
Subject: Network slowness

Posted by [Kai Hendry](#) on Mon, 02 Jul 2007 14:13:43 GMT

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

Hi, I am using a kernel from deb <http://download.openvz.org/debian> etch main

Linux argo 2.6.18-028stab033.1-ovz-smp #1 SMP Thu May 31 04:17:27 CEST 2007 i686 GNU/Linux

```
argo$ sudo vzctl --version
vzctl version 3.0.11
```

I've setup an openvz VE (running Debian Unstable) like so:
<http://faq.database.com/faqw.py?req=show&file=faq01.152.ht p>

And I am quite happy with it. Except I have experiencing some odd slow network behaviour, which I am having trouble debugging.

I first noticed it when my apt-get upgrades hung (timed out).

Now I've been noticing for an application in the VE to connect to a remote mysql DB.

monty is my laptop in the UK
simon is the VE
argo is the host

<http://pastebin.com/940574>

You can see that simon seems to take consistently almost 30 secs to get a couple of records from the DB. Any ideas why?

I have tried stopping and starting simon. I suspected DNS troubles at first, though simon to be able to look up stuff as fast as argo.

Kind regards,

```
simon:~% sudo /sbin/ifconfig
lo      Link encap:Local Loopback
        inet addr:127.0.0.1  Mask:255.0.0.0
        inet6 addr: ::1/128 Scope:Host
        UP LOOPBACK RUNNING MTU:16436 Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
```

```
venet0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
inet addr:127.0.0.1 P-t-P:127.0.0.1 Bcast:0.0.0.0 Mask:255.255.255.255
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
RX packets:1858952 errors:0 dropped:0 overruns:0 frame:0
TX packets:1771614 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:851745664 (812.2 MiB) TX bytes:258658607 (246.6 MiB)
```

```
venet0:0 Link encap:UNSPEC HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
inet addr:192.168.2.66 P-t-P:192.168.2.66 Bcast:0.0.0.0 Mask:255.255.255.255
UP BROADCAST POINTOPOINT RUNNING NOARP MTU:1500 Metric:1
```

```
simon:~% sudo route
```

```
Kernel IP routing table
```

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
191.255.255.1	*	255.255.255.255	UH	0	0	0	venet0
default	191.255.255.1	0.0.0.0	UG	0	0	0	venet0

```
simon:~% sudo cat /etc/resolv.conf
```

```
search nipl.net
```

```
nameserver 72.29.96.250
```

File Attachments

1) [si-network.txt](#), downloaded 606 times

Subject: Re: Network slowness

Posted by [dev](#) on Mon, 02 Jul 2007 14:27:53 GMT

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

Check /proc/user_beancounters limits (post it here).
It can be due to a limit network buffers allowed.

Thanks,
Kirill

Kai Hendry wrote:

```
> Hi, I am using a kernel from deb http://download.openvz.org/debian etch main
>
> Linux argo 2.6.18-028stab033.1-ovz-smp #1 SMP Thu May 31 04:17:27 CEST
> 2007 i686 GNU/Linux
>
> argo$ sudo vzctl --version
> vzctl version 3.0.11
>
>
> I've setup an openvz VE (running Debian Unstable) like so:
> http://faq.dabase.com/faqw.py?req=show&file=faq01.152.ht p
>
```

```

> And I am quite happy with it. Except I have experiencing some odd
> slow network behaviour, which I am having trouble debugging.
>
> I first noticed it when my apt-get upgrades hung (timed out).
>
> Now I've been noticing for an application in the VE to connect to a
> remote mysql DB.
>
> monty is my laptop in the UK
> simon is the VE
> argo is the host
>
> http://pastebin.com/940574
>
> You can see that simon seems to take consistently almost 30 secs to
> get a couple of records from the DB. Any ideas why?
>
> I have tried stopping and starting simon. I suspected DNS troubles at
> first, though simon to be able to look up stuff as fast as argo.
>
> Kind regards,
>
> -----
>
> simon:~% sudo /sbin/ifconfig
> lo      Link encap:Local Loopback
>         inet addr:127.0.0.1  Mask:255.0.0.0
>         inet6 addr: ::1/128 Scope:Host
>         UP LOOPBACK RUNNING  MTU:16436  Metric:1
>         RX packets:0 errors:0 dropped:0 overruns:0 frame:0
>         TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
>         collisions:0 txqueuelen:0
>         RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
>
> venet0  Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
>         inet addr:127.0.0.1  P-t-P:127.0.0.1  Bcast:0.0.0.0  Mask:255.255.255.255
>         UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
>         RX packets:1858952 errors:0 dropped:0 overruns:0 frame:0
>         TX packets:1771614 errors:0 dropped:0 overruns:0 carrier:0
>         collisions:0 txqueuelen:0
>         RX bytes:851745664 (812.2 MiB)  TX bytes:258658607 (246.6 MiB)
>
> venet0:0 Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00
>         inet addr:192.168.2.66  P-t-P:192.168.2.66  Bcast:0.0.0.0  Mask:255.255.255.255
>         UP BROADCAST POINTOPOINT RUNNING NOARP  MTU:1500  Metric:1
>
> simon:~% sudo route

