
Subject: Debian: recommended kernel

Posted by [Roman Haefeli](#) on Wed, 02 May 2012 16:31:22 GMT

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

Hi all

We're running OpenVZ on Debian Squeeze with the kernel shipped by Debian.

Several sources recommend to use RHEL 6 stable kernel. Is it recommended to use it also on Debian stable? If so, how should it be installed? The wiki has links to rpm files only, it seems.

The reason I ask is that we're considering switching kernel, since the Debian OpenVZ kernel seems to have issues. Specifically, we suffer from the same problem as described here:

<http://lists.debian.org/debian-kernel/2011/12/msg00689.html>

Cheers

Roman

Subject: Re: Debian: recommended kernel

Posted by [Timh B](#) on Wed, 02 May 2012 17:39:09 GMT

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

This was linked earlier this week;

<https://github.com/CoolCold/tools/blob/master/opensvz/kernel/ create-ovz-kernel-for-debian.sh>

Might be useful for you if you wish to get a debianized opensvz-kernel.

If you really want to run debian I would suggest looking into proxmox instead since they build the stable rpm-kernels for debian.

Good luck!

//T

On Wed, May 2, 2012 18:31, Roman Haefeli wrote:

> Hi all

>

> We're running OpenVZ on Debian Squeeze with the kernel shipped by

> Debian.

>

> Several sources recommend to use RHEL 6 stable kernel. Is it recommended

> to use it also on Debian stable? If so, how should it be installed? The

> wiki has links to rpm files only, it seems.

>

> The reason I ask is that we're considering switching kernel, since the
> Debian OpenVZ kernel seems to have issues. Specifically, we suffer from
> the same problem as described here:
> <http://lists.debian.org/debian-kernel/2011/12/msg00689.html>
>
> Cheers
> Roman
>
>
-- Timh

Subject: Re: Debian: recommended kernel
Posted by [coolcold](#) on Wed, 02 May 2012 19:18:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, May 2, 2012 at 9:39 PM, Timh B <timh@shiwebs.net> wrote:
> This was linked earlier this week;
> <https://github.com/CoolCold/tools/blob/master/openvz/kernel/create-ovz-kernel-for-debian.sh>
>
> Might be useful for you if you wish to get a debianized openvz-kernel.
>
> If you really want to run debian I would suggest looking into proxmox
> instead since they build the stable rpm-kernels for debian.
May be good, choice, yes. Not tried myself, but reading their maillist.

>
> Good luck!
>
> //T
>
> On Wed, May 2, 2012 18:31, Roman Haefeli wrote:
>> Hi all
>>
>> We're running OpenVZ on Debian Squeeze with the kernel shipped by
>> Debian.
>>
>> Several sources recommend to use RHEL 6 stable kernel. Is it recommended
>> to use it also on Debian stable? If so, how should it be installed? The
>> wiki has links to rpm files only, it seems.
>>
>> The reason I ask is that we're considering switching kernel, since the
>> Debian OpenVZ kernel seems to have issues. Specifically, we suffer from
>> the same problem as described here:
>> <http://lists.debian.org/debian-kernel/2011/12/msg00689.html>
I was forced to switch into RHEL based kernels instead of Debian ones
because of another bug -
http://bugzilla.openvz.org/show_bug.cgi?id=1747 - I had connection

resets while on tcp sockets.

>>
>> Cheers
>> Roman
>>
>>
> -- Timh
>
>
--

Best regards,
[COOLCOLD-RIPN]

Subject: Re: Debian: recommended kernel
Posted by [jjs - mainphrame](#) on Wed, 02 May 2012 19:19:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

I second the proxmox idea - it is an excellent platform, I recently downloaded it to take a look and was quite impressed. It's a more comprehensive solution than openvz+ovz-web-panel. I'm sold on it and will deploy it in future. It does require x86_64 hardware though.

