
Subject: SSH stop responding on HN

Posted by [stec](#) on Wed, 13 Dec 2006 23:21:55 GMT

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

This happend regularly after, especially after a long up-time. When I try to connect to ssh on the host node, it either reply "ssh: connect to host hostname port 22: Operation timed out" or "ssh_exchange_identification: Connection closed by remote host". After shutting down all vm (which are working properly and responding to ssh), the reply is more likely the second one.

Any clue of what could be wrong ?

Denis

PS: This a 2.6.16-026test018 self built kernel and I have no physical access to the host.

Subject: Re: SSH stop responding on host node

Posted by [rickb](#) on Thu, 14 Dec 2006 07:08:42 GMT

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

Hi, do you have ssh access to the HN now? If not, one could only theorize what the problem could be. sshd is not processing your authentication request, theres really no further details beyond that. This could be due to many thing, however without access to the server anything we tell you would be speculation and useless anyway. My recommendation would be to reboot the HN and check /var/log/messages and /var/log/secure on the HN around the times when you could not ssh to the HN..

Rick

Subject: Re: SSH stop responding on host node

Posted by [stec](#) on Thu, 14 Dec 2006 08:12:10 GMT

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

Oh, sorry , I have forgot to tell you that I have an another issue found in the logs, that some process of the host are complaining that fopen report "Too many open files in system". Last time there were nothing about ssh at all, but this time, I have checked again, and here is what I found:

- * A few days ago, sshd has reported "error: fork: Cannot allocate memory"
- * Before I shutdown all vm, ssh report no connection attempt
- * After, these shutdowns, ssh complains:
 sshd[3014]: error: fork: Cannot allocate memory
 Fatal resource shortage: kmemsize, UB 0.

I really do not catch from where these issues may came from, can you suggest me some investigation path ?

Subject: Re: SSH stop responding on host node
Posted by [rickb](#) on Thu, 14 Dec 2006 08:25:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

OK, I see now. When you said:
"When I try to connect to ssh on the host node"
You meant using ssh from the HN to the VE. Metrics numfiles, kmemsize, privvmpages exhaustion would all explain why any software malfunctions in a VE. To verify:

```
cat /proc/user_beancounters
```

and paste here. All you need to do is raise the limits for the failing metrics, understand how much resources your applications need and give them more than that. Have you read the UBC page in the wiki?

Rick

Subject: Re: SSH stop responding on host node
Posted by [stec](#) on Thu, 14 Dec 2006 08:47:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

No, you are wrong, I said that I am trying to connect from home to the HN, the one that is running the VEs (sorry I am still not use to these abbreviation, I am learning now). VEs seems to be working really well, as far as I know, and they fulfill my needs.

This is the HN that complains about file issues and I am pretty sure that this the sshd of the HN that complains. Hope this clarify the situation.

Thanks for helping.

Denis

Subject: Re: SSH stop responding on host node
Posted by [rickb](#) on Thu, 14 Dec 2006 08:59:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

HN has no resource limits. but anyway, send us user_beancounters so we can speak intelligently about the problem. Also check your logs like I mentioned at first. You need to give us specifics, as there is now possible way for us to identify the problem based on errors like "operation timed out"..

Rick

Subject: Re: SSH stop responding on host node
Posted by [stec](#) on Thu, 14 Dec 2006 10:00:13 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dear Rick,

Here is an excerpt of my logs. I have remove many "fopen: Too many open files in system". Note that I have tried to connect using ssh before and after shutting down the VEs.

```
Dec 13 21:24:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:24:25 ns22021 Neighbour table overflow.
Dec 13 21:25:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:25:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:25:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:26:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:26:21 ns22021 IPv6 addrconf: prefix with wrong length 56
Dec 13 21:27:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:28:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:29:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:29:36 ns22021 IPv6 addrconf: prefix with wrong length 56
Dec 13 21:30:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:30:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:30:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:30:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:31:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:32:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:32:29 ns22021 IPv6 addrconf: prefix with wrong length 56
Dec 13 21:33:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:34:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:35:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:35:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:35:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:35:28 ns22021 VPS: 18639: stopped
Dec 13 21:35:38 ns22021 IPv6 addrconf: prefix with wrong length 56
Dec 13 21:35:52 ns22021 VPS: 18638: stopped
Dec 13 21:36:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:37:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:38:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:38:01 ns22021 Uncharging too much 2220 h 1368, res othersockbuf ub 18637
Dec 13 21:38:01 ns22021 Uncharging too much 2220 h 0, res othersockbuf ub 18637
Dec 13 21:38:01 ns22021 Uncharging too much 2220 h 0, res othersockbuf ub 18637
Dec 13 21:38:01 ns22021 Uncharging too much 2220 h 0, res othersockbuf ub 18637
Dec 13 21:38:01 ns22021 Uncharging too much 2220 h 0, res othersockbuf ub 18637
Dec 13 21:38:07 ns22021 VPS: 18637: stopped
Dec 13 21:38:11 ns22021 IPv6 addrconf: prefix with wrong length 56Dec 13 21:38:26 ns22021
VPS: 18636: stopped
Dec 13 21:38:31 ns22021 Fatal resource shortage: kmemsize, UB 0.
```

```

Dec 13 21:38:31 ns22021 sshd[3014]: error: fork: Cannot allocate memory
Dec 13 21:38:35 ns22021 Fatal resource shortage: kmemsize, UB 0.
Dec 13 21:38:35 ns22021 sshd[3014]: error: fork: Cannot allocate memory
Dec 13 21:38:36 ns22021 Fatal resource shortage: kmemsize, UB 0.
Dec 13 21:38:36 ns22021 sshd[3014]: error: fork: Cannot allocate memory
Dec 13 21:38:37 ns22021 Fatal resource shortage: kmemsize, UB 0.
Dec 13 21:38:37 ns22021 sshd[3014]: error: fork: Cannot allocate memory
Dec 13 21:39:01 ns22021 cron[8714]: (CRON) error (can't fork)
Dec 13 21:39:16 ns22021 Ub 18636 holds 13320 in tcpsndbuf on put

