
Subject: occasional high loadavg without any noticeable cpu/memory/io load

Posted by [Rene Dokbua](#) on Mon, 21 May 2012 18:06:53 GMT

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

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

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [svensirk](#) on Tue, 22 May 2012 08:06:15 GMT

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

--

Subject: RE: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Steffan](#) on Tue, 22 May 2012 08:15:40 GMT

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

Sorry dont have the answer for you

But can you tell me what command you used to see all loads on your node ?

Thanxs Steffan

Van: users-bounces@openvz.org [mailto:users-bounces@openvz.org] Namens Rene

Dokbua

Verzonden: maandag 21 mei 2012 20:07

Aan: users@openvz.org

Onderwerp: [Users] occasional high loadavg without any noticeable
cpu/memory/io load

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

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Rene Dokbua](#) on Tue, 22 May 2012 09:06:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Actually I made a small shell script that loops through the list of active containers and outputs the content of each containers /proc/loadavg. It started out as a bit more elaborate script that was intended to provide some of the functionality of a script vzstat, that I used to use with Virtuozzo.

You can download both scripts from
<https://www.ourhelpdesk.net/downloads/z.tgz>

On Tue, May 22, 2012 at 3:15 PM, Steffan <general@ziggo.nl> wrote:

```
> Sorry dont have the answer for you****
>
> But can you tell me what command you used to see all loads on your node ?*
> ***
>
> ** **
>
> Thanxs Steffan****
>
> ** **
>
> *Van:* users-bounces@openvz.org [mailto:users-bounces@openvz.org] *Namens
> *Rene Dokbua
> *Verzonden:* maandag 21 mei 2012 20:07
> *Aan:* users@openvz.org
> *Onderwerp:* [Users] occasional high loadavg without any noticeable
> cpu/memory/io load****
>
```

```

> ** **
>
> 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****
>
> ** **
>

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Rene Dokbua](#) on Tue, 22 May 2012 09:16:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Sirk,

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!

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.

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?

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.

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
>
> --
>
>

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

--

Subject: RE: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Esm](#) on Tue, 22 May 2012 10:00:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Rene,

load could be caused by buffers that are full.

Hope it helps you,

Kind Regards,

Esme de Wolf

Van: users-bounces@openvz.org [mailto:users-bounces@openvz.org] Namens Rene Dokbua
Verzonden: maandag 21 mei 2012 20:07
Aan: users@openvz.org
Onderwerp: [Users] occasional high loadavg without any noticeable cpu/memory/io load

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

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Rene Dokbua](#) on Tue, 22 May 2012 10:27:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Sirk,

> 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).
>

Thanks for the info, much appreciated!

Maybe a little off topic, but I am curious to know: At the moment I find it very convenient to go directly into a containers filesystem from the hardware node - i.e. something like /vz/private/xxx/var/log/... etc - Would I be correct in presuming that by using ploop this will no longer be possible? I know I could just setup a test system and try it out but if you know already it would save me some time ;)

> 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.

>

>

Indeed, this is the only container on that filesystem and that physical drive. This time there were no "spill over" but previous times when load hit 50 or more the load certainly did spill into other containers.

Best,
Rene

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Rene Dokbua](#) on Tue, 22 May 2012 10:49:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Esme,

> Did you check the /proc/user_beancounters of that VPS? Sometime's a high load could be caused by buffers that are full.

Thanks for the suggestion, much appreciated!

I didn't think of checking at the time I'm afraid. I suppose since the container has not been rebooted since, the beancounters should still show any problems encountered at the time right?

Below is the user_beancounters of the problem CT. I notice physpages and dcachesize have maxheld values very close to limits (even if failcnt is zero) could that have been the cause?

| uid | resource | limit | held | maxheld |
|------------|-------------|------------|-----------|------------|
| barrier | | | failcnt | |
| 1407: | kmemsize | | 252703307 | 1124626432 |
| 1932525568 | | 2147483648 | | 0 |
| | lockedpages | | 0 | 15 |
| 524288 | | 524288 | 0 | |

| | | | |
|---------------------|---------------------|------------|---|
| privvmpages | 893372 | 5683554 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| shmpages | 23 | 7399 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| dummy | 0 | 0 | |
| 0 | 0 | 0 | |
| numproc | 136 | 480 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| physpages | 733468 | 1048591 | |
| 0 | 1048576 | 0 | |
| vmguarpages | 0 | 0 | |
| 0 | 9223372036854775807 | 0 | |
| oomguarpages | 137691 | 676209 | |
| 0 | 9223372036854775807 | 0 | |
| numtcpsock | 101 | 459 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| numflock | 7 | 37 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| numpty | 1 | 4 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| numsiginfo | 0 | 66 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| tcpsndbuf | 4024896 | 34884168 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| tcprcvbuf | 1654784 | 7520256 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| othersockbuf | 195136 | 3887232 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| dgramrcvbuf | 0 | 155848 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| numothersock | 130 | 346 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| dcachesize | 222868425 | 1073741824 | |
| 965738496 | 1073741824 | 0 | |
| numfile | 3853 | 12765 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| dummy | 0 | 0 | |
| 0 | 0 | 0 | |
| dummy | 0 | 0 | |
| 0 | 0 | 0 | |
| dummy | 0 | 0 | |
| 0 | 0 | 0 | |
| numiptent | 197 | 197 | |
| 9223372036854775807 | 9223372036854775807 | | 0 |

I'm not that familiar with the nitty-gritties of the beancounters but these are the values I have in the 1407.conf file.

```
PHYSPAGES="0:4096M"
SWAPPAGES="0:8192M"
KMEMSIZE="1843M:2048M"
DCACHESIZE="921M:1024M"
LOCKEDPAGES="2048M"
PRIVVMPAGES="unlimited"
SHMPAGES="unlimited"
NUMPROC="unlimited"
VMGUARPAGES="0:unlimited"
OOMGUARPAGES="0:unlimited"
NUMTCPSOCK="unlimited"
NUMFLOCK="unlimited"
NUMPTY="unlimited"
NUMSIGINFO="unlimited"
TCPSNDBUF="unlimited"
TCPRCVBUF="unlimited"
OTHERSOCKBUF="unlimited"
DGRAMRCVBUF="unlimited"
NUMOTHERSOCK="unlimited"
NUMFILE="unlimited"
NUMIPTENT="unlimited"
```

When user_beancounters physpage limit is 1048576, with PHYSPAGES set to 4GB, then the held value of 733468 should correspond to about 3GB, right?

But top only shows about 1.5GB used at the same time - how is that possible?

dcachesize I think is filesystem stuff? But there seems to be plenty of resources there;

```
# df -i
Filesystem      Inodes  IUsed  IFree IUse% Mounted on
/dev/simfs      20000000 3046139 16953861 16% /
none           524288   109 524179   1% /dev
# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/simfs      492G  156G  312G  34% /
none           2.0G  4.0K  2.0G   1% /dev
```

Best,
Rene

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Rene Dokbua](#) on Tue, 22 May 2012 10:59:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

>

>
> When user_beancounters physpage limit is 1048576, with PHYSPAGES set to
> 4GB, then the held value of 733468 should correspond to about 3GB, right?
> But top only shows about 1.5GB used at the same time - how is that
> possible?

>
>

Actually at the time I cat'ed these beans the memory used according to top was around 2.5GB so that seems right enough. Still doesn't explain how maxheld is so close to limit when top at the time of the trouble showed just around 1.5G memory used.

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Kirill Korotaev](#) on Tue, 22 May 2012 11:05:16 GMT

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

Looks like in your case you've hit physpages limit.

In such situations VPS behaves as a standalone machine - it starts to swap out (though "virtually") and process stuck in D state (swap in / swap out), which contributes to loadavg.

So either increase memory limits for your VPS or kill/tune the memory hungry workload.

Note: loadavg can also increase due to CPU limits as processes are delayed when overuse their CPU.

Thanks,
Kirill

On May 22, 2012, at 14:49 , Rene C. wrote:

Hi Esme,

> Did you check the /proc/user_beancounters of that VPS? Sometime's a high load could be caused by buffers that are full.

Thanks for the suggestion, much appreciated!

I didn't think of checking at the time I'm afraid. I suppose since the container has not been rebooted since, the beancounters should still show any problems encountered at the time right?

Below is the user_beancounters of the problem CT. I notice physpages and dcachesize have maxheld values very close to limits (even if failcnt is zero) could that have been the cause?

| uid | resource | held | maxheld | barrier | limit | failcnt |
|---------------------|--------------|-----------|------------|---------------------|---------------------|---------|
| 1407: | kmemsize | 252703307 | 1124626432 | 1932525568 | | |
| 2147483648 | | 0 | | | | |
| | lockedpages | 0 | 15 | 524288 | 524288 | 0 |
| | privvmpages | 893372 | 5683554 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | shmpages | 23 | 7399 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | dummy | 0 | 0 | 0 | 0 | 0 |
| | numproc | 136 | 480 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | physpages | 733468 | 1048591 | 0 | 1048576 | |
| 0 | | | | | | |
| | vmguarpages | 0 | 0 | 0 | 9223372036854775807 | |
| 0 | | | | | | |
| | oomguarpages | 137691 | 676209 | 0 | 9223372036854775807 | |
| 0 | | | | | | |
| | numtcpsock | 101 | 459 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | numflock | 7 | 37 | 9223372036854775807 | 9223372036854775807 | |
| 0 | | | | | | |
| | numpty | 1 | 4 | 9223372036854775807 | 9223372036854775807 | |
| 0 | | | | | | |
| | numsiginfo | 0 | 66 | 9223372036854775807 | 9223372036854775807 | |
| 0 | | | | | | |
| | tcpsndbuf | 4024896 | 34884168 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | tcprcvbuf | 1654784 | 7520256 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | othersockbuf | 195136 | 3887232 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | dgramrcvbuf | 0 | 155848 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | numothersock | 130 | 346 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | dcachesize | 222868425 | 1073741824 | 965738496 | 1073741824 | |
| 0 | | | | | | |
| | numfile | 3853 | 12765 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | dummy | 0 | 0 | 0 | 0 | 0 |
| | dummy | 0 | 0 | 0 | 0 | 0 |
| | dummy | 0 | 0 | 0 | 0 | 0 |
| | numiptent | 197 | 197 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |

I'm not that familiar with the nitty-gritties of the beancounters but these are the values I have in the 1407.conf file.

```
PHYSPAGES="0:4096M"
SWAPPAGES="0:8192M"
KMEMSIZE="1843M:2048M"
DCACHESIZE="921M:1024M"
LOCKEDPAGES="2048M"
PRIVVMPAGES="unlimited"
SHMPAGES="unlimited"
NUMPROC="unlimited"
VMGUARPAGES="0:unlimited"
OOMGUARPAGES="0:unlimited"
NUMTCPSOCK="unlimited"
NUMFLOCK="unlimited"
NUMPTY="unlimited"
NUMSIGINFO="unlimited"
TCPSNDBUF="unlimited"
TCPRCVBUF="unlimited"
OTHERSOCKBUF="unlimited"
DGRAMRCVBUF="unlimited"
NUMOTHERSOCK="unlimited"
NUMFILE="unlimited"
NUMIPTENT="unlimited"
```

When user_beancounters physpage limit is 1048576, with PHYSPAGES set to 4GB, then the held value of 733468 should correspond to about 3GB, right? But top only shows about 1.5GB used at the same time - how is that possible?

dcachesize I think is filesystem stuff? But there seems to be plenty of resources there;

```
# df -i
Filesystem      Inodes  IUsed  IFree IUse% Mounted on
/dev/simfs      20000000 3046139 16953861  16% /
none           524288   109 524179   1% /dev
# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/simfs      492G  156G  312G  34% /
none           2.0G  4.0K  2.0G   1% /dev
```

Best,
Rene
<ATT00001.c>

Subject: RE: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Esm](#) on Tue, 22 May 2012 11:59:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

I also think that these UBC settings are not consistent. Especially when you have all containers configured with these same UBC settings you will have

soon or later problems.

See: http://wiki.openvz.org/UBC_consistency_check and other pages on the WIKI.

Kind Regards,

Esme

Van: users-bounces@openvz.org [mailto:users-bounces@openvz.org] Namens Kirill Korotaev

Verzonden: dinsdag 22 mei 2012 13:05

Aan: users@openvz.org users@openvz.org; Rene C.

Onderwerp: Re: [Users] occasional high loadavg without any noticeable cpu/memory/io load

Looks like in your case you've hit physpages limit.

In such situations VPS behaves as a standalone machine - it starts to swap out (though "virtually") and process stuck in D state (swap in / swap out),

which contributes to loadavg.

So either increase memory limits for your VPS or kill/tune the memory hungry workload.

Note: loadavg can also increase due to CPU limits as processes are delayed when overuse their CPU.

Thanks,

Kirill

On May 22, 2012, at 14:49 , Rene C. wrote:

Hi Esme,

load could be caused by buffers that are full.

Thanks for the suggestion, much appreciated!

I didn't think of checking at the time I'm afraid. I suppose since the container has not been rebooted since, the beancounters should still show any problems encountered at the time right?

Below is the user_beancounters of the problem CT. I notice physpages and dcachesize have maxheld values very close to limits (even if failcnt is zero) could that have been the cause?

| uid | resource | held | maxheld |
|---------------------|---------------------|-----------|------------|
| barrier | limit | failcnt | |
| 1407: | kmemsize | 252703307 | 1124626432 |
| 1932525568 | 2147483648 | | 0 |
| | lockedpages | 0 | 15 |
| 524288 | 524288 | 0 | |
| | privvmpages | 893372 | 5683554 |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| | shmpages | 23 | 7399 |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| | dummy | 0 | 0 |
| 0 | 0 | 0 | |
| | numproc | 136 | 480 |
| 9223372036854775807 | 9223372036854775807 | | 0 |
| | physpages | 733468 | 1048591 |
| 0 | 1048576 | 0 | |

| | | | |
|---|-----------|------------|---|
| vmguarpages | 0 | 0 | |
| 0 9223372036854775807 | 0 | | |
| oomguarpages | 137691 | 676209 | |
| 0 9223372036854775807 | 0 | | |
| numtcpsock | 101 | 459 | |
| 9223372036854775807 9223372036854775807 | | | 0 |
| numflock | 7 | 37 | |
| 9223372036854775807 9223372036854775807 | | | 0 |
| numpty | 1 | 4 | |
| 9223372036854775807 9223372036854775807 | | | 0 |
| numsiginfo | 0 | 66 | |
| 9223372036854775807 9223372036854775807 | | | 0 |
| tcpsndbuf | 4024896 | 34884168 | |
| 9223372036854775807 9223372036854775807 | | | 0 |
| tcprcvbuf | 1654784 | 7520256 | |
| 9223372036854775807 9223372036854775807 | | | 0 |
| othersockbuf | 195136 | 3887232 | |
| 9223372036854775807 9223372036854775807 | | | 0 |
| dgramrcvbuf | 0 | 155848 | |
| 9223372036854775807 9223372036854775807 | | | 0 |
| numothersock | 130 | 346 | |
| 9223372036854775807 9223372036854775807 | | | 0 |
| dcachesize | 222868425 | 1073741824 | |
| 965738496 1073741824 | | 0 | |
| numfile | 3853 | 12765 | |
| 9223372036854775807 9223372036854775807 | | | 0 |
| dummy | 0 | 0 | |
| 0 0 0 | | | |
| dummy | 0 | 0 | |
| 0 0 0 | | | |
| dummy | 0 | 0 | |
| 0 0 0 | | | |
| numiptent | 197 | 197 | |
| 9223372036854775807 9223372036854775807 | | | 0 |

I'm not that familiar with the nitty-gritties of the beancounters but these are the values I have in the 1407.conf file.

PHYSPAGES="0:4096M"

SWAPPAGES="0:8192M"

KMEMSIZE="1843M:2048M"

DCACHESIZE="921M:1024M"

LOCKEDPAGES="2048M"
PRIVVMPAGES="unlimited"
SHMPAGES="unlimited"
NUMPROC="unlimited"
VMGUARPAGES="0:unlimited"
OOMGUARPAGES="0:unlimited"
NUMTCPSOCK="unlimited"
NUMFLOCK="unlimited"
NUMPTY="unlimited"
NUMSIGINFO="unlimited"
TCPSNDBUF="unlimited"
TCPRCVBUF="unlimited"
OTHERSOCKBUF="unlimited"
DGRAMRCVBUF="unlimited"
NUMOTHERSOCK="unlimited"
NUMFILE="unlimited"
NUMIPTENT="unlimited"

When user_beancounters physpage limit is 1048576, with PHYSPAGES set to 4GB, then the held value of 733468 should correspond to about 3GB, right? But top only shows about 1.5GB used at the same time - how is that possible?

dcachesize I think is filesystem stuff? But there seems to be plenty of resources there;