Joe

On Wed, May 2, 2012 at 10:39 AM, Timh B <timh@shiwebs.net> wrote:

> This was linked earlier this week;
>
> <https://github.com/CoolCold/tools/blob/master/openvz/kernel/create-ovz-kernel-for-debian.sh>
>
> Might be useful for you if you wish to get a debianized openvz-kernel.
>
> If you really want to run debian I would suggest looking into proxmox
> instead since they build the stable rpm-kernels for debian.
>
> Good luck!
>
> //T
>
> On Wed, May 2, 2012 18:31, Roman Haefeli wrote:
> > Hi all
> >
> > We're running OpenVZ on Debian Squeeze with the kernel shipped by
> > Debian.
> >
> > Several sources recommend to use RHEL 6 stable kernel. Is it recommended

> > to use it also on Debian stable? If so, how should it be installed? The
> > wiki has links to rpm files only, it seems.
> >
> > The reason I ask is that we're considering switching kernel, since the
> > Debian OpenVZ kernel seems to have issues. Specifically, we suffer from
> > the same problem as described here:
> > <http://lists.debian.org/debian-kernel/2011/12/msg00689.html>
> >
> > Cheers
> > Roman
> >
> >
> -- Timh
>
>

Subject: Re: Debian: recommended kernel
Posted by [kir](#) on Thu, 03 May 2012 05:14:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 05/02/2012 09:39 PM, Timh B wrote:

> This was linked earlier this week;
> <https://github.com/CoolCold/tools/blob/master/openvz/kernel/create-ovz-kernel-for-debian.sh>
>
> Might be useful for you if you wish to get a debianized openvz-kernel.

Yet another solution is to use alien:

http://wiki.openvz.org/Install_kernel_from_RPM_on_Debian_6.0

The bad thing about it is it's not a "native" way. The good thing is you have the same bit-by-bit binary kernel which we test a lot.

>
> If you really want to run debian I would suggest looking into proxmox
> instead since they build the stable rpm-kernels for debian.
>
> Good luck!
>
> //T
>
> On Wed, May 2, 2012 18:31, Roman Haefeli wrote:
>> Hi all
>>
>> We're running OpenVZ on Debian Squeeze with the kernel shipped by
>> Debian.
>>
>> Several sources recommend to use RHEL 6 stable kernel. Is it recommended

>> to use it also on Debian stable? If so, how should it be installed? The
>> wiki has links to rpm files only, it seems.
>>
>> The reason I ask is that we're considering switching kernel, since the
>> Debian OpenVZ kernel seems to have issues. Specifically, we suffer from
>> the same problem as described here:
>> <http://lists.debian.org/debian-kernel/2011/12/msg00689.html>
>>
>> Cheers
>> Roman
>>

Subject: Re: Debian: recommended kernel
Posted by [coolcold](#) on Thu, 03 May 2012 06:32:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, May 3, 2012 at 9:14 AM, Kir Kolyshkin <kir@openvz.org> wrote:
> On 05/02/2012 09:39 PM, Timh B wrote:
>>
>> This was linked earlier this week;
>>
>> <https://github.com/CoolCold/tools/blob/master/openvz/kernel/create-ovz-kernel-for-debian.sh>
>>
>> Might be useful for you if you wish to get a debianized openvz-kernel.
>
>
> Yet another solution is to use alien:
> http://wiki.openvz.org/Install_kernel_from_RPM_on_Debian_6.0
Well, alien works for binary part of kernel, but when I've tried to
install kernel headers that way, i had files conflicts, that's why
that script above was born.

>
> The bad thing about it is it's not a "native" way. The good thing is you
> have the same bit-by-bit binary kernel which we test a lot.
>
>
>>
>> If you really want to run debian I would suggest looking into proxmox
>> instead since they build the stable rpm-kernels for debian.
>>
>> Good luck!
>>
>> //T
>>
>> On Wed, May 2, 2012 18:31, Roman Haefeli wrote:
>>>

>>> Hi all
>>>
>>> We're running OpenVZ on Debian Squeeze with the kernel shipped by
>>> Debian.
>>>
>>> Several sources recommend to use RHEL 6 stable kernel. Is it recommended
>>> to use it also on Debian stable? If so, how should it be installed? The
>>> wiki has links to rpm files only, it seems.
>>>
>>> The reason I ask is that we're considering switching kernel, since the
>>> Debian OpenVZ kernel seems to have issues. Specifically, we suffer from
>>> the same problem as described here:
>>> <http://lists.debian.org/debian-kernel/2011/12/msg00689.html>
>>>
>>> Cheers
>>> Roman
>>>
>
--
Best regards,
[COOLCOLD-RIPN]

Subject: Re: Debian: recommended kernel
Posted by [Roman Haefeli](#) on Fri, 04 May 2012 09:58:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks for all the responses!

