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

Posted by [Andrey Savochkin](#) on Tue, 06 Dec 2005 14:01:15 GMT

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

On Tue, Dec 06, 2005 at 02:11:20PM +0100, Ingo Molnar wrote:

>

> * 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.

OVZ certainly has room for improvements with respect to swap.

What I want to point out is that swapout management is a complex task.

When a low-priority VPS exceeds its limits, it is not always beneficial for others to make it swap out: swapout wastes disk bandwidth, and to some extent

CPU power. 'Dirty data threshold' could have helped, but it reduces the overall performance of the system, especially if the number of VPSs is small.

Imagine only one VPS running: artificial 'dirty data threshold' would certainly be counter-productive.