```
# df -i
```

| Filesystem | Inodes | IUsed | IFree | IUse% | Mounted on |
|------------|--------|-------|-------|-------|------------|
|------------|--------|-------|-------|-------|------------|

```
/dev/simfs      20000000 3046139 16953861  16% /
```

```
none           524288   109 524179   1% /dev
```

```
# df -h
```

```
Filesystem      Size  Used Avail Use% Mounted on
```

```
/dev/simfs      492G  156G  312G  34% /
```

```
none           2.0G  4.0K  2.0G   1% /dev
```

Best,
Rene

<ATT00001.c>

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [svensirk](#) on Tue, 22 May 2012 12:01:27 GMT

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

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

> Hi Sirk,

>

>>

>> 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).

>

>

> Thanks for the info, much appreciated!

>

> Maybe a little off topic, but I am curious to know: At the moment I find it

> very convenient to go directly into a containers filesystem from the

> hardware node - i.e. something like /vz/private/xxx/var/log/... etc - Would

> I be correct in presuming that by using ploop this will no longer be

> possible? I know I could just setup a test system and try it out but if you

> know already it would save me some time ;)

Only partially correct :-)

You can enter the filesystem of a CT when it's mounted. Meaning - you can enter the root directory when the CT is running.

If the CT is shut down you always have the possibility to mount the ploop file to any directory you desire.

>

>>

>> 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.

>>

>

> Indeed, this is the only container on that filesystem and that physical drive. This time there were no "spill over" but previous times when load hit 50 or more the load certainly did spill into other containers.

>

> Best,

> 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

--

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Rene Dokbua](#) on Tue, 22 May 2012 12:17:29 GMT

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

On Tue, May 22, 2012 at 6:59 PM, Esmé de Wolf <esme@elements.nl> wrote:

> I also think that these UBC settings are not consistent. Especially when
> you have all containers configured with these same UBC settings you will
> have soon or later problems. ****
>
> ** **
>
> See: http://wiki.openvz.org/UBC_consistency_check and other pages on the
> WIKI. ****
>
> ** **
>
> Kind Regards, ****
>
>
> Esme
>

I read that UBC page already and used it to set these values.

No, all my containers do not have the same UBC settings, they were set depending on how much resources each container should have.

Please let me know where any of the values in my conf file conflicts with the UBC recommendations.

I do understand that they may need to be fine tuned in each case, but that's basically what this question is about :)

So basically at this time I have two questions I don't understand:

- 1) how is it possible to have physpages hit the limit when top never shows more than about 75-80% of the memory used?
- 2) how did dcachesize hit limit when both df -i and df -h shows plenty of resources - and haven't been close to limits?

Could the values in the beancounter file be old? Is there a way to reset them (without restarting the CT) ?

Best,
Rene

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Kirill Korotaev](#) on Tue, 22 May 2012 12:35:33 GMT
[View Forum Message](#) <> [Reply to Message](#)

On May 22, 2012, at 16:17 , Rene C. wrote:

On Tue, May 22, 2012 at 6:59 PM, Esmé de Wolf

<esme@elements.nl<mailto:esme@elements.nl>> wrote:

I also think that these UBC settings are not consistent. Especially when you have all containers configured with these same UBC settings you will have soon or later problems.

See: http://wiki.openvz.org/UBC_consistency_check and other pages on the WIKI.

Kind Regards,

Esme

I read that UBC page already and used it to set these values.

No, all my containers do not have the same UBC settings, they were set depending on how much resources each container should have.

Please let me know where any of the values in my conf file conflicts with the UBC recommendations.

I do understand that they may need to be fine tuned in each case, but that's basically what this question is about :)

So basically at this time I have two questions I don't understand:

1) how is it possible to have physpages hit the limit when top never shows more than about 75-80% of the memory used?

once again: top shows current (immediate) values.

maxheld in user_beancounters shows you *maximum* over time.

There is an API for resetting it AFAIR, but no user-space tool in OpenVZ :(((

2) how did dcachesize hit limit when both df -i and df -h shows plenty of resources - and haven't been close to limits?

dcachesize has nothing to do with df.

it's kernel memory used for paths and pinned by opened files and CWD.

You can safely increase it if needed. It's just DoS protection.

Could the values in the beancounter file be old? Is there a way to reset them (without restarting the CT) ?

Best,
Rene

<ATT00001.c>

Subject: RE: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Esm](#) on Tue, 22 May 2012 12:54:09 GMT
[View Forum Message](#) <> [Reply to Message](#)

You could check your <VEID>.conf with vzcfgvalidate. But I think there is

http://wiki.openvz.org/UBC_failcnt_reset There is no need to reset the

Van: users-bounces@openvz.org [mailto:users-bounces@openvz.org] Namens Rene C.

Verzonden: dinsdag 22 mei 2012 14:17

Aan: users@openvz.org

Onderwerp: Re: [Users] occasional high loadavg without any noticeable
cpu/memory/io load

On Tue, May 22, 2012 at 6:59 PM, Esmé de Wolf <esme@elements.nl> wrote:

I also think that these UBC settings are not consistent. Especially when you have all containers configured with these same UBC settings you will have soon or later problems.

See: http://wiki.openvz.org/UBC_consistency_check and other pages on the WIKI.

Kind Regards,

Esme

I read that UBC page already and used it to set these values.

No, all my containers do not have the same UBC settings, they were set depending on how much resources each container should have.

Please let me know where any of the values in my conf file conflicts with the UBC recommendations.

I do understand that they may need to be fine tuned in each case, but that's basically what this question is about :)

So basically at this time I have two questions I don't understand:

1) how is it possible to have physpages hit the limit when top never shows more than about 75-80% of the memory used?

2) how did dcachesize hit limit when both df -i and df -h shows plenty of resources - and haven't been close to limits?

Could the values in the beancounter file be old? Is there a way to reset them (without restarting the CT) ?

Best,

Rene

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Rene Dokbua](#) on Tue, 22 May 2012 17:09:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Tue, May 22, 2012 at 7:54 PM, Esmé de Wolf <esme@elements.nl> wrote:

> You could check your <VEID>.conf with vzcfgvalidate. But I think there is
> quite a risk when giving one of your CT's unlimited resources. If you want
> to read-out the UBC's from the node and see when one fails I could
> recommend you a very good script I'm using myself;
> http://wiki.openvz.org/UBC_failcnt_reset There is no need to reset the

> value's inside your CT.****

>

>

>

Apparently no problems with the file:

```
# vzcfgvalidate -v yes 1407.conf
```

Validation completed: success

Thank you to everyone who provided suggestions, ideas and insight. I've added the user_beancounters to my loadmonitoring script. Next time there is a problem I'll check if any values are hitting the limit and see if increasing them may fix the problem.

Best,
Rene

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Martin Dobrev](#) on Wed, 23 May 2012 07:14:21 GMT

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

Hi,

?? 22.5.2012 ?. 13:27 ?., Rene C. ??????:

> Hi Sirk,

>

>

> 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).

>

>

> Thanks for the info, much appreciated!

>

> Maybe a little off topic, but I am curious to know: At the moment I

> find it very convenient to go directly into a containers filesystem

> from the hardware node - i.e. something like

> /vz/private/xxx/var/log/... etc - Would I be correct in presuming

> that by using ploop this will no longer be possible? I know I could

> just setup a test system and try it out but if you know already it

> would save me some time ;)

>

It's not very practical to access the containers from the VZ/private mount point, as it breaks for example the quota stats of the container. If you still want to do things there better go for the VZ/root mount point. (Advice given to me by one of the now-a-days developer of Viruozzo) And as you already mentioned ploop, as far as I know the ploop-container will be mounted to VZ/root of the CT and you'll still have access to the info in there.

>

> 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.

>

>

> Indeed, this is the only container on that filesystem and that physical drive. This time there were no "spill over" but previous times when load hit 50 or more the load certainly did spill into other containers.

> Best,
> Rene

>

>

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Aleksandar Ivanisevic](#) on Thu, 24 May 2012 14:32:34 GMT

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

"Rene C." <openvz@dokbua.com> writes:

> Thank you to everyone who provided suggestions, ideas and insight. I've added the user_beancounters to my loadmonitoring script. Next time there is a problem I'll check if any values are hitting the limit and see if increasing them may fix the problem.

Well, all your failcnt's were zero so there was nothing hitting the limit.

I'm also interested in what you found since I'm having the same problems on one of my clusters: unexplained high load that goes away as it came: unexplained and suddenly ;)

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load

Posted by [kir](#) on Wed, 30 May 2012 15:07:16 GMT

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

On 05/22/2012 09:09 PM, Rene C. wrote:

>
>
> On Tue, May 22, 2012 at 7:54 PM, Esmé de Wolf <esme@elements.nl
> <mailto:esme@elements.nl>> wrote:
>
> You could check your <VEID>.conf with vzcfgvalidate. But I think
> there is quite a risk when giving one of your CT's unlimited
> resources. If you want to read-out the UBC's from the node and see
> when one fails I could recommend you a very good script I'm using
> myself; http://wiki.openvz.org/UBC_failcnt_reset There is no need
> to reset the value's inside your CT.
>
>
>
> Apparently no problems with the file:
>
> # vzcfgvalidate -v yes 1407.conf
> Validation completed: success
>
> Thank you to everyone who provided suggestions, ideas and insight.
> I've added the user_beancounters to my loadmonitoring script. Next
> time there is a problem I'll check if any values are hitting the limit
> and see if increasing them may fix the problem.
>

Try vzubc -q or something, it might help.

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [kir](#) on Wed, 30 May 2012 15:09:20 GMT

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

On 05/22/2012 01:06 PM, Rene C. wrote:

>
> Actually I made a small shell script that loops through the list of
> active containers and outputs the content of each containers
> /proc/loadavg. It started out as a bit more elaborate script that was
> intended to provide some of the functionality of a script vzstat, that
> I used to use with Virtuozzo.
>
> You can download both scripts from
> <https://www.ourhelpdesk.net/downloads/z.tgz>

vzlist have leverage field that might be of use. I.e.

vzlist -o ctid,lverage

>
>
>
> On Tue, May 22, 2012 at 3:15 PM, Steffan <general@ziggo.nl
> <mailto:general@ziggo.nl> wrote:
>
> Sorry dont have the answer for you
>
> But can you tell me what command you used to see all loads on your
> node ?
>
> Thanxs Steffan
>
> *Van:*users-bounces@openvz.org <mailto:users-bounces@openvz.org>
> [mailto:users-bounces@openvz.org
> <mailto:users-bounces@openvz.org>] *Namens *Rene Dokbua
> *Verzonden:* maandag 21 mei 2012 20:07
> *Aan:* users@openvz.org <mailto:users@openvz.org>
> *Onderwerp:* [Users] occasional high loadavg without any
> noticeable cpu/memory/io load
>
> 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

>
>
>
>
> _____
> Users mailing list
> Users@openvz.org <mailto:Users@openvz.org>
> https://openvz.org/mailman/listinfo/users

>
>

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Rene Dokbua](#) on Wed, 30 May 2012 16:54:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi Kir,

Both the vzubc command and the leverage option to vzlist were new to me
(the leverage options seems undocumented?)

Thanks much, this is VERY useful information!!

Regards,
Rene

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
Posted by [Rene Dokbua](#) on Wed, 04 Jul 2012 09:16:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

Today I again had a VE that went up to a relative high load for no apparent reason.

Below are the details for the hardware node, followed by the high-load container.

I realize it's not the latest kernel, but a reboot takes half an hour (from first VE goes down to last VE is back up, assuming everything goes well and no FCK is forced) so we only reboot into new kernels when there is a really serious reason for it or the server crashes - but I don't see anything in the kernel updates since our current kernel that would address

this issue anyway.

Why does the load in this container suddenly go up like that? Websites hosted by the container becomes very sluggish, so it is a real problem.

It isn't just a problem with this container - or even this hardware node for that reason, I occasionally see it with containers on other hardware nodes as well. One idea I brought up before was that perhaps it's the file system journal, as suggested in <http://wiki.openvz.org/Ploop/Why> - but I think that would affect all containers on that file system, not just a single container?

--- HARDWARE NODE ---