On Thu, 2012-05-03 at 09:14 +0400, Kir Kolyshkin wrote:

> On 05/02/2012 09:39 PM, Timh B wrote:
> > This was linked earlier this week;
> > <https://github.com/CoolCold/tools/blob/master/openvz/kernel/create-ovz-kernel-for-debian.sh>
> >
> > Might be useful for you if you wish to get a debianized openvz-kernel.
>
> Yet another solution is to use alien:
> http://wiki.openvz.org/Install_kernel_from_RPM_on_Debian_6.0
>
> The bad thing about it is it's not a "native" way. The good thing is you
> have the same bit-by-bit binary kernel which we test a lot.

In the meanwhile I have tried two of the suggested solutions:

- * RHEL6-kernel (debianized with alien)
- * proxmox-kernel from their repository (which for the 2.6.32 version says it is based on RHEL6)

(I haven't looked yet into building my own with CoolCold's script)

I didn't have troubles to install either of those. However, both expose the exact same problem. The network bridges vzbr0 and vzbr1 do not work correctly and I need those for the veth devices of the containers. Networking without those bridges seem to work fine, but when the bridges are active, TCP connections fail to work. UDP seems to be not affected, both inbound and outbound connections work. However, with TCP I can only establish outbound connections, but as soon as the connection is established, nothing is received on the other end anymore. Nor can I send from outside to the machine with vzbr0.

When I switch back to Debian's OpenVZ kernel, the problems disappear, but come back, when I boot into the RHEL6 or proxmox kernel. I wonder if really the kernel is the problem or rather some network/bridge configuration. I assume others do run bridges successfully with the aforementioned kernels also run CTs with veths successfully. Is that correct?

Roman

This is my network configuration on said machine:

```
eth2 eth3          physical NICs
|  |
+---+---+
|
bond0
|
+-----+-----+
|         |         |
vlan20  vlan30  vlan31  VLANs
|         |
vzbr0  vzbr1          bridges
```

DMZ local(NFS)

```
# The loopback network interface
auto lo
iface lo inet loopback
```

```
# Fasst die beiden physischen 10Gb-Ports zu
# einem Interface zusammen
auto bond0
```

```
iface bond0 inet manual
    bond-mode 4
    bond-miimon 100
    bond_xmit_hash_policy layer2+3
    bond_lacp_rate slow
    slaves eth2 eth3
```

```
# VLAN für DMZ (über bond0)
auto vlan20
iface vlan20 inet manual
    vlan_raw_device bond0
    address 0.0.0.0
```

```
# VLAN für NFS Netz (über bond0)
auto vlan30
iface vlan30 inet manual
    vlan_raw_device bond0
    address 0.0.0.0
```

```
# VLAN für iSCSI Netz (über bond0)
auto vlan31
iface vlan31 inet static
    vlan_raw_device bond0
    address 10.10.11.92
    network 10.10.11.0
    netmask 255.255.255.0
```

```
# Bridge für VZ (DMZ)
auto vzbr0
iface vzbr0 inet static
    bridge_ports vlan20
    bridge_maxwait 2
    address 195.176.254.92
    network 195.176.254.0
    netmask 255.255.255.128
    broadcast 195.176.254.127
    gateway 195.176.254.1
```

```
# Bridge für VZ (NFS Netz)
auto vzbr1
iface vzbr1 inet static
    bridge_ports vlan30
    bridge_maxwait 2
    address 10.10.10.92
    network 10.10.10.0
    netmask 255.255.255.0
```

Subject: Re: Debian: recommended kernel
Posted by [Roman Haefeli](#) on Wed, 09 May 2012 07:33:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 2012-05-04 at 11:58 +0200, Roman Haefeli wrote:

> Thanks for all the responses!

