
Subject: Outbound networking failure

Posted by [iowissen](#) on Wed, 14 Mar 2012 06:28:41 GMT

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

sorry if this is more of a Xen problem.

we are running openvz in Xen PVs (kernel 2.6.18-194.3.1.el5.028stab069.6xen), and occasionally the network of a PV, as well as the PV's hosted containers, at a virtual interface is lost. sometimes restarting the network can solve the problem but in most cases we have to restart the PV. the PVs another virtual interface was working and other PVs and the HV keeps working.

detailed observation through tcpdump at the trouble time shows:

1. the inbound traffics can be delivered correctly through eth0 of the PV until venet0;
2. it looks the venet0 is also correctly functioned because TCP sync may get response from container and this sync ack can be seen at venet0;
3. but at the eth0 of PV, we only have seen the TCP sync request towards the containers while the sync ack is missing, thus TCP state of the container's application remains SYNC_RECV
4. if ping the PV, at the eth0, we only have seen the ICMP echo requests but the replies are missing.

we didn't encounter the same problem when those openvz stuffs were running over physical machine instead of xen.

does anyone have any hints or ideas? thanks a lot in advance.

- maoke

Subject: Re: Outbound networking failure

Posted by [Andrew Vagin](#) on Wed, 14 Mar 2012 07:12:23 GMT

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

2.6.18-194.3.1.el5.028stab069.6xen is a very old version. Could you update it?

On 03/14/2012 10:28 AM, Maoke wrote:

> sorry if this is more of a Xen problem.

>

> we are running openvz in Xen PVs (kernel

> 2.6.18-194.3.1.el5.028stab069.6xen), and occasionally the network of a

> PV, as well as the PV's hosted containers, at a virtual interface is

> lost. sometimes restarting the network can solve the problem but in

> most cases we have to restart the PV. the PVs another virtual

> interface was working and other PVs and the HV keeps working.

>
> detailed observation through tcpdump at the trouble time shows:
> 1. the inbound traffics can be delivered correctly through eth0 of the
> PV until venet0;
> 2. it looks the venet0 is also correctly functioned because TCP sync
> may get response from container and this sync ack can be seen at venet0;
> 3. but at the eth0 of PV, we only have seen the TCP sync request
> towards the containers while the sync ack is missing, thus TCP state
> of the container's application remains SYNC_RECV
> 4. if ping the PV, at the eth0, we only have seen the ICMP echo
> requests but the replies are missing.
>
> we didn't encounter the same problem when those openvz stuffs were
> running over physical machine instead of xen.
>
> does anyone have any hints or ideas? thanks a lot in advance.
>
> - maoke
>
>

Subject: Re: Outbound networking failure
Posted by [iowissen](#) on Wed, 14 Mar 2012 07:26:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

yes but we still have other troubles with the 2.6.32-042stab044.11. :P so
we are trying to understand the problem first. - maoke

2012/3/14 Andrew Vagin <avagin@parallels.com>

> 2.6.18-194.3.1.el5.028stab069.6xen is a very old version. Could you update
> it?
>
>
> On 03/14/2012 10:28 AM, Maoke wrote:
>
>> sorry if this is more of a Xen problem.
>>
>> we are running openvz in Xen PVs (kernel
>> 2.6.18-194.3.1.el5.028stab069.6xen), and occasionally the network of a PV,
>> as well as the PV's hosted containers, at a virtual interface is lost.
>> sometimes restarting the network can solve the problem but in most cases we
>> have to restart the PV. the PVs another virtual interface was working and
>> other PVs and the HV keeps working.
>>
>> detailed observation through tcpdump at the trouble time shows:
>> 1. the inbound traffics can be delivered correctly through eth0 of the PV

>> until venet0;
>> 2. it looks the venet0 is also correctly functioned because TCP sync may
>> get response from container and this sync ack can be seen at venet0;
>> 3. but at the eth0 of PV, we only have seen the TCP sync request towards
>> the containers while the sync ack is missing, thus TCP state of the
>> container's application remains SYNC_RECV
>> 4. if ping the PV, at the eth0, we only have seen the ICMP echo requests
>> but the replies are missing.
>>
>> we didn't encounter the same problem when those openvz stuffs were
>> running over physical machine instead of xen.
>>
>> does anyone have any hints or ideas? thanks a lot in advance.
>>
>> - maoke
>>
>>

Subject: Re: Outbound networking failure
Posted by [Andrew Vagin](#) on Wed, 14 Mar 2012 07:33:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 03/14/2012 11:26 AM, Maoke wrote:
> yes but we still have other troubles with the 2.6.32-042stab044.11. :P
> so we are trying to understand the problem first.
I don't suggest to move on 2.6.32 kernel. I said about
2.6.18-028stabXXX.Y. You can find a last stable 2.6.18 kernel here:
<http://download.openvz.org/kernel/branches/rhel5-2.6.18/stable/>

This issue may be already fixed.