```

```

> Kernel IP routing table
> Destination Gateway Genmask Flags Metric Ref Use Iface
> 191.255.255.1 * 255.255.255.255 UH 0 0 0 venet0
> default 191.255.255.1 0.0.0.0 UG 0 0 0 venet0
> simon:~% sudo cat /etc/resolv.conf
> search nipl.net
> nameserver 72.29.96.250
>
>
> -----
>

```

Subject: Re: Network slowness
 Posted by [Kai Hendry](#) on Mon, 02 Jul 2007 14:32:12 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 7/2/07, Kirill Korotaev <dev@sw.ru> wrote:
 > Check /proc/user_beancounters limits (post it here).
 > It can be due to a limit network buffers allowed.

Attached.

Version: 2.5

uid resource	held	maxheld	barrier	limit	failcnt	
100: kmemsize	884604	2324396	2752512	2936012		175
lockedpages	0	0	32	32	0	
privvmpages	36369	56976	262144	262144		0
shmpages	640	1328	8192	8192	0	
dummy	0	0	0	0	0	
numproc	27	36	65	65	0	
physpages	4147	9574	0	2147483647		0
vmguarpages	0	0	65536	65536	0	
oomguarpages	4147	9574	6144	2147483647		0
numtcpsock	3	6	80	80	0	
numflock	1	7	100	110	0	
numpty	1	2	16	16	0	
numsiginfo	0	13	256	256	0	
tcpsndbuf	2220	59940	319488	524288		0
tcprcvbuf	0	277500	319488	524288	6502	
othersockbuf	8880	30304	132096	336896		0
dgramrcvbuf	0	8364	132096	132096		0
numothersock	8	18	80	80	0	
dcachesize	0	0	1048576	1097728		0
numfile	606	860	2048	2048	916	
dummy	0	0	0	0	0	
dummy	0	0	0	0	0	
dummy	0	0	0	0	0	

numiptent	10	10	128	128	0	
0: kmemsize	35762556	39840332	2147483647	2147483647	2147483647	0
lockedpages	0	6	2147483647	2147483647	0	
privvmpages	684068	762359	2147483647	2147483647		0
shmpages	154094	161420	2147483647	2147483647		0
dummy	0	0	2147483647	2147483647	0	
numproc	363	416	2147483647	2147483647	0	
physpages	195855	222280	2147483647	2147483647		0
vmguarpages	0	0	2147483647	2147483647	0	
oomguarpages	299011	319377	2147483647	2147483647		1
numtcpsock	228	313	2147483647	2147483647	0	
numflock	39	54	2147483647	2147483647	0	
numpty	19	23	2147483647	2147483647	0	
numsiginfo	10	23	2147483647	2147483647	0	
tcpsndbuf	2824620	3521872	2147483647	2147483647		0
tcprcvbuf	7400004	11458836	2147483647	2147483647		0
othersockbuf	425556	959780	2147483647	2147483647		0
dgramrcvbuf	31752	47628	2147483647	2147483647		0
numothersock	274	331	2147483647	2147483647		0
dcachesize	0	0	2147483647	2147483647	0	
numfile	64362	65677	2147483647	2147483647		0
dummy	0	0	2147483647	2147483647	0	
dummy	0	0	2147483647	2147483647	0	
dummy	0	0	2147483647	2147483647	0	
numiptent	17	17	2147483647	2147483647	0	

