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

Posted by [Ingo Molnar](#) on Tue, 06 Dec 2005 13:11:00 GMT

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

* Ingo Molnar <mingo@elte.hu> wrote:

> 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.

maybe i'm banging on open doors, but the same would be the case not only for userspace-VM overcommit, but also for dirty data. I.e. there should be (already is?) a per-instance 'dirty data threshold', to not force other instances into waiting for writeout/swapout to happen.

or does the OVZ method assume a totally swapless system?

Ingo
