Subject: Re: [ANNOUNCE] first stable release of OpenVZ kernel virtualization solution
Posted by dev on Tue, 06 Dec 2005 11:55:24 GMT
View Forum Message <> Reply to Message

Ingo Molnar wrote:
> * Kirill Korotaev <dev@sw.ru> wrote:
>
>
>>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.
>
>
> interesting. Have you tested the corner case of: 'one lowprio VPS is
> swapping like mad', how it affects highprio VPSs?
Both VPSs have a single page cache. So if one VPS is swapping like hell,
  it's neighbour is swapping as well. This naturally means that node is
out of memory since you created overcommitted configuration.
It is up to you whether:
- to limit the offender
- kill the offender
- migrate the high-prio or low-prio VPS to another node
- add RAM :)

There are many possible actions here...
You can't travel with all your friends/relatives to the sea side in a
car, yeah? You need either to have many cars or to take a bus. Up to you.

Kirill