

Checkpoint-restoring containers with Docker inside

Monday, August 24, 2020 8:40 AM (20 minutes)

CRIU is the most advanced Checkpoint-Restore project on Linux.

But even with CRIU at the moment it is not feasible to checkpoint - restore all possible topologies of processes and namespaces. Even relatively simple case of a process tree with two UTS/IPC namespaces is not supported by CRIU, not mentioning more complex cases like a process tree with more than one PID namespaces.

In OpenVZ and Virtuozzo versions of CRIU these problems were partially solved with introduction of the support for nested PID namespaces, several IPC/UTS namespaces (with respect to USER namespaces) and overlayfs mounts.

These improvements allow us to get basic support of checkpoint-restoring OpenVZ system containers with Docker containers inside.

We have already prepared several upstream kernel patches [4].

I agree to abide by the anti-harassment policy

I agree

Primary authors: MIKHALITSYN, Alexander (Virtuozzo); TIKHOMIROV, Pavel

Presenters: MIKHALITSYN, Alexander (Virtuozzo); TIKHOMIROV, Pavel

Session Classification: Containers and Checkpoint/Restore MC

Track Classification: Containers and Checkpoint/Restore MC