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

Posted by [Ingo Molnar](#) on Tue, 06 Dec 2005 14:22:15 GMT

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

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

> >but what you have right now is an in essence swapless system, correct?

> >Do you support swapping at all in OVZ instances?

> Yes, swap is supported and processes are swapped in/out as in usual

> kernel. The only difference in comparison with std kernel is that UBC

> limits the amount of swappable memory VPS can have.

i mean, the only way to protect a high-prio instance against a low-prio instance doing heavy swapout is by making the low-prio instance swapless, correct? (either by not enabling it to swap at all, or by tweaking the UBC limits in a way that can never lead to swapping).

how about the 'dirty data creator' scenario: an instance filling up all of the RAM with dirty data, at which point a highprio instance is significantly impacted.

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?)

Ingo