```
# uname -a
```

```
Linux server15.hardwarenode.com 2.6.32-042stab049.6 #1 SMP Mon Feb 6  
19:17:43 MSK 2012 x86_64 x86_64 x86_64 GNU/Linux
```

```
# rpm -q sl-release
```

```
sl-release-6.1-2.x86_64
```

```
# top -cbn1 | head -17
```

```
top - 21:00:02 up 123 days, 15:31, 1 user, load average: 0.97, 2.70, 2.37
```

```
Tasks: 886 total, 6 running, 880 sleeping, 0 stopped, 0 zombie
```

```
Cpu(s): 8.4%us, 1.7%sy, 0.0%ni, 86.3%id, 3.5%wa, 0.0%hi, 0.1%si,  
0.0%st
```

```
Mem: 16420716k total, 15566264k used, 854452k free, 1477372k buffers
```

```
Swap: 16777184k total, 623672k used, 16153512k free, 4578176k cached
```

| PID | USER | PR | NI | VIRT | RES | SHR | S | %CPU | %MEM | TIME+ | COMMAND |
|-------------------|--------|----|----|-------|------|------|---|-------|------|----------|------------|
| 94153 | 27 | 20 | 0 | 164m | 41m | 3392 | S | 150.9 | 0.3 | 50575:37 | |
| /usr/libexec/mys | | | | | | | | | | | |
| 9178 | 27 | 20 | 0 | 159m | 29m | 3000 | S | 72.6 | 0.2 | 1284:50 | |
| /usr/libexec/mysq | | | | | | | | | | | |
| 567031 | apache | 20 | 0 | 40296 | 15m | 3588 | S | 17.2 | 0.1 | 0:00.09 | |
| /usr/sbin/httpd | | | | | | | | | | | |
| 567382 | root | 20 | 0 | 15672 | 1820 | 864 | R | 5.7 | 0.0 | 0:00.04 | top -cbn1 |
| 38 | root | 20 | 0 | 0 | 0 | 0 | S | 1.9 | 0.0 | 2:55.25 | [events/3] |
| 41 | root | 20 | 0 | 0 | 0 | 0 | S | 1.9 | 0.0 | 0:29.00 | [events/6] |
| 566362 | apache | 20 | 0 | 43240 | 19m | 4448 | R | 1.9 | 0.1 | 0:01.04 | |
| /usr/sbin/httpd | | | | | | | | | | | |
| 566857 | apache | 20 | 0 | 55248 | 11m | 3456 | R | 1.9 | 0.1 | 0:00.05 | |
| /usr/sbin/httpd | | | | | | | | | | | |
| 566918 | apache | 20 | 0 | 42596 | 17m | 3704 | S | 1.9 | 0.1 | 0:00.15 | |
| /usr/sbin/httpd | | | | | | | | | | | |
| 567033 | apache | 20 | 0 | 39784 | 14m | 3468 | S | 1.9 | 0.1 | 0:00.01 | |
| /usr/sbin/httpd | | | | | | | | | | | |

```
# vzlist -o ctid,lverage
CTID    LAVERAGE
1501 0.00/0.05/0.02
1502 0.00/0.00/0.00
1503 0.08/0.03/0.01
1504 0.00/0.00/0.00
1505 8.29/6.04/3.67
1506 27.11/16.97/7.89
1507 0.00/0.00/0.00
1508 0.19/0.06/0.01
1509 0.07/0.03/0.00
1510 0.02/0.02/0.00
1512 0.00/0.00/0.00
1514 0.00/0.00/0.00
```

```
# iostat -xN
Linux 2.6.32-042stab049.6 (server15.hardwarenode.com) 07/03/12
_x86_64_ (8 CPU)
```

```
avg-cpu: %user  %nice %system %iowait  %steal  %idle
          8.41   0.04   1.75   3.51   0.00  86.28
```

```
Device:      rrqm/s  wrqm/s   r/s    w/s  rsec/s  wsec/s avgrq-sz
avgqu-sz  await  svctm  %util
sdd        0.76  56.58   0.59   0.59   20.27  457.28  402.66
          0.25 211.66  4.03  0.48
sdc        1.72  27.94  17.20  16.16  887.30  336.18  36.68
          0.02 12.71  5.23 17.45
sdb        1.65  27.79  19.48  12.95  975.43  318.64  39.91
          0.09 15.22  3.77 12.23
sda        0.01   0.16   0.10   0.24   1.95   2.79  13.79
          0.00  7.06  4.16  0.14
vg01-swap   0.00   0.00   0.00   0.00   0.00   0.00   8.00
          0.00  3.68  2.22  0.00
vg01-root   0.00   0.00   0.11   0.35   1.94   2.78  10.30
          0.02 38.30  3.12  0.14
vg04-swap   0.00   0.00   1.30   0.22  10.41   1.80   8.00
          0.01  9.28  1.44  0.22
vg04-vz     0.00   0.00   0.05  56.94   9.86  455.49   8.17
          0.01  0.18  0.05  0.27
vg03-swap   0.00   0.00   0.00   0.00   0.00   0.00   8.00
          0.00  6.72  1.10  0.00
vg03-vz     0.00   0.00  18.98  42.41  887.30  336.18  19.93
          0.39  6.33  2.84 17.45
vg02-swap   0.00   0.00   0.00   0.00   0.00   0.00   8.00
          0.00  7.03  0.89  0.00
vg02-vz     0.00   0.00  21.19  39.91  975.43  318.64  21.18
          0.15  8.99  2.00 12.23
```

```
vg01-vz      0.00  0.00  0.00  0.00  0.00  0.00  7.98
  0.00 17.73 17.73  0.00
```

--- CONTAINER ---

top -cbn1 | head -100

top - 21:00:04 up 123 days, 15:25, 0 users, load average: 27.11, 16.97, 7.89

Tasks: 86 total, 2 running, 84 sleeping, 0 stopped, 0 zombie

Cpu(s): 1.4%us, 0.2%sy, 0.0%ni, 98.1%id, 0.1%wa, 0.0%hi, 0.0%si, 0.2%st

Mem: 655360k total, 316328k used, 339032k free, 0k buffers

Swap: 1310720k total, 68380k used, 1242340k free, 58268k cached

| PID | USER | PR | NI | VIRT | RES | SHR | S | %CPU | %MEM | TIME+ | COMMAND |
|---------------------|--------|----|----|-------|------|------|---|------|------|---------|-----------------|
| 916 | mysql | 20 | 0 | 159m | 29m | 3000 | S | 79.3 | 4.6 | 1284:51 | |
| /usr/libexec/mysqld | | | | | | | | | | | |
| 1 | root | 20 | 0 | 2156 | 92 | 64 | S | 0.0 | 0.0 | 0:36.50 | init [3] |
| 2 | root | 20 | 0 | 0 | 0 | 0 | S | 0.0 | 0.0 | 0:00.00 | [kthreadd/1506] |
| 3 | root | 20 | 0 | 0 | 0 | 0 | S | 0.0 | 0.0 | 0:00.00 | [khelper/1506] |
| 97 | root | 16 | -4 | 2244 | 8 | 4 | S | 0.0 | 0.0 | 0:00.00 | /sbin/udevd -d |
| 634 | root | 20 | 0 | 1812 | 212 | 136 | S | 0.0 | 0.0 | 2:39.88 | syslogd -m 0 |
| 667 | root | 20 | 0 | 7180 | 268 | 168 | S | 0.0 | 0.0 | 1:01.55 | /usr/sbin/sshd |
| 676 | root | 20 | 0 | 2832 | 392 | 304 | S | 0.0 | 0.1 | 0:15.13 | xinetd |
| -stayalive - | | | | | | | | | | | |
| 690 | root | 20 | 0 | 6040 | 124 | 72 | S | 0.0 | 0.0 | 0:02.45 | |
| /usr/lib/courier-im | | | | | | | | | | | |
| 693 | root | 20 | 0 | 4872 | 252 | 200 | S | 0.0 | 0.0 | 0:01.94 | |
| /usr/sbin/courierlo | | | | | | | | | | | |
| 701 | root | 20 | 0 | 6040 | 124 | 72 | S | 0.0 | 0.0 | 0:06.34 | |
| /usr/lib/courier-im | | | | | | | | | | | |
| 703 | root | 20 | 0 | 4872 | 256 | 200 | S | 0.0 | 0.0 | 0:03.09 | |
| /usr/sbin/courierlo | | | | | | | | | | | |
| 709 | root | 20 | 0 | 6040 | 128 | 72 | S | 0.0 | 0.0 | 0:18.15 | |
| /usr/lib/courier-im | | | | | | | | | | | |
| 711 | root | 20 | 0 | 4872 | 256 | 200 | S | 0.0 | 0.0 | 0:09.15 | |
| /usr/sbin/courierlo | | | | | | | | | | | |
| 718 | root | 20 | 0 | 6040 | 124 | 72 | S | 0.0 | 0.0 | 0:05.68 | |
| /usr/lib/courier-im | | | | | | | | | | | |
| 720 | root | 20 | 0 | 4872 | 252 | 200 | S | 0.0 | 0.0 | 0:02.54 | |
| /usr/sbin/courierlo | | | | | | | | | | | |
| 730 | qmails | 20 | 0 | 1796 | 224 | 144 | S | 0.0 | 0.0 | 1:27.21 | qmail-send |
| 732 | qmail | 20 | 0 | 1752 | 244 | 192 | S | 0.0 | 0.0 | 0:22.64 | splogger qmail |
| 733 | root | 20 | 0 | 1780 | 140 | 64 | S | 0.0 | 0.0 | 0:07.85 | qmail-lspawn |
| /usr | | | | | | | | | | | |
| 734 | qmailr | 20 | 0 | 1776 | 148 | 76 | S | 0.0 | 0.0 | 0:14.07 | qmail-rspawn |
| 735 | qmailq | 20 | 0 | 1748 | 104 | 68 | S | 0.0 | 0.0 | 0:14.01 | qmail-clean |
| 781 | root | 20 | 0 | 51880 | 4364 | 196 | S | 0.0 | 0.7 | 1:35.02 | /usr/sbin/httpd |

```

828 named    20  0 44104 5708 1112 S  0.0  0.9 10:10.53
/usr/sbin/named -u
866 root     20  0 3708   8   4 S  0.0  0.0  0:00.00 /bin/sh
/usr/bin/my
981 root     20  0 33912 3756  916 S  0.0  0.6 10:55.30 /usr/bin/spamd
--us
1107 xfs      20  0 3392   72  40 S  0.0  0.0  0:00.09 xfs -droppriv
-daem
1115 root     20  0 5672   8   4 S  0.0  0.0  0:00.00
/usr/sbin/saslauthd
1116 root     20  0 5672   8   4 S  0.0  0.0  0:00.00
/usr/sbin/saslauthd
1122 root     20  0 22992 1868 1084 S  0.0  0.3  2:09.79
/usr/bin/sw-engine
1123 root     20  0 27328 1508 1160 S  0.0  0.2  6:06.30
/usr/local/psa/admi
7251 root     20  0 4488  192 136 S  0.0  0.0  0:22.85 crond
9463 apache   20  0 59184 14m 4356 S  0.0  2.3  0:05.10 /usr/sbin/httpd
10512 apache  20  0 42316 2504  84 S  0.0  0.4  0:00.91 /usr/sbin/httpd
12090 apache  20  0 56964 14m 4492 S  0.0  2.2  0:04.48 /usr/sbin/httpd
12682 apache  20  0 61060 17m 4516 S  0.0  2.7  0:02.45 /usr/sbin/httpd
13870 sw-cp-se 20  0 7852 1932  16 S  0.0  0.3  1:19.03
/usr/sbin/sw-cp-ser
17443 apache  20  0 62416 17m 4436 S  0.0  2.7  0:05.27 /usr/sbin/httpd
17461 apache  20  0 52788 10m 4480 S  0.0  1.6  0:02.24 /usr/sbin/httpd
20430 apache  20  0 62164 17m 4356 S  0.0  2.7  0:04.25 /usr/sbin/httpd
23539 popuser  20  0 37612 25m 2328 S  0.0  3.9  0:01.50 spamd child
23924 apache  20  0 58004 15m 5536 S  0.0  2.4  0:01.56 /usr/sbin/httpd
26361 apache  20  0 54496 11m 3864 S  0.0  1.8  0:01.35 /usr/sbin/httpd
26366 apache  20  0 52944 9.8m 3892 S  0.0  1.5  0:01.45 /usr/sbin/httpd
26964 apache  20  0 59184 14m 4316 S  0.0  2.3  0:07.26 /usr/sbin/httpd
27096 apache  20  0 53728 10m 3868 S  0.0  1.6  0:00.33 /usr/sbin/httpd
27102 apache  20  0 54736 11m 3780 S  0.0  1.8  0:00.15 /usr/sbin/httpd
27103 apache  20  0 54480 11m 3784 S  0.0  1.7  0:00.11 /usr/sbin/httpd
27115 apache  20  0 57064 12m 3816 S  0.0  2.0  0:00.32 /usr/sbin/httpd
27118 apache  20  0 53728 10m 3884 S  0.0  1.6  0:01.21 /usr/sbin/httpd
27120 apache  20  0 52184 8376 3120 S  0.0  1.3  0:00.00 /usr/sbin/httpd
27129 apache  20  0 52168 8072 2960 S  0.0  1.2  0:00.00 /usr/sbin/httpd
27139 apache  20  0 53304 9840 3744 S  0.0  1.5  0:01.08 /usr/sbin/httpd
27140 apache  20  0 53000 9.8m 3832 S  0.0  1.5  0:00.66 /usr/sbin/httpd
27144 apache  20  0 52168 8072 2960 S  0.0  1.2  0:00.00 /usr/sbin/httpd
27147 apache  20  0 53252 12m 5536 S  0.0  1.9  0:00.50 /usr/sbin/httpd
27149 apache  20  0 52980 9924 3740 S  0.0  1.5  0:00.17 /usr/sbin/httpd
27153 apache  20  0 53728 10m 3836 S  0.0  1.6  0:00.49 /usr/sbin/httpd
27164 apache  20  0 55224 11m 3812 S  0.0  1.9  0:00.47 /usr/sbin/httpd
27171 apache  20  0 52916 9776 3708 S  0.0  1.5  0:00.16 /usr/sbin/httpd
27172 apache  20  0 52916 9452 3436 S  0.0  1.4  0:00.17 /usr/sbin/httpd
27173 apache  20  0 55340 11m 3720 S  0.0  1.8  0:00.08 /usr/sbin/httpd

```

