

[PATCH] sendfile removal

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-05/msg13684.html>

- *From:* Jens Axboe <jens.axboe@xxxxxxxxxx>
 - *Date:* Thu, 31 May 2007 12:33:16 +0200
-

Hi,

This patch removes the `->sendfile()` hook from the `file_operations` structure, and replaces the `sys_sendfile()` mechanism to be based on `->splice_read()` instead. There should be no functional changes.

Work to be done:

- The ext2 xip support needs a `splice_read()` implementation, currently I just if 0'ed out the `send_xip_file_sendfile()`. CC'ed Carsten, who seems to be the author of this code.
- shmem needs a `splice_read()` implementation. Optimistically CC'ed Hugh.
- nfsd: The `->rq_sendfile_ok` optimization is gone for now. I can't determine the value of it, but I'm assuming it's there for a reason. Any chance this can be converted to splice, or use something else than `->sendfile()`? CC'ed Neil.
- relay: needs a `splice_read()` implementation. I think Tom already has one, CC'ed him.

Apart from that, it was mostly straight forward. Almost everybody uses `generic_file_sendfile()`, which makes the conversion easy. I changed loop to use `do_generic_file_read()` instead of `sendfile`, it works for me...

Patch is against current git.

```
drivers/block/loop.c | 22 ++++++++-----  
fs/adfs/file.c | 2 -  
fs/affs/file.c | 2 -  
fs/afs/file.c | 2 -  
fs/bad_inode.c | 7 -----  
fs/bfs/file.c | 2 -  
fs/block_dev.c | 1  
fs/cifs/cifsfs.c | 8 +++---  
fs/coda/file.c | 11 ++++----  
fs/ecryptfs/file.c | 15 +++++-----  
fs/ext2/file.c | 5 +-+
```

[PATCH] sendfile removal

```
fs/ext3/file.c | 1
fs/ext4/file.c | 1
fs/fat/file.c | 2 -
fs/fuse/file.c | 4 +-
fs/gfs2/ops_file.c | 1
fs/hfs/inode.c | 2 -
fs/hfsplus/inode.c | 2 -
fs/hostfs/hostfs_kern.c | 2 -
fs/hpfs/file.c | 2 -
fs/jffs2/file.c | 2 -
fs/jfs/file.c | 1
fs/minix/file.c | 2 -
fs/nfs/file.c | 15 ++++++-----
fs/nfsd/vfs.c | 19 ++++++-----
fs/ntfs/file.c | 2 -
fs/ocfs2/file.c | 1
fs/qnx4/file.c | 2 -
fs/ramfs/file-mm.c | 2 -
fs/ramfs/file-nommu.c | 2 -
fs/read_write.c | 13 ++++++---
fs/reiserfs/file.c | 1
fs/smbfs/file.c | 9 +++----
fs/sysv/file.c | 2 -
fs/udf/file.c | 2 -
fs/ufs/file.c | 2 -
fs/xfs/linux-2.6/xfs_file.c | 26 -----
fs/xfs/linux-2.6/xfs_linux.h | 1
fs/xfs/linux-2.6/xfs_lrw.c | 44 -----
fs/xfs/linux-2.6/xfs_lrw.h | 3 --
fs/xfs/linux-2.6/xfs_vnode.h | 6 -----
fs/xfs/xfs_vnodeops.c | 3 ---
include/linux/fs.h | 5 ----
include/linux/sunrpc/svc.h | 2 -
kernel/relay.c | 16 -----
mm/filemap.c | 20 -----
mm/filemap_xip.c | 2 +
mm/shmem.c | 32 ++++++++-----
net/sunrpc/auth_gss/svcauth_gss.c | 2 -
net/sunrpc/svc.c | 2 -
50 files changed, 105 insertions(+), 230 deletions(-)
```

```
diff --git a/drivers/block/loop.c b/drivers/block/loop.c
index 5526ead..92bac14 100644
--- a/drivers/block/loop.c
+++ b/drivers/block/loop.c
@@ -435,16 +435,24 @@ do_lo_receive(struct loop_device *lo,
{
struct lo_read_data cookie;
struct file *file;
- int retval;
+ read_descriptor_t desc;
```

[PATCH] sendfile removal

```
+
+ desc.written = 0;
+ desc.count = bvec->bv_len;
+ desc.arg.data = &cookie;
+ desc.error = 0;

cookie.lo = lo;
cookie.page = bvec->bv_page;
cookie.offset = bvec->bv_offset;
cookie.bsize = bsize;
file = lo->lo_backing_file;
- retval = file->f_op->sendfile(file, &pos, bvec->bv_len,
- lo_read_actor, &cookie);
- return (retval < 0)? retval: 0;
+
+ do_generic_file_read(file, &pos, &desc, lo_read_actor);
+ if (desc.written)
+ return 0;
+
+ return desc.error;
}

static int
@@ -679,8 +687,8 @@ static int loop_change_fd(struct loop_device *lo, struct file *lo_file,
if (!S_ISREG(inode->i_mode) && !S_ISBLK(inode->i_mode))
goto out_putf;

- /* new backing store needs to support loop (eg sendfile) */
- if (!inode->i_fop->sendfile)
+ /* new backing store needs to support loop (eg readpage) */
+ if (!file->f_mapping->a_ops->readpage)
goto out_putf;

/* size of the new backing store needs to be the same */
@@ -760,7 +768,7 @@ static int loop_set_fd(struct loop_device *lo, struct file *lo_file,
* If we can't read - sorry. If we only can't write - well,
* it's going to be read-only.
*/
- if (!file->f_op->sendfile)
+ if (!aops->readpage)
goto out_putf;
if (aops->prepare_write && aops->commit_write)
lo_flags |= LO_FLAGS_USE_AOPS;
diff --git a/fs/adfs/file.c b/fs/adfs/file.c
index f544a28..36e381c 100644
--- a/fs/adfs/file.c
+++ b/fs/adfs/file.c
@@ -33,7 +33,7 @@ const struct file_operations adfs_file_operations = {
.fsync = file_fsync,
.write = do_sync_write,
.aio_write = generic_file_aio_write,
```

