

Linux Plumbers Conference 2024



Contribution ID: 175

Type: **not specified**

Isolated user namespaces & cgroupfs

Thursday, 19 September 2024 15:50 (20 minutes)

This talk is about a problem of integration between the concept of an “isolated” ([1], [2], [3], [4]) user namespace and cgroup-v2 delegation model.

The biggest challenge here is that cgroup delegation is based on cgroupfs inodes ownership and cgroupfs superblock is shared between all containers which makes it impossible to deal with cgroupfs as with any other containerized filesystem like procfs or tmpfs.

[1] More flexible user namespaces <https://fosdem.org/2024/schedule/event/fosdem-2024-2987-more-flexible-user-namespaces/>

[2] User namespaces with host-isolated UIDs/GIDs <https://lpc.events/event/17/contributions/1569/>

[3] Isolated dynamic user namespaces <https://lpc.events/event/7/contributions/836/>

[4] Simplified user namespace allocation <https://lpc.events/event/11/contributions/982/>

Primary author: MIKHALITSYN, Aleksandr (Canonical)

Co-author: GRABER, Stéphane (Zabbly)

Presenter: MIKHALITSYN, Aleksandr (Canonical)

Session Classification: Containers and checkpoint/restore MC

Track Classification: Containers and checkpoint/restore MC