

Linux Plumbers Conference 2024



Contribution ID: 216

Type: **not specified**

Memory Allocation Profiling deployment results and future improvements

Friday, 20 September 2024 15:00 (15 minutes)

Memory allocation profiling infrastructure provides a low-overhead mechanism to make all kernel allocations in the system visible. This allows for monitoring memory usage, tracking hotspots, detecting leaks, and identifying regressions.

Unlike previous discussions on the design of this technique, we will now focus on the changes since it was incorporated into the upstream kernel, planned future improvements, and initial discoveries within the Google fleet using Memory Allocation Profiling.

The discussion will cover ongoing improvements to reduce the overhead of this feature (minimizing metadata), enhance support for modules (decrease overhead when allocations persist at unload), improve observability (provide access to certain GFP flags data), adding context capture for select allocations and covering more allocators.

Initial discoveries will be based on our experiences deploying memory allocation profiling on a portion of the Google fleet. We will provide an analysis of the collected data, focusing on reducing kernel memory overheads.

The desired outcome of this discussion is to identify a reduction plan for the top allocation call sites and determine which other call sites to investigate next.

Primary authors: OVERSTREET, Kent; TATASHIN, Pasha; PANDA, Sourav (Google); BAGHDASARYAN, Suren

Presenters: OVERSTREET, Kent; TATASHIN, Pasha; PANDA, Sourav (Google); BAGHDASARYAN, Suren

Session Classification: Kernel Memory Management MC

Track Classification: Kernel Memory Management MC