

On 2/20/07, Sam Vilain <sam@vilain.net> wrote:

>
> The term "segregated group of processes" is too vague. Segregated for
> what? What is the kernel supposed to do with this information?

The generic part of the kernel just keeps track of the fact that they're segregated (and their children, etc).

It's the clients of this subsystem (virtual servers, resource controllers) that can decide to give different per-process behaviour based on the membership of various groups.

> Did you like the names I came up with in my original reply?

>
> - Cpuset namespace for CPU partitioning
> - Resource namespaces:
> - cpusched namespace for CPU
> - ulimit namespace for memory
> - quota namespace for disk space
> - io namespace for disk activity
> - etc

This is a strange abuse of the term "namespace".
http://en.wikipedia.org/wiki/Namespace_%28computer_science%29

For the virtual server work that you're doing, namespace is a good term:

pids name processes, hence a pid namespace lets you have multiple distinct mappings from pids to processes
filenames name files, so a filename (or mount) namespace lets you have multiple distinct mappings from filenames to files.

For resource QoS control, it doesn't really make sense to talk about namespaces. We're not virtualizing resources to rename them for different virtual servers, we're limiting the quality of access to the resources.

But the semantics of the term "namespace" notwithstanding, you're equating a virtual server namespace (pid, ipc, uts, mounts, etc) with a resource controller (memory, I/O, etc) in terms of their place in a hierarchy, which is something I agree with. All of these subsystems can be considered to be units that can be associated with groups of processes; the ultimate grouping of processes is something that we're both ultimately referring to as a container.

>
> Maybe what's missing is a set of helper macros/functions that assist
> with writing new namespaces.

That's pretty much what my containers patch does - provides a generic way to associate state (what you're referring to as a namespace) with groups of processes and a consistent user-space API for manipulating them.

Paul

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