

Energy-Aware Scheduling on Intel Hybrid Topologies

Latency Analysis and Idle CPU selection




LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026




Motivation

- OSPM 2026: drop the schedutil requirement for EAS on x86
 - EAS soon to be active by default
- EAS packs light tasks onto fewer, power-efficient cores
 - Latency cost
- **Question for this talk:**
 - Can latency cost be reduced by using idle CPUs more aggressively?




schbench wakeup latency



LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026

Decomposing wakeup latency




LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026



Dispatch latency without EAS




LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026




Dispatch latency with EAS



LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026

Proposed tweak: cluster-local idle preference




LINUX
PLUMBERS CONFERENCE


PRAGUE, CZECHIA / OCT. 5-7, 2026



Proposed tweak: cluster-local idle preference



Tweak results: no-EAS vs idle-slack settings




LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026



Latency vs EEVDF slice duration




LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026



Idle slack vs EEVDF slice duration




LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026



Idle slacks in the real world




LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026



Idle slacks vs EEVDF slice in the real world




LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026



The cost in energy



LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026



Thanks!



LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026



Backup




LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026



Where does latency go?




LINUX
PLUMBERS CONFERENCE


PRAGUE, CZECHIA / OCT. 5-7, 2026



queue_wait dominates, not selection cost



Who wakes the workers?



LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026




The batching effect

queue_wait latency by position in TTWU_QUEUE batch

Hypothesis: position 0 is picked immediately; higher positions wait behind predecessors


Confirmation: queue_wait should increase monotonically with batch_pos

Batch size  1  2



Share of wakeups in single vs multi-worker TTWU batches

A large multi-worker share supports the packing-induced contention hypothesis



LINUX
PLUMBERS CONFERENCE

PRAGUE, CZECHIA / OCT. 5-7, 2026

