

---

Subject: openvz and SuSE

Posted by [Daniel Bauer](#) on Wed, 01 Feb 2006 13:25:56 GMT

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

---

Hello @all,

I'm tried to work with XEN3 and SuSE but I've the problem with compiling my own kernel (SMP, 64GB). So I've read the artikel about openvz and think it's better for me because I only need linux VPS on i386 machines.

Is it possible to patch a running SuSE Box with openvz? And how to setup a SuSE VPS on this host, are there any good HowTos for this? German is preferred.

Thanks a lot  
Daniel

---

---

Subject: Re: openvz and SuSE

Posted by [kir](#) on Wed, 01 Feb 2006 13:43:03 GMT

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

---

Which suse version do you have?

Specifically, we do not currently provide suse rpms, but the ones made for redhat should be OK. Latest vzctl (an utility to control VPSs) version that fixes initscript to work with SuSE will be released today or tomorrow.

As for the suse VPS, there is a opensuse-10 precreated VPS template available. Unfortunately for its networking to work it requires some stuff from kernel (sysfs support from inside a VPS), and this is off by default in our kernel. The situation will be fixed with the release of 068 (or later) kernel there sysfs can be turned on/off per VPS. For now, you need to change one config setting and recompile the kernel to get opensuse-10 template working.

Our future plans include releasing a custom kernel which is basically SuSE kernel (all the SuSE patches and stuff) with OpenVZ functionality added. We will also be creating more suse templates.

Kir.

Daniel Bauer wrote:

> Hello @all,

>

> I'm tried to work with XEN3 and SuSE but I've the problem with

> compiling my own kernel (SMP, 64GB). So I've read the artikel about  
> openvz and think it's better for me because I only need linux VPS on  
> i386 machines.  
>  
> Is it possible to patch a running SuSE Box with openvz? And how to  
> setup a SuSE VPS on this host, are there any good HowTos for this?  
> German is preferred.  
>  
> Thanks a lot  
> Daniel

---

---

Subject: Re: openvz and SuSE  
Posted by [Daniel Bauer](#) on Wed, 01 Feb 2006 15:06:45 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hi Kir,

From: "Kir Kolyshkin" <[kir@openvz.org](mailto:kir@openvz.org)>  
> Which suse version do you have?

9.3 or 10.0, I prefer 9.3, the VPS should be SuSE 9.3

> Specifically, we do not currently provide suse rpms, but the ones made  
> for redhat should be OK. Latest vzctl (an utility to control VPSs)  
> version that fixes initscript to work with SuSE will be released today  
> or tomorrow.  
>  
> As for the suse VPS, there is a opensuse-10 precreated VPS template  
> available. Unfortunately for its networking to work it requires some  
> stuff from kernel (sysfs support from inside a VPS), and this is off  
> by default in our kernel. The situation will be fixed with the release  
> of 068 (or later) kernel there sysfs can be turned on/off per VPS. For  
> now, you need to change one config setting and recompile the kernel to  
> get opensuse-10 template working.  
>  
> Our future plans include releasing a custom kernel which is basically  
> SuSE kernel (all the SuSE patches and stuff) with OpenVZ functionality  
> added. We will also be creating more suse templates.

so is it possible to patch a 9.3 kernel with the sources of OpenVZ for  
the host and build templates von 9.3 for the clients?

Daniel

---

---

Subject: Re: openvz and SuSE  
Posted by [kir](#) on Wed, 01 Feb 2006 17:00:28 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Daniel Bauer wrote:

> From: "Kir Kolyshkin" <[kir@openvz.org](mailto:kir@openvz.org)>  
>  
>> Daniel Bauer wrote:  
>>  
>>>> Which suse version do you have?  
>>>  
>>>  
>>> 9.3 or 10.0, I prefer 9.3, the VPS should be SuSE 9.3  
>>  
>>  
>> OK. I remember that somebody created some SuSE templates, search in  
>> [forum.openvz.org](http://forum.openvz.org).  
>  
>  
> I've found this message:  
> <http://forum.openvz.org/index.php?t=msg&th=190&start=0&>  
>  
> and his site  
> <http://web150.dd01.profihoster.net/>  
>  
> but there is only a template for SuSE 9.3 64bit server-image for  
> openVZ (155 MB)

Ask the author about x86 (i386) template as well :) if he will have enough requests he may consider doing it :)

>> If you are brave enough you can try doing it all yourself. This is  
>> not very complicated -- basically you need to create a chrooted  
>> environment with all the needed stuff installed. All the needed stuff  
>> are libraries, startup scripts and init, some programs (like ssh)  
>> etc. Then you run this in a VPS.  
>  
>  
> I think this is the best way, maybe I could use my existing real  
> machines or do I have to include something for VPS in the  
> configurations / patches / modification?

