
Subject: Re: A question about Node RAM

Posted by [Martin Dobrev](#) on Fri, 06 Jan 2012 18:23:23 GMT

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

Sounds really massive. I'm not really sure if the I/O will fit, even with RAID 10, but of course it only depends on how I/O intensive your VEs will be. On the other hand, as Kiril already mentioned, OpenVZ kernel is not so different from mainstream kernels and as such there should be no limit in the number of the running VEs.

Martin Dobrev

Sent from iPhonespam SPAMSPAM 4

On 06.01.2012, at 20:08, Quentin MACHU <quentin.machu@gmail.com> wrote:

> Hello,
>
> Thanks for this answer.
> So, we can use 128Gb/256Gb server ? =]
>
> Actually, we're working on Debian 6.
> Do you have any tips on Distro / Kernel ?
>
> Debian 6 + Kernel from Debian repos is really stable ? Debian 5 more maybe ?
>
> We'll have 6*3To HardDrive SAS in RAID 10 to improve I/O
> And Two Opteron 6128 8 cores Magny-Cours 8x 2Ghz.
>
> Do you think it's ok for something like 126 VM with 1Gb of RAM ? =)
>
> Thanks for all :)
>
>
> 2012/1/6 Kirill Korotaev <dev@parallels.com>
> Sure, it's old information and likely it was about 32bit kernels which are limited to 64GB just because CPUs are... :)
> 64bit kernels are not limited anyhow and OpenVZ is not different in this regard from standard Linux.
>
> fixed a couple of places I found with 64GB mentioning:
> [http://wiki.openvz.org/Different_kernel_flavors_\(UP,_SMP,_ENTERPRISE,_ENTNOSPLIT\)](http://wiki.openvz.org/Different_kernel_flavors_(UP,_SMP,_ENTERPRISE,_ENTNOSPLIT))
> http://wiki.openvz.org/FAQ#How_scalable_is_OpenVZ.3F
>
> Thanks,
> Kirill
>
> On Jan 6, 2012, at 20:59 , Quentin MACHU wrote:
>

> > Hello,
> >
> > I've a question for this mailing-list ^^
> >
> > My enterprise is going to order a 128Gb of RAM server.
> > I saw that the OpenVZ Kernel can only support 64Gb.
> >
> > That's because the wiki isn't up to date ?
> > What's about that ?
> > How to bypass this limit ? Can we ?
> > Recompiling the kernel.. ?
> >
> > It's important for us =)
> >
> > Thanks !
> --
> Cordialement,
> MACHU Quentin
>