File Attachments

1) [bean.txt](#), downloaded 620 times

Subject: Re: Network slowness

Posted by [dev](#) on Mon, 02 Jul 2007 15:00:40 GMT

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

It is obvious from this output, that this VE requires much more resources than available for it.

Check whether failcnt column counters are increasing over the time to understand what resource shortages your VEs do experience.

>From this output it is clear that it has shortages of kmemsize, tcprcvbuf and numfile resources.

You can increase those using the command:
vzctl set <VEID> --<resource> <new-value>

(also check for vzsplittool)

Thanks,

Kirill

Kai Hendry wrote:

> On 7/2/07, Kirill Korotaev <dev@sw.ru> wrote:

>

>>Check /proc/user_beancounters limits (post it here).

>>It can be due to a limit network buffers allowed.

>

>

> Attached.

>

>

> -----

>

> Version: 2.5

> uid resource	held	maxheld	barrier	limit	failcnt	
> 100: kmemsize	884604	2324396	2752512	2936012		175
> lockedpages	0	0	32	32	0	
> privvmpages	36369	56976	262144	262144		0
> shmpages	640	1328	8192	8192	0	
> dummy	0	0	0	0	0	
> numproc	27	36	65	65	0	
> physpages	4147	9574	0	2147483647		0
> vmguarpages	0	0	65536	65536	0	
> oomguarpages	4147	9574	6144	2147483647		0
> numtcpsock	3	6	80	80	0	
> numflock	1	7	100	110	0	
> numpty	1	2	16	16	0	
> numsiginfo	0	13	256	256	0	
> tcpsndbuf	2220	59940	319488	524288		0
> tcprcvbuf	0	277500	319488	524288	6502	
> othersockbuf	8880	30304	132096	336896		0
> dgramrcvbuf	0	8364	132096	132096		0
> numothersock	8	18	80	80	0	
> dcachesize	0	0	1048576	1097728		0
> numfile	606	860	2048	2048	916	
> dummy	0	0	0	0	0	
> dummy	0	0	0	0	0	
> dummy	0	0	0	0	0	
> numiptent	10	10	128	128	0	
> 0: kmemsize	35762556	39840332	2147483647	2147483647		0
> lockedpages	0	6	2147483647	2147483647		0
> privvmpages	684068	762359	2147483647	2147483647		0
> shmpages	154094	161420	2147483647	2147483647		0
> dummy	0	0	2147483647	2147483647		0
> numproc	363	416	2147483647	2147483647		0
> physpages	195855	222280	2147483647	2147483647		0

```

> vmguarpages      0      0 2147483647 2147483647      0
> oomguarpages    299011  319377 2147483647 2147483647      1
> numtcpsock      228     313 2147483647 2147483647      0
> numflock        39      54 2147483647 2147483647      0
> numpty          19      23 2147483647 2147483647      0
> numsigninfo     10      23 2147483647 2147483647      0
> tcpsndbuf       2824620 3521872 2147483647 2147483647      0
> tcprcvbuf       7400004 11458836 2147483647 2147483647      0
> othersockbuf    425556  959780 2147483647 2147483647      0
> dgramrcvbuf     31752   47628 2147483647 2147483647      0
> numothersock    274     331 2147483647 2147483647      0
> dcache size     0      0 2147483647 2147483647      0
> numfile         64362  65677 2147483647 2147483647      0
> dummy           0      0 2147483647 2147483647      0
> dummy           0      0 2147483647 2147483647      0
> dummy           0      0 2147483647 2147483647      0
> numiptent       17      17 2147483647 2147483647      0

```

Subject: Re: Network slowness

Posted by [Kai Hendry](#) on Mon, 02 Jul 2007 16:01:04 GMT

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

On 7/2/07, Thorsten Schifferdecker <curx@openvz.org> wrote:

```

> ^- this is a precompiled OpenVZ kernel from repro at debian.systs.org, not
> from download.openvz.org !

