Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!
Posted by Sam Vilain on Thu, 08 Mar 2007 03:32:27 GMT
View Forum Message <> Reply to Message

Paul Menage wrote:
> I made sure to check [...]wikipedia.org[...] when this argument started ... :-)
>

Wikipedia?! That's not a referen[...]

oh bugger it. I've vented enough today and we're on the same page now I
think.

>> This is the classic terminology problem between substance and function.
>> ie, some things share characteristics but does that mean they are the
>> same thing?
>>
>
> Aren't you arguing my side here? My point is that what I'm trying to
> add with "containers" (or whatever name we end up using) can't easily
> be subsumed into the "namespace" concept, and you're arguing that they
> should go into nsproxy because they share some characteristics.
>

Ok, they share this characteristic with namespaces: that they group
processes. So, they conceptually hang off task_struct. But we put them
on ns_proxy because we've got this vague notion that things might be
better that way.

>> about this you still insist on calling this sub-system specific stuff
>> the "container",
>>
> Uh, no. I'm trying to call a *grouping* of processes a container.
>

Ok, so is this going to supplant the namespaces too?

>> and then go screaming that I am wrong and you are right
>> on terminology.
>>
>
> Actually I asked if you/Eric had better suggestions.
>

Cool, let's review them.

Me, 07921311:38+12:

> 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.
Me, 07022110:58+12:
> 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

Ok, there's nothing original or useful there; I'm obviously quite deliberately still punting on the
issue.

Eric, 07030718:32-07:
> Pretty much. For most of the other cases I think we are safe referring
> to them as resource controls or resource limits.   I know that roughly
> covers what cpusets and beancounters and ckrm currently do.

Let's go back in time to the thread I referred to:

Me, 06032209:08+12 and nearby posts
> - "vserver" spelt in full
> - family
> - container
> - jail
> - task_ns (sort for namespace)
> Using the term "box" and ID term "boxid":
> create_space - creates a new space and "hashes" it

Kirill, 06032418:36+03:
> I propose to use "namespace" naming.
> 1. This is already used in fs.
> 2. This is what IMHO suites at least OpenVZ/Eric
> 3. it has good acronym "ns".

Right. So, now I'll also throw into the mix:

  - supply chains (think supply and demand)
  - accounting classes

Do any of those sound remotely close? If not, your turn :)

And do we bother changing IPC namespaces or let that one slide?

Sam.

_____