Well, as a minimum you should

- (1) modify `/etc/fstab` so it will not try mounting your real hard disk partition(s)
- (2) remove `getty` from `/etc/inittab`
- (3) some suse versions requires to comment out one line in `/sbin/rc`  
(will tell you later which one)

Everything else is optimizations. This can include:

- (4) quit using tmpfs for /dev
- (5) moving from udev to static /dev
- (6) remove extra entries from /dev to save disk space and inodes
- (7) remove kernel, grub, kernel module tools - those are not needed (in order for RPM to be happy, you might add a short dummy rpm which virtually 'Provides:' all this stuff)
- (8) modify initscripts to boot faster
- (9) ...and so on.

But those are just optimizations -- they are not strictly required.

> I found this kernel for my machine:

> <http://download.openvz.org/kernel/022stab064.1/ovzkernel-enterprise-2.6.8-022stab064.1.i686.rpm>

>

>

> What about security fixes with this kernel?

We have included a lot of security fixes in our kernel (all that are known to date) -- just read change logs.

---

Subject: Re: openvz and SuSE

Posted by [Daniel Bauer](#) on Thu, 02 Feb 2006 16:27:16 GMT

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

---

From: "Kir Kolyshkin" <[kir@openvz.org](mailto:kir@openvz.org)>

> Daniel Bauer wrote:

>> From: "Kir Kolyshkin" <[kir@openvz.org](mailto:kir@openvz.org)>

>>> Daniel Bauer wrote:

>>>> Which suse version do you have?

>>>>

>>>> 9.3 or 10.0, I prefer 9.3, the VPS should be SuSE 9.3

>>>

>>>

>>> OK. I remember that somebody created some SuSE templates, search in

>>> [forum.openvz.org](http://forum.openvz.org).

>>

>>

>> I've found this message:

>> <http://forum.openvz.org/index.php?t=msg&th=190&start=0&>

>>

>> and his site

>> <http://web150.dd01.profihoster.net/>

>>

>> but there is only a template for SuSE 9.3 64bit server-image for

>> openVZ (155 MB)  
>  
> Ask the author about x86 (i386) template as well :) if he will have  
> enough requests he may consider doing it :)  
>  
>>> If you are brave enough you can try doing it all yourself. This is  
>>> not very complicated -- basically you need to create a chrooted  
>>> environment with all the needed stuff installed. All the needed  
>>> stuff are libraries, startup scripts and init, some programs (like  
>>> ssh) etc. Then you run this in a VPS.  
>>  
>>  
>> I think this is the best way, maybe I could use my existing real  
>> machines or do I have to include something for VPS in the  
>> configurations / patches / modification?  
>  
> Well, as a minimum you should  
> (1) modify /etc/fstab so it will not try mounting your real hard disk  
> partition(s)  
> (2) remove getty from /etc/inittab  
> (3) some suse versions requires to comment out one line in /sbin/rc  
> (will tell you later which one)

ok, that sounds good, I'll have a try with SuSE.

I've started now with installing fedora-core 4 and updated the kernel as mentioned, some prebuild VPS are still running. But how can I access now some physical harddiscs? I need to mount /dev/sdd1 (my data partition). How to enable the second eth interface?

> Everything else is optimizations. This can include:  
>  
> (4) quit using tmpfs for /dev  
> (5) moving from udev to static /dev  
> (6) remove extra entries from /dev to save disk space and inodes  
> (7) remove kernel, grub, kernel module tools - those are not needed  
> (in order for RPM to be happy, you might add a short dummy rpm which  
> virtually 'Provides:' all this stuff)  
> (8) modify initscripts to boot faster  
> (9) ...and so on.  
> But those are just optimizations -- they are not strictly required.  
>  
>> I found this kernel for my machine:  
>> <http://download.openvz.org/kernel/022stab064.1/ovzkernel-entprise-2.6.8-022stab064.1.i686.rpm>  
>>  
>> What about security fixes with this kernel?

>  
> We have included a lot of security fixes in our kernel (all that are  
> known to date) -- just read change logs.

so you update this kernel with actual patches and the OpenVZ Users could  
install the new one everytime it's released?

