
Subject: VPS not isolated

Posted by [Wolfgang Schnerring](#) on Sat, 20 Jan 2007 07:32:08 GMT

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

Hello!

I've encountered a rather strange behaviour of OpenVZ:

The host is a Debian Sarge system running a vanilla 2.6.18 kernel that I've patched with patch-ovz028test010.1-combined. The guest system is a minimal Debian Sarge image, the config file is below.

When I vzctl start the VPS, I had expected that the VPS is started in the background and the command then returns to the shell, but that's not what happens: Instead I am logged out from the console I just used, I see messages from init, just like they appear on boot ("starting sshd... done, starting inetd... done", and so on) and finally get another login prompt for the HN.

When I then vzctl enter the VPS (which sometimes fails with the error "cannot open PTY") and do a "ps aux", I see all processes from the HN, not only those of the VPS.

When I vzctl stop the VPS, I again see messages from init which end in "sending KILL signal to all processes" -- and then the machine freezes.

Examining the syslog is not very enlightening, but there are a few totally garbled entries like this

```
Jan 19 15:22:29 amun2 kernel: 2: I6 rours esent  
(I think that was supposed read "no IPv6 routers present")
```

The same guest image and vz-configuration work fine on another machine (which is Debian Testing instead of Sarge, and runs a Debian kernel instead of a vanilla one).

To me, the whole thing looks like the VPS is not properly isolated from the HN, and its init somehow instead is started on the HN or something. But that's of course just a guess.

I am rather stumped on how to debug this issue and would greatly appreciate any help.

Thanks very much,
Wolfgang

```
# VPS config file, based on ve-vps.basic.conf-sample  
ONBOOT="no"
```

```
# UBC parameters (in form of barrier:limit)  
# Primary parameters  
AVNUMPROC="40:40"
```

```
NUMPROC="65:65"
NUMTCPSOCK="80:80"
NUMOTHERSOCK="80:80"
VMGUARPAGES="6144:2147483647"
# Secondary parameters
KMEMSIZE="2752512:2936012"
TCPSNDBUF="319488:524288"
TCPRCVBUF="319488:524288"
OTHERSOCKBUF="132096:336896"
DGRAMRCVBUF="132096:132096"
OOMGUARPAGES="6144:2147483647"
# Auxiliary parameters
LOCKEDPAGES="32:32"
SHMPAGES="8192:8192"
PRIVVMPAGES="49152:53575"
NUMFILE="2048:2048"
NUMFLOCK="100:110"
NUMPTY="16:16"
NUMSIGINFO="256:256"
DCACHESIZE="1048576:1097728"

PHYSPAGES="0:2147483647"
NUMIPTENT="128:128"

# Disk quota parameters (in form of softlimit:hardlimit)
DISKSPACE="1048576:1153434"
DISKINODES="200000:220000"
QUOTATIME="0"

# CPU fair scheduler parameter
CPUUNITS="1000"

VE_ROOT="/vm/vm25710/root"
VE_PRIVATE="/vm/vm25710/private"
HOSTNAME="vm0.local"
OSTEMPLATE="debian-sarge"
ORIGIN_SAMPLE="vps.basic"
```

Subject: Re: VPS not isolated
Posted by [Thorsten Schifferdeck\[1\]](#) on Sat, 20 Jan 2007 11:28:32 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Wolfgang,

Wolfgang Schnerring wrote:

> Hello!