>

> On Thu, 2012-05-03 at 09:14 +0400, Kir Kolyshev wrote:

> > On 05/02/2012 09:39 PM, Tim B wrote:

> > > This was linked earlier this week;

> > > <https://github.com/CoolCold/tools/blob/master/openvz/kernel/create-ovz-kernel-for-debian.sh>

> > >

> > > Might be useful for you if you wish to get a debianized openvz-kernel.

> >

> > Yet another solution is to use alien:

> > http://wiki.openvz.org/Install_kernel_from_RPM_on_Debian_6.0

> >

> > The bad thing about it is it's not a "native" way. The good thing is you

> > have the same bit-by-bit binary kernel which we test a lot.

>

> In the meanwhile I have tried two of the suggested solutions:

>

> * RHEL6-kernel (debianized with alien)

> * proxmox-kernel from their repository (which for the 2.6.32 version

> says it is based on RHEL6)

>

> (I haven't looked yet into building my own with CoolCold's script)

>

> I didn't have troubles to install either of those. However, both expose

> the exact same problem. The network bridges vzbr0 and vzbr1 do not work
> correctly and I need those for the veth devices of the containers.

> Networking without those bridges seem to work fine, but when the bridges

> are active, TCP connections fail to work. UDP seems to be not affected,

> both inbound and outbound connections work. However, with TCP I can only

> establish outbound connections, but as soon as the connection is

> established, nothing is received on the other end anymore. Nor can I

> send from outside to the machine with vzbr0.

Our problem looks very similar to the one described here:

https://bugzilla.redhat.com/show_bug.cgi?id=623199

Could it be that current stable OpenVZ suffers from the same issue? Or
asked differently: Is the current stable OpenVZ kernel based on RHEL 6.0
which does not contain the fixes as described in:

<http://rhn.redhat.com/errata/RHSA-2011-0542.html>

Roman

Subject: Re: Debian: recommended kernel
Posted by [Roman Haefeli](#) on Wed, 09 May 2012 07:49:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2012-05-02 at 18:31 +0200, Roman Haefeli wrote:

> Hi all
>
> We're running OpenVZ on Debian Squeeze with the kernel shipped by
> Debian.
>
> Several sources recommend to use RHEL 6 stable kernel. Is it recommended
> to use it also on Debian stable? If so, how should it be installed? The
> wiki has links to rpm files only, it seems.
>
> The reason I ask is that we're considering switching kernel, since the
> Debian OpenVZ kernel seems to have issues. Specifically, we suffer from
> the same problem as described here:
> <http://lists.debian.org/debian-kernel/2011/12/msg00689.html>

It turns out that RHEL 6 based OpenVZ kernel causes more severe issues than the one from Debian (see my last mail in this thread).

Actually, we found a simple work-around for our initial problem (see above link): We turn off locking with the mount.nfs option 'nolock'.

We have a HA cluster with corosync/pacemaker running, which manages the CTs which are running on an NFS export shared across all nodes. As the HA layer makes sure to not run duplicate instances of CTs, there seems to be no reason to have NFS locking enabled.

Roman

Subject: Re: Debian: recommended kernel
Posted by [Aleksandar Ivanisevic](#) on Wed, 09 May 2012 10:54:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

Roman Haefeli <reduzent@gmail.com>
writes:

[...]

> We have a HA cluster with corosync/pacemaker running, which manages the
> CTs which are running on an NFS export shared across all nodes. As the
> HA layer makes sure to not run duplicate instances of CTs, there seems
> to be no reason to have NFS locking enabled.

How do you do STONITH in openvz? I hope you are doing it, because thats the only thing that prevents corruption in the split brain situations, now that you've turned off locking (i.e. fencing).

That being said, I'm still struggling to figure out how to do proper clusters in the virtualized environment. If someone has any pointers/tutorials/etc I would appreciate it. I'm specifically interested in situations where you have more than one virtualized clusters running on the same group of physical nodes.