Daniel

---

---

Subject: Re: openvz and SuSE  
Posted by [kir](#) on Thu, 02 Feb 2006 17:06:10 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Daniel Bauer wrote:

> ok, that sounds good, I'll have a try with SuSE.  
>  
> I've started now with installing fedora-core 4 and updated the kernel  
> as mentioned, some prebuild VPS are still running. But how can I  
> access now some physical harddiscs? I need to mount /dev/sdd1 (my data  
> partition).

See pages 66-67 of OpenVZ User's Guide. In short, here it is.

You just enable device access for a given VPS, like this:  
vzctl set NNN --devnodes sdd1:rw --save

Next, you create the device file (you can just cp /dev/sdd1  
/vz/root/NNN/dev)

Next, you mount your partition from inside a VPS as usual.

> How to enable the second eth interface?

You want it to be dedicated to a single VPS, or you just want to access  
some other network via it?

In the first case, see page 68 of OpenVZ User's Guide. In the second  
case, use it as you would do on a normal Linux box, i.e. add IP, proper  
routes and everything should work.

> so you update this kernel with actual patches and the OpenVZ Users could  
> install the new one everytime it's released?

Exactly.

Regards,

Kir.

---

---

Subject: Re: openvz and SuSE  
Posted by [Daniel Bauer](#) on Thu, 02 Feb 2006 18:36:22 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

From: "Kir Kolyshkin" <[kir@openvz.org](mailto:kir@openvz.org)>  
> Daniel Bauer wrote:  
>  
>> ok, that sounds good, I'll have a try with SuSE.  
>>  
>> I've started now with installing fedora-core 4 and updated the kernel  
>> as mentioned, some prebuild VPS are still running. But how can I  
>> access now some physical harddisks? I need to mount /dev/sdd1 (my  
>> data partition).  
>  
> See pages 66-67 of OpenVZ User's Guide. In short, here it is.  
>  
> You just enable device access for a given VPS, like this:  
> `vzctl set NNN --devnodes sdd1:rw --save`  
>  
> Next, you create the device file (you can just `cp /dev/sdd1`  
> `/vz/root/NNN/dev`)  
>  
> Next, you mount your partition from inside a VPS as usual.  
>  
>> How to enable the second eth interface?  
>  
> You want it to be dedicated to a single VPS, or you just want to  
> access some other network via it?  
>  
> In the first case, see page 68 of OpenVZ User's Guide. In the second  
> case, use it as you would do on a normal Linux box, i.e. add IP,  
> proper routes and everything should work.

Hi Kir,

sorry for asking silly questions, I don't see the manual, only the quick install, I'll be back when I've read it ;)

Thanks a lot for all your help. OpenVZ is basicly working now within 1 hours, GREAT! I think I'll do a custom SuSE install at the weekend , if it works and I'm ready I'll give you the precreated tar.gz if you want to.

Daniel

---

---

Subject: Re: openvz and SuSE  
Posted by [Daniel Bauer](#) on Fri, 03 Feb 2006 08:59:50 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

From: "Kir Kolyskin" <[kir@openvz.org](mailto:kir@openvz.org)>  
> Daniel Bauer wrote:  
>> ok, that sounds good, I'll have a try with SuSE.

Is there a howto for creating templates, or should I start with building the directory tree by myself?

Why is the stuff double in private and root, I thought that the root is only the mounted stuff from private?

>> I've started now with installing fedora-core 4 and updated the kernel  
>> as mentioned, some prebuild VPS are still running. But how can I  
>> access now some physical harddisks? I need to mount /dev/sdd1 (my  
>> data partition).  
>  
> See pages 66-67 of OpenVZ User's Guide. In short, here it is.  
>  
> You just enable device access for a given VPS, like this:  
> vzctl set NNN --devnodes sdd1:rw --save  
>  
> Next, you create the device file (you can just cp /dev/sdd1  
> /vz/root/NNN/dev)  
>  
> Next, you mount your partition from inside a VPS as usual.

Great!!!

>> How to enable the second eth interface?  
>  
> You want it to be dedicated to a single VPS, or you just want to  
> access some other network via it?  
>  
> In the first case, see page 68 of OpenVZ User's Guide. In the second  
> case, use it as you would do on a normal Linux box, i.e. add IP,  
> proper routes and everything should work.