>

> I've encountered a rather strange behaviour of OpenVZ:
> The host is a Debian Sarge system running a vanilla 2.6.18 kernel that
> I've patched with patch-ovz028test010.1-combined. The guest system is
> a minimal Debian Sarge image, the config file is below.
>
> When I vzctl start the VPS, I had expected that the VPS is started in
> the background and the command then returns to the shell, but that's
> not what happens: Instead I am logged out from the console I just
> used, I see messages from init, just like they appear on boot
> ("starting sshd... done, starting inetd... done", and so on) and
> finally get another login prompt for the HN.
> When I then vzctl enter the VPS (which sometimes fails with the error
> "cannot open PTY") and do a "ps aux", I see all processes from the HN,
> not only those of the VPS.
> When I vzctl stop the VPS, I again see messages from init which end in
> "sending KILL signal to all processes" -- and then the machine
> freezes.
> Examining the syslog is not very enlightening, but there are a few
> totally garbled entries like this
> Jan 19 15:22:29 amun2 kernel: 2: I6 rours esent
> (I think that was supposed read "no IPv6 routers present")
>
> The same guest image and vz-configuration work fine on another machine
> (which is Debian Testing instead of Sarge, and runs a Debian kernel
> instead of a vanilla one).
> To me, the whole thing looks like the VPS is not properly isolated
> from the HN, and its init somehow instead is started on the HN or
> something. But that's of course just a guess.
>
> I am rather stumped on how to debug this issue and would greatly
> appreciate any help.
>
> Thanks very much,
> Wolfgang
>
>
> # VPS config file, based on ve-vps.basic.conf-sample
> ONBOOT="no"
>
> # UBC parameters (in form of barrier:limit)
> # Primary parameters
> AVNUMPROC="40:40"
> NUMPROC="65:65"
> NUMTCPSOCK="80:80"
> NUMOTHERSOCK="80:80"
> VMGUARPAGES="6144:2147483647"
> # Secondary parameters
> KMEMSIZE="2752512:2936012"

```
> TCPSNDBUF="319488:524288"
> TCPRCVBUF="319488:524288"
> OTHERSOCKBUF="132096:336896"
> DGRAMRCVBUF="132096:132096"
> OOMGUARPAGES="6144:2147483647"
> # Auxiliary parameters
> LOCKEDPAGES="32:32"
> SHMPAGES="8192:8192"
> PRIVVMPAGES="49152:53575"
> NUMFILE="2048:2048"
> NUMFLOCK="100:110"
> NUMPTY="16:16"
> NUMSIGINFO="256:256"
> DCACHESIZE="1048576:1097728"
>
> PHYSPAGES="0:2147483647"
> NUMIPTENT="128:128"
>
> # Disk quota parameters (in form of softlimit:hardlimit)
> DISKSPACE="1048576:1153434"
> DISKINODES="200000:220000"
> QUOTATIME="0"
>
> # CPU fair scheduler parameter
> CPUUNITS="1000"
>
> VE_ROOT="/vm/vm25710/root"
> VE_PRIVATE="/vm/vm25710/private"
> HOSTNAME="vm0.local"
> OSTEMPLATE="debian-sarge"
> ORIGIN_SAMPLE="vps.basic"
```

Did you self-create your debian-sarge os-template ?
Still exist in this ve some tty dev-files (ls -l
/vm/vm25710/private/dev/tty*)

Can you execute a ps in your VE, like

```
# vzctl exec <VEID> ps axf
and only the VE process are shown ?
```

Bye
Thorsten

Subject: Re: VPS not isolated
Posted by [Wolfgang Schnerring](#) on Mon, 22 Jan 2007 09:11:28 GMT

* Thorsten Schifferdecker <ts@systs.org>:
> Wolfgang Schnerring wrote:
>> I've encountered a rather strange behaviour of OpenVZ:
>> The host is a Debian Sarge system running a vanilla 2.6.18 kernel that
>> I've patched with patch-ovz028test010.1-combined. The guest system is
>> a minimal Debian Sarge image, the config file is below.
>>
>> When I vzctl start the VPS, I had expected that the VPS is started in
>> the background and the command then returns to the shell, but that's
>> not what happens: Instead I am logged out from the console I just
>> used, I see messages from init, just like they appear on boot
>> ("starting sshd... done, starting inetd... done", and so on) and
>> finally get another login prompt for the HN.
>
> Did you self-create your debian-sarge os-template ?

Yes, I've created that template myself. And "unfortunately" it works perfectly fine on another machine.

I've just tried to use a precreated Sarge template (from <http://download.openvz.org/template/precreated/debian-3.1-i386-minimal.tar.gz>) but it shows the exact same behaviour.

> Still exist in this ve some tty dev-files (ls -l
> /vm/vm25710/private/dev/tty*)

Yes, there are lots of tty dev-files in the template. Could that be a problem?
As an experiment, I've removed them, but this did not change the behaviour at all.

> Can you execute a ps in your VE, like
> # vzctl exec <VEID> ps axf
> and only the VE process are shown ?

No, this shows all processes.

I'm grateful for any hints on how to debug this issue.

