
Subject: Weird performance problemPosted by [ugob](#) on Wed, 22 Apr 2009 12:27:37 GMT[View Forum Message](#) <> [Reply to Message](#)

Hi,

I'm running a CentOS 4. server and I sometimes face a weird problem. It is a weird performance problem, and here is how I discovered it.

This server runs OpenVZ virtual machines, and one of them is an asterisk server for my personal use. The first symptom of the problem is that the voice quality became flaky. So I logged on the server to see what could be eating cpu cycles, when I ran top, it took almost one minute before top actually showed. Another hint is that when I run dstat (a monitoring utility that is a mix of iostat and vmstat and other stats), I often get a "missed xx ticks", where xx is a number.

Example (current) (sorry for the wrap):

```
----total-cpu-usage---- -dsk/total- -net/total- ---paging-- ---system--
usr sys idl wai hiq siq| read writ| recv send| in  out| int  csw
 3  2 93  2  0  0| 106k 273k|  0   0| 0.2  0.4| 1039 389
 3  6 91  0  0  0|  0 6416k| 276k 275k|  0   0| 2160 6822
missed 55 ticks
 4 10 84  2  0  0| 1200k 1992k| 82k  93k|  0   0| 1188 6275
missed 29 ticks
 1  0 99  0  0  0|  0 1312k| 65k  66k|  0   0| 1050 1114
missed 38 ticks
 2  1 96  0  0  0|  0 1168k| 57k  59k|  0   0|  491  877
missed 13 ticks
 3  1 94  1  0  0|  0 6016k| 181k 176k|  0   0| 2169 5996
missed 50 ticks
 4  2 91  1  0  0| 28k 8744k| 216k 214k|  0   0| 2159 5438
missed 37 ticks
 1  1 98  0  0  0|  0 2632k| 93k  91k|  0   0|  983 1381
missed 34 ticks
 1  1 98  1  0  0|  0 5624k| 113k 110k|  0   0| 1569 2643
missed 52 ticks
 1  1 98  1  0  0|  0 2432k| 29k  28k|  0   0|  679  647
missed 12 ticks
 0  0 100  0  0  0|  0   0| 60B 374B|  0   0|  13  15
 2  3 94  0  0  0|  0 1872k| 209k 210k|  0   0| 1375 3590
missed 30 ticks
```

The problem is currently occurring, but it doesn't seem to be affecting voice quality for now, so I have some time to try to find the cause.

The only solution I've found up to now is to reboot... But hey, this isn't a Windows 98 machine !

I tried restarting the VZ system, which restarts all the VMs, but it didn't solve the problem. I can't tell if the problem occurs on a stock centos kernel, because the server is running production (but non-critical) virtual machines, so it is always running the openVZ kernel.

So here is what I've done for now:

- Top shows a load of about 0.4

- vmstat 1 10 shows this:

```
procs -----memory----- ---swap-- -----io----- --system--
----cpu----
 r b swpd free buff cache si so bi bo in cs us
sy id wa
 0 0 592 191092 381720 537956 0 0 53 68 4 3 3
2 93 2
 0 0 592 190720 381720 537956 0 0 0 0 32 60 1
1 98 0
 0 0 592 191092 381720 537956 0 0 0 0 41 59 0
0 100 0
 1 0 592 191092 381728 537948 0 0 0 2584 311 96 10
4 66 19
 0 0 592 189968 381732 537944 0 0 0 2080 222 174 2
3 79 16
 0 1 592 189968 381732 537944 0 0 0 3244 170 73 10
4 73 12
 0 0 592 190216 381732 537944 0 0 0 136 76 113 1
2 93 4
 0 0 592 189844 381732 537944 0 0 0 0 33 69 1
1 98 0
 0 0 592 189844 381732 537944 0 0 0 0 24 32 0
0 100 0
 0 0 592 190340 381732 537944 0 0 0 0 28 42 0
0 100 0
```

iostat -x 1 (excerpt)

```
Device: rrqm/s wrqm/s r/s w/s rsec/s wsec/s rkB/s wkB/s
avgrq-sz avgqu-sz await svctm %util
sda      0.00 171.00 0.00 124.00 0.00 2368.00 0.00 1184.00
 19.10 0.14 1.13 0.02 0.20
sdb      0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00
 0.00 0.00 0.00 0.00 0.00
sdc      0.00 171.00 0.00 124.00 0.00 2368.00 0.00 1184.00
```

```

19.10  0.17  1.35  0.02  0.30
sdd    0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
0.00  0.00  0.00  0.00  0.00
md0    0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
0.00  0.00  0.00  0.00  0.00
md2    0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
0.00  0.00  0.00  0.00  0.00
md1    0.00  0.00  0.00  294.00  0.00  2352.00  0.00  1176.00
8.00  0.00  0.00  0.00  0.00
dm-0   0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
0.00  0.00  0.00  0.00  0.00
dm-1   0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
0.00  0.00  0.00  0.00  0.00
dm-2   0.00  0.00  0.00  294.00  0.00  2352.00  0.00  1176.00
8.00  0.30  1.01  0.02  0.50
dm-3   0.00  0.00  0.00  294.00  0.00  2352.00  0.00  1176.00
8.00  0.30  1.01  0.02  0.50
dm-4   0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
0.00  0.00  0.00  0.00  0.00
dm-5   0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
0.00  0.00  0.00  0.00  0.00
dm-6   0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
0.00  0.00  0.00  0.00  0.00