I've three nics, but only one venet0.  
i.e. eth0 is connected to intranet for use of VPS  
    (10.x.x.x)  
    eth1 is connected to internet for use of VPS  
    (20.x.x.x)  
    eth2 s connected to intranet only for use of the Host  
    (30.x.x.x)



I'd like to use:

VPS1: intranet (eth0)

VPS2: internet (eth1)

VPS3: intranet (eth0) + internet (eth1)

On the host (eth2) I don't want to have an IP of the intranet or internet, the nics should be used for the VPS.

As I understand it's not possible to share a real nic with Host and diff. VPS, so howto setup this configuration?

Daniel

---

---

Subject: Re: openvz and SuSE

Posted by [kir](#) on Fri, 03 Feb 2006 09:20:29 GMT

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

---

Daniel Bauer wrote:

> From: "Kir Kolyshkin" <[kir@openvz.org](mailto:kir@openvz.org)>

>

>> Daniel Bauer wrote:

>>

>>> ok, that sounds good, I'll have a try with SuSE.

>>

> Is there a howto for creating templates, or should I start with  
> building the directory tree by myself?

No, there is no formal howto, but somehow people manage to find out all the details. It is probably not that hard, and there are two approaches:

(A)

Currently our template tools can use yum-enabled repositories, so the way to go is:

1. Take yum-enabled repo (or create one from a bunch of RPMS using createrepo).
2. Take one of the existing OS templates metadata and find out what is it.
3. Create a metadata for your OS template, based on the above.
4. Run vzpkgcache to create a template cache.

(B)

Alternative approach is just to install a set of packages into chrooted environment (in RPM world you use 'rpm -i --root' for that, in Debian world you use 'debootstrap'), then do some fine tuning and finally tar this environment. The final tarball can be placed to /vz/template/cache and used as a template. Basically, you need to install some basic stuff

- system libraries, init scripts, useful binaries etc.

> Why is the stuff double in private and root, I thought that the root is  
> only the mounted stuff from private?

It is not doubled. You can think of it as a bind mount (although it's a bit more than that).

Private is a source for mount (i.e. what we mount), while root is a mountpoint (i.e. where we mount it).

---

---

Subject: Re: openvz and SuSE

Posted by [dim](#) on Fri, 03 Feb 2006 09:30:03 GMT

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

---

> I've three nics, but only one venet0.

> i.e. eth0 is connected to intranet for use of VPS

> (10.x.x.x)

> eth1 is connected to internet for use of VPS

> (20.x.x.x)

> eth2 s connected to intranet only for use of the Host

> (30.x.x.x)

>

> I'd like to use:

> VPS1: intranet (eth0)

> VPS2: internet (eth1)

> VPS3: intranet (eth0) + internet (eth1)

>

> On the host (eth2) I don't want to have an IP of the intranet or

> internet, the nics should be used for the VPS.

>

> As I understand it's not possible to share a real nic with Host and

> diff. VPS, so howto setup this configuration?

You should assign IP from intranet to VPS1, IP for internet to VPS2 and both such IPs to VPS3.

>

> Daniel

>

--

Thanks,  
Dmitry.

---

---

Subject: Re: openvz and SuSE

Posted by [Daniel Bauer](#) on Fri, 03 Feb 2006 10:40:54 GMT

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

---

From: "Mishin Dmitry" <dim@sw.ru>  
>> I've three nics, but only one venet0.  
>> i.e. eth0 is connected to intranet for use of VPS  
>> (10.x.x.x)  
>> eth1 is connected to internet for use of VPS  
>> (20.x.x.x)  
>> eth2 s connected to intranet only for use of the Host  
>> (30.x.x.x)  
>>  
>> I'd like to use:  
>> VPS1: intranet (eth0)  
>> VPS2: internet (eth1)  
>> VPS3: intranet (eth0) + internet (eth1)  
>>  
>> On the host (eth2) I don't want to have an IP of the intranet or  
>> internet, the nics should be used for the VPS.  
>>  
>> As I understand it's not possible to share a real nic with Host and  
>> diff. VPS, so howto setup this configuration?  
> You should assign IP from intranet to VPS1, IP for internet to VPS2  
> and both  
> such IPs to VPS3.

Sorry that doesn't work here if you don't have a IP on the host.

i.e.

