

---

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load  
Posted by [svensirk](#) on Tue, 22 May 2012 09:50:40 GMT

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

---

2012/5/22 Rene C. <openvz@dokbua.com>:

> Hi Sirk,  
>

Hi Rene,

> Thanks for your reply. I'm so pleased having found this mailing list after  
> having tried the forum, which seem to have very little activity!  
>

True, but this list has helped me a lot as well :-)

> Ploop is a great idea technically, but I'm a little concerned about the "  
> Warning: This is a new feature, not yet ready for production systems. Use  
> with caution." on the OpenVZ Wiki page, so I'm kinda waiting for the  
> green-light that it's ready for production environments.  
>

If you want some practical information on ploop: We are using it in a highly productive environment.

It was either, try ploop and hope it works, or have the systems fail every 2nd day.

So we decided to use ploop and are more than happy.

It even solves a lot of issues we had with the private areas directly on the nfs share.

But of course, thats totally up to you.

I started with only a few "unimportant" CTs and then merged everything after a while (42 CTs).

> It did occur to me that disk-IO could be the cause of the problem, but  
> iostat on the hardware node did not suggest any particular IO problems. I  
> still haven't found a way to see the IO activity within a container - iostat  
> just comes up blank when it's run within a container. Is there a way?  
>

To be honest, I don't know.

iostat ist not working because you do not really have a device.

This ist handled the way with ploop sadly but could be modified I guess.

For ploop you have the ploop-stat command but that dosen't work as expected for me :-)

> We're not using any network storage with this server so that is not the  
> reason.  
>

> The server has 4 SATA-3 drives, with the root partition being on one drive,  
> the problem container alone on a second drive, and the remaining containers  
> on a third.

So you have a different FileSystem for the "problem"-Container that is  
even on a different disk ?

If that is the case, this CT should not affect the others at all in terms of IO.

best regards,

Sirk

>  
> Best,  
> Rene  
>  
> On Tue, May 22, 2012 at 3:06 PM, Sirk Johannsen <s.johannsen@satzmedia.de>  
> wrote:  
>>  
>> Hi Rene,  
>>  
>> Since CPU and MEM are fine it's most likely to be Disk-IO.  
>> I have similar Problems with a Cluster Setup based on OpenVZ.  
>> The problem is that our Storage is way to slow.  
>> We have been accessing the storage via NFS and put all our CTs private  
>> areas on it.  
>> I noticed many times that one CT was doing a lot of disk IO and all  
>> other were suffering from that... that even lead to total system  
>> failures.  
>> This has been solved by converting everything to ploop. Since then our  
>> system is at least in a stable state.  
>> IO Performance is still an issue but does not bring our system down.  
>>  
>> You should give ploop a try :-) I am very happy with it.  
>>  
>> best regards,  
>>  
>> Sirk  
>>  
>> 2012/5/21 Rene Dokbua <openvz@dokbua.com>:  
>> > Hello,  
>> >  
>> > I occasionally get this extreme load on one of our VPS servers. It is  
>> > quite  
>> > large, 4 full E31230 cores, 4 GB RAM and hosting ca. 400 websites +  
>> > parked/addon/subdomains.  
>> >

```

>> > The hardware node has 12 active VPS servers and most of the time things
>> > are
>> > chugging along just fine, something like this.
>> >
>> > 1401: 0.00 0.00 0.00 1/23 4561
>> > 1402: 0.02 0.05 0.05 1/57 16991
>> > 1404: 0.01 0.02 0.00 1/73 18863
>> > 1406: 0.07 0.13 0.06 1/39 31189
>> > 1407: 0.86 1.03 1.14 1/113 31460
>> > 1408: 0.17 0.17 0.18 1/79 32579
>> > 1409: 0.00 0.00 0.02 1/77 21784
>> > 1410: 0.01 0.02 0.00 1/60 7454
>> > 1413: 0.00 0.00 0.00 1/46 18579
>> > 1414: 0.00 0.00 0.00 1/41 23812
>> > 1415: 0.00 0.00 0.00 1/45 9831
>> > 1416: 0.05 0.02 0.00 1/59 11332
>> > 12 active
>> >
>> > The problem VPS is 1407. As you can see below it only uses a bit of the
>> > cpu
>> > and memory.
>> >
>> > top - 17:34:12 up 32 days, 12:21, 0 users, load average: 0.78, 0.95,
>> > 1.09
>> > Tasks: 102 total, 4 running, 90 sleeping, 0 stopped, 8 zombie
>> > Cpu(s): 16.3%us, 2.9%sy, 0.4%ni, 78.5%id, 1.8%wa, 0.0%hi, 0.0%si,
>> > 0.1%st
>> > Mem: 4194304k total, 2550572k used, 1643732k free, 0k buffers
>> > Swap: 8388608k total, 105344k used, 8283264k free, 1793828k cached
>> >
>> > Also iostat and vmstat shows no particular io or swap activity.
>> >
>> > Now for the problem. Every once in a while the loadavg of this
>> > particular
>> > VPS shoots up to like crazy values, 30 or more and it becomes completely
>> > sluggish. The odd thing is load goes up for the VPS server, and starts
>> > spilling into other VPS servers on the same hardware node - but there are
>> > still no particular cpu/memory/io usage going on that I can see. No
>> > particular network activity. In this example load has fallen back to
>> > around 10 but it was much higher earlier.
>> >
>> > 16:19:44 up 32 days, 11:19, 3 users, load average: 12.87, 19.11,
>> > 18.87
>> >
>> > 1401: 0.01 0.03 0.00 1/23 2876
>> > 1402: 0.00 0.11 0.13 1/57 15334
>> > 1404: 0.02 0.20 0.16 1/77 14918
>> > 1406: 0.01 0.13 0.10 1/39 29595

