

cgroups for GPU resources

Currently there is no notion of cgroup accounting for GPU memory and execution. Discuss how we could integrate this with GEM/TTM memory management, including how to correctly account for allocations which are transferred between processes (e.g. Android gralloc-as-a-service), and integrating scheduler/runtime constraints with hardware-based scheduling on newer hardware designs.

I agree to abide by the anti-harassment policy

Primary author: STONE, Daniel (Collabora)

Presenter: STONE, Daniel (Collabora)

Session Classification: GPU/media/AI buffer management and interop MC