Host:

eth0 10.1.0.1/255.255.255.0  
eth1 10.2.0.1/255.255.255.0\*

VPS1:

venet0:0 10.3.0.2/255.255.255.255

VPS2:

venet0:0 10.4.0.3/255.255.255.255\*

VPS3:

venet0:0 10.3.0.4/255.255.255.255  
venet0:1 10.4.0.4/255.255.255.255\*

\* not used in this test

a Client with IP 10.1.0.100 could ping the host

a Client with IP 10.3.0.100 could ping nothing

If I set the host to 10.3.0.1 then I could ping the Host, VPS1 and VPS3.  
So it seems to me that is necessary to have an IP on the host from the  
same subnet.

I see that when I start the VPS a new route for each IP, maybe this is the problem, is it not better to bridge the adapters between Host and VPS?

Bye  
Daniel

---

---

Subject: Re: openvz and SuSE  
Posted by [Daniel Bauer](#) on Fri, 03 Feb 2006 10:43:10 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

From: "Kir Kolyshkin" <kir@openvz.org>  
> Daniel Bauer wrote:  
>> From: "Kir Kolyshkin" <kir@openvz.org>  
>>> Daniel Bauer wrote:  
>>>> ok, that sounds good, I'll have a try with SuSE.  
>>>>  
>> Is there a howto for creating templates, or should I start with  
>> building the directory tree by myself?  
>  
> No, there is no formal howto, but somehow people manage to find out  
> all the details. It is probably not that hard, and there are two  
> approaches:  
>  
> (A)  
> Currently our template tools can use yum-enabled repositories, so the  
> way to go is:  
> 1. Take yum-enabled repo (or create one from a bunch of RPMS using  
> createrepo).  
> 2. Take one of the existing OS templates metadata and find out what is  
> it.  
> 3. Create a metadata for your OS template, based on the above.  
> 4. Run vzpkgcache to create a template cache.  
>  
> (B)  
> Alternative approach is just to install a set of packages into  
> chrooted environment (in RPM world you use 'rpm -i --root' for that,  
> in Debian world you use 'debootstrap'), then do some fine tuning and  
> finally tar this environment. The final tarball can be placed to  
> /vz/template/cache and used as a template. Basically, you need to  
> install some basic stuff - system libraries, init scripts, useful  
> binaries etc.

ok I'll have a try with SuSE ;)

>> Why is the stuff double in private and root, I thought that the root

>> is only the mounted stuff from private?

>

> It is not doubled. You can think of it as a bind mount (although it's  
> a bit more than that).

> Private is a source for mount (i.e. what we mount), while root is a  
> mountpoint (i.e. where we mount it).

ok, that is what I understood, but when I stop the VPS, the root is  
still full of content, shouldn't it be unmounted?

Daniel

---

---

Subject: Re: openvz and SuSE

Posted by [dim](#) on Fri, 03 Feb 2006 10:52:45 GMT

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

---

On Friday 03 February 2006 13:40, Daniel Bauer wrote:

> From: "Mishin Dmitry" <[dim@sw.ru](mailto:dim@sw.ru)>

>

> >> I've three nics, but only one venet0.

> >> i.e. eth0 is connected to intranet for use of VPS

> >> (10.x.x.x)

> >> eth1 is connected to internet for use of VPS

> >> (20.x.x.x)

> >> eth2 s connected to intranet only for use of the Host

> >> (30.x.x.x)

> >>

> >> I'd like to use:

> >> VPS1: intranet (eth0)

> >> VPS2: internet (eth1)

> >> VPS3: intranet (eth0) + internet (eth1)

> >>

> >> On the host (eth2) I don't want to have an IP of the intranet or

> >> internet, the nics should be used for the VPS.

> >>

> >> As I understand it's not possible to share a real nic with Host and

> >> diff. VPS, so howto setup this configuration?

> >

> > You should assign IP from intranet to VPS1, IP for internet to VPS2

> > and both

> > such IPs to VPS3.

>

> Sorry that doesn't work here if you don't have a IP on the host.

> i.e.

> Host:

> eth0 10.1.0.1/255.255.255.0

> eth1 10.2.0.1/255.255.255.0\*

>  
> VPS1:  
> venet0:0 10.3.0.2/255.255.255.255  
>  
> VPS2:  
> venet0:0 10.4.0.3/255.255.255.255\*  
>  
> VPS3:  
> venet0:0 10.3.0.4/255.255.255.255  
> venet0:1 10.4.0.4/255.255.255.255\*  
>  
> \* not used in this test  
>  
> a Client with IP 10.1.0.100 could ping the host  
> a Client with IP 10.3.0.100 could ping nothing  
>  
> If I set the host to 10.3.0.1 then I could ping the Host, VPS1 and VPS3.  
> So it seems to me that is necessary to have an IP on the host from the  
> same subnet.  
You are right.  
>  
> I see that when I start the VPS a new route for each IP, maybe this is  
> the problem, is it not better to bridge the adapters between Host and  
> VPS?  
venet works on IP level, not Ethernet. So, bridging is impossible for now.  
  