[PATCH] sendfile removal

```
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
};

const struct inode_operations adfs_file_inode_operations = {
diff --git a/fs/affs/file.c b/fs/affs/file.c
index c879690..c314a35 100644
--- a/fs/affs/file.c
+++ b/fs/affs/file.c
@@ -35,7 +35,7 @@ const struct file_operations affs_file_operations = {
 .open = affs_file_open,
 .release = affs_file_release,
 .fsync = file_fsync,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
};

const struct inode_operations affs_file_inode_operations = {
diff --git a/fs/afs/file.c b/fs/afs/file.c
index 9c0e721..aede7eb 100644
--- a/fs/afs/file.c
+++ b/fs/afs/file.c
@@ -32,7 +32,7 @@ const struct file_operations afs_file_operations = {
 .aio_read = generic_file_aio_read,
 .aio_write = afs_file_write,
 .mmap = generic_file_readonly_mmap,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
 .fsync = afs_fsync,
};

diff --git a/fs/bad_inode.c b/fs/bad_inode.c
index 329ee47..521ff7c 100644
--- a/fs/bad_inode.c
+++ b/fs/bad_inode.c
@@ -114,12 +114,6 @@ static int bad_file_lock(struct file *file, int cmd, struct file_lock *fl)
return -EIO;
}

-static ssize_t bad_file_sendfile(struct file *in_file, loff_t *ppos,
- size_t count, read_actor_t actor, void *target)
-{
- return -EIO;
-}
-
static ssize_t bad_file_sendpage(struct file *file, struct page *page,
int off, size_t len, loff_t *ppos, int more)
{
@@ -182,7 +176,6 @@ static const struct file_operations bad_file_ops =
.aio_fsync = bad_file_aio_fsync,
.fsync = bad_file_fsync,
```

[PATCH] sendfile removal

[PATCH] sendfile removal

```
.lock = bad_file_lock,
- .sendfile = bad_file_sendfile,
.sendpage = bad_file_sendpage,
.get_unmapped_area = bad_file_get_unmapped_area,
.check_flags = bad_file_check_flags,
diff --git a/fs/bfs/file.c b/fs/bfs/file.c
index ef4d1fa..24310e9 100644
--- a/fs/bfs/file.c
+++ b/fs/bfs/file.c
@@ -24,7 +24,7 @@ const struct file_operations bfs_file_operations = {
.write = do_sync_write,
.aio_write = generic_file_aio_write,
.mmap = generic_file_mmap,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
};

static int bfs_move_block(unsigned long from, unsigned long to, struct super_block *sb)
diff --git a/fs/block_dev.c b/fs/block_dev.c
index ea1480a..b3e9bfa 100644
--- a/fs/block_dev.c
+++ b/fs/block_dev.c
@@ -1346,7 +1346,6 @@ const struct file_operations def_blk_fops = {
#ifdef CONFIG_COMPAT
.compat_ioctl = compat_blkdev_ioctl,
#endif
- .sendfile = generic_file_sendfile,
.splice_read = generic_file_splice_read,
.splice_write = generic_file_splice_write,
};
diff --git a/fs/cifs/cifsfs.c b/fs/cifs/cifsfs.c
index d38c69b..c05fa32 100644
--- a/fs/cifs/cifsfs.c
+++ b/fs/cifs/cifsfs.c
@@ -616,7 +616,7 @@ const struct file_operations cifs_file_ops = {
.fsync = cifs_fsync,
.flush = cifs_flush,
.mmap = cifs_file_mmap,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
.llseek = cifs_llseek,
#ifdef CONFIG_CIFS_POSIX
.ioctl = cifs_ioctl,
@@ -637,7 +637,7 @@ const struct file_operations cifs_file_direct_ops = {
.lock = cifs_lock,
.fsync = cifs_fsync,
.flush = cifs_flush,
- .sendfile = generic_file_sendfile, /* BB removeme BB */
+ .splice_read = generic_file_splice_read,
#ifdef CONFIG_CIFS_POSIX
.ioctl = cifs_ioctl,
```

[PATCH] sendfile removal

```
#endif /* CONFIG_CIFS_POSIX */
@@ -656,7 +656,7 @@ const struct file_operations cifs_file_nobrl_ops = {
 .fsync = cifs_fsync,
 .flush = cifs_flush,
 .mmap = cifs_file_mmap,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
 .llseek = cifs_llseek,
#ifdef CONFIG_CIFS_POSIX
 .ioctl = cifs_ioctl,
@@ -676,7 +676,7 @@ const struct file_operations cifs_file_direct_nobrl_ops = {
 .release = cifs_close,
 .fsync = cifs_fsync,
 .flush = cifs_flush,
- .sendfile = generic_file_sendfile, /* BB removeme BB */
+ .splice_read = generic_file_splice_read,
#ifdef CONFIG_CIFS_POSIX
 .ioctl = cifs_ioctl,
#endif /* CONFIG_CIFS_POSIX */
diff --git a/fs/coda/file.c b/fs/coda/file.c
index 5ef2b60..99dbe86 100644
--- a/fs/coda/file.c
+++ b/fs/coda/file.c
@@ -47,8 +47,9 @@ coda_file_read(struct file *coda_file, char __user *buf, size_t count, loff_t *p
 }

static ssize_t
-coda_file_sendfile(struct file *coda_file, loff_t *ppos, size_t count,
-read_actor_t actor, void *target)
+coda_file_splice_read(struct file *coda_file, loff_t *ppos,
+ struct pipe_inode_info *pipe, size_t count,
+ unsigned int flags)
{
 struct coda_file_info *cfi;
 struct file *host_file;
@@ -57,10 +58,10 @@ coda_file_sendfile(struct file *coda_file, loff_t *ppos, size_t count,
 BUG_ON(!cfi || cfi->cfi_magic != CODA_MAGIC);
 host_file = cfi->cfi_container;

- if (!host_file->f_op || !host_file->f_op->sendfile)
+ if (!host_file->f_op || !host_file->f_op->splice_read)
 return -EINVAL;

- return host_file->f_op->sendfile(host_file, ppos, count, actor, target);
+ return host_file->f_op->splice_read(host_file, ppos, pipe, count, flags);
}

static ssize_t
@@ -295,6 +296,6 @@ const struct file_operations coda_file_operations = {
 .flush = coda_flush,
 .release = coda_release,
```

