

Cgroup v1/v2 Abstraction Layer

Monday, September 20, 2021 7:30 AM (25 minutes)

The Oracle database offers a long-term-stable version that is supported and maintained for many years. But as Linux distributions slowly transition from cgroup v1 to cgroup v2, this creates a challenge for the DB. cgroup v1 and cgroup v2 have different interfaces and best practices.

This talk is to discuss the current status of the cgroup abstraction layer, how applications like the Oracle database plan to use it, and gather/discuss other users' requirements for this layer.

- cgroup v2 support added to libcgroup - DONE [1]
- cgroup v1/v2 abstraction layer - OUT FOR REVIEW [2]
- This layer hides the underlying details of cgroup v1 vs v2 (e.g. cpu.weight vs cpu.shares). The user can make a request in v1 or v2 format, and the layer will do the proper translation. The user still needs to have intimate cgroup and system knowledge
- Higher-level abstraction layer - GATHERING REQUIREMENTS
- The goal of this layer is to abstract away the gory cgroup details and let the user specify their needs (e.g. 2 CPUs that are side channel attack safe). This layer should also handle interactions with systemd (dbus, cgroup delegation, etc.)

[1] <https://github.com/libcgroup/libcgroup/releases/tag/v2.0>

[2] <https://sourceforge.net/p/libcg/mailman/message/37317688/>

I agree to abide by the anti-harassment policy

I agree

Primary author: HROMATKA, Tom

Presenter: HROMATKA, Tom

Session Classification: Containers and Checkpoint/Restore MC

Track Classification: Containers and Checkpoint/Restore MC