```

Another hint is that pings are really slow. Even pinging localhost is very long. The first reply is fast, but the second takes ages to come.

It seems to be blocking here:

```

recvmsg(3, 0xbfbf84b0, MSG_DONTWAIT) = -1 EAGAIN (Resource
temporarily unavailable)
gettimeofday({1239887784, 389347}, NULL) = 0
poll(

```

The rest comes as soon as there is another response:

```

[{{fd=3, events=POLLIN|POLLERR}}, 1, 999) = 0
gettimeofday({1239887903, 119727}, NULL) = 0
gettimeofday({1239887903, 119791}, NULL) = 0
sendmsg(3, {msg_name(16)={sa_family=AF_INET, sin_port=htons(0),
sin_addr=inet_addr("127.0.0.1")},
msg_iov(1)=[{" \10\0\335\2018)\0\4\0370\347i\357\323\1\0\10\t\n\vfr\16\17
\20\21\22\23\24\25\26\27 "...,
64}], msg_controllen=0, msg_flags=0}, MSG_CONFIRM) = 64
recvmsg(3, {msg_name(16)={sa_family=AF_INET, sin_port=htons(0),
sin_addr=inet_addr("127.0.0.1")},
msg_iov(1)=[{" E\0\0T\26\264\0\0@\1e\363\177\0\0\1\177\0\0\1\0\0\345\2018)\ 0\4\0370\347i "...,
192}], msg_controllen=20, {cmsg_len=20, cmsg_level=SOL_SOCKET,

```

```
msg_type=0x1d /* SCM_??? */, ...}, msg_flags=0}, 0) = 84
write(1, "64 bytes from hn01.domain"... , 82) = 82
recvmsg(3, 0xbfbf84b0, MSG_DONTWAIT) = -1 EAGAIN (Resource
temporarily unavailable)
gettimeofday({1239887903, 120785}, NULL) = 0
poll(
```

Then it blocks again...

This confuses Nagios that is running in a VM on this server.

Can the 'gettimeofday' be the problem? 'date' runs w/o delay.

Any insight would be greatly appreciated. It is not critical, but I'd be glad to be able to finally pinpoint and solve the problem.

Hardware: HP Netserver, software raid, SCSI disks, 1.7 GB RAM.

I can provide more information if needed.

Thanks,

Ugo

Subject: Re: Weird performance problem
Posted by [maratrus](#) on Tue, 28 Apr 2009 11:34:47 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

what kernel do you use?

So, the provided vmstat output show that CPU is idle.

The only thing I can assume for now is that somebody disables interrupts and thus "missed * ticks" messages appeared and network packets handling is delayed.

Do you have any assumptions?

BTW, does "service vz stop" resolve the problem or the problem takes place when vz service is stopped?

It may be useful to gather alt-sysrq-* outputs:

Alt-SysRQ-

- 1) "m" - for memory info dump
- 2) "p" - for registers - several times, please, twice the number of CPUs
- 3) "a" - for scheduler stat - 3 times
- 4) "w" - another scheduler info - 3 times
- 5) "t" - for all processes calltraces. Warning - this is a resource consuming operation. At least -

twice.

Subject: Re: Weird performance problem
Posted by [ugob](#) on Tue, 28 Apr 2009 12:53:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

I was using kernel 2.6.18-92.1.18.el5.028stab060.2, now I'm using 2.6.18-92.1.18.el5.028stab060.8. I was previously using PAE kernel, now I'm trying the non-PAE kernel. I only have 1.7 GB RAM anyway. The problem hasn't occurred yet since the last reboot.

Thanks,

Ugo

Subject: Re: Weird performance problem
Posted by [ugob](#) on Tue, 28 Apr 2009 12:55:56 GMT
[View Forum Message](#) <> [Reply to Message](#)

service vz restart didn't solve the problem, the next time it occurs, I'll try a service vz stop.

I don't have a KVM on this server, so it may be hard for me to get the results of alt-sysreq commands.

Ugo