[PATCH] sendfile removal

[PATCH] sendfile removal

```
.fsync = coda_fsync,  
- .sendfile = coda_file_sendfile,  
+ .splice_read = coda_file_splice_read,  
};
```

```
diff --git a/fs/ecryptfs/file.c b/fs/ecryptfs/file.c
```

```
index 59288d8..94f456f 100644
```

```
--- a/fs/ecryptfs/file.c
```

```
+++ b/fs/ecryptfs/file.c
```

```
@@ -338,16 +338,17 @@ static int ecryptfs_fasync(int fd, struct file *file, int flag)  
return rc;  
}
```

```
-static ssize_t ecryptfs_sendfile(struct file *file, loff_t * ppos,  
- size_t count, read_actor_t actor, void *target)  
+static ssize_t ecryptfs_splice_read(struct file *file, loff_t * ppos,  
+ struct pipe_inode_info *pipe, size_t count,  
+ unsigned int flags)  
{  
struct file *lower_file = NULL;  
int rc = -EINVAL;
```

```
lower_file = ecryptfs_file_to_lower(file);  
- if (lower_file->f_op && lower_file->f_op->sendfile)  
- rc = lower_file->f_op->sendfile(lower_file, ppos, count,  
- actor, target);  
+ if (lower_file->f_op && lower_file->f_op->splice_read)  
+ rc = lower_file->f_op->splice_read(lower_file, ppos, pipe,  
+ count, flags);
```

```
return rc;
```

```
}
```

```
@@ -364,7 +365,7 @@ const struct file_operations ecryptfs_dir_fops = {  
.release = ecryptfs_release,  
.fsync = ecryptfs_fsync,  
.fasync = ecryptfs_fasync,  
- .sendfile = ecryptfs_sendfile,  
+ .splice_read = ecryptfs_splice_read,  
};
```

```
const struct file_operations ecryptfs_main_fops = {  
@@ -381,7 +382,7 @@ const struct file_operations ecryptfs_main_fops = {  
.release = ecryptfs_release,  
.fsync = ecryptfs_fsync,  
.fasync = ecryptfs_fasync,  
- .sendfile = ecryptfs_sendfile,  
+ .splice_read = ecryptfs_splice_read,  
};
```

```
static int
```

```
diff --git a/fs/ext2/file.c b/fs/ext2/file.c
```

[PATCH] sendfile removal

[PATCH] sendfile removal

index 566d4e2..a89edc9 100644

--- a/fs/ext2/file.c

+++ b/fs/ext2/file.c

```
@@ -53,7 +53,6 @@ const struct file_operations ext2_file_operations = {
.open = generic_file_open,
.release = ext2_release_file,
.fsync = ext2_sync_file,
- .sendfile = generic_file_sendfile,
.splice_read = generic_file_splice_read,
.splice_write = generic_file_splice_write,
};
@@ -71,7 +70,9 @@ const struct file_operations ext2_xip_file_operations = {
.open = generic_file_open,
.release = ext2_release_file,
.fsync = ext2_sync_file,
- .sendfile = xip_file_sendfile,
+#if 0
+ .splice_read = xip_file_splice_read,
+#endif
};
#endif
```

diff --git a/fs/ext3/file.c b/fs/ext3/file.c

index 1e6f138..acc4913 100644

--- a/fs/ext3/file.c

+++ b/fs/ext3/file.c

```
@@ -120,7 +120,6 @@ const struct file_operations ext3_file_operations = {
.open = generic_file_open,
.release = ext3_release_file,
.fsync = ext3_sync_file,
- .sendfile = generic_file_sendfile,
.splice_read = generic_file_splice_read,
.splice_write = generic_file_splice_write,
};
```

diff --git a/fs/ext4/file.c b/fs/ext4/file.c

index 3c6c1fd..d4c8186 100644

--- a/fs/ext4/file.c

+++ b/fs/ext4/file.c

```
@@ -120,7 +120,6 @@ const struct file_operations ext4_file_operations = {
.open = generic_file_open,
.release = ext4_release_file,
.fsync = ext4_sync_file,
- .sendfile = generic_file_sendfile,
.splice_read = generic_file_splice_read,
.splice_write = generic_file_splice_write,
};
```

diff --git a/fs/fat/file.c b/fs/fat/file.c

index 55d3c74..69a83b5 100644

--- a/fs/fat/file.c

+++ b/fs/fat/file.c

```
@@ -134,7 +134,7 @@ const struct file_operations fat_file_operations = {
```

[PATCH] sendfile removal

[PATCH] sendfile removal

```
.release = fat_file_release,  
.ioctl = fat_generic_ioctl,  
.fsync = file_fsync,  
- .sendfile = generic_file_sendfile,  
+ .splice_read = generic_file_splice_read,  
};
```

```
static int fat_cont_expand(struct inode *inode, loff_t size)
```

```
diff --git a/fs/fuse/file.c b/fs/fuse/file.c
```

```
index adf7995..f79de7c 100644
```

```
--- a/fs/fuse/file.c
```

```
+++ b/fs/fuse/file.c
```

```
@@ -802,7 +802,7 @@ static const struct file_operations fuse_file_operations = {
```

```
.release = fuse_release,
```

```
.fsync = fuse_fsync,
```

```
.lock = fuse_file_lock,
```

```
- .sendfile = generic_file_sendfile,
```

```
+ .splice_read = generic_file_splice_read,
```

```
};
```

```
static const struct file_operations fuse_direct_io_file_operations = {
```

```
@@ -814,7 +814,7 @@ static const struct file_operations fuse_direct_io_file_operations = {
```

```
.release = fuse_release,
```

```
.fsync = fuse_fsync,
```

```
.lock = fuse_file_lock,
```

```
- /* no mmap and sendfile */
```

```
+ /* no mmap and splice_read */
```

```
};
```

```
static const struct address_space_operations fuse_file_aops = {
```

```
diff --git a/fs/gfs2/ops_file.c b/fs/gfs2/ops_file.c
```

```
index 064df88..7dc3be1 100644
```

```
--- a/fs/gfs2/ops_file.c
```

```
+++ b/fs/gfs2/ops_file.c
```

```
@@ -635,7 +635,6 @@ const struct file_operations gfs2_file_fops = {
```

```
.release = gfs2_close,
```

```
.fsync = gfs2_fsync,
```

```
.lock = gfs2_lock,
```

```
- .sendfile = generic_file_sendfile,
```

```
.flock = gfs2_flock,
```

```
.splice_read = generic_file_splice_read,
```

```
.splice_write = generic_file_splice_write,
```

```
diff --git a/fs/hfs/inode.c b/fs/hfs/inode.c
```

```
index 9a934db..bc835f2 100644
```

```
--- a/fs/hfs/inode.c
```

```
+++ b/fs/hfs/inode.c
```

```
@@ -607,7 +607,7 @@ static const struct file_operations hfs_file_operations = {
```

```
.write = do_sync_write,
```

```
.aio_write = generic_file_aio_write,
```

```
.mmap = generic_file_mmap,
```

```
- .sendfile = generic_file_sendfile,
```

