
Subject: Re: [PATCH] containers: define a namespace container subsystem

Posted by [Paul Menage](#) on Fri, 02 Feb 2007 22:25:25 GMT

[View Forum Message](#) <> [Reply to Message](#)

On 1/30/07, Serge E. Hallyn <serue@us.ibm.com> wrote:

>
> It behaves pretty differently from other subsystems implemented
> so far, which could either be seen as evidence that it doesn't
> belong as a subsystem, or, more likely, that the container
> subsystem approach is quite flexible.

The latter - the container system is meant to be able to support more
than just resource controllers.

```
> +           char *name)
> +{
> +    struct dentry *d = container_get_dentry(parent->dentry, name);
> +    return d ? d->d_fsd : NULL;
> +}
> +
> +#define NS_CONT_MODE (S_IFDIR | S_IRUGO | S_IXUGO | S_IWUSR)
> +int container_switch(struct task_struct *tsk)
> +{
> +    int h;
> +    struct container *cur_cont, *new_cont;
> +    char path[20];
> +    struct qstr name;
> +    struct dentry *dentry;
> +    int ret;
> +    char *pathbuf = NULL;
> +    char buffer[20];
> +
> +    /* check if nsproxy subsys is registered */
> +    if (ns_container_subsys_idx == -1)
> +        return 0;
> +
> +    printk(KERN_NOTICE "%s: ns_container subsys registered\n", __FUNCTION__);
> +    /* check if nsproxy subsys is mounted in some hierarchy */
> +    rCU_read_lock();
> +    h = rCU_dereference(subsys[ns_container_subsys_idx]->hierarchy);
> +    rCU_read_unlock();
> +    if (h == 0) {
> +        /* do we mount the nsproxy subsys, or just skip
> +         * creating a container? I think we just skip
> +         * it.
```

I'd say that we should try to create a fresh hierarchy with just the
nsproxy subsystem on it. Otherwise if someone tries to mount the

nsproxy subsystem later, we end up with some namespaces with no container.

This looks great - I'll incorporate it or something like it in my next patch set.

Cheers,

Paul

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
