
Subject: Re: openvz and SuSE
Posted by [dim](#) on Fri, 03 Feb 2006 15:03:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

OpenVZ allows to use firewall both on HN and in VPSs.
And I was completely wrong talking that there is no way!!!
You can set no IP to eth0, but have VPSs accessible from intranet.
Here:

```
ifconfig eth0 0
ip r add 10.0.0.0/8 dev eth0
ip r add default via GW_ADDR
sysctl -w net.ipv4.conf.eth0.proxy_arp=1
ip route add VPS1_IP dev venet0
vzctl start VPS1
there should be some warnings, just skip them.
```

So, the main point is to enable arp_proxying on intranet interface and have added VPS related routes before VPS start (else vzctl will return with error - you may fix this in /usr/lib/vzctl/scripts/vps-functions)

On Friday 03 February 2006 16:37, Daniel Bauer wrote:

> From: "Mishin Dmitry" <dim@sw.ru>

>

> > On Friday 03 February 2006 15:38, Daniel Bauer wrote:

> >> is it possible to take another way to work on Ethernet level, because

> >> I

> >> don't want a official IP for the host.

> >> 1. security

> >> 2. no need for

> >> 3. one official IP less for each block

> >

> > If you have only one or two VPSs, you can use real devices dedicated
> > to each

> > VPS, but this is not your case. For now, we don't work on Ethernet
> > level and

> > you are required to have one more real IP for the block.

> >

> > I suppose, that from security point of view it is a not big deal,

> > while you

> > can use netfilter to protect it and additionally all VPSs, because

> > their

> > traffic goes through HN route tables.

> >

> > If it is still the problem, you can check Virtuozzo's Name Based
> > Hosting

> > feature - it allows to use one real IP for multiple VPSs (pop, smtp,

> > http,

> > ftp)

>
> Hello Dmitry,
>
> thanks for your explanation.
>
> If I understand you right, you do the firewalling on the host, not in
> the VPS. I think it will work and I could afford one more IP for the
> host, but my opinion was to have less as possible on my host and let the
> VPS do the work ;)
>
> Thanks again
> Daniel
>
--
Thanks,
Dmitry.
