

Re: 1352 NUL bytes at the end of a page? (was Re: Assertion `s && s->tree' failed: The saga continues.)

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-05/3310.html>

From: Steven Cole (elenstev_at_mesatop.com)

Date: 05/16/04

To: Andrew Morton <akpm@osdl.org>

Date: Sun, 16 May 2004 07:49:24 -0600

On Sunday 16 May 2004 04:01 am, Andrew Morton wrote:

> Steven Cole <elenstev@mesatop.com> wrote:

>>

>> *The script only reported one*

>> *iteration finished, while I got it to do 36 iterations over several hours earlier*

>> *today (with a 2.6.3-4mdk vendor kernel), so I'm going to add some timing*

>> *tests to the script to see if things are really slowing down with current kernels,*

>> *or if it's just my worried imaginings.*

>

> *I did a bit of testing on a 256MB laptop with a fairly slow disk, ext3.*

> *Three iterations of the test took:*

>

> *2.6.6: 1055.53s user 327.14s system 32% cpu 1:10:06.71 total*

>

> *2.4.27-pre2: 1042.03s user 307.21s system 32% cpu 1:09:46.00 total*

>

> *2.6.3: 1053.23s user 326.16s system 27% cpu 1:22:07.24 total*

>

>

> *So there's nothing particularly wild there. It's possible I guess that the*

> *2.6 VM is very sucky but something else made up for it - possibly the*

> *anticipatory scheduler but more likely the Orlov allocator.*

>

> *You're using reiserfs, yes?*

Yes, but I also have a BK repo on an ext3 fs which saw the exact same initial problem. Well, almost in that I got the Assertion `s && s->tree' failed message, but didn't preserve the possibly null-containing s.ChangeSet file.

>

> *Are you sure the IDE disks are in DMA mode?*

>

>

For the 2.6.3 tests, yes.

Re: 1352 NUL bytes at the end of a page? (was Re: Assertion `s && s->tree' failed: The saga continues.)

Linux-Kernel: Re: 1352 NUL bytes at the end of a page? (was Re: Assertion `s && s->tree' failed: The saga continues.)

For the overnight 2.6.6 tests, apparently not, but it is now.
Without DMA, the clone took about 22 minutes, but even with
DMA back on, 2.6.6-current is much slower than 2.6.3-4mdk.

It's hard to make your wheels fall off when putting around the
race track at 2/3 speed. If I get time, I'll try the patch from Nick
and/or a kernel from around 4/15 when I first saw this problem,
but the weather's nice outside, so it may be later this evening
before I return to testing.

```
/dev/hda:
multcount = 0 (off)
IO_support = 0 (default 16-bit)
unmaskirq = 0 (off)
using_dma = 1 (on)
keepsettings = 0 (off)
readonly = 0 (off)
readahead = 256 (on)
geometry = 65535/16/63, sectors = 78125000, start = 0
```

2.6.3-4mdk:

```
time bk clone -qlr40514130hBbvgP4CvwEVEu27oxm46w testing-2.6 foo
287.93user 59.33system 11:14.03elapsed 51%CPU (0avgtext+0avgdata 0maxresident)k
0inputs+0outputs (41264major+523595minor)pagefaults 0swaps
```

```
(cd foo; time bk pull -q)
397.05user 184.05system 16:38.35elapsed 58%CPU (0avgtext+0avgdata 0maxresident)k
0inputs+0outputs (30287major+2115127minor)pagefaults 0swaps
```

2.6.6-current:

```
time bk clone -qlr40514130hBbvgP4CvwEVEu27oxm46w testing-2.6 foo
290.48user 96.76system 15:00.85elapsed 42%CPU (0avgtext+0avgdata 0maxresident)k
0inputs+0outputs (50893major+515118minor)pagefaults 0swaps
```

```
(cd foo; time bk pull -q)
402.74user 254.98system 23:25.43elapsed 46%CPU (0avgtext+0avgdata 0maxresident)k
0inputs+0outputs (60130major+2089908minor)pagefaults 0swaps
```

Steven

-

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/>

Re: 1352 NUL bytes at the end of a page? (was Re: Assertion `s && s->tree' failed: The saga continues.)