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

	held	maxheld	barrier	limit	failcnt
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>