

---

Subject: Re: [RFC] Virtualization steps

Posted by [Herbert Poetzl](#) on Tue, 28 Mar 2006 15:35:58 GMT

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

---

On Tue, Mar 28, 2006 at 07:15:17PM +1000, Nick Piggin wrote:

> Kirill Korotaev wrote:

> >

> >Nick, will be glad to shed some light on it.

> >

>

> Thanks very much Kirill.

>

> I don't think I'm qualified to make any decisions about this,

> so I don't want to detract from the real discussions, but I

> just had a couple more questions:

>

> >First of all, what it does which low level virtualization can't:

> >- it allows to run 100 containers on 1GB RAM

> > (it is called containers, VE - Virtual Environments,

> > VPS - Virtual Private Servers).

> >- it has no much overhead (<1-2%), which is unavoidable with hardware

> > virtualization. For example, Xen has >20% overhead on disk I/O.

>

> Are any future hardware solutions likely to improve these problems?

not really, but as you know, "640K ought to be enough for anybody", so maybe future hardware developments will make shared resources possible (with different kernels)

> >OS kernel virtualization

> >-----

>

> Is this considered secure enough that multiple untrusted VEs are run

> on production systems?

definitely! there are many, many, hosting providers using exactly this technology to provide Virtual Private Servers for their customers, of course, in production

> What kind of users want this, who can't use alternatives like real

> VMs?

well, the same users who do not want to use Bochs for emulating a PC on a PC, when they can use UML for example, because it's much faster and easier to use ...

aside from that, Linux-VServer for example, is not only designed to create complete virtual servers, it also

works for service separation and increasing security for many applications, like for example:

- test environments (one guest per distro)
- service separation (one service per 'container')
- resource management and accounting

> >Summary of previous discussions on LKML

> >-----

>

> Have their been any discussions between the groups pushing this  
> virtualization, and ...

yes, the discussions are ongoing ... maybe to clarify the situation for the folks not involved (projects in alphabetical order):

FreeVPS (Free Virtual Private Server Solution):

=====

[<http://www.freevps.com/>]

not pushing for inclusion, early Linux-VServer spinoff, partially maintained but they seem to have other interests lately

Alex Lyashkov (FreeVPS kernel maintainer)

[Positive Software Corporation <http://www.freevps.com/>]

BSD Jail LSM (Linux-Jails security module):

=====

[<http://kerneltrap.org/node/3823>]

Serge E. Hallyn (Patch/Module maintainer) [IBM]  
interested in some kind of mainline solution

Dave Hansen (IBM Linux Technology Center)  
interested in virtualization for context/container migration

Linux-VServer (community project, maintained):

=====

[<http://linux-vserver.org/>]

Jacques Gelinas (previous VServer maintainer)  
not pushing for inclusion

Herbert Poetzl (Linux-VServer kernel maintainer)  
not pushing for inclusion, but I want to make damn sure that there does not come bloat into the kernel

and the mainline efforts will be usable for  
Linux-VServer and similar ...

Sam Vilain (Refactoring Linux-VServer patches)  
[Catalyst <http://catalyst.net.nz/>]  
trying hard to provide a simple/minimalistic version  
of Linux-VServer for mainline

many others, not really pushing anything here :)

OpenVZ (open project, maintained, subset of Virtuozzo(tm)):

=====  
[<http://openvz.org/>]

Kir Kolyshkin (OpenVZ maintainer):  
[SWsoft <http://www.swsoft.com> I gues?]  
maybe pushing for inclusion ...

Kirill Korotaev (OpenVZ/Virtuozzo kernel developer?)  
[SWsoft <http://www.swsoft.com>]  
heavily pushing for inclusion ...

Alexey Kuznetsov (Chief Software Engineer)  
[SWsoft <http://www.swsoft.com>]  
not pushing but supporting company interrests

PID Virtualization (kernel branch for inclusion):

=====

Eric W. Biederman (branch developer/maintainer)  
[XMission <http://xmission.com/>]

Virtuozzo(tm) (Commercial solution form SWsoft):

=====

[<http://www.virtuozzo.com/>]

not involved yet, except via OpenVZ

Stanislav Protasov (Director of Engineering)  
[SWsoft <http://www.swsoft.com>]

A ton of IBM and VZ folks are not listed here, but I  
guess you can figure who is who from the email addresses

there are also a bunch of folks from Columbia and  
Princeton university interested and/or involved in  
kernel level virtualization and context migration.

please extend this list where appropriate, I'm pretty sure I forgot at least five important/involved persons

> important kernel developers who are not part of a virtualization  
> effort?

no idea, probably none for now ...

> Ie. is there any consensus about the future of these patches?

what patches? what future?

HTC,  
Herbert

> Thanks,

> Nick

>

> --

> SUSE Labs, Novell Inc.

> Send instant messages to your online friends <http://au.messenger.yahoo.com>

---