```

27179 apache 20 0 52020 7764 2716 S 0.0 1.2 0:00.00 /usr/sbin/httpd
27182 apache 20 0 52020 7764 2716 S 0.0 1.2 0:00.00 /usr/sbin/httpd
27185 apache 20 0 55224 11m 3824 S 0.0 1.9 0:00.30 /usr/sbin/httpd
27186 apache 20 0 53788 10m 3840 S 0.0 1.7 0:00.11 /usr/sbin/httpd
27187 apache 20 0 52916 9448 3436 S 0.0 1.4 0:00.08 /usr/sbin/httpd
27188 apache 20 0 54628 10m 3504 S 0.0 1.7 0:00.05 /usr/sbin/httpd
27196 apache 20 0 53728 10m 3572 S 0.0 1.6 0:00.36 /usr/sbin/httpd
27200 apache 20 0 54628 11m 3796 S 0.0 1.7 0:00.05 /usr/sbin/httpd
27202 apache 20 0 54480 11m 3796 S 0.0 1.7 0:00.10 /usr/sbin/httpd
27204 apache 20 0 53992 10m 3544 S 0.0 1.6 0:00.09 /usr/sbin/httpd
27207 apache 20 0 52168 8084 2960 S 0.0 1.2 0:00.00 /usr/sbin/httpd
27213 apache 20 0 52020 6464 1788 S 0.0 1.0 0:00.00 /usr/sbin/httpd
27214 apache 20 0 54216 10m 3516 S 0.0 1.6 0:00.05 /usr/sbin/httpd
27215 apache 20 0 52020 6456 1788 S 0.0 1.0 0:00.00 /usr/sbin/httpd
27216 apache 20 0 52020 7860 2804 S 0.0 1.2 0:00.00 /usr/sbin/httpd
27218 root 20 0 9400 1900 1408 S 0.0 0.3 0:00.00 crond
27219 root 20 0 2492 956 848 S 0.0 0.1 0:00.00 /bin/sh -c
/usr/loc
27220 root 20 0 2496 1052 920 S 0.0 0.2 0:00.00 /bin/sh
/usr/local/
27233 root 20 0 2540 1016 892 S 0.0 0.2 0:00.00 /bin/bash -c
top -c
27234 root 20 0 2284 952 724 R 0.0 0.1 0:00.00 top -cbn1
27235 root 20 0 1756 420 352 S 0.0 0.1 0:00.00 head -100
27247 root 20 0 2496 452 320 S 0.0 0.1 0:00.00 /bin/sh
/usr/local/
27248 root 20 0 8280 1504 1120 R 0.0 0.2 0:00.00 /usr/bin/mysql
-uad
27249 root 20 0 1800 448 376 S 0.0 0.1 0:00.00 sed -e 1d
27250 root 20 0 2240 640 540 S 0.0 0.1 0:00.00 awk
{printf("%s", $

```

```

# netstat -ptan | grep ESTABLISHED
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:77.87.207.166:21863
ESTABLISHED 23924/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:95.165.204.26:62259
ESTABLISHED 27144/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80
::ffff:193.151.105.100:4059ESTABLISHED 27200/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80
::ffff:109.169.207.68:50087ESTABLISHED 27185/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:31.131.70.135:57017
ESTABLISHED 27179/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:95.165.204.26:62220
ESTABLISHED 27103/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:188.134.61.1:60732
ESTABLISHED 27215/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80

```

```

::ffff:193.151.105.100:4112ESTABLISHED 26964/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80
::ffff:109.169.207.68:50043ESTABLISHED 27164/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:31.131.70.135:56976
ESTABLISHED 27153/httpd

```

```

# cat /proc/user_beancounters
Version: 2.5

```

| uid | resource | limit | held | maxheld | |
|---------------------|--------------|---------------------|----------|-----------|---|
| barrier | | | failcnt | | |
| 1506: | kmemsize | | 27735306 | 179081216 | |
| 304087040 | | 335544320 | | 0 | |
| | lockedpages | | 0 | 0 | |
| 81920 | | 81920 | 0 | | |
| | privvmpages | | 393683 | 430195 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | shmpages | | 823 | 21639 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | dummy | | 0 | 0 | |
| 0 | | 0 | 0 | | |
| | numproc | | 128 | 204 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | physpages | | 79702 | 163840 | |
| 0 | | 163840 | 0 | | |
| | vmguarpages | | 0 | 0 | |
| 0 | | 9223372036854775807 | | 0 | |
| | oomguarpages | | 74734 | 75707 | |
| 0 | | 9223372036854775807 | | 0 | |
| | numtcpsock | | 59 | 153 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | numflock | | 46 | 62 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | numpty | | 0 | 1 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | numsiginfo | | 0 | 33 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | tcpsndbuf | | 1037680 | 11426176 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | tcprcvbuf | | 966656 | 2867584 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | othersockbuf | | 53824 | 838688 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | dgramrcvbuf | | 0 | 502224 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | numothersock | | 114 | 273 | |
| 9223372036854775807 | | 9223372036854775807 | | | 0 |
| | dcachesize | | 10070617 | 167772160 | |
| 150994944 | | 167772160 | | 0 | |

| | | | | |
|---------------------|---------------------|------|------|---|
| numfile | | 1634 | 1865 | |
| 9223372036854775807 | 9223372036854775807 | | | 0 |
| dummy | | 0 | 0 | |
| 0 | 0 | 0 | | |
| dummy | | 0 | 0 | |
| 0 | 0 | 0 | | |
| dummy | | 0 | 0 | |
| 0 | 0 | 0 | | |
| numiptent | | 20 | 20 | |
| 9223372036854775807 | 9223372036854775807 | | | 0 |

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
 Posted by [Rene Dokbua](#) on Tue, 10 Jul 2012 14:40:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

No takers for this one?

If I missed to provide any important information please let me know. The issue happens regularly on several hardware nodes so if I missed anything I can check it next time it happens.

On Wed, Jul 4, 2012 at 4:16 PM, Rene C. <openvz@dokbua.com> wrote:

> Today I again had a VE that went up to a relative high load for no
 > apparent reason.
 >
 > Below are the details for the hardware node, followed by the high-load
 > container.
 >
 > I realize it's not the latest kernel, but a reboot takes half an hour
 > (from first VE goes down to last VE is back up, assuming everything goes
 > well and no FSCK is forced) so we only reboot into new kernels when there
 > is a really serious reason for it or the server crashes - but I don't see
 > anything in the kernel updates since our current kernel that would address
 > this issue anyway.
 >
 > Why does the load in this container suddenly go up like that? Websites
 > hosted by the container becomes very sluggish, so it is a real problem.
 >
 > It isn't just a problem with this container - or even this hardware node
 > for that reason, I occasionally see it with containers on other hardware
 > nodes as well. One idea I brought up before was that perhaps it's the file
 > system journal, as suggested in <http://wiki.openvz.org/Ploop/Why> - but I
 > think that would affect all containers on that file system, not just a
 > single container?
 >
 > --- HARDWARE NODE ---


```

>
> # uname -a
> Linux server15.hardwarenode.com 2.6.32-042stab049.6 #1 SMP Mon Feb 6
> 19:17:43 MSK 2012 x86_64 x86_64 x86_64 GNU/Linux
>
> # rpm -q sl-release
> sl-release-6.1-2.x86_64
>
> # top -cbn1 | head -17
> top - 21:00:02 up 123 days, 15:31, 1 user, load average: 0.97, 2.70, 2.37
> Tasks: 886 total, 6 running, 880 sleeping, 0 stopped, 0 zombie
> Cpu(s): 8.4%us, 1.7%sy, 0.0%ni, 86.3%id, 3.5%wa, 0.0%hi, 0.1%si,
> 0.0%st
> Mem: 16420716k total, 15566264k used, 854452k free, 1477372k buffers
> Swap: 16777184k total, 623672k used, 16153512k free, 4578176k cached
>
>  PID USER    PR  NI  VIRT  RES  SHR  S %CPU %MEM  TIME+  COMMAND
> 94153 27      20   0 164m  41m 3392  S 150.9 0.3 50575:37
> /usr/libexec/mys
> 9178 27      20   0 159m  29m 3000  S 72.6 0.2 1284:50
> /usr/libexec/mysq
> 567031 apache  20   0 40296 15m 3588  S 17.2 0.1 0:00.09
> /usr/sbin/httpd
> 567382 root    20   0 15672 1820 864   R 5.7 0.0 0:00.04 top -cbn1
> 38 root    20   0 0 0 0 S 1.9 0.0 2:55.25 [events/3]
> 41 root    20   0 0 0 0 S 1.9 0.0 0:29.00 [events/6]
> 566362 apache  20   0 43240 19m 4448  R 1.9 0.1 0:01.04
> /usr/sbin/httpd
> 566857 apache  20   0 55248 11m 3456  R 1.9 0.1 0:00.05
> /usr/sbin/httpd
> 566918 apache  20   0 42596 17m 3704  S 1.9 0.1 0:00.15
> /usr/sbin/httpd
> 567033 apache  20   0 39784 14m 3468  S 1.9 0.1 0:00.01
> /usr/sbin/httpd
>
> # vzlist -o ctid,laverage
>  CTID  LAVERAGE
> 1501 0.00/0.05/0.02
> 1502 0.00/0.00/0.00
> 1503 0.08/0.03/0.01
> 1504 0.00/0.00/0.00
> 1505 8.29/6.04/3.67
> 1506 27.11/16.97/7.89
> 1507 0.00/0.00/0.00
> 1508 0.19/0.06/0.01
> 1509 0.07/0.03/0.00
> 1510 0.02/0.02/0.00
> 1512 0.00/0.00/0.00

```



```

> 1514 0.00/0.00/0.00
>
> # iostat -xN
> Linux 2.6.32-042stab049.6 (server15.hardwarenode.com) 07/03/12
> _x86_64_ (8 CPU)
>
> avg-cpu: %user %nice %system %iowait %steal %idle
> 8.41 0.04 1.75 3.51 0.00 86.28
>
> Device: rrqm/s wrqm/s r/s w/s rsec/s wsec/s avgrq-sz
> avgqu-sz await svctm %util
> sdd 0.76 56.58 0.59 0.59 20.27 457.28 402.66
> 0.25 211.66 4.03 0.48
> sdc 1.72 27.94 17.20 16.16 887.30 336.18 36.68
> 0.02 12.71 5.23 17.45
> sdb 1.65 27.79 19.48 12.95 975.43 318.64 39.91
> 0.09 15.22 3.77 12.23
> sda 0.01 0.16 0.10 0.24 1.95 2.79 13.79
> 0.00 7.06 4.16 0.14
> vg01-swap 0.00 0.00 0.00 0.00 0.00 0.00 8.00
> 0.00 3.68 2.22 0.00
> vg01-root 0.00 0.00 0.11 0.35 1.94 2.78 10.30
> 0.02 38.30 3.12 0.14
> vg04-swap 0.00 0.00 1.30 0.22 10.41 1.80 8.00
> 0.01 9.28 1.44 0.22
> vg04-vz 0.00 0.00 0.05 56.94 9.86 455.49 8.17
> 0.01 0.18 0.05 0.27
> vg03-swap 0.00 0.00 0.00 0.00 0.00 0.00 8.00
> 0.00 6.72 1.10 0.00
> vg03-vz 0.00 0.00 18.98 42.41 887.30 336.18 19.93
> 0.39 6.33 2.84 17.45
> vg02-swap 0.00 0.00 0.00 0.00 0.00 0.00 8.00
> 0.00 7.03 0.89 0.00
> vg02-vz 0.00 0.00 21.19 39.91 975.43 318.64 21.18
> 0.15 8.99 2.00 12.23
> vg01-vz 0.00 0.00 0.00 0.00 0.00 0.00 7.98
> 0.00 17.73 17.73 0.00
>
> --- CONTAINER ---
>
> # top -cbn1 | head -100
> top - 21:00:04 up 123 days, 15:25, 0 users, load average: 27.11, 16.97,
> 7.89
> Tasks: 86 total, 2 running, 84 sleeping, 0 stopped, 0 zombie
> Cpu(s): 1.4%us, 0.2%sy, 0.0%ni, 98.1%id, 0.1%wa, 0.0%hi, 0.0%si,
> 0.2%st
> Mem: 655360k total, 316328k used, 339032k free, 0k buffers
> Swap: 1310720k total, 68380k used, 1242340k free, 58268k cached

```

