

---

Subject: RE: lvm and openvz  
Posted by [Mark Olliver](#) on Tue, 27 Mar 2012 15:02:50 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hi All,

I am setting up a new openvz setup and am looking to have each guest store its data on its own lvm partition. Can anyone give me a clue what options I should say to the vzctl create script?

I would guess I tell it the private is the root of the lvm partition for the guest, then root is in the normal place as that is a runtime mount?

Thanks

Mark

---

---

Subject: Re: lvm and openvz  
Posted by [David Brown](#) on Wed, 28 Mar 2012 07:49:07 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On 27/03/2012 17:02, Mark Olliver wrote:

> Hi All,  
>  
> I am setting up a new openvz setup and am looking to have each guest  
> store its data on its own lvm partition. Can anyone give me a clue what  
> options I should say to the vzctl create script?  
>  
> I would guess I tell it the private is the root of the lvm partition for  
> the guest, then root is in the normal place as that is a runtime mount?  
>  
> Thanks  
>  
> Mark  
>

I have an lvm partition for each of my openvz virtual machines. The way

I organise it is to have a base directory /vz, with subdirectories for each machine. I mount the lvm logical disk for server1 on /vz/server1, and then create the virtual machine with "vzctl create --root /vz/server1/root --private /vz/server1/private" (plus other options, obviously).

If you need multiple lvm partitions in the same virtual machine, I guess you mount them within /vz/server1/private before starting the machine - though I haven't needed to do that.

I /really/ wish the openvz developers would move beyond kernel 2.6.32 - kernel 2.6.33 introduced snapshot merging to LVM which would play wonderfully with this setup.

mvh.,

David

---

Subject: Re: Re: lvm and openvz  
Posted by [kir](#) on Thu, 29 Mar 2012 09:04:10 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On 03/28/2012 11:49 AM, David Brown wrote:

>

> I /really/ wish the openvz developers would move beyond kernel 2.6.32 -  
> kernel 2.6.33 introduced snapshot merging to LVM which would play  
> wonderfully with this setup.

I'm not sure why people think that RHEL6 kernel is pure 2.6.32. It is definitely not!

For snapshot merging, I am not an expert here but googling for 'rhel6 lvm snapshot merging' gave me this:

[http://www.linuxtopia.org/online\\_books/rhel6/rhel\\_6\\_lvm\\_admin/rhel\\_6\\_lvm\\_snapshot\\_merge.html](http://www.linuxtopia.org/online_books/rhel6/rhel_6_lvm_admin/rhel_6_lvm_snapshot_merge.html)

and this:

<https://access.redhat.com/knowledge/solutions/58510>

Both articles suggest RHEL6 kernel supports LVM snapshot merging, and so should OpenVZ RHEL6-based kernel.

PS If you are using non-rhel6 openvz kernel, it's definitely time to

switch, and lots of reasons

to do that besides LVM snapshot merging. Notable things are vswap, ploop, stability...

Kir.

---

Subject: Re: Re: lvm and openvz

Posted by [David Brown](#) on Thu, 29 Mar 2012 09:09:59 GMT

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

---

On 29/03/2012 11:04, Kir Kolyshkin wrote:

> On 03/28/2012 11:49 AM, David Brown wrote:

>>

>> I /really/ wish the openvz developers would move beyond kernel

>> 2.6.32 - kernel 2.6.33 introduced snapshot merging to LVM which

>> would play wonderfully with this setup.

>

> I'm not sure why people think that RHEL6 kernel is pure 2.6.32. It is

> definitely not!

>

> For snapshot merging, I am not an expert here but googling for 'rhel6

> lvm snapshot merging' gave me this:

>

> [http://www.linuxtopia.org/online\\_books/rhel6/rhel\\_6\\_lvm\\_admin/rhel\\_6\\_lvm\\_snapshot\\_merge.html](http://www.linuxtopia.org/online_books/rhel6/rhel_6_lvm_admin/rhel_6_lvm_snapshot_merge.html)

>

>

>

> and this:

>

> <https://access.redhat.com/knowledge/solutions/58510>

>

> Both articles suggest RHEL6 kernel supports LVM snapshot merging, and

> so should OpenVZ RHEL6-based kernel.

>

> PS If you are using non-rhel6 openvz kernel, it's definitely time to

> switch, and lots of reasons to do that besides LVM snapshot merging.

> Notable things are vswap, ploop, stability...

>

> Kir.

