
Subject: Re: [ckrm-tech] [RFC] Resource Management - Infrastructure choices
Posted by [Paul Menage](#) on Tue, 31 Oct 2006 16:34:52 GMT

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

On 10/31/06, Pavel Emelianov <xemul@openvz.org> wrote:

>
> That's functionality user may want. I agree that some users
> may want to create some kind of "persistent" beancounters, but
> this must not be the only way to control them. I like the way
> TUN devices are done. Each has TUN_PERSIST flag controlling
> whether or not to destroy device right on closing. I think that
> we may have something similar - a flag BC_PERSISTENT to keep
> beancounters with zero refcounter in memory to reuse them.

How about the cpusets approach, where once a cpuset has no children and no processes, a usermode helper can be executed - this could immediately remove the container/bean-counter if that's what the user wants. My generic containers patch copies this from cpusets.

>
> Moreover, I hope you agree that beancounters can't be made as
> module. If so user will have to built-in configs, and thus
> CONFIG_CONFIGFS_FS essentially becomes "bool", not a "tristate".

How about a small custom filesystem as part of the containers support, then? I'm not wedded to using configfs itself, but I do think that a filesystem interface is much more debuggable and extensible than a system call interface, and the simple filesystem is only a couple of hundred lines.

Paul
