

Scheduler Instrumentation

- Goals:
 - Extract the relevant scheduling details at sched_switch, waking/wakeup, PI, prio change
 - Depending on the scheduling policy of the target process
 - A single event to enable
 - Without breaking existing tools

Current Status

```
# cat /sys/kernel/debug/tracing/events/sched/sched_switch/format
```

```
name: sched_switch
```

```
ID: 273
```

```
Format: prev_comm, prev_pid, prev_prio, prev_state, next_comm, next_pid,  
next_prio
```

(simplified output)

Proposition

```
# echo 2 > /sys/kernel/debug/tracing/events/sched/sched_switch/version
```

```
# cat /sys/kernel/debug/tracing/events/sched/sched_switch/format
```

```
name: sched_switch
```

```
ID: 273
```

```
Format: prev_comm, prev_pid, prev_state, next_comm, next_pid, next_policy, next_nice
```

```
ID: 274
```

```
Format: prev_comm, prev_pid, prev_state, next_comm, next_pid, next_policy,  
next_rt_priority
```

```
ID: 275
```

```
Format: prev_comm, prev_pid, prev_state, next_comm, next_pid, next_policy,  
next_dl_runtime, next_dl_deadline, next_dl_period
```

(simplified output)