I am using a non-rhel6 openvz kernel because I don't use RHEL - I use Debian on my servers. Are you suggesting that I should specifically use the RHEL6 openvz kernel even though I use Debian? That's something I haven't thought of trying, but if it is the recommendation of the OpenVZ developers, then I will give it a shot.

More generally, I would hope that one day OpenVZ will change over to following the current kernel (or perhaps the current long-term support kernels) - there has been a lot of development since 2.6.32, not all of which gets backported by Red Hat. I'd expect that a lot of OpenVZ code can be merged with or replaced by the container support in later kernels. I also think that if OpenVZ doesn't catch up, then people will migrate to other solutions such as Linux VServer or LXC (I know I considered it for the last server I configured).

Of course, I fully appreciate that something like that takes a lot of effort, and that means time, money, people to do the work, testing, etc., etc. But one can still hope!

David

---

---

Subject: Re: Re: lvm and openvz  
Posted by [John Knight](#) on Thu, 29 Mar 2012 09:32:34 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

> I'm not sure why people think that RHEL6 kernel is pure 2.6.32. It is  
> definitely not!

I completely agree, Kir. RHEL 6's 2.6.32 branch is also an attractive target because of the security and feature support coming from Red Hat for a long period of time.

I look back two years ago and there was a plethora of kernels being actively developed (at one point it was 2.6.26, 2.6.27, 2.6.32 vanilla, 2.6.32-el6 testing, 2.6.18-el5, 2.6.18-el5 testing, etc) and look at what's happening now and it's night and day.

Namely I can see the quality of work done by the OpenVZ dev team really shining when they've stopped working on so many branches. I think that was really hurting the project. Now we have awesome stuff like Ploop and vSwap coming out too. It's an exciting time for OpenVZ.

With regard to David Brown, I haven't been a Debian user for nearly 5 years but I do recall other people compiling the openvz el6 branch in Debian without much trouble. It becomes harder to support I'd imagine as you would have to compile each release by hand unless you made your own deb package but I've certainly heard of it being possible. I just checked the OpenVZ wiki and came across this as well:  
[http://wiki.openvz.org/Install\\_kernel\\_from\\_RPM\\_on\\_Debian\\_6.0](http://wiki.openvz.org/Install_kernel_from_RPM_on_Debian_6.0)

\*John Knight\*  
Classic City Telco LLC  
\*Email:\* [john@classiccitytelco.com](mailto:john@classiccitytelco.com) \*|\* \*Main:\* (706) 995-0200

\*Direct:\* (706) 995-0201

On 3/29/2012 5:04 AM, Kir Kolyshkin wrote:

> On 03/28/2012 11:49 AM, David Brown wrote:

>>

>> I /really/ wish the openvz developers would move beyond kernel 2.6.32 -

>> kernal 2.6.33 introduced snapshot merging to LVM which would play

>> wonderfully with this setup.

>

> I'm not sure why people think that RHEL6 kernel is pure 2.6.32. It is

> definitely not!

>

> For snapshot merging, I am not an expert here but googling for 'rhel6

> lvm snapshot merging'

> gave me this:

>

> [http://www.linuxtopia.org/online\\_books/rhel6/rhel\\_6\\_lvm\\_admin/rhel\\_6\\_lvm\\_snapshot\\_merge.html](http://www.linuxtopia.org/online_books/rhel6/rhel_6_lvm_admin/rhel_6_lvm_snapshot_merge.html)

>

>

> and this:

>

> <https://access.redhat.com/knowledge/solutions/58510>

>

> Both articles suggest RHEL6 kernel supports LVM snapshot merging, and

> so should

> OpenVZ RHEL6-based kernel.

>

> PS If you are using non-rhel6 openvz kernel, it's definitely time to

> switch, and lots of reasons

> to do that besides LVM snapshot merging. Notable things are vswap,

> ploop, stability...

>

> Kir.

---

Subject: Re: Re: lvm and openvz

Posted by [kir](#) on Thu, 29 Mar 2012 12:16:47 GMT

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

---

On 03/29/2012 01:09 PM, David Brown wrote:

> On 29/03/2012 11:04, Kir Kolyshkin wrote:

>> On 03/28/2012 11:49 AM, David Brown wrote:

>>> I /really/ wish the openvz developers would move beyond kernel

>>> 2.6.32 - kernal 2.6.33 introduced snapshot merging to LVM which

>>> would play wonderfully with this setup.