> - maoke
>
> 2012/3/14 Andrew Vagin <avagin@parallels.com>
> <<mailto:avagin@parallels.com>>>
>
> 2.6.18-194.3.1.el5.028stab069.6xen is a very old version. Could
> you update it?
>
>
> On 03/14/2012 10:28 AM, Maoke wrote:
>
> sorry if this is more of a Xen problem.
>
> we are running openvz in Xen PVs (kernel
> 2.6.18-194.3.1.el5.028stab069.6xen), and occasionally the
> network of a PV, as well as the PV's hosted containers, at a

> virtual interface is lost. sometimes restarting the network
> can solve the problem but in most cases we have to restart the
> PV. the PVs another virtual interface was working and other
> PVs and the HV keeps working.
>
> detailed observation through tcpdump at the trouble time shows:
> 1. the inbound traffics can be delivered correctly through
> eth0 of the PV until venet0;
> 2. it looks the venet0 is also correctly functioned because
> TCP sync may get response from container and this sync ack can
> be seen at venet0;
> 3. but at the eth0 of PV, we only have seen the TCP sync
> request towards the containers while the sync ack is missing,
> thus TCP state of the container's application remains SYNC_RECV
> 4. if ping the PV, at the eth0, we only have seen the ICMP
> echo requests but the replies are missing.
>
> we didn't encounter the same problem when those openvz stuffs
> were running over physical machine instead of xen.
>
> does anyone have any hints or ideas? thanks a lot in advance.
>
> - maoke

> Users mailing list
> Users@openvz.org <mailto:Users@openvz.org>
> <https://openvz.org/mailman/listinfo/users>
>
>
>

Subject: Re: Outbound networking failure
Posted by [iowissen](#) on Wed, 14 Mar 2012 07:55:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

2012/3/14 Andrew Vagin <avagin@parallels.com>

> On 03/14/2012 11:26 AM, Maoke wrote:
>
>> yes but we still have other troubles with the 2.6.32-042stab044.11. :P so
>> we are trying to understand the problem first.
>>
> I don't suggest to move on 2.6.32 kernel. I said about
> 2.6.18-028stabXXX.Y. You can find a last stable 2.6.18 kernel here:
> <http://download.openvz.org/kernel/branches/rhel5-2.6.18/stable/>

>
> This issue may be already fixed.
>

thanks for the information! then is this a known bug of the old kernel? is there a way to repeat it so that we may test before putting the new kernel to the production environment. thanks again!

maoke

>
> - maoke
>>
>> 2012/3/14 Andrew Vagin <avagin@parallels.com <mailto:avagin@parallels.com
>> >>
>>
>>
>> 2.6.18-194.3.1.el5.028stab069.6xen is a very old version. Could
>> you update it?
>>
>>
>> On 03/14/2012 10:28 AM, Maoke wrote:
>>
>> sorry if this is more of a Xen problem.
>>
>> we are running openvz in Xen PVs (kernel
>> 2.6.18-194.3.1.el5.028stab069.6xen), and occasionally the
>> network of a PV, as well as the PV's hosted containers, at a
>> virtual interface is lost. sometimes restarting the network
>> can solve the problem but in most cases we have to restart the
>> PV. the PVs another virtual interface was working and other
>> PVs and the HV keeps working.
>>
>> detailed observation through tcpdump at the trouble time shows:
>> 1. the inbound traffics can be delivered correctly through
>> eth0 of the PV until venet0;
>> 2. it looks the venet0 is also correctly functioned because
>> TCP sync may get response from container and this sync ack can
>> be seen at venet0;
>> 3. but at the eth0 of PV, we only have seen the TCP sync
>> request towards the containers while the sync ack is missing,
>> thus TCP state of the container's application remains SYNC_RECV
>> 4. if ping the PV, at the eth0, we only have seen the ICMP
>> echo requests but the replies are missing.
>>
>> we didn't encounter the same problem when those openvz stuffs
>> were running over physical machine instead of xen.

>>
>> does anyone have any hints or ideas? thanks a lot in advance.
>>
>> - maoke
>>
>>
>>
>> _____
>> Users mailing list
>> Users@openvz.org <mailto:Users@openvz.org>
>> https://openvz.org/mailman/listinfo/users
>>
>>
>>
>>
>

Subject: separate block device for CT's swap
Posted by [stealth](#) on Thu, 15 Mar 2012 09:38:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

Is it possible to use separate block device as a swap for CT?

Subject: Re: separate block device for CT's swap
Posted by [Tim Small](#) on Thu, 15 Mar 2012 09:58:30 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 15/03/12 09:38, stealth wrote:

> Is it possible to use separate block device as a swap for CT?

AFAIK, no because of the way that the kernel's VM subsystem works. You can of course use multiple block devices as swap for the entire machine, and this will give you better performance overall than separate swaps for each CT would anyway...

Cheers,

Tim.

--

South East Open Source Solutions Limited
Registered in England and Wales with company number 06134732.
Registered Office: 2 Powell Gardens, Redhill, Surrey, RH1 1TQ
VAT number: 900 6633 53 <http://seoss.co.uk/> +44-(0)1273-808309

Subject: Re: separate block device for CT's swap
Posted by [Rick van Rein](#) on Thu, 15 Mar 2012 09:59:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

> Is it possible to use separate block device as a swap for CT?

That conflicts the design of OpenVZ; resources such as memory are shared between containers, so you cannot split the swap.

-Rick
