
Subject: Routing problems using SMP kernel

Posted by [Steve Hodges](#) on Sun, 26 Aug 2007 14:59:57 GMT

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

After getting most of my problems solved I decided to move my test environment onto the production server.

The server is a dual xeon which, with hyperthreading, appears (to Linux) to have 4 processors. So, when I built this machine I decided to use the ovzkernel-2.6.18-smp

The rebuild caused me all sorts of routing problems which I have managed to track down to being caused by the kernel. I just replaced the kernel with ovzkernel-2.6.18

```
aptitude install ovzkernel-2.6.18
aptitude remove ovzkernel-2.6.18-smp
shutdown -r now
```

problem solved!

It seems pretty odd that the smp kernel could cause this, but I really don't know what is different about that kernel.

The symptoms were similar to the ones I had before I set the netmask of the venets correctly, but more extreme. Whereas the netmask issue seemed to cause packets to go out of the wrong interface, this problem seemed to stop packets getting out of the server at all.

If there are any questions about the symptoms, I will be able to swap back to that kernel for the next day or so to test things out.

What will the impact be of running the non-smp kernel on a multi-processor machine? Will I only effectively use a single processor?

Steve

Subject: Re: Routing problems using SMP kernel

Posted by [dev](#) on Mon, 27 Aug 2007 14:56:23 GMT

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

Steve,

Sure, SMP shouldn't affect your routing and it is very strange. I guess >90% of people are running SMP kernels.

>From your report it is totally unclear what OVZ kernel version is (e.g. something like 028stab039)

and where this kernel was got from. Have you built it yourself?
Can you please provide a bit more details on what is working and what not?
Why have you decided that it is routing to blame to?

Thanks,
Kirill

Steve Hodges wrote:

- > After getting most of my problems solved I decided to move my test
- > environment onto the production server.
- >
- > The server is a dual xeon which, with hyperthreading, appears (to Linux)
- > to have 4 processors. So, when I built this machine I decided to use
- > the ovzkernel-2.6.18-smp
- >
- > The rebuild caused me all sorts of routing problems which I have managed
- > to track down to being caused by the kernel. I just replaced the kernel
- > with ovzkernel-2.6.18
- >
- > aptitude install ovzkernel-2.6.18
- > aptitude remove ovzkernel-2.6.18-smp
- > shutdown -r now
- >
- > problem solved!
- >
- > It seems pretty odd that the smp kernel should cause this, but I really
- > don't know what is different about that kernel.
- >
- > The symptoms were similar to the ones I had before I set the netmask of
- > the venets correctly, but more extreme. Whereas the netmask issue
- > seemed to cause packets to go out of the wrong interface, this problem
- > seemed to stop packets getting out of the server at all.
- >
- > If there are any questions about the symptoms, I will be able to swap
- > back to that kernel for the next day or so to test things out.
- >
- > What will the impact be of running the non-smp kernel on a
- > multi-processor machine? Will I only effectively use a single processor?
- >
- > Steve

Subject: Re: Routing problems using SMP kernel
Posted by [Steve Hodges](#) on Mon, 27 Aug 2007 17:38:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 27/08/2007 10:57 PM, Kirill Korotaev wrote:
> Steve,

>
> Sure, SMP shouldn't affect your routing and it is very strange. I guess >90% of people
> are running SMP kernels.
>
> >From your report it is totally unclear what OVZ kernel version is (e.g. something like
028stab039)
> and where this kernel was got from. Have you built it yourself?
> Can you please provide a bit more details on what is working and what not?
> Why have you decided that it is routing to blame to?
>

it's 2.6.18-028stab035.1-ovz-smp obtained from deb
<http://debian.systs.org/stable/ovz/>

when I use the normal kernel I can ping from the VE to the HN and to
other VE's on this HN, to my other HN and to an external site (google.com)

when I use the smp kernel (no other change) I can ping from the VE to
the HN and to other VEs on this HN, but not the other HN or to external
sites

in all cases pinging from the HN is ok.

from the VE, if I try to do a traceroute to the HN it shows the HN as
the first hop (with either smp or normal kernel). If I traceroute to my
other HN, I just get endless * * * lines with the smp kernel (it doesn't
even show the HN as the first hop). With the normal kernel it shows the
HN, then the destination of the ping (the other HN in this case).

Is that a routing issue? dunno? but it looks like it might be. I was
actually leaning toward it being a hardware fault until I noticed the
anomaly in the traceroute.