>> I'm not sure why people think that RHEL6 kernel is pure 2.6.32. It is

>> definitely not!  
>>  
>> For snapshot merging, I am not an expert here but googling for 'rhel6  
>> lvm snapshot merging' gave me this:  
>>  
>> [http://www.linuxtopia.org/online\\_books/rhel6/rhel\\_6\\_lvm\\_admin/rhel\\_6\\_lvm\\_snapshot\\_merge.html](http://www.linuxtopia.org/online_books/rhel6/rhel_6_lvm_admin/rhel_6_lvm_snapshot_merge.html)  
>>  
>>  
>>  
>> and this:  
>>  
>> <https://access.redhat.com/knowledge/solutions/58510>  
>>  
>> Both articles suggest RHEL6 kernel supports LVM snapshot merging, and  
>> so should OpenVZ RHEL6-based kernel.  
>>  
>> PS If you are using non-rhel6 openvz kernel, it's definitely time to  
>> switch, and lots of reasons to do that besides LVM snapshot merging.  
>> Notable things are vswap, ploop, stability...  
>>  
>> Kir.  
> I am using a non-rhel6 openvz kernel because I don't use RHEL - I use  
> Debian on my servers. Are you suggesting that I should specifically use  
> the RHEL6 openvz kernel even though I use Debian? That's something I  
> haven't thought of trying, but if it is the recommendation of the OpenVZ  
> developers, then I will give it a shot.

Yes please. This is the official recommendation, and many people already do this.

We choose RHEL6 kernel as a base not because we are Red Hat fans or we plan to support RHEL distribution only. The reason is RHEL6 is truly a good, well-maintained, stable kernel with lots of Red Hat developer and QA resources invested into it. That is why we use it as a base, but that doesn't mean we only have RHEL in mind.

Now to the practical point: we have modified post-install scripts in kernel rpm to be compatible with Debian as well, so all you need to do is to convert kernel rpm to deb using alien. Some info is provided at [http://wiki.openvz.org/Install\\_kernel\\_from\\_RPM\\_on\\_Debian\\_6.0](http://wiki.openvz.org/Install_kernel_from_RPM_on_Debian_6.0)

> More generally, I would hope that one day OpenVZ will change over to  
> following the current kernel (or perhaps the current long-term support

> kernels) - there has been a lot of development since 2.6.32, not all of  
> which gets backported by Red Hat.

Same as we did before - we do have such plans, although not immediate.

> I'd expect that a lot of OpenVZ code  
> can be merged with or replaced by the container support in later  
> kernels.

That is right, we are actively pushing our stuff upstream, and then we are rebasing our code to what is available in upstream, gradually reducing the size of openvz patchset.

For example, if you will take a look at linux kernel git repo, you will see more than about 150 patches from OpenVZ developers were merged this year. Actually, here's a command to do that:

```
$ git log -E --author='@parallels.com|@openvz.org' --since=2012-01-01 |  
grep -c ^commit
```

Most of the recent patches are CRUI and NFS virtualization.

> I also think that if OpenVZ doesn't catch up, then people will  
> migrate to other solutions such as Linux VServer or LXC (I know I  
> considered it for the last server I configured).

Linux-VServer is totally obsolete from my POV, because they are not interested in pushing the stuff upstream. Of course they benefit from the code that is appearing in mainline (and since a good proportion of that code comes from OpenVZ developers, it is true to say that Linux-VServer benefits from OpenVZ).

As for the LXC, please do not forget that LXC is not something that is opposed to OpenVZ, but rather something that is complementary. I mean, having said that the good proportion of containers code in mainline comes from OpenVZ, it might be true to say that we are probably the biggest contributor to the LXC (kernel code).

> Of course, I fully appreciate that something like that takes a lot of  
> effort, and that means time, money, people to do the work, testing,  
> etc., etc. But one can still hope!

Our current approach is to use RHEL kernels as a base, and push as much stuff to upstream as we can. So far it's working.

---

Subject: Re: Re: lvm and openvz

On 29/03/2012 14:16, Kir Kolyshkin wrote:

> On 03/29/2012 01:09 PM, David Brown wrote:

>> On 29/03/2012 11:04, Kir Kolyshkin wrote:

>>> On 03/28/2012 11:49 AM, David Brown wrote:

>>>> I /really/ wish the openvz developers would move beyond kernel

>>>> 2.6.32 - kernel 2.6.33 introduced snapshot merging to LVM which

>>>> would play wonderfully with this setup.

>>> I'm not sure why people think that RHEL6 kernel is pure 2.6.32. It is

>>> definitely not!

