
Subject: Re: [PATCH 0/7] Profiling sleep times (v3)
Posted by [Araldo Carvalho de M\[2\]](#) on Tue, 06 Dec 2011 13:45:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

Em Tue, Dec 06, 2011 at 11:15:07AM +0400, Andrey Vagin escreveu:

> Hello Arnaldo,

>

> I found, that you commit patches for perf. Could you review this part?

>

> Arun Sharma said, that the second versions of patches works ok for him.

> (Arun is the first user of this functionality after me.)

I'm going to check if the patches still apply and then review,

- Arnaldo

> Thanks.

>

> On 11/28/2011 01:03 PM, Andrew Vagin wrote:

> >Do you want to know where your code waits locks for a long time?

> >Yes! It's for you. This feature helps you to find bottlenecks.

> >

> >It's not artificial task. Once one of my colleague was investigating a

> >scalability problem. He pressed sysrq-t some times and tried to merge

> >call-chains by hand. But perf can do that.

> >

> >Problem:

> >The problem is that events sched_stat_* contain call-chains of

> >non-target tasks.

> >About month ago I sent series of patches:

> >

> >[PATCH 0/3] trace: add ability to collect call chains of non current task.

> >

> >Peter and Frederic explained me, that this solve isn't good and will be

> >better to make it in userspace.

> >

> >Now it's in userspace. This series expands "perf inject" to be able to

> >merge sched_switch events and sched_stat* events. sched_switch events

> >contain correct call-chains and sched_stat contains a correct time

> >slices.

> >

> >v2:

> >* Removed all known issues. Now it works completely.

> >* Improved usability of sched-stat scripts according with Arun's comments.

> >v3: fixed according to comments from David Ahem

> >

> >Andrew Vagin (7):

> > perf: use event_name() to get an event name

```

> > perf: add ability to record event period
> > perf: add ability to change event according to sample (v2)
> > perf: teach "perf inject" to work with files
> > perf: teach perf inject to merge sched_stat_* and sched_switch events
> > perf: add scripts for profiling sleep times (v2)
> > event: add tracepoint for accounting block time
> >
> > include/trace/events/sched.h          | 7 ++
> > kernel/sched_fair.c                  | 2 +
> > tools/perf/builtin-inject.c           | 120 ++++++
> > tools/perf/builtin-record.c           | 5 +
> > tools/perf/scripts/python/bin/sched-stat-record | 65 ++++++
> > tools/perf/scripts/python/bin/sched-stat-report | 5 +
> > tools/perf/util/event.h               | 2 +
> > tools/perf/util/evsel.c               | 74 ++++++
> > tools/perf/util/header.c              | 2 +-
> > tools/perf/util/session.h             | 9 ++
> > 10 files changed, 288 insertions(+), 3 deletions(-)
> > create mode 100644 tools/perf/scripts/python/bin/sched-stat-record
> > create mode 100644 tools/perf/scripts/python/bin/sched-stat-report
> >

```