[PATCH] sendfile removal

[PATCH] sendfile removal

```
+ .splice_read = generic_file_splice_read,
.fsync = file_fsync,
.open = hfs_file_open,
.release = hfs_file_release,
diff --git a/fs/hfsplus/inode.c b/fs/hfsplus/inode.c
index 45dab5d..409ce54 100644
--- a/fs/hfsplus/inode.c
+++ b/fs/hfsplus/inode.c
@@ -288,7 +288,7 @@ static const struct file_operations hfsplus_file_operations = {
.write = do_sync_write,
.aio_write = generic_file_aio_write,
.mmap = generic_file_mmap,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
.fsync = file_fsync,
.open = hfsplus_file_open,
.release = hfsplus_file_release,
diff --git a/fs/hostfs/hostfs_kern.c b/fs/hostfs/hostfs_kern.c
index 8286491..c778620 100644
--- a/fs/hostfs/hostfs_kern.c
+++ b/fs/hostfs/hostfs_kern.c
@@ -390,7 +390,7 @@ int hostfs_fsync(struct file *file, struct dentry *dentry, int datasync)
static const struct file_operations hostfs_file_fops = {
.llseek = generic_file_llseek,
.read = do_sync_read,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
.aio_read = generic_file_aio_read,
.aio_write = generic_file_aio_write,
.write = do_sync_write,
diff --git a/fs/hpfs/file.c b/fs/hpfs/file.c
index b4eafc0..5b53e5c 100644
--- a/fs/hpfs/file.c
+++ b/fs/hpfs/file.c
@@ -129,7 +129,7 @@ const struct file_operations hpfs_file_ops =
.mmap = generic_file_mmap,
.release = hpfs_file_release,
.fsync = hpfs_file_fsync,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
};

const struct inode_operations hpfs_file_iops =
diff --git a/fs/jffs2/file.c b/fs/jffs2/file.c
index 9987127..c253019 100644
--- a/fs/jffs2/file.c
+++ b/fs/jffs2/file.c
@@ -47,7 +47,7 @@ const struct file_operations jffs2_file_operations =
.ioctl = jffs2_ioctl,
.mmap = generic_file_readonly_mmap,
.fsync = jffs2_fsync,
```

[PATCH] sendfile removal

[PATCH] sendfile removal

```
- .sendfile = generic_file_sendfile
+ .splice_read = generic_file_splice_read,
};

/* jffs2_file_inode_operations */
diff --git a/fs/jfs/file.c b/fs/jfs/file.c
index f7f8eff..87eb936 100644
--- a/fs/jfs/file.c
+++ b/fs/jfs/file.c
@@ -108,7 +108,6 @@ const struct file_operations jfs_file_operations = {
 .aio_read = generic_file_aio_read,
 .aio_write = generic_file_aio_write,
 .mmap = generic_file_mmap,
- .sendfile = generic_file_sendfile,
 .splice_read = generic_file_splice_read,
 .splice_write = generic_file_splice_write,
 .fsync = jfs_fsync,
diff --git a/fs/minix/file.c b/fs/minix/file.c
index f92baa1..17765f6 100644
--- a/fs/minix/file.c
+++ b/fs/minix/file.c
@@ -23,7 +23,7 @@ const struct file_operations minix_file_operations = {
 .aio_write = generic_file_aio_write,
 .mmap = generic_file_mmap,
 .fsync = minix_sync_file,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
};

const struct inode_operations minix_file_inode_operations = {
diff --git a/fs/nfs/file.c b/fs/nfs/file.c
index 9eb8eb4..8689b73 100644
--- a/fs/nfs/file.c
+++ b/fs/nfs/file.c
@@ -41,7 +41,9 @@ static int nfs_file_open(struct inode *, struct file *);
static int nfs_file_release(struct inode *, struct file *);
static loff_t nfs_file_llseek(struct file *file, loff_t offset, int origin);
static int nfs_file_mmap(struct file *, struct vm_area_struct *);
- static ssize_t nfs_file_sendfile(struct file *, loff_t *, size_t, read_actor_t, void *);
+ static ssize_t nfs_file_splice_read(struct file *filp, loff_t *ppos,
+ struct pipe_inode_info *pipe,
+ size_t count, unsigned int flags);
static ssize_t nfs_file_read(struct kiocb *, const struct iovec *iov,
unsigned long nr_segs, loff_t pos);
static ssize_t nfs_file_write(struct kiocb *, const struct iovec *iov,
@@ -65,7 +67,7 @@ const struct file_operations nfs_file_operations = {
 .fsync = nfs_fsync,
 .lock = nfs_lock,
 .flock = nfs_flock,
- .sendfile = nfs_file_sendfile,
+ .splice_read = nfs_file_splice_read,
```

[PATCH] sendfile removal

[PATCH] sendfile removal

```
.check_flags = nfs_check_flags,  
};
```

```
@@ -224,20 +226,21 @@ nfs_file_read(struct kiocb *iocb, const struct iovec *iov,  
}
```

```
static ssize_t  
-nfs_file_sendfile(struct file *filp, loff_t *ppos, size_t count,  
- read_actor_t actor, void *target)  
+nfs_file_splice_read(struct file *filp, loff_t *ppos,  
+ struct pipe_inode_info *pipe, size_t count,  
+ unsigned int flags)  
{  
struct dentry *dentry = filp->f_path.dentry;  
struct inode *inode = dentry->d_inode;  
ssize_t res;
```

```
- dfprintk(VFS, "nfs: sendfile(%s/%s, %lu@%Lu)\n",  
+ dfprintk(VFS, "nfs: splice_read(%s/%s, %lu@%Lu)\n",  
dentry->d_parent->d_name.name, dentry->d_name.name,  
(unsigned long) count, (unsigned long long) *ppos);
```

```
res = nfs_revalidate_mapping(inode, filp->f_mapping);  
if (!res)  
- res = generic_file_sendfile(filp, ppos, count, actor, target);  
+ res = generic_file_splice_read(filp, ppos, pipe, count, flags);  
return res;  
}
```

```
diff --git a/fs/nfsd/vfs.c b/fs/nfsd/vfs.c
```

```
index 7e6aa24..1676e34 100644
```

```
--- a/fs/nfsd/vfs.c
```

```
+++ b/fs/nfsd/vfs.c
```

```
@@ -805,6 +805,7 @@ found:
```

```
* so that they can be passed to the network sendmsg/sendpage routines
```

```
* directly. They will be released after the sending has completed.
```

```
*/
```

```
+#if 0
```

```
static int
```

```
nfsd_read_actor(read_descriptor_t *desc, struct page *page, unsigned long offset, unsigned long size)
```

```
{
```

```
@@ -836,6 +837,7 @@ @ nfsd_read_actor(read_descriptor_t *desc, struct page *page, unsigned long offset
```

```
desc->written += size;
```

```
return size;
```

```
}
```

```
+#endif
```

```
static __be32
```

```
nfsd_vfs_read(struct svc_rqst *rqstp, struct svc_fh *fhp, struct file *file,
```

```
@@ -861,16 +863,13 @@ @ nfsd_vfs_read(struct svc_rqst *rqstp, struct svc_fh *fhp, struct file *file,
```

```
if (ra && ra->p_set)
```

