
Subject: Re: [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by [Herbert Poetzl](#) on Sat, 03 Mar 2007 17:45:53 GMT

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

On Fri, Mar 02, 2007 at 06:45:06PM +0300, Kirill Korotaev wrote:

> 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... interesting comparison.

> as for me, I can't see much difference between virtualization

> namespaces and resource namespaces.

I agree here, there is not much difference for the following aspects:

- resource accounting, limits and namespaces apply to a group of processes
- they isolate those processes in some way from other groups of processes
- they apply a virtual view and/or limitation to those processes

> 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.

IMHO invisibility only applies to the pid space :)

but as I said, the processes are isolated in some way, might it be pids, networking, ipc, uts or filesystem, similar can be said for resource limits and resource accounting, where you are only focusing on a certain group of processes, applying an artificial limit and ideally virtualizing all kernel interfaces

in such a way, that it looks like the artificial limit
is a real physical limitation

> 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.

yes, but they do it in a way a normal Linux system
would do, so no 'new' limits there, unless you
disallow changing those limits from inside a space

best,
Herbert

> 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>

Containers mailing list

Containers@lists.osdl.org

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