```

```

>> > 1407: 10.95 15.71 15.05 1/128 13950
>> > 1408: 0.36 0.52 0.57 1/81 27167
>> > 1409: 0.09 0.26 0.43 1/78 17851
>> > 1410: 0.09 0.17 0.18 1/61 4344
>> > 1413: 0.00 0.03 0.00 1/46 16539
>> > 1414: 0.01 0.01 0.00 1/41 22372
>> > 1415: 0.00 0.01 0.00 1/45 8404
>> > 1416: 0.05 0.10 0.11 1/58 9292
>> > 12 active
>> >
>> > top - 16:20:02 up 32 days, 11:07, 0 users, load average: 9.14, 14.97,
>> > 14.82
>> > Tasks: 135 total, 1 running, 122 sleeping, 0 stopped, 12 zombie
>> > Cpu(s): 16.3%us, 2.9%sy, 0.4%ni, 78.5%id, 1.8%wa, 0.0%hi, 0.0%si,
>> > 0.1%st
>> > Mem: 4194304k total, 1173844k used, 3020460k free, 0k buffers
>> > Swap: 8388608k total, 115576k used, 8273032k free, 725144k cache
>> >
>> > Notice how cpu is plenty idle, and only 1/4 of the available memory is
>> > being
>> > used.
>> >
>> > http://wiki.openvz.org/Ploop/Why explains "One such property that
>> > deserves a
>> > special item in this list is file system journal. While journal is a
>> > good
>> > thing to have, because it helps to maintain file system integrity and
>> > improve reboot times (by eliminating fsck in many cases), it is also a
>> > bottleneck for containers. If one container will fill up in-memory
>> > journal
>> > (with lots of small operations leading to file metadata updates, e.g.
>> > file
>> > truncates), all the other containers I/O will block waiting for the
>> > journal
>> > to be written to disk. In some extreme cases we saw up to 15 seconds of
>> > such
>> > blockage.". The problem I noticed last much longer than 15 seconds
>> > though
>> > - typically 15-30 minutes, then load goes back where it should be.
>> >
>> > Any suggestions where I could look for the cause of this? It's not like
>> > it
>> > happens everyday, maybe once or twice per month, but it's enough to
>> > cause
>> > customers to complain.
>> >
>> > Regards,
>> > Rene

```

>> >  
>> >  
>> --  
>> Satzmedia GmbH  
>>  
>> Altonaer Poststraße 9  
>> 22767 Hamburg  
>> Tel: +49 (0) 40 - 1 888 969 - 140  
>> Fax: +49 (0) 40 - 1 888 969 - 200  
>> E-Mail: s.johannsen@satzmedia.de  
>> E-Business-Lösungen: <http://www.satzmedia.de>  
>> Amtsgericht Hamburg, HRB 71729  
>> Ust-IDNr. DE201979921  
>> Geschäftsführer:  
>> Dipl.-Kfm. Christian Satz  
>> Dipl.-Inform. Markus Meyer-Westphal  
>>  
>> --  
>>  
>>  
>>  
--  
Satzmedia GmbH

Altonaer Poststraße 9  
22767 Hamburg  
Tel: +49 (0) 40 - 1 888 969 - 140  
Fax: +49 (0) 40 - 1 888 969 - 200  
E-Mail: s.johannsen@satzmedia.de  
E-Business-Lösungen: <http://www.satzmedia.de>  
Amtsgericht Hamburg, HRB 71729  
Ust-IDNr. DE201979921  
Geschäftsführer:  
Dipl.-Kfm. Christian Satz  
Dipl.-Inform. Markus Meyer-Westphal

--

---