



Contribution ID: 44

Type: **not specified**

Launching new processes with ‘io_uring_spawn’ for fast builds

Monday, September 12, 2022 10:45 AM (45 minutes)

io_uring allows running a batch of operations fast, on behalf of the current process. As the name suggests, this works exceptionally well for I/O workloads. However, one of the most prominent workloads in software development involves executing other processes: make and other build systems launch many other processes over the course of a build. How can we launch those processes faster?

What if we could launch other processes, and give them initial work to do using io_uring, ending with an exec? What if we could handle the pre-exec steps for a new process *entirely* in the kernel, with no userspace required, eliminating the need for fork or even vfork, and eliminating page-table CoW overhead?

In this talk, I'll introduce io_uring_spawn, a mechanism for launching empty new processes with an associated io_uring. I'll show how the kernel can launch a blank process, with no initial copy-on-write page tables, and initialize all of its resources from an io_uring. I'll walk through both the successful path and the error-handling path, and show how to get information about the launched process. Finally, I'll show how existing userspace can take advantage of io_uring_spawn to speed up posix_spawn, and provide performance numbers for common workloads, including kernel compilation.

I agree to abide by the anti-harassment policy

Yes

Primary author: TRIPLETT, Josh

Presenter: TRIPLETT, Josh

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track