```

>
> PID USER   PR NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
> 916 mysql   20  0 159m 29m 3000 S 79.3  4.6   1284:51
> /usr/libexec/mysqld
>   1 root    20  0 2156  92  64 S  0.0  0.0   0:36.50 init [3]
>   2 root    20  0   0   0   0 S  0.0  0.0   0:00.00
> [kthreadd/1506]
>   3 root    20  0   0   0   0 S  0.0  0.0   0:00.00 [khelper/1506]
>  97 root    16 -4 2244   8   4 S  0.0  0.0   0:00.00 /sbin/udevd -d
> 634 root    20  0 1812 212 136 S  0.0  0.0   2:39.88 syslogd -m 0
> 667 root    20  0 7180 268 168 S  0.0  0.0   1:01.55 /usr/sbin/sshd
> 676 root    20  0 2832 392 304 S  0.0  0.1   0:15.13 xinetd
> -stayalive -
> 690 root    20  0 6040 124  72 S  0.0  0.0   0:02.45
> /usr/lib/courier-im
> 693 root    20  0 4872 252 200 S  0.0  0.0   0:01.94
> /usr/sbin/courierlo
> 701 root    20  0 6040 124  72 S  0.0  0.0   0:06.34
> /usr/lib/courier-im
> 703 root    20  0 4872 256 200 S  0.0  0.0   0:03.09
> /usr/sbin/courierlo
> 709 root    20  0 6040 128  72 S  0.0  0.0   0:18.15
> /usr/lib/courier-im
> 711 root    20  0 4872 256 200 S  0.0  0.0   0:09.15
> /usr/sbin/courierlo
> 718 root    20  0 6040 124  72 S  0.0  0.0   0:05.68
> /usr/lib/courier-im
> 720 root    20  0 4872 252 200 S  0.0  0.0   0:02.54
> /usr/sbin/courierlo
> 730 qmails   20  0 1796 224 144 S  0.0  0.0   1:27.21 qmail-send
> 732 qmail    20  0 1752 244 192 S  0.0  0.0   0:22.64 splogger qmail
> 733 root    20  0 1780 140  64 S  0.0  0.0   0:07.85 qmail-lspawn
> | /usr
> 734 qmailr   20  0 1776 148  76 S  0.0  0.0   0:14.07 qmail-rspawn
> 735 qmailq   20  0 1748 104  68 S  0.0  0.0   0:14.01 qmail-clean
> 781 root    20  0 51880 4364 196 S  0.0  0.7   1:35.02
> /usr/sbin/httpd
> 828 named    20  0 44104 5708 1112 S  0.0  0.9  10:10.53
> /usr/sbin/named -u
> 866 root    20  0 3708   8   4 S  0.0  0.0   0:00.00 /bin/sh
> /usr/bin/my
> 981 root    20  0 33912 3756 916 S  0.0  0.6  10:55.30
> /usr/bin/spamd --us
> 1107 xfs     20  0 3392   72  40 S  0.0  0.0   0:00.09 xfs -droppriv
> -daem
> 1115 root    20  0 5672   8   4 S  0.0  0.0   0:00.00
> /usr/sbin/saslauthd
> 1116 root    20  0 5672   8   4 S  0.0  0.0   0:00.00

```

```

> /usr/sbin/saslauthd
> 1122 root    20  0 22992 1868 1084 S  0.0  0.3  2:09.79
> /usr/bin/sw-engine
> 1123 root    20  0 27328 1508 1160 S  0.0  0.2  6:06.30
> /usr/local/psa/admi
> 7251 root    20  0 4488  192  136 S  0.0  0.0  0:22.85 crond
> 9463 apache  20  0 59184 14m 4356 S  0.0  2.3  0:05.10
> /usr/sbin/httpd
> 10512 apache 20  0 42316 2504  84 S  0.0  0.4  0:00.91
> /usr/sbin/httpd
> 12090 apache 20  0 56964 14m 4492 S  0.0  2.2  0:04.48
> /usr/sbin/httpd
> 12682 apache 20  0 61060 17m 4516 S  0.0  2.7  0:02.45
> /usr/sbin/httpd
> 13870 sw-cp-se 20  0 7852 1932  16 S  0.0  0.3  1:19.03
> /usr/sbin/sw-cp-ser
> 17443 apache 20  0 62416 17m 4436 S  0.0  2.7  0:05.27
> /usr/sbin/httpd
> 17461 apache 20  0 52788 10m 4480 S  0.0  1.6  0:02.24
> /usr/sbin/httpd
> 20430 apache 20  0 62164 17m 4356 S  0.0  2.7  0:04.25
> /usr/sbin/httpd
> 23539 popuser 20  0 37612 25m 2328 S  0.0  3.9  0:01.50 spamd child
> 23924 apache 20  0 58004 15m 5536 S  0.0  2.4  0:01.56
> /usr/sbin/httpd
> 26361 apache 20  0 54496 11m 3864 S  0.0  1.8  0:01.35
> /usr/sbin/httpd
> 26366 apache 20  0 52944 9.8m 3892 S  0.0  1.5  0:01.45
> /usr/sbin/httpd
> 26964 apache 20  0 59184 14m 4316 S  0.0  2.3  0:07.26
> /usr/sbin/httpd
> 27096 apache 20  0 53728 10m 3868 S  0.0  1.6  0:00.33
> /usr/sbin/httpd
> 27102 apache 20  0 54736 11m 3780 S  0.0  1.8  0:00.15
> /usr/sbin/httpd
> 27103 apache 20  0 54480 11m 3784 S  0.0  1.7  0:00.11
> /usr/sbin/httpd
> 27115 apache 20  0 57064 12m 3816 S  0.0  2.0  0:00.32
> /usr/sbin/httpd
> 27118 apache 20  0 53728 10m 3884 S  0.0  1.6  0:01.21
> /usr/sbin/httpd
> 27120 apache 20  0 52184 8376 3120 S  0.0  1.3  0:00.00
> /usr/sbin/httpd
> 27129 apache 20  0 52168 8072 2960 S  0.0  1.2  0:00.00
> /usr/sbin/httpd
> 27139 apache 20  0 53304 9840 3744 S  0.0  1.5  0:01.08
> /usr/sbin/httpd
> 27140 apache 20  0 53000 9.8m 3832 S  0.0  1.5  0:00.66

```

```

> /usr/sbin/httpd
> 27144 apache 20 0 52168 8072 2960 S 0.0 1.2 0:00.00
> /usr/sbin/httpd
> 27147 apache 20 0 53252 12m 5536 S 0.0 1.9 0:00.50
> /usr/sbin/httpd
> 27149 apache 20 0 52980 9924 3740 S 0.0 1.5 0:00.17
> /usr/sbin/httpd
> 27153 apache 20 0 53728 10m 3836 S 0.0 1.6 0:00.49
> /usr/sbin/httpd
> 27164 apache 20 0 55224 11m 3812 S 0.0 1.9 0:00.47
> /usr/sbin/httpd
> 27171 apache 20 0 52916 9776 3708 S 0.0 1.5 0:00.16
> /usr/sbin/httpd
> 27172 apache 20 0 52916 9452 3436 S 0.0 1.4 0:00.17
> /usr/sbin/httpd
> 27173 apache 20 0 55340 11m 3720 S 0.0 1.8 0:00.08
> /usr/sbin/httpd
> 27179 apache 20 0 52020 7764 2716 S 0.0 1.2 0:00.00
> /usr/sbin/httpd
> 27182 apache 20 0 52020 7764 2716 S 0.0 1.2 0:00.00
> /usr/sbin/httpd
> 27185 apache 20 0 55224 11m 3824 S 0.0 1.9 0:00.30
> /usr/sbin/httpd
> 27186 apache 20 0 53788 10m 3840 S 0.0 1.7 0:00.11
> /usr/sbin/httpd
> 27187 apache 20 0 52916 9448 3436 S 0.0 1.4 0:00.08
> /usr/sbin/httpd
> 27188 apache 20 0 54628 10m 3504 S 0.0 1.7 0:00.05
> /usr/sbin/httpd
> 27196 apache 20 0 53728 10m 3572 S 0.0 1.6 0:00.36
> /usr/sbin/httpd
> 27200 apache 20 0 54628 11m 3796 S 0.0 1.7 0:00.05
> /usr/sbin/httpd
> 27202 apache 20 0 54480 11m 3796 S 0.0 1.7 0:00.10
> /usr/sbin/httpd
> 27204 apache 20 0 53992 10m 3544 S 0.0 1.6 0:00.09
> /usr/sbin/httpd
> 27207 apache 20 0 52168 8084 2960 S 0.0 1.2 0:00.00
> /usr/sbin/httpd
> 27213 apache 20 0 52020 6464 1788 S 0.0 1.0 0:00.00
> /usr/sbin/httpd
> 27214 apache 20 0 54216 10m 3516 S 0.0 1.6 0:00.05
> /usr/sbin/httpd
> 27215 apache 20 0 52020 6456 1788 S 0.0 1.0 0:00.00
> /usr/sbin/httpd
> 27216 apache 20 0 52020 7860 2804 S 0.0 1.2 0:00.00
> /usr/sbin/httpd
> 27218 root 20 0 9400 1900 1408 S 0.0 0.3 0:00.00 crond

```

```

> 27219 root    20   0 2492 956 848 S 0.0 0.1 0:00.00 /bin/sh -c
> /usr/loc
> 27220 root    20   0 2496 1052 920 S 0.0 0.2 0:00.00 /bin/sh
> /usr/local/
> 27233 root    20   0 2540 1016 892 S 0.0 0.2 0:00.00 /bin/bash -c
> top -c
> 27234 root    20   0 2284 952 724 R 0.0 0.1 0:00.00 top -cbn1
> 27235 root    20   0 1756 420 352 S 0.0 0.1 0:00.00 head -100
> 27247 root    20   0 2496 452 320 S 0.0 0.1 0:00.00 /bin/sh
> /usr/local/
> 27248 root    20   0 8280 1504 1120 R 0.0 0.2 0:00.00
> /usr/bin/mysql -uad
> 27249 root    20   0 1800 448 376 S 0.0 0.1 0:00.00 sed -e 1d
> 27250 root    20   0 2240 640 540 S 0.0 0.1 0:00.00 awk
> {printf("%s", $
>
> # netstat -ptan | grep ESTABLISHED
> tcp      0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:77.87.207.166:21863 ESTABLISHED 23924/httpd
> tcp      0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:95.165.204.26:62259 ESTABLISHED 27144/httpd
> tcp      0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:193.151.105.100:4059 ESTABLISHED 27200/httpd
> tcp      0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:109.169.207.68:50087 ESTABLISHED 27185/httpd
> tcp      0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:31.131.70.135:57017 ESTABLISHED 27179/httpd
> tcp      0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:95.165.204.26:62220 ESTABLISHED 27103/httpd
> tcp      0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:188.134.61.1:60732
> ESTABLISHED 27215/httpd
> tcp      0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:193.151.105.100:4112 ESTABLISHED 26964/httpd
> tcp      0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:109.169.207.68:50043 ESTABLISHED 27164/httpd
> tcp      0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:31.131.70.135:56976 ESTABLISHED 27153/httpd
>
> # cat /proc/user_beancounters
> Version: 2.5
>      uid resource          held          maxheld
>      barrier      limit      failcnt
> 1506: kmemsize      27735306      179081216
> 304087040      335544320      0
>      lockedpages      0      0
> 81920      81920      0
>      privvmpages      393683      430195
> 9223372036854775807 9223372036854775807      0
>      shmpages      823      21639
> 9223372036854775807 9223372036854775807      0
>      dummy      0      0
> 0      0      0
>      numproc      128      204
> 9223372036854775807 9223372036854775807      0
>      physpages      79702      163840
> 0      163840      0
>      vmguarpages      0      0

```

```

> 0 9223372036854775807 0
> oomguarpages 74734 75707
> 0 9223372036854775807 0
> numtcpsock 59 153
> 9223372036854775807 9223372036854775807 0
> numflock 46 62
> 9223372036854775807 9223372036854775807 0
> numpty 0 1
> 9223372036854775807 9223372036854775807 0
> numsiginfo 0 33
> 9223372036854775807 9223372036854775807 0
> tcpsndbuf 1037680 11426176
> 9223372036854775807 9223372036854775807 0
> tcprcvbuf 966656 2867584
> 9223372036854775807 9223372036854775807 0
> othersockbuf 53824 838688
> 9223372036854775807 9223372036854775807 0
> dgramrcvbuf 0 502224
> 9223372036854775807 9223372036854775807 0
> numothersock 114 273
> 9223372036854775807 9223372036854775807 0
> dcachesize 10070617 167772160
> 150994944 167772160 0
> numfile 1634 1865
> 9223372036854775807 9223372036854775807 0
> dummy 0 0
> 0 0 0
> dummy 0 0
> 0 0 0
> dummy 0 0
> 0 0 0
> numiptent 20 20
> 9223372036854775807 9223372036854775807 0
>

```

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
 Posted by [Kirill Korotaev](#) on Tue, 10 Jul 2012 16:34:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

I can take a look if you give me access to node.
 If agree - send it privately, w/o users@ on CC.

Kirill

On Jul 10, 2012, at 18:40 , Rene C. wrote:

No takers for this one?

If I missed to provide any important information please let me know. The issue happens regularly on several hardware nodes so if I missed anything I can check it next time it happens.

On Wed, Jul 4, 2012 at 4:16 PM, Rene C. <openvz@dokbua.com<mailto:openvz@dokbua.com>> wrote:

Today I again had a VE that went up to a relative high load for no apparent reason.

Below are the details for the hardware node, followed by the high-load container.

I realize it's not the latest kernel, but a reboot takes half an hour (from first VE goes down to last VE is back up, assuming everything goes well and no FSCK is forced) so we only reboot into new kernels when there is a really serious reason for it or the server crashes - but I don't see anything in the kernel updates since our current kernel that would address this issue anyway.

Why does the load in this container suddenly go up like that? Websites hosted by the container becomes very sluggish, so it is a real problem.

It isn't just a problem with this container - or even this hardware node for that reason, I occasionally see it with containers on other hardware nodes as well. One idea I brought up before was that perhaps it's the file system journal, as suggested in <http://wiki.openvz.org/Ploop/Why> - but I think that would affect all containers on that file system, not just a single container?

--- HARDWARE NODE ---

```
# uname -a
```

```
Linux server15.hardwarenode.com<http://server15.hardwarenode.com/> 2.6.32-042stab049.6 #1  
SMP Mon Feb 6 19:17:43 MSK 2012 x86_64 x86_64 x86_64 GNU/Linux
```

```
# rpm -q sl-release
```

```
sl-release-6.1-2.x86_64
```

```
# top -cbn1 | head -17
```

```
top - 21:00:02 up 123 days, 15:31, 1 user, load average: 0.97, 2.70, 2.37
```

