mnt/smb/0540

777 mnt/smb/0777

1777 mnt/smb/1777

2777 mnt/smb/2777

3777 mnt/smb/3777

1777 mnt/smb/d01777

2777 mnt/smb/d02777

444 mnt/smb/d0444

4777 mnt/smb/d04777

540 mnt/smb/d0540

777 mnt/smb/d0777

644 mnt/smb/emptyfile

stat: cannot stat 'mnt/smb/fifo': Operation not supported

644 mnt/smb/file-as-sfrench

- What about the Linux Kernel?
 - New API changes added about once a year to the VFS (minor global changes added more often, but not all could affect what we need to send on the wire in perfect world ...)
 - Need to quickly update protocol when not possible to do over SMB3
 - Need better interaction with key communities (containers, databases and many others) about what they would like to see added to the protocol or emulated better

SMB3 POSIX Extensions

- Negotiate Protocol
 - SMB3.1.1 (or later required)
 - POSIX Negotiate Context 0x100
 - Version is included in the context by including the GUID of the supported POSIX open context(s) – currently only one supported
 - If POSIX open contexts not supported, negotiate context must be ignored
 - If POSIX open contexts supported for some files then negotiate context is returned, but server must fail opens with POSIX contexts for files where POSIX is not supported (rather than ignoring the POSIX context)
- Tree Connect – in future dialects tree connect contexts may allow more granularity in allowing servers to tell clients which shares they can't use POSIX opens on
- Case sensitivity yes/no can be exposed via existing QFS Info call

POSIX Extension Requirements

- If server returns a POSIX create context on an open:
 - It supports case sensitive names on this path
 - It supports POSIX unlink/rename semantics on this file
 - It supports advisory (POSIX) locking on this file.
 - Actually they are “OFD” not “POSIX” locks (see e.g. <https://gavv.github.io/blog/file-locks/#emulating-open-file-description-locks>)
 - PATH names are not remapped (no SFU remap needed for * and \ and > and < and : ...). UCS2 converted directly to UTF-8 and server supports POSIX pathnames
 - It is permitted to have POSIX opens and non-POSIX opens on the same file and also the server is allowed to forbid posix extensions on some files (or across the whole share) but allow it on others

We Leverage Existing SMB3 features

- Hardlinks use Windows setinfo call (long ago implemented)
- ~~Symlinks, mkfifo, mknod use “nfs reparse point” (MS-FSCC 2.1.2.6)~~ Distinct reparse point tags for each type of special file (ala what “WSL” does) - allows us to better optimize readdir
- ACE with special SID (with mode at end) ala “NFS ACL” mapping can be used to set mode (SID: S-1-5-88-3) see <http://people.redhat.com/steved/Bakeathon-2010/SDC2010-NFS-Windows.jbiseda.20100921.pdf>
- Other linux extensions, e.g. fallocation are mapped to existing SMB3 operations where possible

Format of the POSIX owner and mid information in the ACL

<NTSecurityAuthority>-<SECURITY_NFS_ID_BASE_RID>-
<NfsSidType>-<NfsSidValue>

✈ Owner SID for UID: "S-1-5-88-1-<uid>"

✈ Group SID for GID: "S-1-5-88-2-<gid>"

✈ Mode SID: "S-1-5-88-3-<mode>"

✈ Everyone: "S-1-5-88-4"

POSIX Infolevels

- Query/SetInfo and Query_DIR (and also FSInfo)
 - Level 0x64 SMB2_FIND_POSIX_INFORMATION
 - Payload variable (Max = 216 bytes)
 - Timestamps
 - File size
 - Dos attributes
 - U64 Inode number
 - U32 device id
 - U32 zero
 - Struct posix_create_context_response

Wireshark


- See Aurelien's dissector improvements
 - <https://github.com/aaptel/wireshark/commits/smb3unix>
 - And Pike sample test code
 - <https://github.com/aaptel/pike/tree/smb3unix>

Next Steps

- Continue debugging test implementations (cifs.ko and JRAs Samba POSIX test branch). Current focus: enhancing smb3 client to better handle POSIX stat (getattr)
- Continue to add xfstests to the 'jraposix' test group in the buildbot (to regression test the client against Samba server with POSIX extensions)
- Examine EVERY xfstest skip and every xfstest fail for potential match to features in (or to add to) SMB3.1.1 POSIX Extensions
 - e.g. enabling special files (fifos, blkdevs etc.) as reparse points enables five tests.
- Continue extending the wireshark dissectors (see Aurelien)
- Continue updating the wiki with details:
<https://wiki.samba.org/index.php/SMB3-Linux>
- Questions/comments welcome: samba-technical and linux-cifs lists

Thank you for your time

- This is a very exciting time for ...



