
Subject: Re: [PATCH] event: fix TP_printk() argument in sched_switch
Posted by [Steven Rostedt](#) on Thu, 03 Nov 2011 14:19:37 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2011-11-02 at 19:07 -0400, Steven Rostedt wrote:

> 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.

>

Can you try this patch on perf. It's untested (not even compiled tested)

-- Steve

```
diff --git a/tools/perf/util/trace-event-parse.c b/tools/perf/util/trace-event-parse.c
index 0a7ed5b..6c164dc 100644
--- a/tools/perf/util/trace-event-parse.c
+++ b/tools/perf/util/trace-event-parse.c
@@ -1537,6 +1537,8 @@ process_flags(struct event *event, struct print_arg *arg, char **tok)
    field = malloc_or_die(sizeof(*field));

    type = process_arg(event, field, &token);
+   while (type == EVENT_OP)
+       type = process_op(event, field, &token);
    if (test_type_token(type, token, EVENT_DELIM, ","))
        goto out_free;
```