```

Sorry, you're right. Should I hurry to upgrade to -35?

Subject: Re: Network slowness

Posted by [Kai Hendry](#) on Mon, 02 Jul 2007 17:49:30 GMT

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

I've ramped up these limits, though it is still slow. ./

```

python db-test.py 0.46s user 0.01s system 1% cpu 28.852 total
python db-test.py 0.46s user 0.01s system 1% cpu 28.791 total

```

Version: 2.5

uid	resource	held	maxheld	barrier	limit	failcnt
100:	kmemsize	1061955	2324396	2147483647	2147483647	175
	lockedpages	0	0	32	32	0
	privvmpages	37094	56976	262144	262144	0
	shmpages	640	1328	8192	8192	0
	dummy	0	0	0	0	0
	numproc	30	39	65	65	0

physpages	4636	9779	0	2147483647	0
vmguarpages	0	0	65536	65536	0
oomguarpages	4636	9779	6144	2147483647	0
numtcpsock	3	6	80	80	0
numflock	1	7	100	110	0
numpty	2	3	16	16	0
numsiginfo	0	13	256	256	0
tcpsndbuf	2220	59940	319488	524288	0
tcprcvbuf	49152	277500	2147483647	2147483647	6502
othersockbuf	8880	30304	132096	336896	0
dgramrcvbuf	0	8364	132096	132096	0
numothersock	8	18	80	80	0
dcachesize	0	0	1048576	1097728	0
numfile	658	860	2147483647	2147483647	916
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	10	10	128	128	0
0: kmemsize	36178920	39840332	2147483647	2147483647	0
lockedpages	0	6	2147483647	2147483647	0
privvmpages	705250	762359	2147483647	2147483647	0
shmpages	154094	161420	2147483647	2147483647	0
dummy	0	0	2147483647	2147483647	0
numproc	359	416	2147483647	2147483647	0
physpages	204693	226136	2147483647	2147483647	0
vmguarpages	0	0	2147483647	2147483647	0
oomguarpages	307863	329313	2147483647	2147483647	1
numtcpsock	225	313	2147483647	2147483647	0
numflock	42	54	2147483647	2147483647	0
numpty	21	23	2147483647	2147483647	0
numsiginfo	10	23	2147483647	2147483647	0
tcpsndbuf	2569320	3712620	2147483647	2147483647	0
tcprcvbuf	7332912	11458836	2147483647	2147483647	0
othersockbuf	425556	959780	2147483647	2147483647	0
dgramrcvbuf	31752	47628	2147483647	2147483647	0
numothersock	274	331	2147483647	2147483647	0
dcachesize	0	0	2147483647	2147483647	0
numfile	66572	66890	2147483647	2147483647	0
dummy	0	0	2147483647	2147483647	0
dummy	0	0	2147483647	2147483647	0
dummy	0	0	2147483647	2147483647	0
numiptent	17	17	2147483647	2147483647	0

File Attachments

1) [bean.txt](#), downloaded 583 times

Subject: Re: Network slowness
Posted by [Leslie Watter](#) on Mon, 02 Jul 2007 19:44:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Kai,

Try to increase these limits again, until you don't get an increase in the failcnt column.

When you get a failcnt it is telling you that you don't have the right amount of resources, and things can (and probably will) go wrong.

Cheers,

LEslie

2007/7/2, Kai Hendry <kai.hendry@gmail.com>:
> I've ramped up these limits, though it is still slow. :/
>
> python db-test.py 0.46s user 0.01s system 1% cpu 28.852 total
> python db-test.py 0.46s user 0.01s system 1% cpu 28.791 total
>
--
Leslie H. Watter
