
Subject: *SOLVED* No resolve inside one VE
Posted by [3molo](#) on Fri, 11 May 2007 07:41:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

One of my vps users has had the same problem although destroyed->created new VE. No resolve, but only for that specific users VE. All VEs are debian at me casa. This particular VE does not run iptables rules. Tried 4-5 different nameservers, all of which I can ping but none of which I can resolve from. Isolated problem to, better believe me, postfix(!).

```
sputnik:/# ping -c1 ping.sunet.se
<nothing happens for minutes>
sputnik:/# /etc/init.d/postfix stop
Stopping Postfix Mail Transport Agent: postfix.
sputnik:/# ping -c1 ping.sunet.se
PING ping.sunet.se (130.242.80.31) 56(84) bytes of data.
64 bytes from stockholm1.sunet.se (130.242.80.31): icmp_seq=1 ttl=249 time=10.0 ms
--- ping.sunet.se ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 10.057/10.057/10.057/0.000 ms
```

Nameservers.

Tried one of each, some of them and all of them in resolv.conf. All of them works from other VEs.

```
# cat /etc/resolv.conf
nameserver 84.246.88.10
nameserver 207.69.188.185
nameserver 207.69.188.186
nameserver 207.69.188.187
nameserver 24.136.181.182
```

Iptables.

```
sputnik:/# iptables -L
Chain INPUT (policy ACCEPT)
target    prot opt source      destination
Chain FORWARD (policy ACCEPT)
target    prot opt source      destination
Chain OUTPUT (policy ACCEPT)
target    prot opt source      destination
```

Iptables on HN.

```
# iptables -L
Chain INPUT (policy ACCEPT)
target    prot opt source      destination
Chain FORWARD (policy ACCEPT)
target    prot opt source      destination
Chain OUTPUT (policy ACCEPT)
target    prot opt source      destination
```

```
# iptables -L -t nat
Chain PREROUTING (policy ACCEPT)
target    prot opt source                destination
Chain POSTROUTING (policy ACCEPT)
target    prot opt source                destination
Chain OUTPUT (policy ACCEPT)
target    prot opt source                destination
```

Other information:

Debian 4.0. 2.6.18-028stab027 SMP i686.

```
ii postfix                2.4.0-3
ii popa3d                 1.0.2-3
```

Summary:

Problem only occurs when postfix is started inside VE.

VE can ping all, for me, pingable internet hosts.

No iptables inside VE, default policy of HN is ACCEPT.

This specifk VE, this specifk user and his two applications popa3d + postfix is the only thing to cause the problem.

How to debug in order to file this as a bug?

Subject: Re: No resolve inside one VE

Posted by [morik](#) on Fri, 11 May 2007 07:52:41 GMT

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

please execute and show output:

```
sputnik:/# echo "-----Before postfix start-----"
sputnik:/# ip a l
sputnik:/# ip r l
sputnik:/# /etc/init.d/postfix start
sputnik:/# echo "-----After postfix start-----"
sputnik:/# ip a l
sputnik:/# ip r l
sputnik:/# /etc/init.d/postfix stop
sputnik:/# echo "-----After postfix stop-----"
sputnik:/# ip a l
sputnik:/# ip r l
```

Subject: Re: No resolve inside one VE

Posted by [Vasily Tarasov](#) on Fri, 11 May 2007 07:53:30 GMT

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

The very strange problem
I have two ideas

1. postfix init script changes some network configuration
2. Some resources shortage: can you look at /proc/user_beancounters after unfortunate ping, please.

Thanks,
Vasilyu

Subject: Re: No resolve inside one VE

Posted by [3molo](#) on Fri, 11 May 2007 07:59:57 GMT

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

Thanks Morik and Vasily!

failcnt for tcprcvbuf and othersuckbuf is the problem, so Ill just raise them alittle

resource	held	maxheld	barrier	limit	failcnt
tcprcvbuf	0	454392	319488	524288	238
othersockbuf	130980	133032	132096	336896	16061

This one I should have figured out. Thanks!

Subject: Re: No resolve inside one VE

Posted by [3molo](#) on Fri, 11 May 2007 08:38:44 GMT

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

Ok so something has happened. What I did was change the othersockbuf, tcprcvbuf & tcpsndbuf parameters (set VEID --parameter value:value --save) and upon stoping the VE:

```
flop:/etc/vz/conf# vzctl stop 107
```

```
Stopping VE ...
```

```
flop kernel: Oops: 0000 [#1]
```

```
flop kernel: SMP
```

```
flop kernel: CPU: 1, VCPU: 0.1
```

```
flop kernel: EIP is at ipt_unregister_table+0xb/0x1b0 [ip_tables]
```

```
flop kernel: eax: 00000000 ebx: 00104007 ecx: f8bff770 edx: f6f61e80
```

```
flop kernel: esi: 00000000 edi: 00104007 ebp: 00000000 esp: e7d99f44
```

```
flop kernel: ds: 007b es: 007b ss: 0068
```

```
flop kernel: Process vzmond/107 (pid: 24906, veid: 0, ti=e7d99000 task=f62cb950
```

```
task.ti=e7d99000)
```

```
flop kernel: Stack: 00000282 c013f797 f6295600 00000000 c2337800 c22fc2c0 f8c00d28
```

