
Subject: Re: [ckrm-tech] [PATCH 7/7] containers (V7): Container interface to nsproxy subsystem

Posted by [serue](#) on Mon, 02 Apr 2007 14:09:39 GMT

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

Quoting Srivatsa Vaddagiri (vatsa@in.ibm.com):

> On Mon, Feb 12, 2007 at 12:15:28AM -0800, menage@google.com wrote:

> > +int ns_container_clone(struct task_struct *tsk)

> > +{

> > + return container_clone(tsk, &ns_subsys);

> > +}

>

> This function is a no-op if ns hierarchy is not mounted at this point.

> This would mean that we will miss out on some directories in ns

> hierarchy if it happened to be mounted later. It would be nice to

> recreate such missing directories upon mount. However I suspect it would

> not be easy ..Maybe we need to scan the task list and (re-)invoke

> ns_container_clone() for every new task->nsproxy we find in the list.

> Alternately perhaps we could auto mount (kern_mount) ns hierarchy very early at

> bootup? On the flip side that would require remount support so that additional

> controllers (like cgroup, mem) can be bound to (non-empty) ns hierarchy after

> bootup.

Losing the directory isn't a big deal though. And both unsharing a namespace (which causes a ns_container_clone) and mounting the hierarchy are done by userspace, so if for some installation it is a big deal, the init scripts on initrd can mount the hierarchy before doing any unsharing.

Can you think of a reason why losing a few directories matters?

-serge
