
Subject: [PATCH 3/16] Introduce MS_KERNMOUNT flag
Posted by [Pavel Emelianov](#) on Fri, 06 Jul 2007 08:04:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

This flag tells the .get_sb callback that this is a kern_mount() call so that it can trust *data pointer to be valid in-kernel one.

Running a few steps forward - this will be needed for proc to create the superblock and store a valid pid namespace on it during the namespace creation. The reason, why the namespace cannot live without proc mount is described in the appropriate patch.

Signed-off-by: Pavel Emelianov <xemul@openvz.org>

```
fs/namespace.c | 3 ++-
fs/super.c     | 6 +++---
include/linux/fs.h | 4 +++-
3 files changed, 8 insertions(+), 5 deletions(-)
```

```
diff -upr linux-2.6.22-rc4-mm2.orig/fs/namespace.c linux-2.6.22-rc4-mm2-2/fs/namespace.c
--- linux-2.6.22-rc4-mm2.orig/fs/namespace.c 2007-06-14 12:00:06.000000000 +0400
+++ linux-2.6.22-rc4-mm2-2/fs/namespace.c 2007-07-04 19:00:39.000000000 +0400
@@ -1558,7 +1558,8 @@ long do_mount(char *dev_name, char *dir_
     mnt_flags |= MNT_NOMNT;
```

```
    flags &= ~(MS_NOSUID | MS_NOEXEC | MS_NODEV | MS_ACTIVE |
-    MS_NOATIME | MS_NODIRATIME | MS_RELATIME | MS_NOMNT);
+    MS_NOATIME | MS_NODIRATIME | MS_RELATIME |
+    MS_NOMNT | MS_KERNMOUNT);
```

```
/* ... and get the mountpoint */
retval = path_lookup(dir_name, LOOKUP_FOLLOW, &nd);
diff -upr linux-2.6.22-rc4-mm2.orig/fs/super.c linux-2.6.22-rc4-mm2-2/fs/super.c
--- linux-2.6.22-rc4-mm2.orig/fs/super.c 2007-06-07 15:37:30.000000000 +0400
+++ linux-2.6.22-rc4-mm2-2/fs/super.c 2007-07-04 19:00:39.000000000 +0400
@@ -942,9 +942,9 @@ do_kern_mount(const char *fstype, int fl
    return mnt;
}
```

```
-struct vfsmount *kern_mount(struct file_system_type *type)
+struct vfsmount *kern_mount_data(struct file_system_type *type, void *data)
{
-    return vfs_kern_mount(type, 0, type->name, NULL);
+    return vfs_kern_mount(type, MS_KERNMOUNT, type->name, data);
}
```

```

-EXPORT_SYMBOL(kern_mount);
+EXPORT_SYMBOL_GPL(kern_mount_data);
diff -upr linux-2.6.22-rc4-mm2.orig/include/linux/fs.h linux-2.6.22-rc4-mm2-2/include/linux/fs.h
--- linux-2.6.22-rc4-mm2.orig/include/linux/fs.h 2007-06-14 12:00:06.000000000 +0400
+++ linux-2.6.22-rc4-mm2-2/include/linux/fs.h 2007-07-04 19:00:39.000000000 +0400
@@ -130,6 +130,7 @@ extern int dir_notify_enable;
#define MS_NO_LEASES (1<<22) /* fs does not support leases */
#define MS_SETUSER (1<<23) /* set mnt_uid to current user */
#define MS_NOMNT (1<<24) /* don't allow unprivileged submounts */
+#define MS_KERNMOUNT (1<<25) /* this is a kern_mount call */
#define MS_ACTIVE (1<<30)
#define MS_NOUSER (1<<31)

@@ -1490,7 +1491,8 @@ void unnamed_dev_init(void);

extern int register_filesystem(struct file_system_type *);
extern int unregister_filesystem(struct file_system_type *);
-extern struct vfsmount *kern_mount(struct file_system_type *);
+extern struct vfsmount *kern_mount_data(struct file_system_type *, void *data);
+#define kern_mount(type) kern_mount_data(type, NULL)
extern int may_umount_tree(struct vfsmount *);
extern int may_umount(struct vfsmount *);
extern void umount_tree(struct vfsmount *, int, struct list_head *);

```

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