```

Here is the beancounters as of now. Of course, I cannot produce those when the issue occurs since I cannot ssh anymore.

Version: 2.5

uid resource	held	maxheld	barrier	limit	failcnt	
0: kmemsize	109296870	109841586	2147483647	2147483647		0
lockedpages	0	0	2147483647	2147483647	0	
privvmpages	3262	14897	2147483647	2147483647		0
shmpages	679	695	2147483647	2147483647		0
dummy	0	0	2147483647	2147483647	0	
numproc	61	82	2147483647	2147483647		0
physpages	1560	6760	2147483647	2147483647		0
vmguarpages	0	0	2147483647	2147483647		0
oomguarpages	1560	6760	2147483647	2147483647		0
numtcpsock	3	5	2147483647	2147483647		0
numflock	1	2	2147483647	2147483647		0
numpty	1	1	2147483647	2147483647		0
numsiginfo	0	2	2147483647	2147483647		0
tcpsndbuf	31080	35520	2147483647	2147483647		0
tcprcvbuf	87564	87564	2147483647	2147483647		0
othersockbuf	11100	16224	2147483647	2147483647		0
dgramrcvbuf	0	8364	2147483647	2147483647		0
numothersock	23	27	2147483647	2147483647		0
dcachesize	0	0	2147483647	2147483647		0
numfile	525860	526009	2147483647	2147483647		0
dummy	0	0	2147483647	2147483647		0
dummy	0	0	2147483647	2147483647		0
dummy	0	0	2147483647	2147483647		0
numiptent	0	0	2147483647	2147483647		0
18636: kmemsize	2162620	2702410	18486804	20335484		0
lockedpages	0	0	902	902	0	
privvmpages	63043	69808	152556	167811		0
shmpages	14	350	15255	15255	0	
dummy	0	0	0	0	0	
numproc	46	58	800	800		0
physpages	39341	42286	0	2147483647		0
vmguarpages	0	0	152556	2147483647		0

oomguarpages	39341	42286	152556	2147483647	0
numtcpsock	7	19	800	800	0
numflock	6	7	720	792	0
numpty	0	0	80	80	0
numsiginfo	0	2	1024	1024	0
tcpsndbuf	62160	417360	2885468	6162268	0
tcprcvbuf	114688	147456	2885468	6162268	0
othersockbuf	133012	376720	1442734	4719534	0
dgramrcvbuf	0	4268	1442734	1442734	0
numothersock	95	100	800	800	0
dcachesize	0	0	4026407	4147200	0
numfile	1339	1736	7200	7200	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	0	0	200	200	0
18637: kmemsize	890020	1425116	18486804	20335484	0
lockedpages	0	0	902	902	0
privvmpages	12696	19885	152556	167811	0
shmpages	14	350	15255	15255	0
dummy	0	0	0	0	0
numproc	18	31	800	800	0
physpages	3081	6911	0	2147483647	0
vmguarpages	0	0	152556	2147483647	0
oomguarpages	3081	6911	152556	2147483647	0
numtcpsock	11	25	800	800	0
numflock	2	3	720	792	0
numpty	0	0	80	80	0
numsiginfo	0	2	1024	1024	0
tcpsndbuf	97680	426240	2885468	6162268	0
tcprcvbuf	180224	368984	2885468	6162268	0
othersockbuf	15540	19128	1442734	4719534	0
dgramrcvbuf	0	606056	1442734	1442734	0
numothersock	7	11	800	800	0
dcachesize	0	0	4026407	4147200	0
numfile	619	820	7200	7200	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	0	0	200	200	0
18638: kmemsize	454989	1046206	18486804	20335484	0
lockedpages	0	0	902	902	0
privvmpages	1253	8677	152556	167811	0
shmpages	0	336	15255	15255	0
dummy	0	0	0	0	0
numproc	8	20	800	800	0
physpages	837	1734	0	2147483647	0
vmguarpages	0	0	152556	2147483647	0

