Re: running Linux with no swap space (but lots of RAM)

Source: http://linux.derkeiler.com/Newsgroups/comp.os.linux.development.system/2007−11/msg00216.html

- *From*: Rainer Weikusat <rweikusat@xxxxxxxxxxx>
- *Date*: Thu, 29 Nov 2007 15:55:33 +0100

phil−news−nospam@xxxxxxxx writes:

[...]

> | e> In that situation it is quite likely that you don't have "some
> | e> other device" to swap to. That's probably why you are booting
> | e> from flash in the first place.
> |
> | The problem in your new hypothetical is then that you have no device
> | capable of tolerating paging I/O, not that the system is paging to
> | disc.
>
> Please explain what you mean by "tolerating paging I/O".

In plain words, this means 'only devices whose actual behaviour and limitations the OP understands even less well than he believes to understand disks'.

[...]

> A frequent scenario I see happens when my need for memory by user space programs is below the capacity is available, and would not have even begun to swap anything. A program is run that will be doing a large amount of I/O output, such as copying 100+ GB of files between filesystems. The I/O buffering goes beyond just the pages that are free. The buffering logic tries to buffer far more of those 100+ GB than needed to keep the writing drive continuously busy or even to minimize head seeks.

The purpose of the page cache is neither 'to keep the drive continously busy' nor 'to minimze head seeks'. Both would be tasks the elevator (or I/O−scheduler) is supposed to accomplish.

> The end result is that other programs that were not immediately in use (for example Firefox, Thunderbird, and even Xorg itself), get substantially swapped out.

'programs' (ie text pages corresponding to running processes) are never 'swapped out', because the corresponding memory pages are usually clean, so their contents can be discarded and later re−read from disk.

Just for an informal test, I have just created an archive of all of my filesystem to /dev/null (~16G). The amount of memory allocated to 'buffers' and 'cached' (vmstat) peaked well below 95000K and 45000K, respectively. No paging activities occured during creation of the archive.

BTW, the by−and−far easiest path to personal happiness for you in this respect is to just misconfigure your system to your hearts content (its yours, after all) instead of talking about hypothesises you have about situations which − for some strange reason − are not that generally reproducible than the generality of your inferences would require.

It would still be more sensible to add RAM until you don't experience regular paging activities occuring for some unknown reason on your system and leave the virtual memory configuration as−is to deal with non−regular situations. Or consider reducing your working set.
.