```
Tasks: 886 total, 6 running, 880 sleeping, 0 stopped, 0 zombie
```

```
Cpu(s): 8.4%us, 1.7%sy, 0.0%ni, 86.3%id, 3.5%wa, 0.0%hi, 0.1%si, 0.0%st
```

```
Mem: 16420716k total, 15566264k used, 854452k free, 1477372k buffers
```

```
Swap: 16777184k total, 623672k used, 16153512k free, 4578176k cached
```

| PID | USER | PR | NI | VIRT | RES | SHR | S | %CPU | %MEM | TIME+ | COMMAND |
|--------|--------|----|----|-------|------|------|---|-------|------|----------|-------------------|
| 94153 | 27 | 20 | 0 | 164m | 41m | 3392 | S | 150.9 | 0.3 | 50575:37 | /usr/libexec/mys |
| 9178 | 27 | 20 | 0 | 159m | 29m | 3000 | S | 72.6 | 0.2 | 1284:50 | /usr/libexec/mysq |
| 567031 | apache | 20 | 0 | 40296 | 15m | 3588 | S | 17.2 | 0.1 | 0:00.09 | /usr/sbin/httpd |
| 567382 | root | 20 | 0 | 15672 | 1820 | 864 | R | 5.7 | 0.0 | 0:00.04 | top -cbn1 |
| 38 | root | 20 | 0 | 0 | 0 | 0 | S | 1.9 | 0.0 | 2:55.25 | [events/3] |
| 41 | root | 20 | 0 | 0 | 0 | 0 | S | 1.9 | 0.0 | 0:29.00 | [events/6] |
| 566362 | apache | 20 | 0 | 43240 | 19m | 4448 | R | 1.9 | 0.1 | 0:01.04 | /usr/sbin/httpd |


```
566857 apache 20 0 55248 11m 3456 R 1.9 0.1 0:00.05 /usr/sbin/httpd
566918 apache 20 0 42596 17m 3704 S 1.9 0.1 0:00.15 /usr/sbin/httpd
567033 apache 20 0 39784 14m 3468 S 1.9 0.1 0:00.01 /usr/sbin/httpd
```

```
# vzlist -o ctid,lavrage
CTID    LAVERAGE
1501 0.00/0.05/0.02
1502 0.00/0.00/0.00
1503 0.08/0.03/0.01
1504 0.00/0.00/0.00
1505 8.29/6.04/3.67
1506 27.11/16.97/7.89
1507 0.00/0.00/0.00
1508 0.19/0.06/0.01
1509 0.07/0.03/0.00
1510 0.02/0.02/0.00
1512 0.00/0.00/0.00
1514 0.00/0.00/0.00
```

```
# iostat -xN
Linux 2.6.32-042stab049.6 (server15.hardwarenode.com<http://server15.hardwarenode.com/>)
07/03/12    _x86_64_    (8 CPU)
```

```
avg-cpu:  %user  %nice %system %iowait  %steal   %idle
           8.41   0.04   1.75   3.51   0.00  86.28
```

| Device: | rrqm/s | wrqm/s | r/s | w/s | rsec/s | wsec/s | avgrq-sz | avgqu-sz | await | svctm | %util |
|-----------|--------|--------|-------|-------|--------|--------|----------|----------|--------|-------|-------|
| sdd | 0.76 | 56.58 | 0.59 | 0.59 | 20.27 | 457.28 | 402.66 | 0.25 | 211.66 | 4.03 | 0.48 |
| sdc | 1.72 | 27.94 | 17.20 | 16.16 | 887.30 | 336.18 | 36.68 | 0.02 | 12.71 | 5.23 | 17.45 |
| sdb | 1.65 | 27.79 | 19.48 | 12.95 | 975.43 | 318.64 | 39.91 | 0.09 | 15.22 | 3.77 | 12.23 |
| sda | 0.01 | 0.16 | 0.10 | 0.24 | 1.95 | 2.79 | 13.79 | 0.00 | 7.06 | 4.16 | 0.14 |
| vg01-swap | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 8.00 | 0.00 | 3.68 | 2.22 | 0.00 |
| vg01-root | 0.00 | 0.00 | 0.11 | 0.35 | 1.94 | 2.78 | 10.30 | 0.02 | 38.30 | 3.12 | 0.14 |
| vg04-swap | 0.00 | 0.00 | 1.30 | 0.22 | 10.41 | 1.80 | 8.00 | 0.01 | 9.28 | 1.44 | 0.22 |
| vg04-vz | 0.00 | 0.00 | 0.05 | 56.94 | 9.86 | 455.49 | 8.17 | 0.01 | 0.18 | 0.05 | 0.27 |
| vg03-swap | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 8.00 | 0.00 | 6.72 | 1.10 | 0.00 |
| vg03-vz | 0.00 | 0.00 | 18.98 | 42.41 | 887.30 | 336.18 | 19.93 | 0.39 | 6.33 | 2.84 | 17.45 |
| vg02-swap | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 8.00 | 0.00 | 7.03 | 0.89 | 0.00 |
| vg02-vz | 0.00 | 0.00 | 21.19 | 39.91 | 975.43 | 318.64 | 21.18 | 0.15 | 8.99 | 2.00 | 12.23 |
| vg01-vz | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 7.98 | 0.00 | 17.73 | 17.73 | 0.00 |

```
--- CONTAINER ---
```

```
# top -cbn1 | head -100
top - 21:00:04 up 123 days, 15:25, 0 users, load average: 27.11, 16.97, 7.89
Tasks: 86 total, 2 running, 84 sleeping, 0 stopped, 0 zombie
Cpu(s): 1.4%us, 0.2%sy, 0.0%ni, 98.1%id, 0.1%wa, 0.0%hi, 0.0%si, 0.2%st
Mem: 655360k total, 316328k used, 339032k free, 0k buffers
```


Swap: 1310720k total, 68380k used, 1242340k free, 58268k cached

| PID | USER | PR | NI | VIRT | RES | SHR | S | %CPU | %MEM | TIME+ | COMMAND |
|-------|----------|----|----|-------|------|------|---|------|------|----------|---------------------|
| 916 | mysql | 20 | 0 | 159m | 29m | 3000 | S | 79.3 | 4.6 | 1284:51 | /usr/libexec/mysqld |
| 1 | root | 20 | 0 | 2156 | 92 | 64 | S | 0.0 | 0.0 | 0:36.50 | init [3] |
| 2 | root | 20 | 0 | 0 | 0 | 0 | S | 0.0 | 0.0 | 0:00.00 | [kthreadd/1506] |
| 3 | root | 20 | 0 | 0 | 0 | 0 | S | 0.0 | 0.0 | 0:00.00 | [khelper/1506] |
| 97 | root | 16 | -4 | 2244 | 8 | 4 | S | 0.0 | 0.0 | 0:00.00 | /sbin/udevd -d |
| 634 | root | 20 | 0 | 1812 | 212 | 136 | S | 0.0 | 0.0 | 2:39.88 | syslogd -m 0 |
| 667 | root | 20 | 0 | 7180 | 268 | 168 | S | 0.0 | 0.0 | 1:01.55 | /usr/sbin/sshd |
| 676 | root | 20 | 0 | 2832 | 392 | 304 | S | 0.0 | 0.1 | 0:15.13 | xinetd -stayalive - |
| 690 | root | 20 | 0 | 6040 | 124 | 72 | S | 0.0 | 0.0 | 0:02.45 | /usr/lib/courier-im |
| 693 | root | 20 | 0 | 4872 | 252 | 200 | S | 0.0 | 0.0 | 0:01.94 | /usr/sbin/courierlo |
| 701 | root | 20 | 0 | 6040 | 124 | 72 | S | 0.0 | 0.0 | 0:06.34 | /usr/lib/courier-im |
| 703 | root | 20 | 0 | 4872 | 256 | 200 | S | 0.0 | 0.0 | 0:03.09 | /usr/sbin/courierlo |
| 709 | root | 20 | 0 | 6040 | 128 | 72 | S | 0.0 | 0.0 | 0:18.15 | /usr/lib/courier-im |
| 711 | root | 20 | 0 | 4872 | 256 | 200 | S | 0.0 | 0.0 | 0:09.15 | /usr/sbin/courierlo |
| 718 | root | 20 | 0 | 6040 | 124 | 72 | S | 0.0 | 0.0 | 0:05.68 | /usr/lib/courier-im |
| 720 | root | 20 | 0 | 4872 | 252 | 200 | S | 0.0 | 0.0 | 0:02.54 | /usr/sbin/courierlo |
| 730 | qmails | 20 | 0 | 1796 | 224 | 144 | S | 0.0 | 0.0 | 1:27.21 | qmail-send |
| 732 | qmaill | 20 | 0 | 1752 | 244 | 192 | S | 0.0 | 0.0 | 0:22.64 | splogger qmail |
| 733 | root | 20 | 0 | 1780 | 140 | 64 | S | 0.0 | 0.0 | 0:07.85 | qmail-lspawn /usr |
| 734 | qmailr | 20 | 0 | 1776 | 148 | 76 | S | 0.0 | 0.0 | 0:14.07 | qmail-rspawn |
| 735 | qmailq | 20 | 0 | 1748 | 104 | 68 | S | 0.0 | 0.0 | 0:14.01 | qmail-clean |
| 781 | root | 20 | 0 | 51880 | 4364 | 196 | S | 0.0 | 0.7 | 1:35.02 | /usr/sbin/httpd |
| 828 | named | 20 | 0 | 44104 | 5708 | 1112 | S | 0.0 | 0.9 | 10:10.53 | /usr/sbin/named -u |
| 866 | root | 20 | 0 | 3708 | 8 | 4 | S | 0.0 | 0.0 | 0:00.00 | /bin/sh /usr/bin/my |
| 981 | root | 20 | 0 | 33912 | 3756 | 916 | S | 0.0 | 0.6 | 10:55.30 | /usr/bin/spamd --us |
| 1107 | xfs | 20 | 0 | 3392 | 72 | 40 | S | 0.0 | 0.0 | 0:00.09 | xfs -droppriv -daem |
| 1115 | root | 20 | 0 | 5672 | 8 | 4 | S | 0.0 | 0.0 | 0:00.00 | /usr/sbin/saslauthd |
| 1116 | root | 20 | 0 | 5672 | 8 | 4 | S | 0.0 | 0.0 | 0:00.00 | /usr/sbin/saslauthd |
| 1122 | root | 20 | 0 | 22992 | 1868 | 1084 | S | 0.0 | 0.3 | 2:09.79 | /usr/bin/sw-engine |
| 1123 | root | 20 | 0 | 27328 | 1508 | 1160 | S | 0.0 | 0.2 | 6:06.30 | /usr/local/psa/admi |
| 7251 | root | 20 | 0 | 4488 | 192 | 136 | S | 0.0 | 0.0 | 0:22.85 | crond |
| 9463 | apache | 20 | 0 | 59184 | 14m | 4356 | S | 0.0 | 2.3 | 0:05.10 | /usr/sbin/httpd |
| 10512 | apache | 20 | 0 | 42316 | 2504 | 84 | S | 0.0 | 0.4 | 0:00.91 | /usr/sbin/httpd |
| 12090 | apache | 20 | 0 | 56964 | 14m | 4492 | S | 0.0 | 2.2 | 0:04.48 | /usr/sbin/httpd |
| 12682 | apache | 20 | 0 | 61060 | 17m | 4516 | S | 0.0 | 2.7 | 0:02.45 | /usr/sbin/httpd |
| 13870 | sw-cp-se | 20 | 0 | 7852 | 1932 | 16 | S | 0.0 | 0.3 | 1:19.03 | /usr/sbin/sw-cp-ser |
| 17443 | apache | 20 | 0 | 62416 | 17m | 4436 | S | 0.0 | 2.7 | 0:05.27 | /usr/sbin/httpd |
| 17461 | apache | 20 | 0 | 52788 | 10m | 4480 | S | 0.0 | 1.6 | 0:02.24 | /usr/sbin/httpd |
| 20430 | apache | 20 | 0 | 62164 | 17m | 4356 | S | 0.0 | 2.7 | 0:04.25 | /usr/sbin/httpd |
| 23539 | popuser | 20 | 0 | 37612 | 25m | 2328 | S | 0.0 | 3.9 | 0:01.50 | spamd child |
| 23924 | apache | 20 | 0 | 58004 | 15m | 5536 | S | 0.0 | 2.4 | 0:01.56 | /usr/sbin/httpd |
| 26361 | apache | 20 | 0 | 54496 | 11m | 3864 | S | 0.0 | 1.8 | 0:01.35 | /usr/sbin/httpd |
| 26366 | apache | 20 | 0 | 52944 | 9.8m | 3892 | S | 0.0 | 1.5 | 0:01.45 | /usr/sbin/httpd |
| 26964 | apache | 20 | 0 | 59184 | 14m | 4316 | S | 0.0 | 2.3 | 0:07.26 | /usr/sbin/httpd |
| 27096 | apache | 20 | 0 | 53728 | 10m | 3868 | S | 0.0 | 1.6 | 0:00.33 | /usr/sbin/httpd |

