

GPU Recovery

Rodrigo Siqueira, Timur Kristóf

XDC 2025



Topics Overview

Sharing an overview of the current status of GPU recovery. The overall idea follows these steps:

1. When a job cause a hang, call driver specific handler.
2. Stop the scheduler.
3. Handle the bad job, and keep the good jobs.
4. Reset the ring buffer.
5. Resume the scheduler.

Discussion around `drm_sched_resubmit_jobs()`

- Echoing the idea that `drm_sched_resubmit_jobs()` is now deprecated.
- Generally, drivers should have their own specific code for handling resubmission jobs and not use `drm_sched_resubmit_jobs()`.
- We had some discussion about abstracting some of the driver-specific implementation to the scheduler.

GPU Reset debug

- Discussion around GPU reset limitations, such as the impossibility of confidently pinpointing the instruction that caused the issue.
- Get hints from userspace to improve the devcoredump output.

Avoid hang regressions

- Introduce IGT tests that exercise driver behaviour when dealing with bad instructions.
- Find new test cases to expand the HangTest suite.