
Subject: Re: [PATCH 0/7] containers (V7): Generic Process Containers

Posted by [Sam Vilain](#) on Mon, 12 Feb 2007 22:38:32 GMT

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

menage@google.com wrote:

> Generic Process Containers

> -----

>

> There have recently been various proposals floating around for
> resource management/accounting and other task grouping subsystems in
> the kernel, including ResGroups, User BeanCounters, NSProxy
> containers, and others. These all need the basic abstraction of being
> able to group together multiple processes in an aggregate, in order to
> track/limit the resources permitted to those processes, or control
> other behaviour of the processes, and all implement this grouping in
> different ways.

>

I know I'm a bit out of touch, but AUI the NSProxy *is* the container.

We decided a long time ago that a container was basically just a set of namespaces, which includes all of the subsystems you mention.

This would suggesting re-write this patchset, part 2 as a "CPUSet namespace", part 4 as a "CPU scheduling namespace", parts 5 and 6 as "Resource Limits Namespace" (drop this "BeanCounter" brand), and of course part 7 falls away.

Sam.

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