S
+
M
B
3

POSIX Path Names Work

File Edit View Search Terminal Help

```
root@smf-Thinkpad-P51:~# ls /posix-extensions-mount/
d0754      fileacolon:      'fileexclamation!'      hello
'file!'    'fileasterisk*'  'filequestion?'         newfile

root@smf-Thinkpad-P51:~# ls /scratch
d0754      fileacolon:      'fileexclamation!'      hello
'file!'    'fileasterisk*'  'filequestion?'         newfile

root@smf-Thinkpad-P51:~# uname -a
Linux smf-Thinkpad-P51 5.0.0-rc4+ #67 SMP Sun Jan 27 20:49:32
019 x86_64 x86_64 x86_64 GNU/Linux

root@smf-Thinkpad-P51:~# ~/posix/bin/smbd -V
Version 4.10.0pre1-DEVELOPERBUILD

root@smf-Thinkpad-P51:~#
```

Mode bits on create and case sensitivity work!

```
root@Ubuntu-17-Virtual-Machine:/mnt# ~/create-4-files-with-mode-test
root@Ubuntu-17-Virtual-Machine:/mnt# cd /mnt1
root@Ubuntu-17-Virtual-Machine:/mnt1# ~/create-4-files-with-mode-test
root@Ubuntu-17-Virtual-Machine:/mnt1# ls /test /test-no-posix -la
/test:
total 12
drwxrwxrwx  3 root    root    4096 May 31 16:55 █
drwxr-xr-x 32 root    root    4096 May 31 16:46 ..
-rwx----- 1 testuser testuser  0 May 31 16:55 0700
-rwxrwx---  1 testuser testuser  0 May 31 16:55 0770
-rwxrwxr-x  1 testuser testuser  0 May 31 16:55 0775
drwxr-xr-x  2 sfrench sfrench 4096 Mar 24 10:34 tmp

/test-no-posix:
total 8
drwxrwxrwx  2 root    root    4096 May 31 16:55 █
drwxr-xr-x 32 root    root    4096 May 31 16:46 ..
-rwxrw-r--  1 testuser testuser  0 May 31 16:55 0700
-rwxrw-r--  1 testuser testuser  0 May 31 16:55 0770
-rwxrw-r--  1 testuser testuser  0 May 31 16:55 0775
root@Ubuntu-17-Virtual-Machine:/mnt1# mkdir UPPER
root@Ubuntu-17-Virtual-Machine:/mnt1# touch upper
root@Ubuntu-17-Virtual-Machine:/mnt1# cd /mnt
root@Ubuntu-17-Virtual-Machine:/mnt# mkdir UPPER
root@Ubuntu-17-Virtual-Machine:/mnt# touch upper
root@Ubuntu-17-Virtual-Machine:/mnt# ls /test /test-no-posix
/test:
0700 0770 0775 tmp upper UPPER

/test-no-posix:
0700 0770 0775 UPPER
```

