

>> Hello everybody!

>>

>> I'm having problems with OpenVZ memory management. I have (probably)
>> read all the Wiki articles there are that touch that subject, and
>> browsed through the mailing list archives to no avail.

>>

>> Most of my OpenVZ VPS'es work just fine after a bit of fiddling, but one
>> of them is misbehaving constantly. It always runs out of memory, no
>> matter how much I give it. It is running 10 instances of a same server
>> to give better interactive responsiveness. I was just wondering if the
>> server program is leaking memory and causing this erratic behavior, or
>> if there is something wrong with my OpenVZ VPS's configuration.

>>

>

>> The physical server runs only this one VPS. The hardware node has 2GB of
>> RAM plus 2GB of swap. The VPS is given roughly 3.5GB of that if
>> available (privvmpages limit). It is guaranteed 2.5GB (vmguarpages
>> barrier). This is FAR more than the server software in question needs,
>> but still it occasionally (and predictably) runs out of memory.

>>

>> The strange thing about the VPS is that the HELD values in oomguarpages
>> and privvmpages are _much_ lower than the MAXHELD values - almost
>> triple. There should be no usage peaks that should cause this kind of
>> asymmetry, unless an instance of the server software goes amok.

>

> Yes, privvmpages maxheld shows that software tried to allocate lots of memory in reality.
> And this VE used 874046 pages of RAM+swap (oomguarpages) in the peak, i.e. ~3.4Gb !!!
> It looks like accounting doesn't lie and your software really tries
> to allocate that much from time to time.

Yes, oomguarpages maxheld is suspiciously high, and probably slows
things down considerably when the server starts swapping.

>

> Have you seen yourself when this was happening and failcounters were increasing?
> If you can observe this VE being near the limit,
> run top in host system and sort processes by RSS usage.
> This will show you which processes consume most of the memory (RSS column).

I'll have to monitor the VPS when it's approaching the privvmpages
limit. Luckily monit will tell me when that is happening :). I'll let
you know what I find out. It starts to seem like the servers in the VPS
are just misbehaving.

> BTW, have you did something to make VE use less memory (like server restart)
> or the memory usage drops that low itself after some time of being high?

No, I've done nothing special after the latest failure. Sometimes one of the server instances dies and has to be restarted, but it's only one instance, and there are several others still around.

>> The VPS's resource information is shown below. The parameters are not
>> optimized, as you can see, but that not my biggest problem right now :).
>> So can you see anything wrong with these settings, or should I take a
>> look at the server software that is running on the VPS?

>>

>> [root@VPS_NODE ~]# cat /proc/user_beancounters

>> Version: 2.5

| uid resource | held | maxheld | barrier | limit | failcnt | |
|---------------|----------|----------|-----------|------------|---------|--|
| 103: kmemsize | 19265157 | 40765738 | 183079731 | 201387704 | 0 | |
| lockedpages | 0 | 0 | 8939 | 8939 | 0 | |
| privvmpages | 383368 | 930126 | 917504 | 930000 | 153 | |
| shmpages | 21647 | 24239 | 31099 | 31099 | 0 | |
| dummy | 0 | 0 | 0 | 0 | 0 | |
| numproc | 229 | 524 | 8000 | 8000 | 0 | |
| physpages | 236810 | 485688 | 0 | 2147483647 | 0 | |
| vmguarpages | 0 | 0 | 655360 | 2147483647 | 0 | |
| oomguarpages | 382472 | 874046 | 310999 | 2147483647 | 0 | |
| numtcpsock | 532 | 793 | 8000 | 8000 | 0 | |
| numflock | 18 | 594 | 1000 | 1100 | 0 | |
| numpty | 1 | 4 | 512 | 512 | 0 | |
| numsiginfo | 0 | 39 | 1024 | 1024 | 0 | |
| tcpsndbuf | 3524644 | 4236356 | 28258577 | 61026577 | 0 | |
| tcprcvbuf | 3530300 | 6528712 | 28258577 | 61026577 | 0 | |
| othersockbuf | 190060 | 1737340 | 14129288 | 46897288 | 0 | |
| dgramrcvbuf | 0 | 41836 | 14129288 | 14129288 | 0 | |
| numothersock | 204 | 1169 | 8000 | 8000 | 0 | |
| dcachesize | 0 | 0 | 39977755 | 41177088 | 0 | |
| numfile | 11106 | 22346 | 71488 | 71488 | 0 | |
| dummy | 0 | 0 | 0 | 0 | 0 | |
| dummy | 0 | 0 | 0 | 0 | 0 | |
| dummy | 0 | 0 | 0 | 0 | 0 | |
| numiptent | 10 | 10 | 200 | 200 | 0 | |

>>

>> [root@HOST_NODE ~]# free

| | total | used | free | shared | buffers | cached |
|--------------------|---------|---------|---------|--------|---------|--------|
| Mem: | 2073344 | 2017716 | 55628 | 0 | 66928 | 954828 |
| -/+ buffers/cache: | 995960 | 1077384 | | | | |
| Swap: | 2031608 | 586508 | 1445100 | | | |

>

>> Anyways, thanks for a great Open Source virtualization project!

> you are welcome :@)
>
> Thanks,
> Kirill