[PATCH] sendfile removal

[PATCH] sendfile removal

```
file->f_ra = ra->p_ra;

- if (file->f_op->sendfile && rqstp->rq_sendfile_ok) {
- rqstp->rq_resused = 1;
- host_err = file->f_op->sendfile(file, &offset, *count,
- nfsd_read_actor, rqstp);
- } else {
- oldfs = get_fs();
- set_fs(KERNEL_DS);
- host_err = vfs_readv(file, (struct iovec __user *)vec, vlen, &offset);
- set_fs(oldfs);
- }
+ /*
+ * Can we use ->splice_read() for this?
+ */
+ oldfs = get_fs();
+ set_fs(KERNEL_DS);
+ host_err = vfs_readv(file, (struct iovec __user *)vec, vlen, &offset);
+ set_fs(oldfs);

/* Write back readahead params */
if (ra) {
diff --git a/fs/ntfs/file.c b/fs/ntfs/file.c
index 7ed5639..ffcc504 100644
--- a/fs/ntfs/file.c
+++ b/fs/ntfs/file.c
@@ -2276,7 +2276,7 @@ const struct file_operations ntfs_file_ops = {
mounted filesystem. */
.mmap = generic_file_mmap, /* Mmap file. */
.open = ntfs_file_open, /* Open file. */
- .sendfile = generic_file_sendfile, /* Zero-copy data send with
+ .splice_read = generic_file_splice_read /* Zero-copy data send with
the data source being on
the ntfs partition. We do
not need to care about the
diff --git a/fs/ocfs2/file.c b/fs/ocfs2/file.c
index ac6c964..84e2b1f 100644
--- a/fs/ocfs2/file.c
+++ b/fs/ocfs2/file.c
@@ -1817,7 +1817,6 @@ const struct inode_operations ocfs2_special_file_iops = {
const struct file_operations ocfs2_fops = {
.read = do_sync_read,
.write = do_sync_write,
- .sendfile = generic_file_sendfile,
.mmap = ocfs2_mmap,
.fsync = ocfs2_sync_file,
.release = ocfs2_file_release,
diff --git a/fs/qnx4/file.c b/fs/qnx4/file.c
index 4464998..867f42b 100644
--- a/fs/qnx4/file.c
+++ b/fs/qnx4/file.c
```

[PATCH] sendfile removal

[PATCH] sendfile removal

```
@@ -25,7 +25,7 @@ const struct file_operations qnx4_file_operations =
.read = do_sync_read,
.aio_read = generic_file_aio_read,
.mmap = generic_file_mmap,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
#ifdef CONFIG_QNX4FS_RW
.write = do_sync_write,
.aio_write = generic_file_aio_write,
diff --git a/fs/ramfs/file-mmuc.c b/fs/ramfs/file-mmuc.c
index 2f14774..97bdc0b 100644
--- a/fs/ramfs/file-mmuc.c
+++ b/fs/ramfs/file-mmuc.c
@@ -41,7 +41,7 @@ const struct file_operations ramfs_file_operations = {
.aio_write = generic_file_aio_write,
.mmap = generic_file_mmap,
.fsync = simple_sync_file,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
.llseek = generic_file_llseek,
};

diff --git a/fs/ramfs/file-nommu.c b/fs/ramfs/file-nommu.c
index 9345a46..69ff5f9 100644
--- a/fs/ramfs/file-nommu.c
+++ b/fs/ramfs/file-nommu.c
@@ -42,7 +42,7 @@ const struct file_operations ramfs_file_operations = {
.write = do_sync_write,
.aio_write = generic_file_aio_write,
.fsync = simple_sync_file,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
.llseek = generic_file_llseek,
};

diff --git a/fs/read_write.c b/fs/read_write.c
index 4d03008..cb4f456 100644
--- a/fs/read_write.c
+++ b/fs/read_write.c
@@ -15,6 +15,7 @@
#include <linux/module.h>
#include <linux/syscalls.h>
#include <linux/pagemap.h>
+#include <linux/pipe_fs_i.h>
#include "read_write.h"

#include <asm/uaccess.h>
@@ -25,7 +26,7 @@ const struct file_operations generic_ro_fops = {
.read = do_sync_read,
.aio_read = generic_file_aio_read,
.mmap = generic_file_readonly_mmap,
```

[PATCH] sendfile removal

```
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
};

EXPORT_SYMBOL(generic_ro_fops);
@@ -708,7 +709,7 @@ static ssize_t do_sendfile(int out_fd, int in_fd, loff_t *ppos,
struct inode * in_inode, * out_inode;
loff_t pos;
ssize_t retval;
- int fput_needed_in, fput_needed_out;
+ int fput_needed_in, fput_needed_out, fl;

/*
* Get input file, and verify that it is ok..
@@ -723,7 +724,7 @@ static ssize_t do_sendfile(int out_fd, int in_fd, loff_t *ppos,
in_inode = in_file->f_path.dentry->d_inode;
if (!in_inode)
goto fput_in;
- if (!in_file->f_op || !in_file->f_op->sendfile)
+ if (!in_file->f_op || !in_file->f_op->splice_read)
goto fput_in;
retval = -ESPIPE;
if (!ppos)
@@ -776,7 +777,11 @@ static ssize_t do_sendfile(int out_fd, int in_fd, loff_t *ppos,
count = max - pos;
}

- retval = in_file->f_op->sendfile(in_file, ppos, count, file_send_actor, out_file);
+ fl = 0;
+ if (in_file->f_flags & O_NONBLOCK)
+ fl = SPLICE_F_NONBLOCK;
+
+ retval = do_splice_direct(in_file, ppos, out_file, count, fl);

if (retval > 0) {
add_rchar(current, retval);
diff --git a/fs/reiserfs/file.c b/fs/reiserfs/file.c
index 9e451a6..30eebf3 100644
--- a/fs/reiserfs/file.c
+++ b/fs/reiserfs/file.c
@@ -1531,7 +1531,6 @@ const struct file_operations reiserfs_file_operations = {
.open = generic_file_open,
.release = reiserfs_file_release,
.fsync = reiserfs_sync_file,
- .sendfile = generic_file_sendfile,
.aio_read = generic_file_aio_read,
.aio_write = generic_file_aio_write,
.splice_read = generic_file_splice_read,
diff --git a/fs/smbfs/file.c b/fs/smbfs/file.c
index aea3f8a..c5d78a7 100644
--- a/fs/smbfs/file.c
```