Rename works with POSIX extensions!

```
root@Ubuntu-17-Virtual-Machine: ~  
File Edit View Search Terminal Help  
root@Ubuntu-17-Virtual-Machine:~# ls /mnt-rename-test -la  
total 2052  
drwxr-xr-x  2 root root    0 May 31 18:19 .  
drwxr-xr-x 34 root root 4096 May 31 18:13 ..  
-rwxr-xr-x  1 root root    0 May 31 18:18 emptyfile  
-rwxr-xr-x  1 root root    0 May 31 18:19 emptyfile-posix  
-rwxr-xr-x  1 root root   16 May 31 18:17 targetfile  
-rwxr-xr-x  1 root root   16 May 31 18:19 targetfile-posix  
root@Ubuntu-17-Virtual-Machine:~# mount | grep rename  
//localhost/rename-test on /mnt-rename-test type cifs (rw,relatime,vers=3.1.1,cache=strict,username=testuser,doma,uid=0,noforceuid,gid=0,noforcegid,addr=127.0.0.1,file_mode=0755,dir_mode=0755,soft,posix,posixpaths,serverino,mapposix,rsz,echo_interval=60,actimeo=1)  
root@Ubuntu-17-Virtual-Machine:~# mv /mnt-rename-test/emptyfile /mnt-rename-test/targetfile  
mv: cannot move '/mnt-rename-test/emptyfile' to '/mnt-rename-test/targetfile': Permission denied
```

```
root@Ubuntu-17-Virtual-Machine: ~  
File Edit View Search Terminal Help  
root@Ubuntu-17-Virtual-Machine:~# mount | grep rename  
//localhost/rename-test on /mnt-rename-test type cifs (rw,relatime,vers=3.1.1,cache=strict,username=testuser,doma,uid=0,noforceuid,gid=0,noforcegid,addr=127.0.0.1,file_mode=0755,dir_mode=0755,soft,posix,posixpaths,serverino,mapposix,rsz,echo_interval=60,actimeo=1)  
root@Ubuntu-17-Virtual-Machine:~# mv /mnt-rename-test/emptyfile-posix /mnt-rename-test/targetfile-posix  
root@Ubuntu-17-Virtual-Machine:~#
```

Statfs ("stat -f") without POSIX extensions:

```
root@smf-Thinkpad-P51:~/cifs-2.6# cat /proc/mounts | grep cifs
//localhost/scratch /mnt cifs rw,relatime,vers=3.0,cache=strict,username=testuser,domain=,uid=0,noforceuid,gid=0,noforcegid,addr=127.0.0.1,file_mode=0755,dir_mode=0755,soft,nounix,serverino,mfsymlinks,noperm,rsize=1048576,wsiz=1048576,echo_interval=60,actimeo=1 0
root@smf-Thinkpad-P51:~/cifs-2.6# stat -f /mnt
  File: "/mnt"
    ID: 0          Namelen: 4096    Type: smb2
Block size: 1024      Fundamental block size: 1024
Blocks: Total: 234804176  Free: 28323720   Available: 28323720
Inodes: Total: 0        Free: 0
root@smf-Thinkpad-P51:~/cifs-2.6# stat -f /scratch
  File: "/scratch"
    ID: e94471edc7140504 Namelen: 255    Type: ext2/ext3
Block size: 4096      Fundamental block size: 4096
Blocks: Total: 58701044  Free: 10080212   Available: 7080929
Inodes: Total: 14983168  Free: 13901548
```

Statfs ("stat -f") with POSIX extensions – works!

```
root@smf-Thinkpad-P51:~/cifs-2.6# cat /proc/mounts | grep smb3
//127.0.0.1/scratch /mnt1 smb3 rw,relatime,vers=3.1.1,cache=strict,username=testuser,doma
in=,uid=0,noforceuid,gid=0,noforcegid,addr=127.0.0.1,file_mode=0755,dir_mode=0755,soft,p
six,posixpaths,serverino,mapposix,noperm,rsize=1048576,wsiz=1048576,echo_interval=60,act
imeo=1 0 0
root@smf-Thinkpad-P51:~/cifs-2.6# stat -f /mnt1
File: "/mnt1"
  ID: 0          Namelen: 4096      Type: smb2
Block size: 4096      Fundamental block size: 4096
Blocks: Total: 58701044   Free: 10080249   Available: 7080966
Inodes: Total: 14983168   Free: 13901538
root@smf-Thinkpad-P51:~/cifs-2.6# stat -f /scratch
File: "/scratch"
  ID: e94471edc7140504 Namelen: 255      Type: ext2/ext3
Block size: 4096      Fundamental block size: 4096
Blocks: Total: 58701044   Free: 10080127   Available: 7080844
Inodes: Total: 14983168   Free: 13901536
```