00104007

flop kernel: 00000000 00104007 00000000 f8bff3fb 00000000 f8bff89d f8c00c80 00000006

flop kernel: f8b9158f f6143000 00000000 f6143000 f6143028 c0520700 00000000 f8b93ab6

flop kernel: Call Trace:

flop kernel: Code: 04 89 d0 5a c3 c7 04 24 02 00 00 00 e8 5f b0 ff ff 89 c2 89 d0 5a c3 89 f6 8d
bc 27 00 00 00 55 57 56 53 83 ec 1c 8b 44 24 30 <8b> 50 34 89 04 24 89 54 24 14 e8 c6 ad ff
ff 89 44 24 0c 89 e0

flop kernel: EIP: [<f89521ab>] ipt_unregister_table+0xb/0x1b0 [ip_tables] SS:ESP
0068:e7d99f44

Unable to stop VE: operation timed out

flop:~# cat /var/log/kernel

<snip>

May 11 10:16:28 flop kernel: CPU: 1, VCPU: 0.1

May 11 10:16:28 flop kernel: EIP: 0060:[<f89521ab>] Not tainted VLI

May 11 10:16:28 flop kernel: EFLAGS: 00010296 (2.6.18-028stab027 #1)

May 11 10:16:28 flop kernel: EIP is at ipt_unregister_table+0xb/0x1b0 [ip_tables]

May 11 10:16:28 flop kernel: eax: 00000000 ebx: 00104007 ecx: f8bff770 edx: f6f61e80

May 11 10:16:28 flop kernel: esi: 00000000 edi: 00104007 ebp: 00000000 esp: e7d99f44

May 11 10:16:28 flop kernel: ds: 007b es: 007b ss: 0068

May 11 10:16:28 flop kernel: Process vzmond/107 (pid: 24906, veid: 0, ti=e7d99000
task=f62cb950 task.ti=e7d99000)

May 11 10:16:28 flop kernel: Stack: 00000282 c013f797 f6295600 00000000 c2337800 c22fc2c0
f8c00d28 00104007

May 11 10:16:28 flop kernel: 00000000 00104007 00000000 f8bff3fb 00000000 f8bff89d f8c00c80
00000006

May 11 10:16:28 flop kernel: f8b9158f f6143000 00000000 f6143000 f6143028 c0520700
00000000 f8b93ab6

May 11 10:16:28 flop kernel: Call Trace:

May 11 10:16:28 flop kernel: [<c013f797>] ub_slab_uncharge+0x87/0xb0

May 11 10:16:28 flop kernel: [<f8bff3fb>] ip_nat_rule_cleanup+0x3b/0x80 [iptable_nat]

May 11 10:16:28 flop kernel: [<f8bff89d>] fini_iptable_nat+0x1d/0x60 [iptable_nat]

May 11 10:16:28 flop kernel: [<f8b9158f>] do_ve_ip_tables+0x3cf/0x14a0 [vzmon]

May 11 10:16:28 flop kernel: [<f8b93ab6>] env_cleanup+0xb6/0x180 [vzmon]

May 11 10:16:28 flop kernel: [<f8b93bc3>] vzmond_helper+0x43/0x60 [vzmon]

May 11 10:16:28 flop kernel: [<f8b93b80>] vzmond_helper+0x0/0x60 [vzmon]

May 11 10:16:28 flop kernel: [<c0100f35>] kernel_thread_helper+0x5/0x10

May 11 10:16:28 flop kernel: Code: 04 89 d0 5a c3 c7 04 24 02 00 00 00 e8 5f b0 ff ff 89 c2 89 d0
5a c3 89 f6 8d bc 27 00 00 00 55 57 56 53 83 ec 1c 8b 4 4 24 30

<8b> 50 34 89 04 24 89 54 24 14 e8 c6 ad ff ff 89 44 24 0c 89 e0

May 11 10:16:28 flop kernel: EIP: [<f89521ab>] ipt_unregister_table+0xb/0x1b0 [ip_tables]
SS:ESP 0068:e7d99f44

/proc/user_beancounters

107: kmemsize	211264	2444162	10485760	15728640	0
lockedpages	0	0	32	32	0

privvmpages	1	98451	98816	131072	5
shmpages	1	977	8192	8192	0
dummy	0	0	0	0	0
numproc	0	69	120	120	0
physpages	1	9761	0	2147483647	0
vmguarpages	0	0	65536	131072	0
oomguarpages	1	9761	6144	2147483647	0
numtcpsock	0	47	512	512	0
numflock	0	8	100	110	0
numpty	0	5	16	16	0
numsiginfo	0	13	256	256	0
tcpsndbuf	0	35520	319488	524288	0
tcprcvbuf	0	454392	319488	524288	238
othersockbuf	0	133032	132096	336896	16123
dgramrcvbuf	0	8364	132096	132096	0
numothersock	0	110	512	512	0
dcachesize	0	0	1048576	1097728	0
numfile	0	1272	2048	2048	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	0	10	128	128	0

Do I have to stop vz itself? Rather not since other VEs still run smooth.

Subject: Re: No resolve inside one VE

Posted by [Vasily Tarasov](#) on Fri, 11 May 2007 08:56:03 GMT

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

Hello,

It is defenetely a kernel BUG! Please, fill it in bugzilla.openvz.org with detailed description and the link to this post. Thank you!

Quote:Do I have to stop vz itself? Rather not since other VEs still run smooth.

It's recomendado to reboot your node after such Oops.
