

Scheduler Microconference

Monday, September 20, 2021 7:00 AM (5 minutes)

Title: Scheduler Microconference

The scheduler is an important functionality of the Linux kernel, deciding what process gets to run when and for how long. With different topologies and workloads, it is no easy task to give the user the best experience possible. Schedulers are one of the most discussed topics at the Linux Kernel Mailing List, but many of these topics need further discussion in a conference format. Indeed, the scheduler microconference is responsible for many topics to make progress.

For example, at last year's Scheduler MC, we discussed core scheduling which is now on its way to being merged [1]. The scheduling fairness patches were merged [2], NUMA topology limitations fixes were added to the kernel [3]. Not only some progress was made in the direction of accepting patches, but also some topics were proved to be not feasible, like "Flattening the CFS runqueue," and this was facilitated by the conference format.

This year, we think the following topics will lead to a productive microconference:

- Cgroup interface and other updates for core-scheduling [1]
- Cgroup and SCHED_DEADLINE [4]
- Capacity Awareness –For busy systems
- Interrupt Awareness
- Load Balancing
 - Wakeup [5] [6] [7]
 - Periodic [5] [6]
 - NUMA load balancing

Come and join us in the discussion of controlling what tasks get to run on your machine and when. We hope to see you there!

Attendees list:

- Peter Zijlstra peterz@infradead.org
- Thomas Gleixner tglx@linutronix.de
- Steven Rostedt rostedt@goodmis.org
- Vincent Guittot vincent.guittot@linaro.org
- Ingo Molnar mingo@redhat.com
- Juri Lelli juri.elli@redhat.com
- Daniel Bristot de Oliveira bristot@redhat.com
- Dietmar Eggemann dietmar.eggemann@arm.com
- Sebastian Andrzej Siewior bigeasy@linutronix.de
- Valentin Schneider valentin.schneider@arm.com
- Clark Williams williams@redhat.com
- Paul E. McKenney paulmck@kernel.org
- Ben Segall bsegall@google.com
- Mel Gorman mgorman@suse.de
- Joel Fernandes joel@joelfernandes.org
- Quentin Perret qperret@google.com
- Aubrey Li aubrey.li@intel.com
- Paul Turner pjt@google.com
- Rik van Riel riel@surriel.com

- Ricardo Neri ricardo.neri-calderon@linux.intel.com
- Catalin Marinas catalin.marinas@arm.com
- Qais Yousef qais.yousef@arm.com
- Patrick Bellasi patrick.bellasi@matbug.net
- Morten Rasmussen morten.rasmussen@arm.com
- Viresh Kumar viresh.kumar@linaro.org
- Phil Auld pauld@redhat.com
- Waiman Long longman@redhat.com
- Josef Bacik josef@toxicpanda.com
- Frederic Weisbecker fweisbec@gmail.com

Links:

[1] <https://lore.kernel.org/lkml/20210422120459.447350175@infradead.org/T/>

[2] scheduling fairness commits:

- 6e7499135db7 ("sched/fair: Reduce busy load balance interval")
- e4d32e4d5444 ("sched/fair: Minimize concurrent LBs between domain level")
- 2208cdaa56c9 ("sched/fair: Reduce minimal imbalance threshold")
- 5a7f55590467 ("sched/fair: Relax constraint on task's load during load balance")

[3] numa topology commits:

- 620a6dc40754 ("sched/topology: Make sched_init_numa() use a set for the deduplicating sort")
- 585b6d2723dc ("sched/topology: fix the issue groups don't span domain->span for NUMA diameter > 2")

[4] <https://lore.kernel.org/lkml/cover.1610463999.git.bristot@redhat.com/>

[5] <https://lore.kernel.org/linux-arm-kernel/20210420001844.9116-5-song.bao.hua@hisilicon.com/T/>

[6] <https://www.spinics.net/lists/kernel/msg3894298.html>

[7] <https://www.spinics.net/lists/kernel/msg3914884.html>

I agree to abide by the anti-harassment policy

I agree

Primary authors: GIANI, Dhaval (Oracle); BRISTOT DE OLIVEIRA, Daniel (Red Hat, Inc.); HYSER, chris; LELLI, Juri (Red Hat); GUITTOT, Vincent (Linaro)

Presenters: GIANI, Dhaval (Oracle); BRISTOT DE OLIVEIRA, Daniel (Red Hat, Inc.); HYSER, chris; LELLI, Juri (Red Hat); GUITTOT, Vincent (Linaro)

Session Classification: Scheduler MC

Track Classification: Scheduler MC