Subject: Re: [PATCH 05/10] Containers(V10): Add container_clone() interface
Posted by serue on Thu, 31 May 2007 19:56:31 GMT

View Forum Message <> Reply to Message

Quoting Andrew Morton (akpm@linux-foundation.org):

```
> On Tue, 29 May 2007 06:01:09 -0700 menage@google.com wrote:
>
> > This patch adds support for container_clone(), a speculative interface
> > to creating new containers intended to be used for systems such as
> > namespace unsharing.
> >
> > ...
> >
> > +
> > +static atomic_t namecnt;
> > +static void get_unused_name(char *buf)
> > +{
> > + sprintf(buf, "node%d", atomic_inc_return(&namecnt));
> > +}
>
> A stupid thing, but a sufficiently determined attacker could cause this to
> wrap.
```

Yeah, this was very consciously done as a "just make it work for now" naming system. If we want to stick with this naming, then I suppose we could do a global bitmap.

But imo this naming is not very convenient - it would be nicer if we

a) allowed users to specify a name (not sure how that would work logistically)

b) made the namecnt variable for automatically named containers be per-directory. I'd much rather see

```
 /containers/node1/node1
 /containers/node2
than
 /containers/node1/node3
 /containers/node2
```

(assuming /node2 was created between /node1 and /node1/node1 or /node1/node3)

thanks,
-serge