I'm not sure if having 2 nics in the box has any bearing on it.

with the smp kernel I also note checksum errors when I do a ping -R. I
don't get those errors using the non-smp kernel.

OK, this gets extremely weird. I just checked the kernel I'm running and
it is still the smp version. and that is after I executed:

```
aptitude install ovzkernel-2.6.18  
aptitude remove ovzkernel-2.6.18-smp  
shutdown -r now
```

I am now concerned that this problem will recur if I am forced to reboot. It can't be as simple as
the reboot fixing it as I rebooted several times while I was having the problem and it didn't go
away.

I wonder if I have just entered the twilight zone?

Steve

> Thanks,

> Kirill

>

> Steve Hodges wrote:

>

>> After getting most of my problems solved I decided to move my test
>> environment onto the production server.

>>

>> The server is a dual xeon which, with hyperthreading, appears (to Linux)
>> to have 4 processors. So, when I built this machine I decided to use
>> the ovzkernel-2.6.18-smp

>>

>> The rebuild caused me all sorts of routing problems which I have managed
>> to track down to being caused by the kernel. I just replaced the kernel
>> with ovzkernel-2.6.18

>>

>> aptitude install ovzkernel-2.6.18
>> aptitude remove ovzkernel-2.6.18-smp
>> shutdown -r now

>>

>> problem solved!

>>

>> It seems pretty odd that the smp kernel could cause this, but I really
>> don't know what is different about that kernel.

>>

>> The symptoms were similar to the ones I had before I set the netmask of
>> the venets correctly, but more extreme. Whereas the netmask issue
>> seemed to cause packets to go out of the wrong interface, this problem
>> seemed to stop packets getting out of the server at all.

>>

>> If there are any questions about the symptoms, I will be able to swap
>> back to that kernel for the next day or so to test things out.

>>

>> What will the impact be of running the non-smp kernel on a
>> multi-processor machine? Will I only effectively use a single processor?

>>

>> Steve

Subject: Re: Routing problems using SMP kernel
Posted by [kir](#) on Mon, 27 Aug 2007 19:47:23 GMT

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

I guess it makes much sense to diagnose the hardware at this point. Some

info is available at http://wiki.openvz.org/Hardware_testing

Steve Hodges wrote:

> On 27/08/2007 10:57 PM, Kirill Korotaev wrote:

>> Steve,

>>

>> Sure, SMP shouldn't affect your routing and it is very strange. I

>> guess >90% of people

>> are running SMP kernels.

>>

>> >From your report it is totally unclear what OVZ kernel version is

>> (e.g. something like 028stab039)

>> and where this kernel was got from. Have you built it yourself?

>> Can you please provide a bit more details on what is working and what

>> not?

>> Why have you decided that it is routing to blame to?

>>

>

> it's 2.6.18-028stab035.1-ovz-smp obtained from deb

> [http://debian.systs.org/stable openvz](http://debian.systs.org/stable/openvz)

>

> when I use the normal kernel I can ping from the VE to the HN and to

> other VE's on this HN, to my other HN and to an external site

> (google.com)

>

> when I use the smp kernel (no other change) I can ping from the VE to

> the HN and to other VEs on this HN, but not the other HN or to

> external sites

>

> in all cases pinging from the HN is ok.

>

> from the VE, if I try to do a traceroute to the HN it shows the HN as

> the first hop (with either smp or normal kernel). If I traceroute to

> my other HN, I just get endless * * * lines with the smp kernel (it

> doesn't even show the HN as the first hop). With the normal kernel it

> shows the HN, then the destination of the ping (the other HN in this

> case).

>

> Is that a routing issue? dunno? but it looks like it might be. I

> was actually leaning toward it being a hardware fault until I noticed

> the anomaly in the traceroute.

>

> I'm not sure if having 2 nics in the box has any bearing on it.

>

> with the smp kernel I also note checksum errors when I do a ping -R. I

> don't get those errors using the non-smp kernel.

>

> OK, this gets extremely weird. I just checked the kernel I'm running

> and it is still the smp version. and that is after I executed:
>
> aptitude install ovzkernel-2.6.18
> aptitude remove ovzkernel-2.6.18-smp
> shutdown -r now
>
> I am now concerned that this problem will recur if I am forced to
> reboot. It can't be as simple as the reboot fixing it as I rebooted
> several times while I was having the problem and it didn't go away.
>
> I wonder if I have just entered the twilight zone?
