# Subject: \*SOLVED\* Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by tchipman on Wed, 13 Sep 2006 19:33:52 GMT

View Forum Message <> Reply to Message

Hi All,

I'm not positive this is an openVz issue, but wanted to post here in case others observe the same thing (?)

I've got a CentOS 4.4 machine with OpenVZ installed and a few OpenVZ virtual hosts running on it. The hardware / software is very solid, no problems at all there.

One ODD thing that has been happening: The clock on the system is drifting at an absurd rate -- upwards of 5 seconds every 5 minutes -- and when I try to use ntp to keep it in synch with a local time source, it refuses to do so.

Right now I've got a crontab hitting "ntpdate" command every 5 minutes, which is working "OK" but clearly it is a band-aid solution.

I know from other reading that there are issues with VMWare hosting boxes which result in time-skew issues.. and it seemed not-inconcievable that this might also be an issue for OpenVZ.

However, it could be purely a problem with CentOS (I've found some threads in that vein already, but nothing conclusive / to my satisfaction yet).

So .. if anyone has any comments .. they certainly would be appreciated..

Sorry if this is a bit off topic.. clearly if/when I find an answer, I'll post it here just for reference.

--Tim Chipman

Subject: Re: Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by Vasily Tarasov on Thu, 14 Sep 2006 07:27:18 GMT View Forum Message <> Reply to Message

Can you check, please, if this behaviour persist if boot kernel with "acpi=off"?

Thanks.

Subject: Re: Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by tchipman on Thu, 14 Sep 2006 21:24:18 GMT

View Forum Message <> Reply to Message

Many thanks for the suggestion. I won't be able to reboot the machine until Monday morning, but

will do so at that time and then post a followup to the thread here after I've observed behaviour through the day.

--Tim

Subject: Re: Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by John Kelly on Fri, 15 Sep 2006 02:58:02 GMT View Forum Message <> Reply to Message

Chrony works fine on my OpenVZ HN. Haven't tried ntpd yet. Don't really need it, since chrony does the job.

My HN is debian; maybe there is no chrony for CentOS.

Subject: Re: Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by tchipman on Mon, 25 Sep 2006 19:39:12 GMT

View Forum Message <> Reply to Message

Hi all,

so I rebooted the machine today (finally) with acpi turned off, as suggested above. At first, it seemed to help. Alas, after a few hours it seems things are no better than before. I can reach my local NTP servers, but the offset and jitter are crazy, the clock is drifting again (was off more than 140seconds after a few hours)

So. I'll try to build and run chrony as my next test, I guess. Somewhat frustrating though.

I'll keep posting as things progress.

--Tim

Sample problem output below:

---paste---

[root@hydra etc]# ntpq -p

remote refid st t when poll reach delay offset jitter

\_\_\_\_\_

\_\_\_\_\_

NTP\_SERVER1 128.233.150.93 2 u 35 64 3 1.118 -290409 8552.76 NTP\_SERVER2 132.246.168.9 2 u 35 64 3 0.931 -290409 8553.04 LOCAL(0) LOCAL(0) 10 I 33 64 3 0.000 0.000 0.001 Subject: Re: Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by HaroldB on Mon, 25 Sep 2006 22:12:17 GMT

View Forum Message <> Reply to Message

Does the clock drift under the CentOS kernel? It seems odd that the openvz patch could cause this...

Subject: Re: Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by tchipman on Tue, 26 Sep 2006 12:15:29 GMT View Forum Message <> Reply to Message

BTW, I compiled Chrony from source and tried that last night. No Joy. It complained of various things, including indications that NTP time sources drifted too much; and then core-dump exits by chronyd, with such things logged to messages as shown below:

---paste---

Sep 25 17:11:52 hydra chronyd[31170]: Unexpected condition [adjtimex failed for set\_frequency, freq\_ppm=1.1522e+05 scaled\_freq=-1.8268e+01 required\_tick=8848] at sys\_linux.c:453, core dumped

Sep 25 17:13:15 hydra chronyd[31439]: chronyd version 1.21 starting

Sep 25 17:13:15 hydra chronyd[31439]: Initial txc.tick=10000 txc.freq=0 (0.00000000) txc.offset=0 => hz=100 shift\_hz=7

