

# [PATCH 00/20] Mount writer count and read-only bind mounts (v3)

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-06/msg09294.html>

---

- *From:* Dave Hansen <[haveblue@xxxxxxxxxx](mailto:haveblue@xxxxxxxxxx)>
  - *Date:* Tue, 27 Jun 2006 15:14:36 -0700
- 

I think these are ready for -mm. They've gone through a few revisions and a week or two of normal burn-in testing.

---

The following series implements read-only bind mounts. This feature allows a read-only view into a read-write filesystem. In the process of doing that, it also provides infrastructure for keeping track of the number of writers to any given mount. New in this version is that if that number is non-zero, remounts from r/w to r/o are not allowed.

This set does not take the previously tried approach of pushing down the vfstmount structure deeply into call paths, such that it might be checked in functions like permission(), may\_create() and may\_open(). Instead, it does checks near the entry points in the kernel, bumping a reference count in the vfstmount structure. I've also eliminated the use of the MNT\_RDONLY flag. It was redundant since we have the reference count.

This set also makes no attempt to keep the return codes for these r/o bind mounts the same as for a real r/o filesystem or device. It would require significantly more code and be quite a bit more invasive. Unless there is a very strong reason to do so, I believe it isn't worth the trouble.

One note: the previous patches all worked this way:

```
mount --bind -o ro /source /dest
```

These patches have changed that behavior. It now requires two steps:

```
mount --bind /source /dest
mount -o remount,ro /dest
```

Since the last revision, the locking in faccessat() and mnt\_is\_readonly() has been changed to fix a race which might have caused a false-negative mount-is-readonly return when faccessat() is called while another two processes are racing to make a mount

[PATCH 00/20] Mount writer count and read-only bind mounts (v3)

readonly.

Signed-off-by: Dave Hansen <haveblue@xxxxxxxxxx>

–

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/>