

Re: [RFC] sys_setrlimit() in 2.6.16

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-02/msg05804.html>

- *From:* Ingo Molnar <mingo@xxxxxxx>
 - *Date:* Thu, 16 Feb 2006 10:47:31 +0100
-

* Andrew Morton <akpm@xxxxxxx> wrote:

This has to be considered a bug. The spec certainly implies that a limit of zero should be honoured and, probably more importantly, that's how it works in 2.4.

Problem is, the code in there all assumes that an it_prof_expires of zero means "it was never set", and changing that (add a yes-it-has flag?) would be less than trivial.

So I think the path of least resistance here is to just convert the caller's zero seconds into one second. That in fact gives the same behaviour as 2.4: you get whacked after one second or more CPU time.

(This is not a final patch – that revolting expression in sys_setrlimit() needs help first).

your approach looks good to me. It doesn't make much sense anyway to have a task whacked right after startup ... so adding a common-sense "the user must have meant some really small value" thing doesn't look all that wrong.

Acked-by: Ingo Molnar <mingo@xxxxxxx>

Ingo

–

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>