Subject: Re: Re: Debian: recommended kernel
Posted by [Roman Haefeli](#) on Wed, 09 May 2012 12:43:51 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2012-05-09 at 12:54 +0200, Aleksandar Ivanisevic wrote:

> Roman Haefeli <reduzent@gmail.com>

> writes:

>

>

> [...]

>

> > We have a HA cluster with corosync/pacemaker running, which manages the
> > CTs which are running on an NFS export shared across all nodes. As the
> > HA layer makes sure to not run duplicate instances of CTs, there seems
> > to be no reason to have NFS locking enabled.

>

> How do you do STONITH in openvz? I hope you are doing it, because
> thats the only thing that prevents corruption in the split brain
> situations,

We use SBD (Storage Based Death)

> now that you've turned off locking (i.e. fencing).

Turning on locking wouldn't prevent data corruption in case of a split brain situation as for most files no lock is acquired.

Roman

Subject: Re: Debian: recommended kernel
Posted by [Aleksandar Ivanisevic](#) on Thu, 10 May 2012 09:02:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

Roman Haefeli <reduzent@gmail.com>

writes:

[...]

>> now that you've turned off locking (i.e. fencing).
>
> Turning on locking wouldn't prevent data corruption in case of a split
> brain situation as for most files no lock is acquired.

Heh, I always thought openvz locks a file in \$LOCKDIR to signal that the container is running, but now that I look in the straces it looks like it uses a custom ioctl, which is local kernel scope only.

Would probably be a nice feature request to use a flock or at least leave a lockfile somewhere that can be mounted as a poor man's fencing mechanism.

Subject: Re: Debian: recommended kernel
Posted by [kir](#) on Thu, 10 May 2012 16:06:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 05/09/2012 11:33 AM, Roman Haefeli wrote:
> On Fri, 2012-05-04 at 11:58 +0200, Roman Haefeli wrote:
>> Thanks for all the responses!
>>
>> On Thu, 2012-05-03 at 09:14 +0400, Kir Kolyshkin wrote:
>>> On 05/02/2012 09:39 PM, Timh B wrote:
>>>> This was linked earlier this week;
>>>> <https://github.com/CoolCold/tools/blob/master/openvz/kernel/create-ovz-kernel-for-debian.sh>
>>>>
>>>> Might be useful for you if you wish to get a debianized openvz-kernel.
>>> Yet another solution is to use alien:
>>> http://wiki.openvz.org/Install_kernel_from_RPM_on_Debian_6.0
>>>
>>> The bad thing about it is it's not a "native" way. The good thing is you
>>> have the same bit-by-bit binary kernel which we test a lot.
>> In the meanwhile I have tried two of the suggested solutions:
>>
>> * RHEL6-kernel (debianized with alien)
>> * proxmox-kernel from their repository (which for the 2.6.32 version
>> says it is based on RHEL6)
>>
>> (I haven't looked yet into building my own with CoolCold's script)
>>
>> I didn't have troubles to install either of those. However, both expose
>> the exact same problem. The network bridges vzbr0 and vzbr1 do not work

>> correctly and I need those for the veth devices of the containers.
>> Networking without those bridges seem to work fine, but when the bridges
>> are active, TCP connections fail to work. UDP seems to be not affected,
>> both inbound and outbound connections work. However, with TCP I can only
>> establish outbound connections, but as soon as the connection is
>> established, nothing is received on the other end anymore. Nor can I
>> send from outside to the machine with vzbr0.
> Our problem looks very similar to the one described here:
> https://bugzilla.redhat.com/show_bug.cgi?id=623199
>
> Could it be that current stable OpenVZ suffers from the same issue? Or
> asked differently: Is the current stable OpenVZ kernel based on RHEL 6.0
> which does not contain the fixes as described in:
> <http://rhn.redhat.com/errata/RHSA-2011-0542.html>

If you go to <http://wiki.openvz.org> and put RHSA 2011 0542 into a search box you will get an answer, which is kernel 042stab013.1 released a bit less than a year ago, end of May 2011.

So I am pretty sure the bug you are referring to is fixed.

Anyway, please file a new bug to bugzilla.openvz.org, and provide a simple scenario of how to reproduce it (please write a complete report).
