
Subject: Re: Limits for low memory

Posted by [Wolfgang Schnerring](#) on Wed, 25 Apr 2007 16:14:57 GMT

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

Hello Kirill,

thank you very much for your insights!

* Kirill Korotaev <dev@sw.ru> [2007-04-25 18:19]:

> 1. 230 VEs correspond to kir@ results when he was running fine ~120
> VEs on 1Gb notebook.

That's good to hear, so the number seems to be about right.

> 2. can you please attach your kernel panic
> (if it was a really panic and not an OOM kill message)?

Yes, sorry, it was an OOM kill message. The machine was dead either way, though. ;-) But since that's my fault, I don't mind -- after all, I pretty much told it to kill itself by asking for more memory than there is physically available.

The traceback goes over several screens, I've uploaded them to <<http://wosc.de/tmp/oom/>> if anybody is interested (please pardon my bad photography skills).

> 3. what kernel version do you use?
> is it i686 or x86-64 kernel?

It's 2.6.18-028test19 on i686.

> 1. You also need to take into account dcachesize
> (which is not immediately accounted by default for optimization purposes,
> but can be enabled with dentry_watermark sysctl and
> usually dcachesize is around ~2Mb per VE, so 460Mb of low memory in your case)

Ah I see, so dcachesize seems to be the missing piece.

I tried to enable dcachesize accounting by "echo 0 0 > /proc/sys/ubc/dentry_watermark", is that correct?

I reran the 220 VE experiment, and the numbers it yielded were
kmemsize 127537452
sockbuf 156184224
dcachesize 16297263
which again is only 290MB... strange.

> physpages can go both to low memory and highmemory and you don't

> know how it was spread from UBC counters.

I understand. But I would have thought that the kernel tries to put everything that's possible to high memory first, especially before starting OOM-killing.

> > -- Question: Where do the 832 MB come from?

> The number depends on kernel and configuration:

> [dev@dev ~]\$ cat /proc/meminfo | grep Low

> LowTotal: 903348 kB

> the most impact on the size of low memory have vmalloc reserved area.

You see, this is the kind of insight you'll rarely find in any books or something. Either, you read the source (which is a little daunting given the size of the kernel) or you are lucky -- like I am right now -- and are able to ask an expert. Thanks very much for your help!

> > -- Question: Where does the 0.4 come from?

I'm afraid I still haven't understood what the meaning of this factor is in the formula on the wiki-page.

> plus you need to take into account fragmentation issues.

> i.e. it's not that easy like a sum of counters == total memory.

That makes sense, but I'd still expect that the sum of the counters to be at least in the same general area -- or is that a wrong expectation?

Wolfgang
