

Re: 2.6.0-test9 – poor swap performance on low end machines

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2003-11/0479.html>

From: Chris Vine (chris_at_cvine.freemove.co.uk)

Date: 11/04/03

To: Con Kolivas <kernel@kolivas.org>, Rik van Riel <riel@redhat.com>

Date: Tue, 4 Nov 2003 22:08:05 +0000

On Tuesday 04 November 2003 2:55 am, Con Kolivas wrote:

- > *That's pretty much what I expected. Overall I'm happier with this later*
- > *version as it doesn't impact on the noticeable improvement on systems that*
- > *are not overloaded, yet keeps performance at least that of the untuned*
- > *version. I can tune it to be better for this work load but it would be to*
- > *the detriment of the rest.*
- >
- > *Ultimately this is the problem I see with 2.6 ; there is no way for the vm*
- > *to know that "all the pages belonging to the currently running tasks should*
- > *try their best to fit into the available space by getting an equal share".*
- > *It seems the 2.6 vm gives nice emphasis to the most current task, but at*
- > *the detriment of other tasks that are on the runqueue and still need ram.*
- > *The original design of the 2.6 vm didn't even include this last ditch*
- > *effort at taming swappiness with the "knob", and behaved as though the*
- > *swappiness was always set at 100. Trying to tune this further with just*
- > *the swappiness value will prove futile as can be seen by the "best" setting*
- > *of 20 in your test case still taking 4 times longer to compile the kernel.*
- >
- > *This is now a balance tradeoff of trying to set a value that works for your*
- > *combination of the required ram of the applications you run concurrently,*
- > *the physical ram and the swap ram. As you can see from your example, in*
- > *your workload it seems there would be no point having more swap than your*
- > *physical ram since even if it tries to use say 40Mb it just drowns in a*
- > *swapstorm. Clearly this is not the case in a machine with more ram in*
- > *different circumstances, as swapping out say openoffice and mozilla while*
- > *it's not being*