
Subject: Re: [ANNOUNCE] first stable release of OpenVZ kernel virtualization solution

Posted by [Ingo Molnar](#) on Tue, 06 Dec 2005 12:01:11 GMT

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

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

> >interesting. Have you tested the corner case of: 'one lowprio VPS is
> >swapping like mad', how it affects highprio VPSs?

> Both VPSs have a single page cache. So if one VPS is swapping like hell,
> it's neighbour is swapping as well. This naturally means that node is
> out of memory since you created overcommitted configuration.
> It is up to you whether:
> - to limit the offender
> - kill the offender
> - migrate the high-prio or low-prio VPS to another node
> - add RAM :)

well, the other solution is to let certain instances overcommit
userspace RAM, and to just use a swap device for that scenario. An admin
can fix or shut down the offender, but other, more important instances
would still be up and running. Hard limits have other problems: they are
hard failures, instead of graceful failures.

by 'swapping madly' i dont mean lowmem pressure, but plain userspace VM
pressure. I fear it's not flexible enough to not allow for that. I.e. it
would be nice to extend the beancounters to let userspace to _swap_
instead of exposing it to a hard limit.

Ingo