>>>

>>> For snapshot merging, I am not an expert here but googling for 'rhel6

>>> lvm snapshot merging' gave me this:

>>>

>>> [http://www.linuxtopia.org/online\\_books/rhel6/rhel\\_6\\_lvm\\_admin/rhel\\_6\\_lvm\\_snapshot\\_merge.html](http://www.linuxtopia.org/online_books/rhel6/rhel_6_lvm_admin/rhel_6_lvm_snapshot_merge.html)

>>>

>>>

>>>

>>>

>>> and this:

>>>

>>> <https://access.redhat.com/knowledge/solutions/58510>

>>>

>>> Both articles suggest RHEL6 kernel supports LVM snapshot merging, and

>>> so should OpenVZ RHEL6-based kernel.

>>>

>>> PS If you are using non-rhel6 openvz kernel, it's definitely time to

>>> switch, and lots of reasons to do that besides LVM snapshot merging.

>>> Notable things are vswap, ploop, stability...

>>>

>>> Kir.

>> I am using a non-rhel6 openvz kernel because I don't use RHEL - I use

>> Debian on my servers. Are you suggesting that I should specifically use

>> the RHEL6 openvz kernel even though I use Debian? That's something I

>> haven't thought of trying, but if it is the recommendation of the OpenVZ

>> developers, then I will give it a shot.

>

> Yes please. This is the official recommendation, and many people already  
> do this.

>

> We choose RHEL6 kernel as a base not because we are Red Hat fans or we  
> plan to support

> RHEL distribution only. The reason is RHEL6 is truly a good,

> well-maintained, stable kernel with

> lots of Red Hat developer and QA resources invested into it. That is why

> we use it as a base,



> but that doesn't mean we only have RHEL in mind.  
>  
> Now to the practical point: we have modified post-install scripts in  
> kernel rpm to be compatible  
> with Debian as well, so all you need to do is to convert kernel rpm to  
> deb using alien. Some  
> info is provided at  
> [http://wiki.openvz.org/Install\\_kernel\\_from\\_RPM\\_on\\_Debian\\_6.0](http://wiki.openvz.org/Install_kernel_from_RPM_on_Debian_6.0)  
>

Thank you for that pointer. I will read through the information, and try it out. I don't know when I will get the chance on my main servers - playing with the kernel on the host machine means taking all the guest systems off-line for a little while - but I'll find a spare system somewhere to test it.

>> More generally, I would hope that one day OpenVZ will change over to  
>> following the current kernel (or perhaps the current long-term support  
>> kernels) - there has been a lot of development since 2.6.32, not all of  
>> which gets backported by Red Hat.  
>  
> Same as we did before - we do have such plans, although not immediate.  
>  
>> I'd expect that a lot of OpenVZ code  
>> can be merged with or replaced by the container support in later  
>> kernels.  
>  
> That is right, we are actively pushing our stuff upstream, and then we  
> are rebasing our code to what is available in upstream, gradually  
> reducing the size of openvz patchset.  
>  
> For example, if you will take a look at linux kernel git repo, you will  
> see more than about 150 patches from OpenVZ developers were merged this  
> year. Actually, here's a command to do that:  
>  
> \$ git log -E --author='@parallels.com|@openvz.org' --since=2012-01-01 |  
> grep -c ^commit  
>  
> Most of the recent patches are CRUI and NFS virtualization.  
>  
>> I also think that if OpenVZ doesn't catch up, then people will  
>> migrate to other solutions such as Linux VServer or LXC (I know I  
>> considered it for the last server I configured).  
>  
> Linux-VServer is totally obsolete from my POV, because they are not  
> interested in pushing the stuff upstream. Of course they benefit from  
> the code that is appearing in mainline (and since a good proportion of  
> that code comes from OpenVZ developers, it is true to say that

> Linux-VServer benefits from OpenVZ).

>

OK. Linux VServer struck me as being more limited, and less active than OpenVZ when I first looked at virtualisation solutions several years ago - from your comments, it looks like I made the right decision.

> As for the LXC, please do not forget that LXC is not something that is  
> opposed to OpenVZ, but rather something that is complementary. I mean,  
> having said that the good proportion of containers code in mainline  
> comes from OpenVZ, it might be true to say that we are probably the  
> biggest contributor to the LXC (kernel code).

>

Yes, I understand that a lot of LXC is based on OpenVZ ideas and code moved into the mainline. To my uninformed mind, it looks like LXC has many of the basic features of OpenVZ, while OpenVZ provides more detailed control of the virtual machines and more useful tools and utilities for creating and controlling the guests.