```

27102 apache 20 0 54736 11m 3780 S 0.0 1.8 0:00.15 /usr/sbin/httpd
27103 apache 20 0 54480 11m 3784 S 0.0 1.7 0:00.11 /usr/sbin/httpd
27115 apache 20 0 57064 12m 3816 S 0.0 2.0 0:00.32 /usr/sbin/httpd
27118 apache 20 0 53728 10m 3884 S 0.0 1.6 0:01.21 /usr/sbin/httpd
27120 apache 20 0 52184 8376 3120 S 0.0 1.3 0:00.00 /usr/sbin/httpd
27129 apache 20 0 52168 8072 2960 S 0.0 1.2 0:00.00 /usr/sbin/httpd
27139 apache 20 0 53304 9840 3744 S 0.0 1.5 0:01.08 /usr/sbin/httpd
27140 apache 20 0 53000 9.8m 3832 S 0.0 1.5 0:00.66 /usr/sbin/httpd
27144 apache 20 0 52168 8072 2960 S 0.0 1.2 0:00.00 /usr/sbin/httpd
27147 apache 20 0 53252 12m 5536 S 0.0 1.9 0:00.50 /usr/sbin/httpd
27149 apache 20 0 52980 9924 3740 S 0.0 1.5 0:00.17 /usr/sbin/httpd
27153 apache 20 0 53728 10m 3836 S 0.0 1.6 0:00.49 /usr/sbin/httpd
27164 apache 20 0 55224 11m 3812 S 0.0 1.9 0:00.47 /usr/sbin/httpd
27171 apache 20 0 52916 9776 3708 S 0.0 1.5 0:00.16 /usr/sbin/httpd
27172 apache 20 0 52916 9452 3436 S 0.0 1.4 0:00.17 /usr/sbin/httpd
27173 apache 20 0 55340 11m 3720 S 0.0 1.8 0:00.08 /usr/sbin/httpd
27179 apache 20 0 52020 7764 2716 S 0.0 1.2 0:00.00 /usr/sbin/httpd
27182 apache 20 0 52020 7764 2716 S 0.0 1.2 0:00.00 /usr/sbin/httpd
27185 apache 20 0 55224 11m 3824 S 0.0 1.9 0:00.30 /usr/sbin/httpd
27186 apache 20 0 53788 10m 3840 S 0.0 1.7 0:00.11 /usr/sbin/httpd
27187 apache 20 0 52916 9448 3436 S 0.0 1.4 0:00.08 /usr/sbin/httpd
27188 apache 20 0 54628 10m 3504 S 0.0 1.7 0:00.05 /usr/sbin/httpd
27196 apache 20 0 53728 10m 3572 S 0.0 1.6 0:00.36 /usr/sbin/httpd
27200 apache 20 0 54628 11m 3796 S 0.0 1.7 0:00.05 /usr/sbin/httpd
27202 apache 20 0 54480 11m 3796 S 0.0 1.7 0:00.10 /usr/sbin/httpd
27204 apache 20 0 53992 10m 3544 S 0.0 1.6 0:00.09 /usr/sbin/httpd
27207 apache 20 0 52168 8084 2960 S 0.0 1.2 0:00.00 /usr/sbin/httpd
27213 apache 20 0 52020 6464 1788 S 0.0 1.0 0:00.00 /usr/sbin/httpd
27214 apache 20 0 54216 10m 3516 S 0.0 1.6 0:00.05 /usr/sbin/httpd
27215 apache 20 0 52020 6456 1788 S 0.0 1.0 0:00.00 /usr/sbin/httpd
27216 apache 20 0 52020 7860 2804 S 0.0 1.2 0:00.00 /usr/sbin/httpd
27218 root 20 0 9400 1900 1408 S 0.0 0.3 0:00.00 crond
27219 root 20 0 2492 956 848 S 0.0 0.1 0:00.00 /bin/sh -c /usr/loc
27220 root 20 0 2496 1052 920 S 0.0 0.2 0:00.00 /bin/sh /usr/local/
27233 root 20 0 2540 1016 892 S 0.0 0.2 0:00.00 /bin/bash -c top -c
27234 root 20 0 2284 952 724 R 0.0 0.1 0:00.00 top -cbn1
27235 root 20 0 1756 420 352 S 0.0 0.1 0:00.00 head -100
27247 root 20 0 2496 452 320 S 0.0 0.1 0:00.00 /bin/sh /usr/local/
27248 root 20 0 8280 1504 1120 R 0.0 0.2 0:00.00 /usr/bin/mysql -uad
27249 root 20 0 1800 448 376 S 0.0 0.1 0:00.00 sed -e 1d
27250 root 20 0 2240 640 540 S 0.0 0.1 0:00.00 awk {printf("%s", $

```

```
# netstat -ptan | grep ESTABLISHED
```

```

tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:77.87.207.166:21863<http://77.87.207.166:21863/>
ESTABLISHED 23924/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:95.165.204.26:62259<http://95.165.204.26:62259/>
ESTABLISHED 27144/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:193.151.105.100:4059<http://193.151.105.100:4059/>

```

```

ESTABLISHED 27200/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:109.169.207.68:50087<http://109.169.207.68:50087/>
ESTABLISHED 27185/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:31.131.70.135:57017<http://31.131.70.135:57017/>
ESTABLISHED 27179/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:95.165.204.26:62220<http://95.165.204.26:62220/>
ESTABLISHED 27103/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:188.134.61.1:60732<http://188.134.61.1:60732/>
ESTABLISHED 27215/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:193.151.105.100:4112<http://193.151.105.100:4112/>
ESTABLISHED 26964/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:109.169.207.68:50043<http://109.169.207.68:50043/>
ESTABLISHED 27164/httpd
tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:31.131.70.135:56976<http://31.131.70.135:56976/>
ESTABLISHED 27153/httpd

```

```
# cat /proc/user_beancounters
```

```
Version: 2.5
```

| uid | resource | held | maxheld | barrier | limit | failcnt |
|---------------------|--------------|----------|-----------|---------------------|---------------------|-----------|
| 1506: | kmemsize | 27735306 | 179081216 | 304087040 | | 335544320 |
| | 0 | | | | | |
| | lockedpages | 0 | 0 | 81920 | 81920 | 0 |
| | privvmpages | 393683 | 430195 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | shmpages | 823 | 21639 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | dummy | 0 | 0 | 0 | 0 | 0 |
| | numproc | 128 | 204 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | physpages | 79702 | 163840 | 0 | 163840 | |
| 0 | | | | | | |
| | vmguarpages | 0 | 0 | 0 | 9223372036854775807 | |
| 0 | | | | | | |
| | oomguarpages | 74734 | 75707 | 0 | 9223372036854775807 | |
| | 0 | | | | | |
| | numtcpsock | 59 | 153 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | numflock | 46 | 62 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | numpty | 0 | 1 | 9223372036854775807 | 9223372036854775807 | |
| 0 | | | | | | |
| | numsiginfo | 0 | 33 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | tcpsndbuf | 1037680 | 11426176 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | tcprcvbuf | 966656 | 2867584 | 9223372036854775807 | | |
| 9223372036854775807 | | 0 | | | | |
| | othersockbuf | 53824 | 838688 | 9223372036854775807 | | |

| | | | | | |
|---------------------|----------|-----------|---------------------|-----------|---|
| 9223372036854775807 | 0 | | | | |
| dgramrcvbuf | 0 | 502224 | 9223372036854775807 | | |
| 9223372036854775807 | 0 | | | | |
| numothersock | 114 | 273 | 9223372036854775807 | | |
| 9223372036854775807 | 0 | | | | |
| dcachesize | 10070617 | 167772160 | 150994944 | 167772160 | |
| 0 | | | | | |
| numfile | 1634 | 1865 | 9223372036854775807 | | |
| 9223372036854775807 | 0 | | | | |
| dummy | 0 | 0 | 0 | 0 | 0 |
| dummy | 0 | 0 | 0 | 0 | 0 |
| dummy | 0 | 0 | 0 | 0 | 0 |
| numiptent | 20 | 20 | 9223372036854775807 | | |
| 9223372036854775807 | 0 | | | | |

<ATT00001.c>

Subject: Re: occasional high loadavg without any noticeable cpu/memory/io load
 Posted by [Rene Dokbua](#) on Tue, 10 Jul 2012 18:36:29 GMT

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

Thanks, that'd be very cool. Access to the hardware node is limited by IP but if you send me (privately if you prefer) the IP address you will use to access I'll add that to allowed hosts and reply with the login coordinates.

Rene

On Tue, Jul 10, 2012 at 11:34 PM, Kirill Korotaev <dev@parallels.com> wrote:

> I can take a look if you give me access to node.
 > If agree - send it privately, w/o users@ on CC.
 >
 > Kirill
 >
 >
 > On Jul 10, 2012, at 18:40 , Rene C. wrote:
 >
 > No takers for this one?
 >
 > If I missed to provide any important information please let me know. The
 > issue happens regularly on several hardware nodes so if I missed anything I
 > can check it next time it happens.
 >
 > On Wed, Jul 4, 2012 at 4:16 PM, Rene C. <openvz@dokbua.com> wrote:
 >
 >> Today I again had a VE that went up to a relative high load for no
 >> apparent reason.

```

>>
>> Below are the details for the hardware node, followed by the high-load
>> container.
>>
>> I realize it's not the latest kernel, but a reboot takes half an hour
>> (from first VE goes down to last VE is back up, assuming everything goes
>> well and no FSCK is forced) so we only reboot into new kernels when there
>> is a really serious reason for it or the server crashes - but I don't see
>> anything in the kernel updates since our current kernel that would address
>> this issue anyway.
>>
>> Why does the load in this container suddenly go up like that? Websites
>> hosted by the container becomes very sluggish, so it is a real problem.
>>
>> It isn't just a problem with this container - or even this hardware node
>> for that reason, I occasionally see it with containers on other hardware
>> nodes as well. One idea I brought up before was that perhaps it's the file
>> system journal, as suggested in http://wiki.openvz.org/Ploop/Why - but I
>> think that would affect all containers on that file system, not just a
>> single container?
>>
>> --- HARDWARE NODE ---
>>
>> # uname -a
>> Linux server15.hardwarenode.com 2.6.32-042stab049.6 #1 SMP Mon Feb 6
>> 19:17:43 MSK 2012 x86_64 x86_64 x86_64 GNU/Linux
>>
>> # rpm -q sl-release
>> sl-release-6.1-2.x86_64
>>
>> # top -cbn1 | head -17
>> top - 21:00:02 up 123 days, 15:31, 1 user, load average: 0.97, 2.70,
>> 2.37
>> Tasks: 886 total, 6 running, 880 sleeping, 0 stopped, 0 zombie
>> Cpu(s): 8.4%us, 1.7%sy, 0.0%ni, 86.3%id, 3.5%wa, 0.0%hi, 0.1%si,
>> 0.0%st
>> Mem: 16420716k total, 15566264k used, 854452k free, 1477372k buffers
>> Swap: 16777184k total, 623672k used, 16153512k free, 4578176k cached
>>
>>  PID USER    PR  NI  VIRT  RES  SHR  S %CPU %MEM  TIME+  COMMAND
>> 94153 27      20   0 164m 41m 3392 S 150.9 0.3 50575:37
>> /usr/libexec/mys
>> 9178 27      20   0 159m 29m 3000 S 72.6 0.2 1284:50
>> /usr/libexec/mysq
>> 567031 apache  20   0 40296 15m 3588 S 17.2 0.1 0:00.09
>> /usr/sbin/httpd
>> 567382 root    20   0 15672 1820 864 R 5.7 0.0 0:00.04 top -cbn1
>> 38 root    20   0 0 0 0 S 1.9 0.0 2:55.25 [events/3]

```

```

>> 41 root 20 0 0 0 0 S 1.9 0.0 0:29.00 [events/6]
>> 566362 apache 20 0 43240 19m 4448 R 1.9 0.1 0:01.04
>> /usr/sbin/httpd
>> 566857 apache 20 0 55248 11m 3456 R 1.9 0.1 0:00.05
>> /usr/sbin/httpd
>> 566918 apache 20 0 42596 17m 3704 S 1.9 0.1 0:00.15
>> /usr/sbin/httpd
>> 567033 apache 20 0 39784 14m 3468 S 1.9 0.1 0:00.01
>> /usr/sbin/httpd
>>
>> # vzlist -o ctid,lverage
>> CTID LAVERAGE
>> 1501 0.00/0.05/0.02
>> 1502 0.00/0.00/0.00
>> 1503 0.08/0.03/0.01
>> 1504 0.00/0.00/0.00
>> 1505 8.29/6.04/3.67
>> 1506 27.11/16.97/7.89
>> 1507 0.00/0.00/0.00
>> 1508 0.19/0.06/0.01
>> 1509 0.07/0.03/0.00
>> 1510 0.02/0.02/0.00
>> 1512 0.00/0.00/0.00
>> 1514 0.00/0.00/0.00
>>
>> # iostat -xN
>> Linux 2.6.32-042stab049.6 (server15.hardwarenode.com) 07/03/12
>> _x86_64_ (8 CPU)
>>
>> avg-cpu: %user %nice %system %iowait %steal %idle
>> 8.41 0.04 1.75 3.51 0.00 86.28
>>
>> Device: rrqm/s wrqm/s r/s w/s rsec/s wsec/s
>> avgrq-sz avgqu-sz await svctm %util
>> sdd 0.76 56.58 0.59 0.59 20.27 457.28
>> 402.66 0.25 211.66 4.03 0.48
>> sdc 1.72 27.94 17.20 16.16 887.30 336.18
>> 36.68 0.02 12.71 5.23 17.45
>> sdb 1.65 27.79 19.48 12.95 975.43 318.64
>> 39.91 0.09 15.22 3.77 12.23
>> sda 0.01 0.16 0.10 0.24 1.95 2.79
>> 13.79 0.00 7.06 4.16 0.14
>> vg01-swap 0.00 0.00 0.00 0.00 0.00 0.00
>> 8.00 0.00 3.68 2.22 0.00
>> vg01-root 0.00 0.00 0.11 0.35 1.94 2.78
>> 10.30 0.02 38.30 3.12 0.14
>> vg04-swap 0.00 0.00 1.30 0.22 10.41 1.80
>> 8.00 0.01 9.28 1.44 0.22

```



