Subject: Re: [ANNOUNCE] first stable release of OpenVZ kernel virtualization solution
Posted by dev on Tue, 06 Dec 2005 11:48:38 GMT
View Forum Message <> Reply to Message

Ingo, I added Andrey note, added our maillist to CC.

> in general, i think the biggest resource-isolation issue for big servers
> is lowmem. Do you have any ideas how to handle that within OVZ? The
> scenario is: one (relatively low-prio) instance starves another
> (high-prio) OVZ instance of lowmem, getting the system into a kswapd
> storm.
> Ingo

We have per VPS UBC (user beancounters) parameters called "kmemsize"
(almost all kernel structures are accounted into this - page tables,
vmas, etc.), tcprcvbuf, tcpsndbuf and others which allows to control VPS
usage of low-mem.

i.e. if OpenVZ is configured appropriately such situation should not
happen (provided host system is not highly overcommited).
Situation with overcommit happened on i386 >4GB RAM in 2.4 kernels. But
as you remember 4GB split helped a lot in this case. In 2.6 kernels
situation is really much better and out tests with high number of VPSs
work correctly even without 4GB split.

Additional note from Andrey Savochkin:
some resources such as disk space can be highly overcommited. Low-mem is
special resource and correctly configured OpenVZ server should not have
much overcommit for lowmem. This eliminates starvation on lowmem and
kswapd storms.

Kirill