[PATCH] sendfile removal

```
+++ b/fs/smbfs/file.c
@@ -262,8 +262,9 @@ out:
}

static ssize_t
-smb_file_sendfile(struct file *file, loff_t *ppos,
- size_t count, read_actor_t actor, void *target)
+smb_file_splice_read(struct file *file, loff_t *ppos,
+ struct pipe_inode_info *pipe, size_t count,
+ unsigned int flags)
{
struct dentry *dentry = file->f_path.dentry;
ssize_t status;
@@ -277,7 +278,7 @@ smb_file_sendfile(struct file *file, loff_t *ppos,
DENTRY_PATH(dentry), status);
goto out;
}
- status = generic_file_sendfile(file, ppos, count, actor, target);
+ status = generic_file_splice_read(file, ppos, pipe, count, flags);
out:
return status;
}
@@ -416,7 +417,7 @@ const struct file_operations smb_file_operations =
.open = smb_file_open,
.release = smb_file_release,
.fsync = smb_fsync,
- .sendfile = smb_file_sendfile,
+ .splice_read = smb_file_splice_read,
};

const struct inode_operations smb_file_inode_operations =
diff --git a/fs/sysv/file.c b/fs/sysv/file.c
index 0732ddb..589be21 100644
--- a/fs/sysv/file.c
+++ b/fs/sysv/file.c
@@ -27,7 +27,7 @@ const struct file_operations sysv_file_operations = {
.aio_write = generic_file_aio_write,
.mmap = generic_file_mmap,
.fsync = sysv_sync_file,
- .sendfile = generic_file_sendfile,
+ .splice_read = generic_file_splice_read,
};

const struct inode_operations sysv_file_inode_operations = {
diff --git a/fs/udf/file.c b/fs/udf/file.c
index 51b5764..df070be 100644
--- a/fs/udf/file.c
+++ b/fs/udf/file.c
@@ -261,7 +261,7 @@ const struct file_operations udf_file_operations = {
.aio_write = udf_file_aio_write,
.release = udf_release_file,
```

[PATCH] sendfile removal

```
.fsync = udf_fsync_file,  
- .sendfile = generic_file_sendfile,  
+ .splice_read = generic_file_splice_read,  
};  
  
const struct inode_operations udf_file_inode_operations = {  
diff --git a/fs/ufs/file.c b/fs/ufs/file.c  
index 1e09632..6705d74 100644  
--- a/fs/ufs/file.c  
+++ b/fs/ufs/file.c  
@@ -60,5 +60,5 @@ const struct file_operations ufs_file_operations = {  
.mmap = generic_file_mmap,  
.open = generic_file_open,  
.fsync = ufs_sync_file,  
- .sendfile = generic_file_sendfile,  
+ .splice_read = generic_file_splice_read,  
};  
diff --git a/fs/xfs/linux-2.6/xfs_file.c b/fs/xfs/linux-2.6/xfs_file.c  
index cb51dc9..8c43cd2 100644  
--- a/fs/xfs/linux-2.6/xfs_file.c  
+++ b/fs/xfs/linux-2.6/xfs_file.c  
@@ -124,30 +124,6 @@ xfs_file_aio_write_invis(  
}
```

```
STATIC ssize_t  
-xfs_file_sendfile(  
- struct file *filp,  
- loff_t *pos,  
- size_t count,  
- read_actor_t actor,  
- void *target)  
- {  
- return bhv_vop_sendfile(vn_from_inode(filp->f_path.dentry->d_inode),  
- filp, pos, 0, count, actor, target, NULL);  
- }  
-  
-STATIC ssize_t  
-xfs_file_sendfile_invis(  
- struct file *filp,  
- loff_t *pos,  
- size_t count,  
- read_actor_t actor,  
- void *target)  
- {  
- return bhv_vop_sendfile(vn_from_inode(filp->f_path.dentry->d_inode),  
- filp, pos, IO_INVIS, count, actor, target, NULL);  
- }  
-  
-STATIC ssize_t  
xfs_file_splice_read(  
struct file *infilp,
```


[PATCH] sendfile removal

```
- XFS_STATS_INC(xs_read_calls);
- if (XFS_FORCED_SHUTDOWN(mp))
- return -EIO;
-
- xfs_ilock(ip, XFS_IOLOCK_SHARED);
-
- if (DM_EVENT_ENABLED(BHV_TO_VNODE(bdp)->v_vfsp, ip, DM_EVENT_READ) &&
- (!(ioflags & IO_INVIS))) {
- bhv_vrwlock_t locktype = VRWLOCK_READ;
- int error;
-
- error = XFS_SEND_DATA(mp, DM_EVENT_READ, BHV_TO_VNODE(bdp),
- *offset, count,
- FILP_DELAY_FLAG(filp), &locktype);
- if (error) {
- xfs_iunlock(ip, XFS_IOLOCK_SHARED);
- return -error;
- }
- }
- xfs_rw_enter_trace(XFS_SENDFILE_ENTER, &ip->i_iocore,
- (void *) (unsigned long)target, count, *offset, ioflags);
- ret = generic_file_sendfile(filp, offset, count, actor, target);
- if (ret > 0)
- XFS_STATS_ADD(xs_read_bytes, ret);
-
- xfs_iunlock(ip, XFS_IOLOCK_SHARED);
- return ret;
-}
-
-ssize_t
xfs_splice_read(
bhv_desc_t *bdp,
struct file *infilp,
diff --git a/fs/xfs/linux-2.6/xfs_lrw.h b/fs/xfs/linux-2.6/xfs_lrw.h
index 7ac51b1..7c60a1e 100644
--- a/fs/xfs/linux-2.6/xfs_lrw.h
+++ b/fs/xfs/linux-2.6/xfs_lrw.h
@@ -90,9 +90,6 @@ extern ssize_t xfs_read(struct bhv_desc *, struct kiocb *,
extern ssize_t xfs_write(struct bhv_desc *, struct kiocb *,
const struct iovec *, unsigned int,
loff_t *, int, struct cred *);
-extern ssize_t xfs_sendfile(struct bhv_desc *, struct file *,
- loff_t *, int, size_t, read_actor_t,
- void *, struct cred *);
extern ssize_t xfs_splice_read(struct bhv_desc *, struct file *, loff_t *,
struct pipe_inode_info *, size_t, int, int,
struct cred *);
diff --git a/fs/xfs/linux-2.6/xfs_vnode.h b/fs/xfs/linux-2.6/xfs_vnode.h
index d1b2d01..013048a 100644
--- a/fs/xfs/linux-2.6/xfs_vnode.h
+++ b/fs/xfs/linux-2.6/xfs_vnode.h
```

