Linux Plumbers Conference 2024



Contribution ID: 14

Type: not specified

Title: Multi-sized THP performance benchmarks and analysis on ARM64

Friday, 20 September 2024 15:45 (45 minutes)

The Linux kernel has supported multi-sized THP since v6.8 allowing the use of intermediate sized huge pages less than 2M. ARM64 supports contiguous PTEs where multiple PTE entries can be coalesced one TLB entry. This will increase the size of memory covered by the TLB entries and avoid page table walks to create TLB entries.

We ran a series of benchmarks on Ampere Altra using some popular workloads in cloud: In-memory databases, kernel compilation etc using different sized huge pages: 2M, 128K, 64K and others.

This presentation will also cover how multi-sized THP work and includes hardware details on the operation of contiguous PTE and variable page size on ARM64.

We conclude the multi-sized THP may not boost all kind of workloads. The overhead of page table walk is a significant contributing factor for some workloads. The reduced page faults help performance for some workloads. We would recommend the use of a kernel with 16K page size as an optimal solution that has most performance gains but does not significantly increase the memory footprint.

Primary author: SHI, Yang (Ampere Computing)
Co-author: SINGLA, Olivier (Ampere Computing)
Presenters: SINGLA, Olivier (Ampere Computing); SHI, Yang (Ampere Computing)
Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track