
Subject: Re: first stable release of OpenVZ kernel virtualization solution

Posted by [Ingo Molnar](#) on Tue, 06 Dec 2005 17:12:58 GMT

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

* Kirill Korotaev <dev@sw.ru> wrote:

> >my point is, that such a swap or writeout related slowdown of a highprio
> >instance can be just as bad as a real DoS - and it brings us essentially
> >back to where we started with vserver. (and writeout related slowdowns
> >of unrelated instances cannot be avoided even with the most conservative
> >UBC settings, correct?)

> We plan to use CFQv2 in some near future, but currently writeout is
> not controlled by UBC anyhow. The only note is that currently used
> disk I/O scheduler (anticipatory) behaves quite well when one VPS is
> doing massive writes... Disk I/O is a kind of problem for any of
> existing virtualization solutions and OpenVZ is not different here...

it's not just massive disk IO and IO starvation of another instance
(which can be mitigated by isolating the disks of instances), it's also
the starvation of RAM of another instance.

Or is that solved already? I.e. can high-prio instances have a
guaranteed amount of RAM set aside for them - even if they do not make
use of that guaranteed amount at the moment?

this is where the Xen approach still seems to differ so much: there the
RAM assigned to an instance truly belongs to that instance. I.e. you can
achieve guaranteed service, no matter what another instance does.

Ingo