>  
> Bye  
> Daniel  
>  
--  
Thanks,  
Dmitry.

---

---

Subject: Re: openvz and SuSE  
Posted by [kir](#) on Fri, 03 Feb 2006 11:04:24 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Daniel Bauer wrote:

>>> Why is the stuff double in private and root, I thought that the root  
>>> is only the mounted stuff from private?  
>>  
>>  
>> It is not doubled. You can think of it as a bind mount (although it's  
>> a bit more than that).  
>> Private is a source for mount (i.e. what we mount), while root is a

>> mountpoint (i.e. where we mount it).  
>  
> ok, that is what I understood, but when I stop the VPS, the root is  
> still full of content, shouldn't it be unmounted?

vzctl stop does unmounting of the root area. If you stop VPS from within a VPS itself, the mount is still there -- and you can unmount manually (vzctl umount NNN).

Creating a cron job for that is problematic - you can not distinguish a case there you stopped VPS from inside so it needs to be unmounted from the case there you just mounted a VPS (say, to do some maintenance). So it is left as is, since there is no problem to left it mounted - it does not consume much resources etc.

---

Subject: Re: openvz and SuSE  
Posted by [Daniel Bauer](#) on Fri, 03 Feb 2006 12:35:57 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

From: "Kir Kolyshkin" <kir@openvz.org>  
> Daniel Bauer wrote:  
>>>> Why is the stuff double in private and root, I thought that the root  
>>>> is only the mounted stuff from private?  
>>>  
>>>  
>>> It is not doubled. You can think of it as a bind mount (although  
>>> it's a bit more than that).  
>>> Private is a source for mount (i.e. what we mount), while root is a  
>>> mountpoint (i.e. where we mount it).  
>>  
>> ok, that is what I understood, but when I stop the VPS, the root is  
>> still full of content, shouldn't it be unmounted?  
>  
> vzctl stop does unmounting of the root area. If you stop VPS from  
> within a VPS itself, the mount is still there -- and you can unmount  
> manually (vzctl umount NNN).

ok, now I see ;)

> Creating a cron job for that is problematic - you can not distinguish  
> a case there you stopped VPS from inside so it needs to be unmounted  
> from the case there you just mounted a VPS (say, to do some  
> maintenance). So it is left as is, since there is no problem to left  
> it mounted - it does not consume much resources etc.

is it better to stop a machine inside by "poweroff" or using vzctl stop?

Daniel

---

---

Subject: Re: openvz and SuSE

Posted by [Daniel Bauer](#) on Fri, 03 Feb 2006 12:38:59 GMT

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

---

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

> On Friday 03 February 2006 13:40, Daniel Bauer wrote:

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

>>

>> >> I've three nics, but only one venet0.

>> >> i.e. eth0 is connected to intranet for use of VPS

>> >> (10.x.x.x)

>> >> eth1 is connected to internet for use of VPS

>> >> (20.x.x.x)

>> >> eth2 s connected to intranet only for use of the Host

>> >> (30.x.x.x)

>> >>

>> >> I'd like to use:

>> >> VPS1: intranet (eth0)

>> >> VPS2: internet (eth1)

>> >> VPS3: intranet (eth0) + internet (eth1)

>> >>

>> >> On the host (eth2) I don't want to have an IP of the intranet or

>> >> internet, the nics should be used for the VPS.

>> >>

>> >> As I understand it's not possible to share a real nic with Host

>> >> and

>> >> diff. VPS, so howto setup this configuration?

>> >>

>> > You should assign IP from intranet to VPS1, IP for internet to VPS2

>> > and both

>> > such IPs to VPS3.

>>

>> Sorry that doesn't work here if you don't have a IP on the host.

>> i.e.

