

---

Subject: [patch 03/22] r/o bind mounts: stub functions  
Posted by [Cedric Le Goater](#) on Thu, 07 Jun 2007 15:25:30 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

From: Dave Hansen <[hansendc@us.ibm.com](mailto:hansendc@us.ibm.com)>

This patch adds two function mnt\_want\_write() and mnt\_drop\_write(). These are used like a lock pair around and fs operations that might cause a write to the filesystem.

Before these can become useful, we must first cover each place in the VFS where writes are performed with a want/drop pair. When that is complete, we can actually introduce code that will safely check the counts before allowing r/w<->r/o transitions to occur.

Note that we put the linux/fs.h #include as far down in mount.h as possible. This is to keep the mount.h->fs.h->sched.h->mount.h include dependency from biting us.

Signed-off-by: Dave Hansen <[hansendc@us.ibm.com](mailto:hansendc@us.ibm.com)>

---

---

```
fs/namespace.c      | 13 ++++++
include/linux/mount.h | 13 ++++++
2 files changed, 26 insertions(+)
```

Index: 2.6.22-rc4-mm2-robindmount/fs/namespace.c

```
=====
--- 2.6.22-rc4-mm2-robindmount.orig/fs/namespace.c
+++ 2.6.22-rc4-mm2-robindmount/fs/namespace.c
@@ -79,6 +79,19 @@ struct vfsmount *alloc_vfsmnt(const char
    return mnt;
}

+int mnt_want_write(struct vfsmount *mnt)
+{
+ if (__mnt_is_readonly(mnt))
+     return -EROFS;
+ return 0;
+}
+EXPORT_SYMBOL_GPL(mnt_want_write);
+
+void mnt_drop_write(struct vfsmount *mnt)
+{
+}
+EXPORT_SYMBOL_GPL(mnt_drop_write);
```

```

+
int simple_set_mnt(struct vfsmount *mnt, struct super_block *sb)
{
    mnt->mnt_sb = sb;
Index: 2.6.22-rc4-mm2-robindmount/include/linux/mount.h
=====
--- 2.6.22-rc4-mm2-robindmount.orig/include/linux/mount.h
+++ 2.6.22-rc4-mm2-robindmount/include/linux/mount.h
@@ @ -74,6 +74,19 @@ static inline struct vfsmount *mntget(st
    return mnt;
}

+#include <linux/fs.h>
+
+/*
+ * This shouldn't be used directly outside of the VFS,
+ * use mnt_want/drop_write() instead.
+ */
+static inline int __mnt_is_READONLY(struct vfsmount *mnt)
+{
+    return (mnt->mnt_sb->s_flags & MS_RDONLY);
+}
+
+extern int mnt_want_write(struct vfsmount *mnt);
+extern void mnt_drop_write(struct vfsmount *mnt);
extern void mntput_no_expire(struct vfsmount *mnt);
extern void mnt_pin(struct vfsmount *mnt);
extern void mnt_unpin(struct vfsmount *mnt);

--
```

---

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

---