```

>> vg04-vz      0.00  0.00  0.05 56.94  9.86 455.49
>> 8.17  0.01  0.18 0.05 0.27
>> vg03-swap     0.00  0.00  0.00  0.00  0.00  0.00
>> 8.00  0.00  6.72 1.10 0.00
>> vg03-vz      0.00  0.00 18.98 42.41 887.30 336.18
>> 19.93  0.39  6.33 2.84 17.45
>> vg02-swap     0.00  0.00  0.00  0.00  0.00  0.00
>> 8.00  0.00  7.03 0.89 0.00
>> vg02-vz      0.00  0.00 21.19 39.91 975.43 318.64
>> 21.18  0.15  8.99 2.00 12.23
>> vg01-vz      0.00  0.00  0.00  0.00  0.00  0.00
>> 7.98  0.00 17.73 17.73 0.00
>>
>> --- CONTAINER ---
>>
>> # top -cbn1 | head -100
>> top - 21:00:04 up 123 days, 15:25, 0 users, load average: 27.11, 16.97,
>> 7.89
>> Tasks: 86 total, 2 running, 84 sleeping, 0 stopped, 0 zombie
>> Cpu(s): 1.4%us, 0.2%sy, 0.0%ni, 98.1%id, 0.1%wa, 0.0%hi, 0.0%si,
>> 0.2%st
>> Mem: 655360k total, 316328k used, 339032k free, 0k buffers
>> Swap: 1310720k total, 68380k used, 1242340k free, 58268k cached
>>
>> PID USER    PR NI VIRT RES  SHR S %CPU %MEM  TIME+ COMMAND
>> 916 mysql   20  0 159m 29m 3000 S 79.3 4.6 1284:51
>> /usr/libexec/mysqld
>>  1 root    20  0 2156 92  64 S 0.0 0.0 0:36.50 init [3]
>>  2 root    20  0  0  0  0 S 0.0 0.0 0:00.00
>> [kthreadd/1506]
>>  3 root    20  0  0  0  0 S 0.0 0.0 0:00.00
>> [khelper/1506]
>> 97 root    16 -4 2244 8  4 S 0.0 0.0 0:00.00 /sbin/udev
>> -d
>> 634 root    20  0 1812 212 136 S 0.0 0.0 2:39.88 syslogd -m 0
>> 667 root    20  0 7180 268 168 S 0.0 0.0 1:01.55
>> /usr/sbin/sshd
>> 676 root    20  0 2832 392 304 S 0.0 0.1 0:15.13 xinetd
>> -stayalive -
>> 690 root    20  0 6040 124 72 S 0.0 0.0 0:02.45
>> /usr/lib/courier-im
>> 693 root    20  0 4872 252 200 S 0.0 0.0 0:01.94
>> /usr/sbin/courierlo
>> 701 root    20  0 6040 124 72 S 0.0 0.0 0:06.34
>> /usr/lib/courier-im
>> 703 root    20  0 4872 256 200 S 0.0 0.0 0:03.09
>> /usr/sbin/courierlo
>> 709 root    20  0 6040 128 72 S 0.0 0.0 0:18.15

```

```

>> /usr/lib/courier-im
>> 711 root    20  0 4872 256 200 S 0.0 0.0  0:09.15
>> /usr/sbin/courierlo
>> 718 root    20  0 6040 124  72 S 0.0 0.0  0:05.68
>> /usr/lib/courier-im
>> 720 root    20  0 4872 252 200 S 0.0 0.0  0:02.54
>> /usr/sbin/courierlo
>> 730 qmails   20  0 1796 224 144 S 0.0 0.0  1:27.21 qmail-send
>> 732 qmaill   20  0 1752 244 192 S 0.0 0.0  0:22.64 splogger
>> qmail
>> 733 root     20  0 1780 140  64 S 0.0 0.0  0:07.85 qmail-lspawn
>> | /usr
>> 734 qmailr   20  0 1776 148  76 S 0.0 0.0  0:14.07 qmail-rspawn
>> 735 qmailq   20  0 1748 104  68 S 0.0 0.0  0:14.01 qmail-clean
>> 781 root     20  0 51880 4364 196 S 0.0 0.7  1:35.02
>> /usr/sbin/httpd
>> 828 named     20  0 44104 5708 1112 S 0.0 0.9 10:10.53
>> /usr/sbin/named -u
>> 866 root     20  0 3708   8   4 S 0.0 0.0  0:00.00 /bin/sh
>> /usr/bin/my
>> 981 root     20  0 33912 3756 916 S 0.0 0.6 10:55.30
>> /usr/bin/spamd --us
>> 1107 xfs      20  0 3392  72  40 S 0.0 0.0  0:00.09 xfs
>> -droppriv -daem
>> 1115 root     20  0 5672   8   4 S 0.0 0.0  0:00.00
>> /usr/sbin/saslauthd
>> 1116 root     20  0 5672   8   4 S 0.0 0.0  0:00.00
>> /usr/sbin/saslauthd
>> 1122 root     20  0 22992 1868 1084 S 0.0 0.3  2:09.79
>> /usr/bin/sw-engine
>> 1123 root     20  0 27328 1508 1160 S 0.0 0.2  6:06.30
>> /usr/local/psa/admi
>> 7251 root     20  0 4488 192 136 S 0.0 0.0  0:22.85 crond
>> 9463 apache   20  0 59184 14m 4356 S 0.0 2.3  0:05.10
>> /usr/sbin/httpd
>> 10512 apache  20  0 42316 2504  84 S 0.0 0.4  0:00.91
>> /usr/sbin/httpd
>> 12090 apache  20  0 56964 14m 4492 S 0.0 2.2  0:04.48
>> /usr/sbin/httpd
>> 12682 apache  20  0 61060 17m 4516 S 0.0 2.7  0:02.45
>> /usr/sbin/httpd
>> 13870 sw-cp-se 20  0 7852 1932  16 S 0.0 0.3  1:19.03
>> /usr/sbin/sw-cp-ser
>> 17443 apache  20  0 62416 17m 4436 S 0.0 2.7  0:05.27
>> /usr/sbin/httpd
>> 17461 apache  20  0 52788 10m 4480 S 0.0 1.6  0:02.24
>> /usr/sbin/httpd
>> 20430 apache  20  0 62164 17m 4356 S 0.0 2.7  0:04.25

```



```

>> /usr/sbin/httpd
>> 23539 popuser 20 0 37612 25m 2328 S 0.0 3.9 0:01.50 spamd child
>> 23924 apache 20 0 58004 15m 5536 S 0.0 2.4 0:01.56
>> /usr/sbin/httpd
>> 26361 apache 20 0 54496 11m 3864 S 0.0 1.8 0:01.35
>> /usr/sbin/httpd
>> 26366 apache 20 0 52944 9.8m 3892 S 0.0 1.5 0:01.45
>> /usr/sbin/httpd
>> 26964 apache 20 0 59184 14m 4316 S 0.0 2.3 0:07.26
>> /usr/sbin/httpd
>> 27096 apache 20 0 53728 10m 3868 S 0.0 1.6 0:00.33
>> /usr/sbin/httpd
>> 27102 apache 20 0 54736 11m 3780 S 0.0 1.8 0:00.15
>> /usr/sbin/httpd
>> 27103 apache 20 0 54480 11m 3784 S 0.0 1.7 0:00.11
>> /usr/sbin/httpd
>> 27115 apache 20 0 57064 12m 3816 S 0.0 2.0 0:00.32
>> /usr/sbin/httpd
>> 27118 apache 20 0 53728 10m 3884 S 0.0 1.6 0:01.21
>> /usr/sbin/httpd
>> 27120 apache 20 0 52184 8376 3120 S 0.0 1.3 0:00.00
>> /usr/sbin/httpd
>> 27129 apache 20 0 52168 8072 2960 S 0.0 1.2 0:00.00
>> /usr/sbin/httpd
>> 27139 apache 20 0 53304 9840 3744 S 0.0 1.5 0:01.08
>> /usr/sbin/httpd
>> 27140 apache 20 0 53000 9.8m 3832 S 0.0 1.5 0:00.66
>> /usr/sbin/httpd
>> 27144 apache 20 0 52168 8072 2960 S 0.0 1.2 0:00.00
>> /usr/sbin/httpd
>> 27147 apache 20 0 53252 12m 5536 S 0.0 1.9 0:00.50
>> /usr/sbin/httpd
>> 27149 apache 20 0 52980 9924 3740 S 0.0 1.5 0:00.17
>> /usr/sbin/httpd
>> 27153 apache 20 0 53728 10m 3836 S 0.0 1.6 0:00.49
>> /usr/sbin/httpd
>> 27164 apache 20 0 55224 11m 3812 S 0.0 1.9 0:00.47
>> /usr/sbin/httpd
>> 27171 apache 20 0 52916 9776 3708 S 0.0 1.5 0:00.16
>> /usr/sbin/httpd
>> 27172 apache 20 0 52916 9452 3436 S 0.0 1.4 0:00.17
>> /usr/sbin/httpd
>> 27173 apache 20 0 55340 11m 3720 S 0.0 1.8 0:00.08
>> /usr/sbin/httpd
>> 27179 apache 20 0 52020 7764 2716 S 0.0 1.2 0:00.00
>> /usr/sbin/httpd
>> 27182 apache 20 0 52020 7764 2716 S 0.0 1.2 0:00.00
>> /usr/sbin/httpd

```

```

>> 27185 apache 20 0 55224 11m 3824 S 0.0 1.9 0:00.30
>> /usr/sbin/httpd
>> 27186 apache 20 0 53788 10m 3840 S 0.0 1.7 0:00.11
>> /usr/sbin/httpd
>> 27187 apache 20 0 52916 9448 3436 S 0.0 1.4 0:00.08
>> /usr/sbin/httpd
>> 27188 apache 20 0 54628 10m 3504 S 0.0 1.7 0:00.05
>> /usr/sbin/httpd
>> 27196 apache 20 0 53728 10m 3572 S 0.0 1.6 0:00.36
>> /usr/sbin/httpd
>> 27200 apache 20 0 54628 11m 3796 S 0.0 1.7 0:00.05
>> /usr/sbin/httpd
>> 27202 apache 20 0 54480 11m 3796 S 0.0 1.7 0:00.10
>> /usr/sbin/httpd
>> 27204 apache 20 0 53992 10m 3544 S 0.0 1.6 0:00.09
>> /usr/sbin/httpd
>> 27207 apache 20 0 52168 8084 2960 S 0.0 1.2 0:00.00
>> /usr/sbin/httpd
>> 27213 apache 20 0 52020 6464 1788 S 0.0 1.0 0:00.00
>> /usr/sbin/httpd
>> 27214 apache 20 0 54216 10m 3516 S 0.0 1.6 0:00.05
>> /usr/sbin/httpd
>> 27215 apache 20 0 52020 6456 1788 S 0.0 1.0 0:00.00
>> /usr/sbin/httpd
>> 27216 apache 20 0 52020 7860 2804 S 0.0 1.2 0:00.00
>> /usr/sbin/httpd
>> 27218 root 20 0 9400 1900 1408 S 0.0 0.3 0:00.00 crond
>> 27219 root 20 0 2492 956 848 S 0.0 0.1 0:00.00 /bin/sh -c
>> /usr/loc
>> 27220 root 20 0 2496 1052 920 S 0.0 0.2 0:00.00 /bin/sh
>> /usr/local/
>> 27233 root 20 0 2540 1016 892 S 0.0 0.2 0:00.00 /bin/bash -c
>> top -c
>> 27234 root 20 0 2284 952 724 R 0.0 0.1 0:00.00 top -cbn1
>> 27235 root 20 0 1756 420 352 S 0.0 0.1 0:00.00 head -100
>> 27247 root 20 0 2496 452 320 S 0.0 0.1 0:00.00 /bin/sh
>> /usr/local/
>> 27248 root 20 0 8280 1504 1120 R 0.0 0.2 0:00.00
>> /usr/bin/mysql -uad
>> 27249 root 20 0 1800 448 376 S 0.0 0.1 0:00.00 sed -e 1d
>> 27250 root 20 0 2240 640 540 S 0.0 0.1 0:00.00 awk
>> {printf("%s", $
>>
>> # netstat -ptan | grep ESTABLISHED
>> tcp 0 0 ::ffff:xx.xx.xx.xx:80 ::ffff:77.87.207.166:21863 ESTABLISHED 23924/httpd
>> tcp 0 0 ::ffff:xx.xx.xx.xx:80 ::ffff:95.165.204.26:62259 ESTABLISHED 27144/httpd
>> tcp 0 0 ::ffff:xx.xx.xx.xx:80 ::ffff:193.151.105.100:4059ESTABLISHED 27200/httpd
>> tcp 0 0 ::ffff:xx.xx.xx.xx:80 ::ffff:109.169.207.68:50087ESTABLISHED 27185/httpd

```

```

>> tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:31.131.70.135:57017 ESTABLISHED 27179/httpd
>> tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:95.165.204.26:62220 ESTABLISHED 27103/httpd
>> tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:188.134.61.1:60732
>> ESTABLISHED 27215/httpd
>> tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:193.151.105.100:4112 ESTABLISHED 26964/httpd
>> tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:109.169.207.68:50043 ESTABLISHED 27164/httpd
>> tcp    0    0 ::ffff:xx.xx.xx.xx:80 ::ffff:31.131.70.135:56976 ESTABLISHED 27153/httpd
>>
>> # cat /proc/user_beancounters
>> Version: 2.5
>>      uid resource          held          maxheld
>>      barrier          limit          failcnt
>> 1506: kmemsize          27735306          179081216
>> 304087040          335544320          0
>>      lockedpages          0          0
>> 81920          81920          0
>>      privvmpages          393683          430195
>> 9223372036854775807 9223372036854775807          0
>>      shmpages          823          21639
>> 9223372036854775807 9223372036854775807          0
>>      dummy          0          0
>> 0          0          0
>>      numproc          128          204
>> 9223372036854775807 9223372036854775807          0
>>      physpages          79702          163840
>> 0          163840          0
>>      vmguarpages          0          0
>> 0 9223372036854775807          0
>>      oomguarpages          74734          75707
>> 0 9223372036854775807          0
>>      numtcpsock          59          153
>> 9223372036854775807 9223372036854775807          0
>>      numflock          46          62
>> 9223372036854775807 9223372036854775807          0
>>      numpty          0          1
>> 9223372036854775807 9223372036854775807          0
>>      numsiginfo          0          33
>> 9223372036854775807 9223372036854775807          0
>>      tcpsndbuf          1037680          11426176
>> 9223372036854775807 9223372036854775807          0
>>      tcprcvbuf          966656          2867584
>> 9223372036854775807 9223372036854775807          0
>>      othersockbuf          53824          838688
>> 9223372036854775807 9223372036854775807          0
>>      dgramrcvbuf          0          502224
>> 9223372036854775807 9223372036854775807          0
>>      numothersock          114          273
>> 9223372036854775807 9223372036854775807          0

```

```
>>      dcache size      10070617      167772160
>> 150994944      167772160      0
>>      numfile      1634      1865
>> 9223372036854775807 9223372036854775807      0
>>      dummy      0      0
>>      0      0      0
>>      dummy      0      0
>>      0      0      0
>>      dummy      0      0
>>      0      0      0
>>      numiptent      20      20
>> 9223372036854775807 9223372036854775807      0
>>
>
> <ATT00001.c>
>
>
>
```
