
Subject: Re: [PATCH 1/4] Virtualization/containers: introduction

Posted by [dev](#) on Tue, 07 Feb 2006 16:16:36 GMT

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

>>I can't think of any real use cases where you would specifically want A)
>>without B).

> You misrepresent my approach.
[...]

> Second I am not trying to just implement a form of virtualizing PIDs.
> Heck I don't intend to virtualize anything. The kernel has already
> virtualized everything I require. I want to implement multiple
> instances of the current kernel global namespaces. All I want is
> to be able to use the same name twice in user space and not have
> a conflict.

if you want not virtualize anything, what is this discussion about? :)

can you provide an URL to your sources? you makes lot's of statements
about that your network virtualization solution is better/more complete,
so I'd like to see your solution in whole rather than only words.

Probably this will help.

> I disagree with a struct container simply because I do not see what
> value it happens to bring to the table. I have yet to see a problem
> that it solves that I have not solved yet.

again, source would help to understand your solution and problem you
solved and not solved yet.

> In addition I depart from vserver and other implementations in another
> regard. It is my impression a lot of their work has been done so
> those projects are maintainable outside of the kernel, which makes
> sense as that is where those code bases live. But I don't think that
> gives the best solution for an in kernel implementation, which is
> what we are implementing.

These solutions are in kernel implementations actually.

Kirill
