
Subject: dcachesize possible issue?

Posted by [daniel_vz](#) on Mon, 25 Nov 2013 12:41:29 GMT

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

Hello everyone,

I noticed today a possible issue with one of my VPSes.

It shows an abnormal value for dcachesize.

Current limits are as follows:

```
-----
CT 103   | HELD Bar% Lim%| MAXH Bar% Lim%| BAR | LIM | FAIL
-----+-----+-----+-----+-----
  kmemsize|10.5G  -  - |15.5G  -  - | - | - | -
 lockedpages| -  -  - | 20K  -  - | - | - | -
 privvmpages|249M  -  - |1.28G  -  - | - | - | -
  shmpages|92.9M  -  - |151M  -  - | - | - | -
  numproc| 32  -  - | 75  -  - | - | - | -
  physpages|11.2G  - 70%| 16G  - 100%| - | 16G| -
 vmguarpages| -  -  - | -  -  - | - | - | -
 oomguarpages|147M  -  - |262M  -  - | - | - | -
 numtcpsock| 14  -  - |117  -  - | - | - | -
  numflock| 2  -  - | 39  -  - | - | - | -
  numpty| -  -  - | 3  -  - | - | - | -
 numsignifo| -  -  - | 27  -  - | - | - | -
 tcpsndbuf|553K  -  - |4.23M  -  - | - | - | -
 tcprcvbuf|2.16M  -  - |26.6M  -  - | - | - | -
 othersockbuf|27.1K  -  - |1.05M  -  - | - | - | -
 dgramrcvbuf| -  -  - |9.03K  -  - | - | - | -
 numothersock| 59  -  - |129  -  - | - | - | -
 dcachesize|10.5G  -  - |15.5G  -  - | - | - | -
  numfile|1.53K  -  - |1.94K  -  - | - | - | -
 numiptent| 24  -  - | 24  -  - | - | - | -
 swappages|3.81M  - 0.2%|126M  - 6%| - | 2G| -
-----
```

There are no failcnts and the machine is up and running for a month now. Restarting the web and php-fpm servers didn't do much to lower the value. Is this metric supposed to stay that way all the time or it does decrease automatically?

Looking at open files usage I don't see something which accounts for that high of a value.

Currently running the following kernel on the host:

Mon Nov 25 15:31:41 2013

ftop 1.0

Processes: 24 total, 1 unreadable

Press h for help, o for options

Open Files: 35 regular, 0 dir, 53 chr, 0 blk, 22 pipe, 67 sock, 17 misc

_	PID	#FD	USER	COMMAND
--	10882	17	root	nginx: master process /usr/sbin/nginx -c /etc/nginx/nginx.confnginx: worker process
	+-	err	--W	-- 0/0 /var/log/nginx/error.log (fd 4 for PID 10882)
	+-	4	--W	-- 0/0 /var/log/nginx/error.log (stderr from PID 10882)
	+-	5	--W	>- 141626/141626 /var/log/nginx/access.log (fd 5 for PID 10883)
	+-	6	--W	-- 442029/442029 /var/log/nginx/website/access.log (fd 6 for PID 10883)
	+-	7	--W	-- 40.3M/0 /var/log/nginx/website/error.log (fd 7 for PID 10883)
--	10883	15	nginx	nginx: worker process
	+-	err	--W	-- 0/0 /var/log/nginx/error.log (stderr from PID 10882)
	+-	4	--W	-- 0/0 /var/log/nginx/error.log (stderr from PID 10882)
	+-	5	--W	>- 141626/141626 /var/log/nginx/access.log (fd 5 for PID 10882)
	+-	6	--W	-- 442029/442029 /var/log/nginx/website/access.log (fd 6 for PID 10882)
	+-	7	--W	-- 40.3M/0 /var/log/nginx/website/error.log (fd 7 for PID 10882)
--	10884	15	nginx	nginx: worker process
	+-	err	--W	-- 0/0 /var/log/nginx/error.log (stderr from PID 10882)
	+-	4	--W	-- 0/0 /var/log/nginx/error.log (stderr from PID 10882)
	+-	5	--W	>- 141626/141626 /var/log/nginx/access.log (fd 5 for PID 10882)
	+-	6	--W	-- 442029/442029 /var/log/nginx/website/access.log (fd 6 for PID 10882)
	+-	7	--W	-- 40.3M/0 /var/log/nginx/website/error.log (fd 7 for PID 10882)
--	10885	15	nginx	nginx: worker process
	+-	err	--W	-- 0/0 /var/log/nginx/error.log (stderr from PID 10882)
	+-	4	--W	-- 0/0 /var/log/nginx/error.log (stderr from PID 10882)
	+-	5	--W	>- 141626/141626 /var/log/nginx/access.log (fd 5 for PID 10882)
	+-	6	--W	-- 442029/442029 /var/log/nginx/website/access.log (fd 6 for PID 10882)
	+-	7	--W	-- 40.3M/0 /var/log/nginx/website/error.log (fd 7 for PID 10882)
--	10887	15	nginx	nginx: worker process
	+-	err	--W	-- 0/0 /var/log/nginx/error.log (stderr from PID 10882)
	+-	4	--W	-- 0/0 /var/log/nginx/error.log (stderr from PID 10882)
	+-	5	--W	>- 141626/141626 /var/log/nginx/access.log (fd 5 for PID 10882)
	+-	6	--W	-- 442029/442029 /var/log/nginx/website/access.log (fd 6 for PID 10882)
	+-	7	--W	-- 40.3M/0 /var/log/nginx/website/error.log (fd 7 for PID 10882)
--	119	10	root	/sbin/udev -d
	+-	10	-rw	-- 1053/1053 /dev/.udev/queue.bin
--	401	5	root	/sbin/rsyslogd -i /var/run/syslogd.pid -c 5
	+-	out	--W	-- 206/206 /var/log/messages
	+-	err	--W	-- 9927/9927 /var/log/cron

```
| +- 3 -r- -- 0/0 /proc/kmsg
| +- 4 --W -- 6103/6103 /var/log/secure
-- 438 6 root sendmail: accepting connectionssendmail: Queue runner@01:00:00 for
/var/spool/clientmqueuecrond
| +- 5 --w -- 32/32 /var/run/sendmail.pid
-- 446 5 smmsp sendmail: Queue runner@01:00:00 for /var/spool/clientmqueuecrond
| +- 4 --w -- 48/48 /var/run/sm-client.pid
-- 473 6 root crond
| +- 3 -rw -- 4/4 /var/run/crond.pid
-- 10904 11 root php-fpm: master process (/etc/php-fpm.conf)sshd: app [priv]
+- err --W -- 2052/2052 /var/log/php-fpm/error.log (fd 5 for PID 10904)
+- 5 --W -- 2052/2052 /var/log/php-fpm/error.log (stderr from PID 10904)
```

No zombie processes that hold open files.

```
[root@app /]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/ploop46926p1 20G  8.3G  11G  44% /
none            3.0G  4.0K  3.0G   1% /dev
[root@app /]# df -i
Filesystem      Inodes  IUsed  IFree IUse% Mounted on
/dev/ploop46926p1 1282560 144098 1138462  12% /
none            786432   151  786281   1% /dev
```

I'm nowhere near the available free space in the VPS.

```
[root@CentOS-64-64-minimal conf]# vzcfgvalidate 103.conf
Validation completed: success
```

VPS is running CentOS 6.4 just like the host system. Is this an OpenVZ matter or the guest system is caching more files than I would like?

Any help is appreciated.

Subject: Re: dcachesize possible issue?
Posted by [pavel.odintsov](#) on Wed, 27 Nov 2013 09:47:46 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello!

Please check directories with big number of files.

You could try this: https://openvz.org/Page_cache_isolation but I'm not sure about it because dcachesize is directory/inode cache and can't be so big.

Subject: Re: dcachesize possible issue?

Posted by [daniel_vz](#) on Wed, 27 Nov 2013 16:02:00 GMT

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

Page cache isolation doesn't seem to be related to dcachesize. I'll try enabling page cache isolation. However, since I don't have any failcnts yet, is this something to worry about? I did look for CentOS explanations and this seems to be something built-in. The O/S will keep this cache and free the memory once a program asks for it. I'm worrying it will cause an OOM condition though. In the past it was a bug related to this in OpenVZ kernels, I wondered if this is a repeat of the same issue.

Don't know what debug data to collect, I can get more if you can tell me what to collect. My main concern is that I don't hit an OOM and have the kernel killing processes in the VPS. The host has 32Gb of memory, although the VPS seems to consume way less than I've allocated for it.

Thanks!