Sep 25 17:13:15 hydra chronyd[31439]: set\_config\_hz=0 hz=100 shift\_hz=7

basic\_freq\_scale=1.28000000 nominal\_tick=10000 slew\_delta\_tick=833 max\_tick\_bias=1000

Sep 25 17:13:15 hydra chronyd[31439]: Linux kernel major=2 minor=6 patch=9

Sep 25 17:13:15 hydra chronyd[31439]: calculated\_freq\_scale=0.99902439

freq scale=0.99902439

Sep 25 17:16:29 hydra chronyd[31439]: Selected source 129.1X3.5.100

Sep 25 17:16:31 hydra chronyd[31439]: Can't synchronise: no majority

Sep 25 17:18:41 hydra chronyd[31439]: Selected source 129.1X3.1.100

Sep 25 17:18:41 hydra chronyd[31439]: Unexpected condition [adjtimex failed for set\_frequency, freq\_ppm=1.0022e+05 scaled\_freq=-1.5904e+01 required\_tick=8998] at sys\_linux.c:453, core dumped

---endpaste---

Possibly the nature of this dump/crash by chrony is a hint, but alas I'm not sure what it is trying to tell me.

# Subject: Re: Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by tchipman on Tue, 26 Sep 2006 12:16:46 GMT

View Forum Message <> Reply to Message

Hi.

You are right -- I'll probably have to try this for a little while. I'm not keen because it means killing my vzhosts for an hour or so, and I've got semi-production services inside some of them now. Suppose I could migrate the vzhosts temporarily for the test period to my backup-test-vzhostbox. Whoo hoo. Fun never ends.

I'll post the results once I've finished the test.

Т

Subject: Re: Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by tchipman on Tue, 26 Sep 2006 13:20:13 GMT

View Forum Message <> Reply to Message

Hi All,

so, some results with a semi-happy-ending. Can't say I understand the mess fully, but at least we have a stopgap that is less hideous.

**Basic Summary:** 

- (0) Starting point: was running with unreliable NTP, Booted on "RHEL-Unsupported" OpenVZ pre-complied Kernel.
- (1) rebooted the machine using a stock nonVZ kernel for CentOS. Let ntp try to keep the clock synched. After 5 minutes it established a reliable synch that seemed to hold.
- (2) Rebooted the machine to Stock Production VZ Kernel (not the RHEL "unsupported" VZ kernel). Started NTP. It refused to synch, held only to local pseudo-clock and then began to drift again.
- (3) Applied some tweaks to the ntp config, and was able to force NTP to synch.

If anyone has thoughts about why this behaviour is here, I'm all ears. It does appear to suggest an issue with this OpenVZ kernel on my hardware being part of the problem, since a stock CentOS kernel wasn't giving any problems.

---Tim

Exact notes follow below:

### (0) STARTING POINT KERNEL:

Grub entry describes it thus:

title CentOS (2.6.9-023stab016.2-smp) OpenVZ RHEL Kernel Aug-26-06 root (hd0,1)

kernel /vmlinuz-2.6.9-023stab016.2-smp ro root=/dev/md3 console=ttyS0,38400 acpi=off initrd /initrd-2.6.9-023stab016.2-smp.img

This is pulled via yum-installation using OpenVZ RPM Repo source. With this kernel booted, the NTP drifts terribly (1 second per minute approx)

(1) Reboot to stock kernel, described by grub entry:

title CentOS (2.6.9-42.0.2.ELsmp) CentOS NO OpenVZ Kernel Aug-24-06 root (hd0,1) kernel /vmlinuz-2.6.9-42.0.2.ELsmp ro root=/dev/md3 console=ttyS0,38400 acpi=off initrd /initrd-2.6.9-42.0.2.ELsmp.img

With this kernel booted (vz hosts down, vz not running) the NTP synch works great.

(2) Reboot using Official VZ Kernel:

title CentOS (2.6.8-022stab078.14-smp) OpenVZ Official Aug-24-06 root (hd0,1)

kernel /vmlinuz-2.6.8-022stab078.14-smp ro root=/dev/md3 console=ttyS0,38400 acpi=off initrd /initrd-2.6.8-022stab078.14-smp.img

With this kernel booted, VZ hosts running, the NTP drift is back.

