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

Posted by [ebiederm](#) on Wed, 08 Feb 2006 21:04:14 GMT

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

Dave Hansen <haveblue@us.ibm.com> writes:

> On Wed, 2006-02-08 at 18:36 +0300, Kirill Korotaev wrote:
>> - full isolation can be inconvenient from containers management point of
>> view. You will need to introduce new modified tools such as top/ps/kill
>> and many many others. You won't be able to strace/gdb processes from the
>> host also.
>
> I'd like to put a theory out there: the more isolation we perform, the
> easier checkpointing and migration become to guarantee.
>
> Agree? Disagree?

Agree. But that does not address the reasons OpenVZ and Vserver exist.

> But, full isolation is hard to code.

Disagree. If you limit your self to just changing the code that translates from names to objects it is a very narrow slice of code, and there are very few surprises. There is a lot of grunt work involved but it is easy to tell if you got everything and did it correctly.

Other approaches are more adhoc, take short cuts, and seem prone to missing the corner cases.

> The right approach is very likely
> somewhere in the middle where we require some things to happen
> underneath us. For instance, requiring that the filesystem be made
> consistent if a container is moved across systems.

Possibly. That is very out from where we are at the moment.
Let's get the isolation and see where we are at.

Eric
