

---

Subject: Re: [PATCH 01/10] Containers(V10): Basic container framework  
Posted by [Andrew Morton](#) on Wed, 30 May 2007 16:00:31 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Wed, 30 May 2007 07:02:00 -0700 "Paul Menage" <menage@google.com> wrote:

>  
> >  
> > People have hit unpleasant problems before now running iput() against  
> > partially-constructed inodes.  
>  
> What kinds of problems? Are there bits of state that I should fully  
> construct even if I'm going to iput() it, or is there a better  
> function to call? fs/ext3/super.c seems to do the same thing.

I don't recall, actually. But it crashed.

I guess the fault-injection code could be used to trigger errors here.

```
> > >  
> > > +static inline void get_first_subsys(const struct container *cont,  
> > > +      struct container_subsys_state **css,  
> > > +      int *subsys_id) {  
> > > +    const struct containerfs_root *root = cont->root;  
> > > +    const struct container_subsys *test_ss;  
> > > +    BUG_ON(list_empty(&root->subsys_list));  
> > > +    test_ss = list_entry(root->subsys_list.next,  
> > > +      struct container_subsys, sibling);  
> > > +    if (css) {  
> > > +      *css = cont->subsys[test_ss->subsys_id];  
> > > +      BUG_ON(!*css);  
> > > +    }  
> > > +    if (subsys_id)  
> > > +      *subsys_id = test_ss->subsys_id;  
> > > +}  
> >  
> >
```

> > This ends up having several callers and its too large to inline.

>  
> Two large from a compiler PoV or from a style PoV? It's basically just  
> six dereferences and two comparisons, plus the BUG\_ON(s).

It will end up generating more .text this way. We figure that this makes  
it slower, due to increased icache footprint.

---