

Linux: Tuning CFS

By *srlinuxx*

Created *05/08/2007 - 7:32am*

Submitted by srlinuxx on Sunday 5th of August 2007 07:32:23 AM Filed under [Linux](#) [1]

Nick Piggin used 'git bisect' to track a lmbench regression to the main CFS commit, leading to an interesting discussion between Nick and Ingo Molnar. Ultimately the regression was tracked down to the temporary configurability of the scheduler while it is tuned for optimal performance, "one reason for the extra overhead is the current tunability of CFS, but that is not fundamental, it's caused by the many knobs that CFS has at the moment." The solution, already coded but not yet merged in the mainline kernel "changes those knobs to constants, allowing the compiler to optimize the math better and reduce code size," and as a result result, "CFS can be faster at micro-context-switching than 2.6.22."

Ingo described the lmbench configuration in question as a "micro-benchmark", and noted that with a macro-benchmark better performance was more pronounced, "because with CFS the `_quality_` of scheduling decisions has increased.

[More Here](#) [2]

[Linux](#)

Source URL: <http://www.tuxmachines.org/node/18771>

Links:

[1] <http://www.tuxmachines.org/taxonomy/term/63>

[2] <http://kerneltrap.org/node/14055>