Contribution ID: **65** Type: **not specified**

# Linux Kernel Scheduling and split-LLC architectures: Overview, Challenges and Opportunities

*Tuesday, 13 September 2022 16:45 (25 minutes)*

Linux Task Scheduler has seen several enhancements to make task scheduling better and smarter for split last level cache (split-LLC) environments. With wider adoption of the chiplet-like technology in current and future processors, these continued efforts become key to squeeze the most out of the silicon.

Work has already gone in to accurately model the domain topology for split-LLC architectures: Optimizing task wakeups to target cache-hot LLCs, reducing cross-LLC communication. NUMA imbalance metrics have been reworked to enable better task distribution across NUMA nodes with multiple LLCs. These enhancements have enabled several workloads to benefit from architectural advantages of split-LLCs. That being said, there is still lot of performance left on the table.

In this talk we provide an overview of recent scheduler changes that have benefitted workloads in a split-LLC environment. We will describe challenges, opportunities and some ambitious ideas to make the Linux Scheduler more performant on split-LLC architectures.

## I agree to abide by the anti-harassment policy

Yes

**Primary authors:** SHENOY, Gautham R (AMD Inc.); NAYAK, Prateek (AMD Inc. )

**Presenters:** SHENOY, Gautham R (AMD Inc.); NAYAK, Prateek (AMD Inc. )

**Session Classification:** Real-time and Scheduling MC