THFN:

(3) as per hints pulled from the URL:

http://www.djack.com.pl/modules.php?name=FAQ&myfaq=yes&xmyfaq=yes&id\_cat=3&id=132

Specifically, I added to my ntp config:

server 13X.6X.1.1X3 burst iburst

the burst and iburst settings

Bingo, NTP was synched in <5 minutes.

So, For now I'm running the official stock OpenVZ kernel (I don't need the extra features of the RHEL version, namely MAC Addresses for VZHosts - for now -- which is why I had tested that kernel in the recent past) -- and I've got NTP working. For now.

MAYBE the problem somehow relates to the fact, I'm running a 32-bit kernel on a 64-bit platform (AMD 64-bit Athlon 64x2 CPU). I've seen brief reference via my google searches to some context where NTP gets pissed off when the kernel is built on platform somewhat dissimilar from your actual. However, this could be just smoke-and-mirrors thoughts on my part. Could be something totally unrelated. however, the fact that a vanilla CentOS kernel was able to run with NTP happily on this machine without any kludges ... is somewhat pointing-the-finger towards the OpenVZ kernel, one way or another ?!

So. That's the end of the story for now.

Subject: Re: \*SOLVED\* Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by tchipman on Thu, 05 Oct 2006 13:06:32 GMT

View Forum Message <> Reply to Message

Hi All,

Thought I should post another followup here, since the issue has evolved a fair bit since my last post. Since making the tweak with "burst / iburst" to my ntp config, it appeared to do the trick .. time wasn't ever so far out.

However, I was still seeing not-great values reported via "ntpq -p", ie, offset values in the thousands. Observing my /var/log/messages file, it was also clear that the ntp daemon was struggling, with frequent shifts of the NTP time source being used, and occasional "time warp" clock resets (ie, 120second adjustments being made periodically), approx as follows:

---PASTE---

from /var/log/messages ...

Oct 1 04:20:42 hydra ntpd[6350]: synchronized to 129.1X3.1.100, stratum 2

Oct 1 04:20:47 hydra ntpd[6350]: no servers reachable

Oct 1 04:23:20 hydra ntpd[6350]: synchronized to 129.1X3.5.100, stratum 2

Oct 1 04:24:18 hydra ntpd[6350]: synchronized to 129.1X3.105.63, stratum 3

Oct 1 04:24:21 hydra ntpd[6350]: no servers reachable

Oct 1 04:24:35 hydra ntpd[6350]: synchronized to 129.1X3.105.63, stratum 3

Oct 1 04:25:31 hydra ntpd[6350]: no servers reachable

Oct 1 04:27:03 hydra ntpd[6350]: synchronized to 129.1X3.1.100, stratum 2

Oct 1 04:25:16 hydra ntpd[6350]: time reset -106.823489 s

Oct 1 04:26:35 hydra ntpd[6350]: synchronized to 129.1X3.1.100, stratum 2

```
Oct 1 04:26:39 hydra ntpd[6350]: no servers reachable
```

- Oct 1 04:27:54 hydra ntpd[6350]: synchronized to 129.1X3.1.100, stratum 2
- Oct 1 04:28:51 hydra ntpd[6350]: synchronized to 129.1X3.105.63, stratum 3
- Oct 1 04:28:53 hydra ntpd[6350]: no servers reachable
- Oct 1 04:29:05 hydra ntpd[6350]: synchronized to 129.1X3.1.100, stratum 2

#### ---ENDPASTE----

Yesterday I was puttering around with this a bit more. Some reading I came across suggested that a boot-flag option, "clock=pmtmr" might be of some use. I rebooted with this, but it was clear from Dmesg that it was being ignored, telling me:

Warning: clock= override failed. Defaulting to PIT

So, I tweaked my boot parameters again and removed my boot parameter, "acpi=off", since I hadn't tested behaviour thus with my "new ntp config" (using burst / iburst).

