

Linux-Kernel: [Patch 10/10]: ext3 online resize: make group-add asynchronous.

[Patch 10/10]: ext3 online resize: make group-add asynchronous.

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-09/9041.html>

From: Stephen Tweedie (sct_at_redhat.com)

Date: 09/30/04

Date: Thu, 30 Sep 2004 14:24:06 +0100

To: linux-kernel@vger.kernel.org, Andrew Morton <akpm@osdl.org>, Andreas Dilger <[There's not much to be gained from making the resize group-add operation fully synchronous. It's already atomic by virtue of the journal, and a resize operation usually doesn't add just one group, but many. The real need is to sync things up at the end of the *entire* resize, not after each individual group-add; and by syncing each group separately, the whole resize operation is made much, much slower if we're running on a live filesystem.](mailto:adilger@clusterf</p></div><div data-bbox=)

Signed-off-by: Stephen Tweedie <sct@redhat.com>

--- linux-2.6.9-rc2-mm4/fs/ext3/resize.c.=K0009=.orig

+++ linux-2.6.9-rc2-mm4/fs/ext3/resize.c

@@ -872,7 +872,6 @@ int ext3_group_add(struct super_block *s

exit_journal:

```
    unlock_super(sb);
- handle->h_sync = 1;
    if ((err2 = ext3_journal_stop(handle)) && !err)
        err = err2;
    if (!err) {
-
```

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

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

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