Linux Plumbers Conference 2022



Contribution ID: 316

Type: not specified

Io_uring command/nvme passthrough: current status and plans

Wednesday, 14 September 2022 15:00 (1h 30m)

Io_uring command is a new async-ioctl-like facility to attach io_uring capabilities to any arbitrary command implemented by the underlying provider (driver, filesystem, etc.). The first use case of the construct is to implement a new passthrough path to NVMe. This path guarantees both availability and scalability. It helps both the early adopters of NVMe and the kernel community as emerging storage interfaces can be consumed and user-space stack can evolve before cementing changes to other mature parts of the kernel (e.g. syscall, filesystem, block-layer, etc.)

While initial support got merged into 5.19 Kernel, we are receiving a bunch of new voices from users. In this BoF we go over to cover

- Io_uring command: current capabilities, elements (big SQE, big CQE) and generic communication-model
- · Modern nvme passthrough: how it compares against the block-path and legacy passthrough path
- Bunch of currently missing capabilities, either at the io_uring command level or at nvme-level. While patches are out for a few, some are under development or at the planning stage. A discussion can help in establishing a path forward for newer elements in this path e.g. multipath, iopoll, fixedbufs, admin-only constraint, io-size limit etc.

I agree to abide by the anti-harassment policy

Yes

Primary author: JOSHI, kanchan (Samsung semiconductor (SSIR))Presenter: JOSHI, kanchan (Samsung semiconductor (SSIR))Session Classification: BOFs Session

Track Classification: Birds of a Feather (BoF)