[PATCH] sendfile removal

```
@@ -139,9 +139,6 @@ typedef ssize_t (*vop_read_t)(bhv_desc_t *, struct kiocb *,
typedef ssize_t (*vop_write_t)(bhv_desc_t *, struct kiocb *,
const struct iovec *, unsigned int,
loff_t *, int, struct cred *);
- typedef ssize_t (*vop_sendfile_t)(bhv_desc_t *, struct file *,
- loff_t *, int, size_t, read_actor_t,
- void *, struct cred *);
typedef ssize_t (*vop_splice_read_t)(bhv_desc_t *, struct file *, loff_t *,
struct pipe_inode_info *, size_t, int, int,
struct cred *);
@@ -206,7 +203,6 @@ typedef struct bhv_vnodeops {
vop_close_t vop_close;
vop_read_t vop_read;
vop_write_t vop_write;
- vop_sendfile_t vop_sendfile;
vop_splice_read_t vop_splice_read;
vop_splice_write_t vop_splice_write;
vop_ioctl_t vop_ioctl;
@@ -254,8 +250,6 @@ typedef struct bhv_vnodeops {
VOP(vop_read, vp)(VNHEAD(vp),file,iiov,segs,offset,iioflags,cr)
#define bhv_vop_write(vp,file,iiov,segs,offset,iioflags,cr) \
VOP(vop_write, vp)(VNHEAD(vp),file,iiov,segs,offset,iioflags,cr)
-#define bhv_vop_sendfile(vp,f,off,iioflags,cnt,act,targ,cr) \
- VOP(vop_sendfile, vp)(VNHEAD(vp),f,off,iioflags,cnt,act,targ,cr)
#define bhv_vop_splice_read(vp,f,o,pipe,cnt,fl,iiofl,cr) \
VOP(vop_splice_read, vp)(VNHEAD(vp),f,o,pipe,cnt,fl,iiofl,cr)
#define bhv_vop_splice_write(vp,f,o,pipe,cnt,fl,iiofl,cr) \
diff --git a/fs/xfs/xfs_vnodeops.c b/fs/xfs/xfs_vnodeops.c
index de17aed..70bc82f 100644
--- a/fs/xfs/xfs_vnodeops.c
+++ b/fs/xfs/xfs_vnodeops.c
@@ -4680,9 +4680,6 @@ bhv_vnodeops_t xfs_vnodeops = {
.vop_open = xfs_open,
.vop_close = xfs_close,
.vop_read = xfs_read,
-#ifdef HAVE_SENDFILE
- .vop_sendfile = xfs_sendfile,
-#endif
#ifdef HAVE_SPLICE
.vop_splice_read = xfs_splice_read,
.vop_splice_write = xfs_splice_write,
diff --git a/include/linux/fs.h b/include/linux/fs.h
index 7cf0c54..4c73cd6 100644
--- a/include/linux/fs.h
+++ b/include/linux/fs.h
@@ -1104,7 +1104,6 @@ struct file_operations {
int (*aio_fsync) (struct kiocb *, int datasync);
int (*fasync) (int, struct file *, int);
int (*lock) (struct file *, int, struct file_lock *);
- ssize_t (*sendfile) (struct file *, loff_t *, size_t, read_actor_t, void *);
ssize_t (*sendpage) (struct file *, struct page *, int, size_t, loff_t *, int);
```

[PATCH] sendfile removal

```
unsigned long (*get_unmapped_area)(struct file *, unsigned long, unsigned long, unsigned long, unsigned
long);
int (*check_flags)(int);
@@ -1734,7 +1733,6 @@ extern ssize_t generic_file_buffered_write(struct kiocb *, const struct iovec *,
unsigned long, loff_t, loff_t *, size_t, ssize_t);
extern ssize_t do_sync_read(struct file *filp, char __user *buf, size_t len, loff_t *ppos);
extern ssize_t do_sync_write(struct file *filp, const char __user *buf, size_t len, loff_t *ppos);
-extern ssize_t generic_file_sendfile(struct file *, loff_t *, size_t, read_actor_t, void *);
extern void do_generic_mapping_read(struct address_space *mapping,
struct file_ra_state *, struct file *,
loff_t *, read_descriptor_t *, read_actor_t);
@@ -1764,9 +1762,6 @@ extern int nonseekable_open(struct inode * inode, struct file * filp);
#ifdef CONFIG_FS_XIP
extern ssize_t xip_file_read(struct file *filp, char __user *buf, size_t len,
loff_t *ppos);
-extern ssize_t xip_file_sendfile(struct file *in_file, loff_t *ppos,
- size_t count, read_actor_t actor,
- void *target);
extern int xip_file_mmap(struct file * file, struct vm_area_struct * vma);
extern ssize_t xip_file_write(struct file *filp, const char __user *buf,
size_t len, loff_t *ppos);
diff --git a/include/linux/sunrpc/svc.h b/include/linux/sunrpc/svc.h
index 4a7ae8a..129d50f 100644
--- a/include/linux/sunrpc/svc.h
+++ b/include/linux/sunrpc/svc.h
@@ -253,7 +253,7 @@ struct svc_rqst {
 * determine what device number
 * to report (real or virtual)
 */
- int rq_sendfile_ok; /* turned off in gss privacy
+ int rq_splice_ok; /* turned off in gss privacy
 * to prevent encrypting page
 * cache pages */
wait_queue_head_t rq_wait; /* synchronization */
diff --git a/kernel/relay.c b/kernel/relay.c
index 4311101..86a6579 100644
--- a/kernel/relay.c
+++ b/kernel/relay.c
@@ -1060,21 +1060,6 @@ static ssize_t relay_file_read(struct file *filp,
NULL, &desc);
}

-static ssize_t relay_file_sendfile(struct file *filp,
- loff_t *ppos,
- size_t count,
- read_actor_t actor,
- void *target)
- {
- read_descriptor_t desc;
- desc.written = 0;
- desc.count = count;
```

