
Subject: OOM killed process, although there is enough memory
Posted by [Roman Haefeli](#) on Wed, 10 Oct 2012 16:25:48 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi all

We're having issues with processes in a container being killed by OOM, although the hostnode does not have even half of its memory used. How can it be that there is a failcnt for oomguarpages, although the barrier wasn't even reached?

This is on Debian Squeeze with the OpenVZ Kernel shipped by Debian.

```
uname -r  
2.6.32-5-openvz-amd64
```

from dmesg:

```
[498530.406463] OOM killed process java (pid=4143, ve=2018) exited, free=11882227 gen=17.
```

from /proc/user_beancounters

uid	resource	held	maxheld	barrier	limit	failcnt
2018:	kmemsize	12790390	34388183	34379782		38379782
	1255					
	lockedpages	0	239	1024	1200	6
	privvmpages	416902	1559453	2097152		
9223372036854775807		0				
	shmpages	7929	8954	21504	21504	
0						
	dummy	0	0	0	0	0
	numproc	184	372	1000	1200	0
	physpages	71964	529494	0	9223372036854775807	
0						
	vmguarpages	0	0	2097152	9223372036854775807	
0						
	oomguarpages	71964	529494	2097152		
9223372036854775807		1				
	numtcpsock	43	117	360	360	0
	numflock	7	29	188	206	0
	numpty	0	1	16	16	0
	numsiginfo	0	7	256	256	0
	tcpsndbuf	213904	955224	1720320	2703360	
0						
	tcprcvbuf	212992	282952	1720320	2703360	
0						

0	othersockbuf	18496	43600	1126080	2097152	
	dgramrcvbuf	0	4360	262144	262144	0
	numothersock	214	253	360	360	0
0	dcachesize	611724	1746754	3409920	3624960	
	numfile	3745	7966	9312	9312	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	numiptent	10	10	128	128	0

Roman

Subject: Re: OOM killed process, although there is enough memory

Posted by [Corin Langosch](#) on Wed, 10 Oct 2012 17:16:09 GMT

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

On 10.10.2012 at 18:25 +0200, Roman Haefeli <reduzent@gmail.com> wrote:

> We're having issues with processes in a container being killed by OOM,
 > although the hostnode does not have even half of its memory used.
 > How can it be that there is a failcnt for oomguarpages, although the
 > barrier wasn't even reached?
 >
 > This is on Debian Squeeze with the OpenVZ Kernel shipped by Debian.

I can only highly recommend against using the debian openvz kernels. I always had stability problems with them and a few months ago I discovered a major bug (<http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=655385>), causing the whole system to get instable/ frozen.

For a few months now I'm using the official rhel6 openvz kernels and convert them into debian packages as described here: http://wiki.openvz.org/Install_kernel_from_RPM_on_Debian_6.0 Since then I never had any problems again so far :)

Corin

Subject: Re: OOM killed process, although there is enough memory

Posted by [Roman Haefeli](#) on Wed, 10 Oct 2012 20:07:22 GMT

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

On Mit, 2012-10-10 at 19:16 +0200, Corin Langosch wrote:

> On 10.10.2012 at 18:25 +0200, Roman Haefeli
> <reduzent@gmail.com> wrote:
> > We're having issues with processes in a container being killed by OOM,
> > although the hostnode does not have even half of its memory used.
> > How can it be that there is a failcnt for oomguarpages, although the
> > barrier wasn't even reached?
> >
> > This is on Debian Squeeze with the OpenVZ Kernel shipped by Debian.
>
> I can only highly recommend against using the debian openvz kernels. I
> always had stability problems with them and a few months ago I
> discovered a major bug
> (<http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=655385>), causing
> the whole system to get instable/ freezed.
>
> For a few months now I'm using the official rhel6 openvz kernels and
> convert them into debian packages as described here:
> http://wiki.openvz.org/Install_kernel_from_RPM_on_Debian_6.0 Since
> then I never had any problems again so far :)

Thanks for your comments. It's not the first time I experienced oddities with the Debian OpenVZ kernel. I'll definitely check out the RHEL6 kernel.

Also, knowing that what I reported is not expected behavior is valuable for me to know. I started to think that I might have misconfigured the CTs due to a lack of understanding the mem limits/guarantees.

Roman