The ideal situation from my viewpoint would be to continue the upstream pushes until all the kernel code for OpenVZ is in the mainline. That would be the most flexible for users, giving them OpenVZ with whatever kernel they wanted. But I guess it's a two-edged sword for the openvz developers - it would mean less effort supporting new kernels, but maybe more work since the kernel changes all the time, and more work for support and testing. I don't know if such a merge would be possible, practical, or desirable (from your viewpoint or from the mainline viewpoint).

>> Of course, I fully appreciate that something like that takes a lot of  
>> effort, and that means time, money, people to do the work, testing,  
>> etc., etc. But one can still hope!

>

> Our current approach is to use RHEL kernels as a base, and push as much  
> stuff to upstream as we can. So far it's working.

I'll give the RHEL kernel a try as soon as I get the chance.

Thank you very much for your explanations and advice. (And of course, thank you for your work on openvz - it is a fantastic system and has made my job much easier.)

Best regards,

David

---

---

Subject: Re: Re: lvm and openvz

Posted by [jjs - mainphrame](#) on Thu, 29 Mar 2012 16:51:46 GMT

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

---

I'm running the rhel-based kernel on my debian 6 openvz server with good results. Better performance and more features.

<https://plus.google.com/114658067490332530482/posts/dbHMM22nRv1>

Joe

On Thu, Mar 29, 2012 at 2:09 AM, David Brown <david@westcontrol.com> wrote:

>  
>> Kir.  
>>  
>  
> I am using a non-rhel6 openvz kernel because I don't use RHEL - I use  
> Debian on my servers. Are you suggesting that I should specifically use  
> the RHEL6 openvz kernel even though I use Debian? That's something I  
> haven't thought of trying, but if it is the recommendation of the OpenVZ  
> developers, then I will give it a shot.  
>  
> More generally, I would hope that one day OpenVZ will change over to  
> following the current kernel (or perhaps the current long-term support  
> kernels) - there has been a lot of development since 2.6.32, not all of  
> which gets backported by Red Hat. I'd expect that a lot of OpenVZ code can  
> be merged with or replaced by the container support in later kernels. I  
> also think that if OpenVZ doesn't catch up, then people will migrate to  
> other solutions such as Linux VServer or LXC (I know I considered it for  
> the last server I configured).  
>  
> Of course, I fully appreciate that something like that takes a lot of  
> effort, and that means time, money, people to do the work, testing, etc.,  
> etc. But one can still hope!  
>  
> David  
>  
> \_\_\_\_\_ \*\* \_\_\_\_\_  
> Users mailing list  
> Users@openvz.org  
> [https://openvz.org/mailman/\\*\\*listinfo/users](https://openvz.org/mailman/**listinfo/users)<<https://openvz.org/mailman/listinfo/users>>  
>

---

---

Subject: Re: Re: lvm and openvz

Posted by [John Knight](#) on Thu, 29 Mar 2012 21:44:24 GMT

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

---

That makes perfect sense. Running debian containers on the el6 kernel on a centos node before has shown me that there are no fundamental issues with the debian toolchain or command line binaries running on the el6-kernel-compiled-on-el6, but it makes sense that there would be a lot more possibilities of bugs compiling it on debian directly.

\*John Knight\*

Classic City Telco LLC

\*Email:\* john@classiccitytelco.com \*|\* \*Main:\* (706) 995-0200

\*Direct:\* (706) 995-0201 \*|\* \*Mobile:\* (706) 255-9203

CCT Enterprise Linux 6 is released! Click here to learn more.

<[http://www.classiccitytelco.com/?page\\_id=488](http://www.classiccitytelco.com/?page_id=488)>

On 3/29/2012 5:36 PM, Kir Kolyshkin wrote:

> On 03/29/2012 01:32 PM, John Knight wrote:

>> With regard to David Brown, I haven't been a Debian user for nearly 5  
>> years but I do recall other people compiling the openvz el6 branch in  
>> Debian without much trouble.

> While theoretically it's probably the best and most natural way,  
> practically I do not recommend doing it (at least unless you will also  
> use toolchain/gcc from RHEL6). Kernel is big and complex, as well as  
> gcc, and the improper combination of two could lead to bad results.  
> This is not paranoia, I have seen a number of times when kernel code  
> was miscompiled because of older/newer gcc version used, with a weird  
> runtime effects.

>

> This is why we recommend taking binary rpm and using alien on it.  
> While not a good thing from purist point of view\*, practical result is  
> the very same (ie bit by bit) kernel and modules, tried and trusted,  
> tested and working.

>

> \* being a purist, I do not like it. But I also know that this way it  
> works, and compiler and toolchain indeed make a difference.

---