[PATCH] sendfile removal

```
- desc.arg.data = target;
- desc.error = 0;
- return relay_file_read_subbufs(filp, ppos, subbuf_send_actor,
- actor, &desc);
-}
-
const struct file_operations relay_file_operations = {
.open = relay_file_open,
.poll = relay_file_poll,
@@ -1082,7 +1067,6 @@ const struct file_operations relay_file_operations = {
.read = relay_file_read,
.llseek = no_llseek,
.release = relay_file_release,
- .sendfile = relay_file_sendfile,
};
EXPORT_SYMBOL_GPL(relay_file_operations);

diff --git a/mm/filemap.c b/mm/filemap.c
index edb1b0b..14ad45a 100644
--- a/mm/filemap.c
+++ b/mm/filemap.c
@@ -1245,26 +1245,6 @@ int file_send_actor(read_descriptor_t * desc, struct page *page, unsigned long o
return written;
}

-ssize_t generic_file_sendfile(struct file *in_file, loff_t *ppos,
- size_t count, read_actor_t actor, void *target)
-{
- read_descriptor_t desc;
-
- if (!count)
- return 0;
-
- desc.written = 0;
- desc.count = count;
- desc.arg.data = target;
- desc.error = 0;
-
- do_generic_file_read(in_file, ppos, &desc, actor);
- if (desc.written)
- return desc.written;
- return desc.error;
-}
-EXPORT_SYMBOL(generic_file_sendfile);
-
static ssize_t
do_readahead(struct address_space *mapping, struct file *filp,
unsigned long index, unsigned long nr)
diff --git a/mm/filemap_xip.c b/mm/filemap_xip.c
index fa360e5..e85bade 100644
--- a/mm/filemap_xip.c
```

[PATCH] sendfile removal

```
+++ b/mm/filemap_xip.c
@@ -159,6 +159,7 @@ xip_file_read(struct file *filp, char __user *buf, size_t len, loff_t *ppos)
}
EXPORT_SYMBOL_GPL(xip_file_read);

+#if 0
ssize_t
xip_file_sendfile(struct file *in_file, loff_t *ppos,
size_t count, read_actor_t actor, void *target)
@@ -180,6 +181,7 @@ xip_file_sendfile(struct file *in_file, loff_t *ppos,
return desc.error;
}
EXPORT_SYMBOL_GPL(xip_file_sendfile);
+#endif

/*
 * __xip_unmap is invoked from xip_unmap and
diff --git a/mm/shmem.c b/mm/shmem.c
index e537317..71f480b 100644
--- a/mm/shmem.c
+++ b/mm/shmem.c
@@ -1483,6 +1483,17 @@ static const struct inode_operations shmem_symlink_inode_operations;
static const struct inode_operations shmem_symlink_inline_operations;

/*
+ * FIXME: need a real implementation :-)
+ */
+static ssize_t
+shmem_file_splice_read(struct file *in, loff_t *ppos,
+ struct pipe_inode_info *pipe, size_t len,
+ unsigned int flags)
+{
+ return -EINVAL;
+}
+
+/*
 * Normally tmpfs makes no use of shmem_prepare_write, but it
 * lets a tmpfs file be used read-write below the loop driver.
 */
@@ -1709,25 +1720,6 @@ static ssize_t shmem_file_read(struct file *filp, char __user *buf, size_t count
return desc.error;
}

-static ssize_t shmem_file_sendfile(struct file *in_file, loff_t *ppos,
- size_t count, read_actor_t actor, void *target)
- {
- read_descriptor_t desc;
-
- if (!count)
- return 0;
-

```

[PATCH] sendfile removal

```
- desc.written = 0;
- desc.count = count;
- desc.arg.data = target;
- desc.error = 0;
-
- do_shmem_file_read(in_file, ppos, &desc, actor);
- if (desc.written)
- return desc.written;
- return desc.error;
-}
-
static int shmem_statfs(struct dentry *dentry, struct kstatfs *buf)
{
struct shmem_sb_info *sbinfo = SHMEM_SB(dentry->d_sb);
@@ -2397,7 +2389,7 @@ static const struct file_operations shmem_file_operations = {
.read = shmem_file_read,
.write = shmem_file_write,
.fsync = simple_sync_file,
- .sendfile = shmem_file_sendfile,
+ .splice_read = shmem_file_splice_read,
#endif
};

diff --git a/net/sunrpc/auth_gss/svcauth_gss.c b/net/sunrpc/auth_gss/svcauth_gss.c
index 099a983..c094583 100644
--- a/net/sunrpc/auth_gss/svcauth_gss.c
+++ b/net/sunrpc/auth_gss/svcauth_gss.c
@@ -853,7 +853,7 @@ @@@@ unwrap_priv_data(struct svc_rqst *rqstp, struct xdr_buf *buf, u32 seq, struct gs
u32 priv_len, maj_stat;
int pad, saved_len, remaining_len, offset;

- rqstp->rq_sendfile_ok = 0;
+ rqstp->rq_splice_ok = 0;

priv_len = svc_getnl(&buf->head[0]);
if (rqstp->rq_deferred) {
diff --git a/net/sunrpc/svc.c b/net/sunrpc/svc.c
index e673ef9..55ea6df 100644
--- a/net/sunrpc/svc.c
+++ b/net/sunrpc/svc.c
@@ -814,7 +814,7 @@ @@@@ svc_process(struct svc_rqst *rqstp)
rqstp->rq_res.tail[0].iov_base = NULL;
rqstp->rq_res.tail[0].iov_len = 0;
/* Will be turned off only in gss privacy case: */
- rqstp->rq_sendfile_ok = 1;
+ rqstp->rq_splice_ok = 1;
/* tcp needs a space for the record length... */
if (rqstp->rq_prot == IPPROTO_TCP)
svc_putnl(resv, 0);
--
```

[PATCH] sendfile removal

Jens Axboe

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>