

[PATCH] reiserfs: handle cnode allocation failure gracefully

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-11/9520.html>

From: Jeff Mahoney (jeffm_at_suse.com)

Date: 11/29/05

Date: Tue, 29 Nov 2005 16:44:03 -0500

To: Andrew Morton <akpm@osdl.org>, Linus Torvalds <torvalds@osdl.org>, Linux Kernel Mailing List

If an external device is used for a journal, by default it will use the entire device. The reiserfs journal code allocates structures per journal block when it mounts the file system. If the journal device is too large, and memory cannot be allocated for the structures, it will continue and ultimately panic when it can't pull one off the free list.

This patch handles the allocation failure gracefully and prints an error message at mount time.

Changes: Updated error message to be more descriptive to the user.

Discussed and approved on ReiserFS Mailing List, Nov 28.

Signed-off-by: Jeff Mahoney <jeffm@suse.com>

```
diff -ruNpX dontdiff a/fs/reiserfs/journal.c b/fs/reiserfs/journal.c
--- a/fs/reiserfs/journal.c 2005-09-16 11:42:58.000000000 -0400
+++ b/fs/reiserfs/journal.c 2005-11-23 19:14:17.000000000 -0500
@@ -2757,6 +2757,15 @@ int journal_init(struct super_block *p_s
     journal->j_cnode_used = 0;
     journal->j_must_wait = 0;

+ if (journal->j_cnode_free == 0) {
+ reiserfs_warning(p_s_sb, "journal-2004: Journal cnode memory "
+ "allocation failed (%ld bytes). Journal is "
+ "too large for available memory. Usually "
+ "this is due to a journal that is too large.",
+ sizeof (struct reiserfs_journal_cnode) * num_cnodes);
+ goto free_and_return;
+ }
+
     init_journal_hash(p_s_sb);
     jl = journal->j_current_jl;
     jl->j_list_bitmap = get_list_bitmap(p_s_sb, jl);
```

Linux-Kernel: [PATCH] reiserfs: handle cnode allocation failure gracefully

--

Jeff Mahoney
SuSE Labs

-

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/>