Subject: Re: first stable release of OpenVZ kernel virtualization solution
Posted by Ingo Molnar on Tue, 06 Dec 2005 14:04:16 GMT
View Forum Message <> Reply to Message

* Andrey Savochkin <saw@sawoct.com> wrote:

> > 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 benefitial 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.

but what you have right now is an in essence swapless system, correct?
Do you support swapping at all in OVZ instances?

Ingo