Thanks,
Wolfgang

Subject: Re: Re: VPS not isolated
Posted by [kir](#) on Mon, 22 Jan 2007 11:19:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

Wolfgang Schnerring wrote:

```
>> Can you execute a ps in your VE, like
>> # vzctl exec <VEID> ps axf
>> and only the VE process are shown ?
>>
```

```
>
> No, this shows all processes.
```

```
>
> The only reason I know why it could happen is if somebody mounted VE's
> /proc from the host system.
```

Try to do the following:

```
umount /path/to/VEroot/proc [usually it is /vz/root/VEID/proc]
vzctl exec VEID mount /proc
```

and see what the following command will show you

```
vzctl exec VEID ps ax
```

Subject: Re: VPS not isolated

Posted by [Wolfgang Schnerring](#) on Tue, 23 Jan 2007 15:58:16 GMT

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

* Kir Kolyshkin <kir@openvz.org>:

```
> Wolfgang Schnerring wrote:
```

```
>>> Can you execute a ps in your VE, like
```

```
>>> # vzctl exec <VEID> ps axf
```

```
>>> and only the VE process are shown ?
```

```
>>
```

```
>> No, this shows all processes.
```

```
>>
```

```
> The only reason I know why it could happen is if somebody mounted VE's /proc
> from the host system.
```

I don't actually understand what's going on, but the reason for this odd behaviour is that the VPS private area was lying on an NFS-mount.

I don't know whether that should work or whether that's a known issue, anyways it doesn't work, and it fails in colorful ways, as I've described.

I've now moved the private area to a "real" disk, and everything's fine.

Thanks,
Wolfgang

Subject: Re: Re: VPS not isolated

Posted by [dev](#) on Tue, 23 Jan 2007 16:07:59 GMT

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

Wolfgang,

> I don't actually understand what's going on, but the reason for this odd
> behaviour is that the VPS private area was lying on an NFS-mount.
> I don't know whether that should work or whether that's a known issue, anyways
> it doesn't work, and it fails in colorful ways, as I've described.
> I've now moved the private area to a "real" disk, and everything's fine.
Ouch, this is a good important piece of additional information!
Can you please post a bug in bugzilla with this information,
or even better provide an access (via private email)
for onsite quick debug/resolving?

Thanks,
Kirill

Subject: Re: VPS not isolated

Posted by [Wolfgang Schnerring](#) on Wed, 24 Jan 2007 09:41:05 GMT

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

* Kirill Korotaev <dev@sw.ru>:

>> I don't actually understand what's going on, but the reason for this odd
>> behaviour is that the VPS private area was lying on an NFS-mount.
>> I don't know whether that should work or whether that's a known issue, anyways
>> it doesn't work, and it fails in colorful ways, as I've described.
>> I've now moved the private area to a "real" disk, and everything's fine.
> Ouch, this is a good important piece of additional information!
> Can you please post a bug in bugzilla with this information,

Sure, it's http://bugzilla.openvz.org/show_bug.cgi?id=456

> or even better provide an access (via private email)
> for onsite quick debug/resolving?

I'm afraid that won't be possible, those are university machines and the regulations are strict. But it is easy to reproduce, just do a NFS mount of the private area from localhost to localhost.

Wolfgang

Subject: Re: Re: VPS not isolated

Posted by [dev](#) on Wed, 24 Jan 2007 09:51:29 GMT

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

>>>I don't actually understand what's going on, but the reason for this odd
>>>behaviour is that the VPS private area was lying on an NFS-mount.
>>>I don't know whether that should work or whether that's a known issue, anyways
>>>it doesn't work, and it fails in colorful ways, as I've described.
>>>I've now moved the private area to a "real" disk, and everything's fine.
>>
>>Ouch, this is a good important piece of additional information!
>>Can you please post a bug in bugzilla with this information,
>
>
> Sure, it's http://bugzilla.openvz.org/show_bug.cgi?id=456
>
>
>>or even better provide an access (via private email)
>>for onsite quick debug/resolving?
>
>
> I'm afraid that won't be possible, those are university machines and the
> regulations are strict. But it is easy to reproduce, just do a NFS mount of the
> private area from localhost to localhost.
ok. let's communicate through the bugzilla for better tracking.
I proposed you to try the same w/o vzquota.

Thanks,
Kirill