>> Host:

>> eth0 10.1.0.1/255.255.255.0

>> eth1 10.2.0.1/255.255.255.0\*

>>

>> VPS1:

>> venet0:0 10.3.0.2/255.255.255.255

>>

>> VPS2:

>> venet0:0 10.4.0.3/255.255.255.255\*

>>



>> VPS3:  
>> venet0:0 10.3.0.4/255.255.255.255  
>> venet0:1 10.4.0.4/255.255.255.255\*  
>>  
>> \* not used in this test  
>>  
>> a Client with IP 10.1.0.100 could ping the host  
>> a Client with IP 10.3.0.100 could ping nothing  
>>  
>> If I set the host to 10.3.0.1 then I could ping the Host, VPS1 and  
>> VPS3.  
>> So it seems to me that is necessary to have an IP on the host from  
>> the  
>> same subnet.  
> You are right.  
>>  
>> I see that when I start the VPS a new route for each IP, maybe this  
>> is  
>> the problem, is it not better to bridge the adapters between Host and  
>> VPS?  
> venet works on IP level, not Ethernet. So, bridging is impossible for  
> now.

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

Daniel

---

---

Subject: Re: openvz and SuSE  
Posted by [kir](#) on Fri, 03 Feb 2006 12:48:22 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Daniel Bauer wrote:

> is it better to stop a machine inside by "poweroff" or using vzctl stop?

It is a bit better to stop it using vzctl stop -- this way vzctl will unmount the area, delete network routes etc.

In case you run poweroff from inside a VPS, network routes will be cleaned up by vpsnetclean script executed each 5 minutes from cron.

---

---

Subject: Re: openvz and SuSE  
Posted by [dim](#) on Fri, 03 Feb 2006 13:00:15 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

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)

>  
> Daniel  
>  
--  
Thanks,  
Dmitry.

---

Subject: Re: openvz and SuSE  
Posted by [Daniel Bauer](#) on Fri, 03 Feb 2006 13:37:13 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

From: "Mishin Dmitry" <[dim@sw.ru](mailto: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

---

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](mailto: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.

---

---

Subject: Re: openvz and SuSE  
Posted by [jbond007](#) on Wed, 08 Feb 2006 15:43:02 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

problem SUSE 10

i try make work the suse 10  
installation was perfect  
Yast work perfect  
but ntwork card don't show  
and the network harware "  
i try add but don't work  
ifconfig -a  
please help to fix this issue

```
lo      Link encap:Local Loopback
        LOOPBACK MTU:16436 Metric:1
        RX packets:36 errors:0 dropped:0 overruns:0 frame:0
        TX packets:36 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:3584 (3.5 Kb) TX bytes:3584 (3.5 Kb)
```

```
venet0  Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
        BROADCAST POINTOPOINT NOARP MTU:1500 Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

localhost:/ # ifconfig -a

```
lo      Link encap:Local Loopback
        LOOPBACK MTU:16436 Metric:1
        RX packets:36 errors:0 dropped:0 overruns:0 frame:0
        TX packets:36 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:3584 (3.5 Kb) TX bytes:3584 (3.5 Kb)
```

```
venet0  Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
        BROADCAST POINTOPOINT NOARP MTU:1500 Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

---

Subject: Re: Re: openvz and SuSE  
Posted by [dim](#) on Wed, 08 Feb 2006 15:52:29 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Could you provide /etc/sysconfig/network/ifcfg-eth\* content?

On Wednesday 08 February 2006 18:43, jbond007 wrote:  
> problem SUSE 10

```
>
> i try make work the suse 10
> installation was perfect
> Yast work perfect
> but network card don't show
> and the network hardware "
> i try add but don't work
> ifconfig -a
> please help to fix this issue
> lo      Link encap:Local Loopback
>         LOOPBACK MTU:16436 Metric:1
>         RX packets:36 errors:0 dropped:0 overruns:0 frame:0
>         TX packets:36 errors:0 dropped:0 overruns:0 carrier:0
>         collisions:0 txqueuelen:0
>         RX bytes:3584 (3.5 Kb) TX bytes:3584 (3.5 Kb)
>
> venet0  Link encap:UNSPEC HWaddr
> 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00 BROADCAST POINTOPOINT NOARP
> MTU:1500 Metric:1
>         RX packets:0 errors:0 dropped:0 overruns:0 frame:0
>         TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
>         collisions:0 txqueuelen:0
>         RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
>
> localhost:/ # ifconfig -a
> lo      Link encap:Local Loopback
>         LOOPBACK MTU:16436 Metric:1
>         RX packets:36 errors:0 dropped:0 overruns:0 frame:0
>         TX packets:36 errors:0 dropped:0 overruns:0 carrier:0
>         collisions:0 txqueuelen:0
>         RX bytes:3584 (3.5 Kb) TX bytes:3584 (3.5 Kb)
>
> venet0  Link encap:UNSPEC HWaddr
> 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00 BROADCAST POINTOPOINT NOARP
> MTU:1500 Metric:1
>         RX packets:0 errors:0 dropped:0 overruns:0 frame:0
>         TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
>         collisions:0 txqueuelen:0
>         RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
>
--
Thanks,
Dmitry.
```

---

Subject: Re: Re: openvz and SuSE  
Posted by [jbond007](#) on Wed, 08 Feb 2006 17:07:50 GMT

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

---

from my host or my Opensuse Vps ?

