
Subject: Re: [PATCH] event: fix TP_printk() argument in sched_switch
Posted by [Steven Rostedt](#) on Wed, 02 Nov 2011 23:07:57 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, 2011-11-03 at 01:40 +0300, Andrew Vagin wrote:
> process_arg(...) can't parse "__entry->prev_state & (TASK_STATE_MAX-1)",
> because a complicated argument should be within brackets.

No it is fine, the userspace tool is broken.

>
> Without this patch "perf report" prints following errors:
> \$./perf record -ag -e sched:sched_switch
> ...
> \$./perf report
> Warning: Error: expected type 5 but read 4
> Warning: Error: expected type 4 but read 0
> Fatal: bad op token {
>
> Signed-off-by: Andrew Vagin <avagin@openvz.org>

NACK!

This is a perf userspace bug, not a kernel one. Please fix the userspace tool instead.

Note, the new version of libparseevent handles this case without issue.
Perf just needs to be updated.

-- Steve

> ---
> include/trace/events/sched.h | 2 +-
> 1 files changed, 1 insertions(+), 1 deletions(-)
>
> diff --git a/include/trace/events/sched.h b/include/trace/events/sched.h
> index 959ff18..1838237 100644
> --- a/include/trace/events/sched.h
> +++ b/include/trace/events/sched.h
> @@ -140,7 +140,7 @@ TRACE_EVENT(sched_switch,
> TP_printk("prev_comm=%s prev_pid=%d prev_prio=%d prev_state=%s%s ==>
next_comm=%s next_pid=%d next_prio=%d",
> __entry->prev_comm, __entry->prev_pid, __entry->prev_prio,
> __entry->prev_state & (TASK_STATE_MAX-1) ?
> - __print_flags(__entry->prev_state & (TASK_STATE_MAX-1), "|",
> + __print_flags((__entry->prev_state & (TASK_STATE_MAX-1)), "|",
> { 1, "S" }, { 2, "D" }, { 4, "T" }, { 8, "t" },
> { 16, "Z" }, { 32, "X" }, { 64, "x" },

```
> { 128, "W" }) : "R",
```