Since the change, NTP behaviour is \*exactly\* what I would have wanted - the time source is no longer having trouble, offset values are reasonable, and the clock is solid. ie:

\*sxt-np-1.XXXX.D 128.XXX.150.93 2 u 111 1024 377 6.370 -0.550 1.929 +KIL-NP-1.XXXX.D XXX.246.168.9 2 u 133 1024 377 3.472 0.198 5.219 Router.Phys.OCE XXX.173.5.100 3 u 100 1024 357 704.997 350.089 111.705

and additionally in /var/log/messages,

```
Oct 4 17:44:44 hydra ntpd[10375]: ntpd 4.2.0a@1.1190-r Sun Aug 13 01:49:12 CDT 2006 (1)
```

Oct 4 17:44:44 hydra ntpd: ntpd startup succeeded

Oct 4 17:44:44 hydra ntpd[10375]: precision = 4.000 usec

Oct 4 17:44:44 hydra ntpd[10375]: Listening on interface wildcard, 0.0.0.0#123

Oct 4 17:44:44 hydra ntpd[10375]: Listening on interface lo, 127.0.0.1#123

Oct 4 17:44:44 hydra ntpd[10375]: Listening on interface eth0, XXX.173.23.235#123

Oct 4 17:44:44 hydra ntpd[10375]: Listening on interface eth2, XXX.168.111.235#123

Oct 4 17:44:44 hydra ntpd[10375]: kernel time sync status 0040

Oct 4 17:44:44 hydra ntpd[10375]: frequency initialized -86.648 PPM from /var/lib/ntp/drift

Oct 4 17:44:51 hydra ntpd[10375]: synchronized to XXX.173.5.100, stratum 2

Oct 4 17:44:59 hydra ntpd[10375]: kernel time sync disabled 0041

Oct 4 17:46:04 hydra ntpd[10375]: kernel time sync enabled 0001

Oct 4 22:16:37 hydra sshd(pam\_unix)[4782]: session opened for user chipmant by (uid=0)

ie, note that once sync was enabled at 17:46:04, it hasn't been lost / nor has it drifted / the clock is solid now.

So - things are infinitely better now as they stand as compared to the original state. I'm not sure if this mess of a history will help anyone else, ever, but I thought I should post the end-of-story here... just in case.

--Tim Chipman

Subject: Re: \*SOLVED\* Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by dahas on Mon, 07 Apr 2008 20:43:08 GMT

View Forum Message <> Reply to Message

hi

I use CentOS 5.1 x64

and openvz kernel 2.6.24-ovz004.1 and got same problem with clock.

I dont see the solution or isnt very clear for me what have to do to get lost of that drift.

Since now i use ntpdate to ajust my clock.

non openvz kernel works ok no drift at all.

Thanks in advance

Subject: Re: \*SOLVED\* Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by adobriyan on Tue, 08 Apr 2008 13:39:58 GMT View Forum Message <> Reply to Message

- > I use CentOS 5.1 x64 and openvz kernel 2.6.24-ovz004.1 and
- > got same problem with clock.
- > I dont see the solution or isnt very clear for me what have
- > to do to get lost of that drift.
- > Since now i use ntpdate to ajust my clock.
- > non openvz kernel works ok no drift at all.

Interesting...

First, is vanilla 2.6.24 kernel OK? Or 2.6.24.4, even.

You can extract .config from ovz kernel in /proc/config.gz and recompile vanilla kernel with it (ignore all missing options messages), then collect "dmesg" output from working and openvz kernel. Differences between them should sched light on this issue.

If you can't or won't recompile kernel, post dmesg and .config from any kernel that doesn't drift.

Attach both dmesgs here, or in bugzilla.

P.S.: acpi=off is crude hammer, something in openvz should be fixed.

Subject: Re: \*SOLVED\* Clock Drift / NTPd - CentOS OpenVZ Host - Problems Posted by dahas on Tue, 08 Apr 2008 17:15:59 GMT

View Forum Message <> Reply to Message

HI

Nice to see some 1 still can give me a hand

Today I come from work and see a strange things

only 7.5334 sec drifted in about 10h which is ok in my opinion or called normal.

I run without acpi=off the kernel

I will atach the config.gz and dmesg.

My plan is to do 2 things 1 to do as you say to compile a vanila kernel 2.6.24 and second to try to run with acpi-off.

Server do nothing now have only 2 VE inside and from now till 2AM I will try to port a old server inside VE.

I dont know what will be ther order but at least I will try to do all of this.

The problem what I sow when sistem boot and ignored is hwclock who say to run --debug because he dont find any way to access clock.

P.S. In attach is the drifting config and dmesg openvz kernel. I will try to recompile 2.6.24.4 with the config of openvz.

### File Attachments

- 1) config.gz, downloaded 920 times
- 2) dmesg, downloaded 923 times