---

---

Subject: Re: Re: openvz and SuSE  
Posted by [Daniel Bauer](#) on Wed, 08 Feb 2006 18:14:48 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hi ?

From: "jbond007" <max007@bellsouth.net>  
> problem SUSE 10  
>  
> i try make work the suse 10  
> installation was perfect  
> Yast work perfect  
> but ntwork card don't show  
> and the network harware "  
> i try add but don't work  
> [...]

did this problem appear in the standard installation of SuSE or when the  
OpenVZ Kernel is booted?  
which hardware?

Daniel

---

---

Subject: Re: Re: openvz and SuSE  
Posted by [jbond007](#) on Thu, 09 Feb 2006 01:47:17 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

sorry you not right !  
i have image Opensuse 10 for vps  
my host run centos4.2

sample  
after i crete the vzctl create 151 --ostemplate open-suse-\*\*\*\*  
vps create perfect  
vzctl set 150 --ipadd \*\*\*\*\*  
i try ping  
the ip not respond !  
bad  
ok so i try using vzctl enter 151  
to find the issue  
i use yast

not network device  
the main issue here Opensuse is not ready to support vps network device not show  
like fedora or centos image  
Please Help me ASAP Regards !!!!!!!!!!!

---

Subject: Re: Re: Re: openvz and SuSE  
Posted by [dim](#) on Thu, 09 Feb 2006 09:23:22 GMT  
[View Forum Message](#) <> [Reply to Message](#)

Correct me if I'm wrong.

You have hardware node with working network, but can't establish network for  
your suse VPS?

If so, I need your ifconfig -a, netstat -rn (and probably iptables -L)  
commands output both in hardware node and guest

If you have hardware node with not-working network card, I need previous  
mentioned info ( /etc/sysconfig/network/ifcfg-eth\*) from hardware node.

On Wednesday 08 February 2006 20:07, [jbond007](#) wrote:

> from my host or my Opensuse Vps ?

>

--

Thanks,  
Dmitry.

---

Subject: Re: Re: Re: openvz and SuSE  
Posted by [jbond007](#) on Thu, 16 Feb 2006 03:11:48 GMT  
[View Forum Message](#) <> [Reply to Message](#)

dear please I love your help but please don't respond Bu\*\*\* respond  
have suse 10 image i create new vps but the NIC is not working inside of vps only for suse 10  
how i make work this ?

thank you  
try !

---

Subject: Re: Re: Re: openvz and SuSE  
Posted by [dev](#) on Thu, 16 Feb 2006 06:38:14 GMT  
[View Forum Message](#) <> [Reply to Message](#)



my 2 cents:

As Kir told already openSUSE10 template requires some kernel support. You need to rebuilt your kernel with CONFIG\_VE\_SYSFS=y.

In next kernel updates this option will be made user-configurable and vzctl will turn it on automatically for SUSE10 template.

---

Subject: Re: Re: Re: openvz and SuSE  
Posted by [jbond007](#) on Thu, 16 Feb 2006 14:50:49 GMT  
[View Forum Message](#) <> [Reply to Message](#)

rebuild the kernel form vps or host ?

1 how you buil the kernel please send me step by step

Urgent  
Best regards  
max

---

Subject: Re: Re: Re: Re: openvz and SuSE  
Posted by [dev](#) on Fri, 17 Feb 2006 06:17:49 GMT  
[View Forum Message](#) <> [Reply to Message](#)

Hello,

> rebuild the kernel form vps or host ?  
VPS has no kernel. in host system.

> 1 how you buil the kernel please send me step by step

1. rpm -ihv vzkernel-XXX.src.rpm
2. cd /usr/src/redhat/SOURCES (or whatever directory you have for RPM builds)
3. edit all i686 config files, replace line:  
# CONFIG\_VE\_SYSFS is not set  
with line:  
CONFIG\_VE\_SYSFS=y
4. rpmbuild -ba --target=i686 kernel-ovz.spec

WBR,  
Kirill