
Subject: Re: [PATCH 0/9] namespaces: Introduction

Posted by [serue](#) on Fri, 19 May 2006 19:38:02 GMT

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

Quoting Hua Zhong (hzhong@gmail.com):

> > snapshot/restart/migration worry me. If they require
> > complete serialisation of complex kernel data structures then
> > we have a problem, because it means that any time anyone
> > changes such a structure they need to update (and test) the
> > serialisation.
>
> Checkpoint/Restart/Migration could be very complicated if done at OS level (per
process/process group/or any subset of an OS). But
> it is much simpler if done on virtual machine level (VMWare/Xen) because there is a natural and
clear boundary, and doesn't get
> affected if the OS kernel internal changes.
>
> It's good to see some progress in supporting virtualization in Linux, but as Andrew put it, some
big decisions need to be made
> up-front. One big question is actually how many virtualization technologies Linux should
support? Particularly, does it need to
> support both OS-level virtualization and VM-level virtualization? And why? And to what degree?

Because migration can be used for more than one purpose. One such
purpose is load-balancing large numbers of jobs. If you have large
numbers of jobs, you do not want the resource overhead of a full OS for
each migrateable job.

The reason it is deemed simpler at the vm level, as you point out, is
that resources are naturally isolated. The same work which will prepare
the kernel for vserver/openvz functionality will isolate kernel resources
between vserver/containers, making c/r and migration at that level level
much simpler.

thanks,
-serge