oomguarpages	837	1734	152556	2147483647	0
numtcpsock	4	6	800	800	0
numflock	2	3	720	792	0
numpty	0	0	80	80	0
numsiginfo	0	2	1024	1024	0
tcpsndbuf	35520	35520	2885468	6162268	0
tcprcvbuf	65536	65536	2885468	6162268	0
othersockbuf	86580	90168	1442734	4719534	0
dgramrcvbuf	0	8536	1442734	1442734	0
numothersock	41	45	800	800	0
dcachesize	0	0	4026407	4147200	0
numfile	255	473	7200	7200	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	0	0	200	200	0
18639: kmemsize	260964	857866	18486804	20335484	0
lockedpages	0	0	902	902	0
privvmpages	709	7954	152556	167811	0
shmpages	0	336	15255	15255	0
dummy	0	0	0	0	0
numproc	5	18	800	800	0
physpages	414	1234	0	2147483647	0
vmguarpages	0	0	152556	2147483647	0
oomguarpages	414	1234	152556	2147483647	0
numtcpsock	2	4	800	800	0
numflock	2	3	720	792	0
numpty	0	0	80	80	0
numsiginfo	0	2	1024	1024	0
tcpsndbuf	17760	17760	2885468	6162268	0
tcprcvbuf	32768	32768	2885468	6162268	0
othersockbuf	2220	5808	1442734	4719534	0
dgramrcvbuf	0	4268	1442734	1442734	0
numothersock	1	5	800	800	0
dcachesize	0	0	4026407	4147200	0
numfile	122	382	7200	7200	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	0	0	200	200	0

Hope this helps diagnosing better.

Denis

Subject: Re: SSH stop responding on host node
Posted by [rickb](#) on Thu, 14 Dec 2006 10:44:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

500,000 open files..

this is probably related to a ulimit limit and not due to an openvz resource limitation. I suspect you are doing something zany in your HN node to have 500k open files. Move that application into a VE so that your HN cannot crash.

Rick

Subject: Re: SSH stop responding on host node
Posted by [stec](#) on Thu, 14 Dec 2006 11:13:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

Nice, but how do I determine which application is allocating so much files ? I do not run anything on the HN, except some monitoring stuff of my hosting company.

Thanks again for your help.

Denis

Subject: Re: SSH stop responding on host node
Posted by [stec](#) on Fri, 15 Dec 2006 00:26:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

I do not understand why numfile report this increasing number. Using lsof | wc -l do not report a large number. What should I conclude ?

Subject: Re: SSH stop responding on HN
Posted by [xemul](#) on Fri, 15 Dec 2006 09:23:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

We had some BUGS in numfile and kmemsize accounting in 026test018 kernel due to ported optimizations from stable OpenVZ branch. They were fixed in 026test020 kernel and in 028 kernels.

Obviously you're experiencing them, so update will definitely help.

But note, that if you want to stay on 026 kernels then you'll have to download the sources and add top patch from GIT - it wasn't included into any 026 release as we dropped this branch for a while and started with 028.

Thank you.

Subject: Re: SSH stop responding on HN
Posted by [stec](#) on Fri, 15 Dec 2006 10:33:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks for this information, this is really good news.
I am currently in the process of building a new server that will replace the old one, and I expect to use the latest build on 2.6.18 kernel. I am really happy to know that my effort will solve that issue.

Anyway, great job guys ! Many thanks !

Subject: Re: SSH stop responding on HN
Posted by [John Kelly](#) on Fri, 15 Dec 2006 14:40:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

xemul wrote on Fri, 15 December 2006 04:23 But note, that if you want to stay on 026 kernels then you'll have to download the sources and add top patch from GIT

OK. But there are many patches in GIT. Which one is it?

Subject: Re: SSH stop responding on HN
Posted by [dim](#) on Tue, 19 Dec 2006 13:38:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

This one:
<http://git.openvz.org/?p=linux-2.6.16-openvz;a=commit;h=46d1e25bce7440b23652caf7f463670a5360890a>.

Subject: Re: SSH stop responding on HN
Posted by [John Kelly](#) on Tue, 19 Dec 2006 14:56:34 GMT
[View Forum Message](#) <> [Reply to Message](#)

OK. I was looking in the 2.6.18 tree for a "top" patch. But now I see that "top" meant, not the command "top", but the latest patch in the 2.6.16 tree, which was not included in the final tarball.

BTW, your raw URL does not work, you have to edit it and extract the portion between the double quotes.

Subject: Re: SSH stop responding on HN
Posted by [dim](#) on Tue, 19 Dec 2006 15:05:08 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hmm, it works for me. Ok, I'll replace link.
