

Hello Peter,

> On 11/09/2011 03:55 PM, Peter Zijlstra wrote:

>>

>> On Mon, 2011-11-07 at 15:54 +0300, Andrew Vagin wrote:

>>>

>>> This patch solves the following problem:

>>>

>>> Now some samples may be lost due to throttling. The number of samples is
>>> restricted by `sysctl_perf_event_sample_rate/HZ`. A trace event is
>>> divided on some samples according to event's period. I don't sure, that
>>> we should generate more than one sample on each trace event. I think the
>>> better way to use `SAMPLE_PERIOD`.

>>

>> It would be yes, but this code predates that, also it needs to work even
>> if the user doesn't provide `SAMPLE_PERIOD`.

I have not understood exactly what I should do now.

I'm going to send the third version of this patches. New version contains only a small fix according with the comment for path 4/7. In new version I am not going to fix the problem about which we discussed early. I have some reasons for it:

- * It's another task. My decision is more effective for my task and I believe that exists many cases where my decision may be suitable. I want to say, that it may be in kernel, event if the problem would not exist at all.

- * `__perf_count()` has been broken for a long time and nobody reports this problem, so we can say, that the problem isn't urgent.

- * Only few events `sched:sched_stat_*` are affected by this problem.

- * I am not sure, that we should to solve this problem. The "problem" may be not a problem at all. We have ability to customize `sample_period` and we have events about that some events were throttled.

If you have any objection, send them to me.
