

Re: Linux machines dieing in swap storms

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-10/msg09750.html>

- *From:* Rik van Riel <riel@xxxxxxxxxx>
 - *Date:* Fri, 26 Oct 2007 00:03:58 -0400
-

On Fri, 26 Oct 2007 05:56:49 +0200
Bodo Eggert <7eggert@xxxxxx> wrote:

Rik van Riel <riel@xxxxxxxxxx> wrote:

On Thu, 25 Oct 2007 16:20:41 +0100
Richard Purdie <rpurdie@xxxxxxxxxx> wrote:

Advice on solving this welcome preferably in mainline but
I'll
happily hack my kernels with a workaround if need be.

I can't see any easy hacks or workarounds to fix the issue in the
current MM, except maybe activate the OOM killer if the amount of
page cache and buffer cache is really low and swap is full...

In the longer run, I'm working on:

<http://linux-mm.org/PageReplacementDesign>

What about only reclaimimn cache if the cache has grown beyond a
watermark and only reclaimimn non-cache if it's below another
watermark? I can imagine it will solve my
diskcache-pushes-out-mousehandler problem, and I'm pretty sure having
very low file cache is bad for performance, too.

There are much better ways to determine such thresholds than
requiring the sysadmin to set them by hand. I have described
one on the page linked above.

Another thing I can imagine is to detect thrashing conditions and to
change scheduling in order to increase the likelihood of cache hits and

Re: Linux machines dieing in swap storms

thereby progress: If an application just got a page, keep it running for a while (accumulating negative credits).

If the process needs another page after the page it just got (very likely), you cannot "keep it running".

--

"Debugging is twice as hard as writing the code in the first place. Therefore, if you write the code as cleverly as possible, you are, by definition, not smart enough to debug it." – Brian W. Kernighan

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>