Subject: Re: [PATCH 0/2] resource control file system - aka containers on top of nsproxy!
Posted by Kirill Korotaev on Fri, 02 Mar 2007 15:45:06 GMT
View Forum Message <> Reply to Message

Paul,

>>I suspect we can make cpusets also work
>>on top of this very easily.
>
>
> I'm skeptical, and kinda worried.
>
> ... can you show me the code that does this?
don't worry. we are not planning to commit any code breaking cpusets...
I will be the first one against it.

> Namespaces are not the same thing as actual resources
> (memory, cpu cycles, ...). Namespaces are fluid mappings;
> Resources are scarce commodities.
hm... interesing comparison.
as for me, I can't see much difference between virtualization namespaces
and resource namespaces.

Both have some impact on what the task in the namespace can do and what can't.
The only difference is that virtualization namespaces usually also
make one user to be invisible to another. That's the only difference imho.

Also if you take a look at IPC namespace you'll note that IPC
can also limit IPC resources in question.
So it is kinda of virtualization + resource namespace.

> I'm wagering you'll break either the semantics, and/or the
> performance, of cpusets doing this.
I like Paul's containers patch. It looks good and pretty well.
After some of the context issues are resolved it's fine.
Maybe it is even the best way of doing things.

Thanks,
Kirill

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