

Mount-v2 CRIU migration engine: status update

Monday, 20 September 2021 08:10 (25 minutes)

Mount checkpoint/restore is an important part of CRIU, it is responsible for consistency of the file system view of dumped processes. In current state we can only restore simple mount configurations, something more complex would either make us fail or, which is even worse, make us creating wrong file system view for restored processes.

In CRIU we only see the final state, the result of probably multiple kernel API calls inside a container, and on restore we need to recreate the sequence of calls which would lead to the exact same state, in general this task can be very complex. So sometimes the only way is to simplify the API so that it become easier to restore all possible configurations.

Last year [1] we discussed a variety of problems CRIU faces with mounts, most important ones related to mount propagation and how to simplify mount propagation configuration so that even complex setups can be re-created simply and correctly.

This talk will start with showing more complex mount configurations to demonstrate that that we still need an API change. Then there will be a status update on kernel patch progress and changes that were done during the last year followed by the discussion on how to make the patch [2] mergeable to the upstream kernel.

Here are links:

- 1 <https://www.linuxplumbersconf.org/event/7/contributions/640/>
- 2 <https://lore.kernel.org/linux-api/20210715100714.120228-1-ptikhomirov@virtuozzo.com/>
 - mount-v2 draft for criu - <https://github.com/Snorch/criu/commits/mount-v2-poc>
 - mount-v2 for VZ criu - <https://src.openvz.org/projects/OVZ/repos/criu/browse/criu/mount-v2.c>

Thanks to Andrei Vagin and Christian Brauner for a great help with it!

I agree to abide by the anti-harassment policy

I agree

Primary author: TIKHOMIROV, Pavel (Virtuozzo)

Presenter: TIKHOMIROV, Pavel (Virtuozzo)

Session Classification: Containers and Checkpoint/Restore MC

Track Classification: Containers and Checkpoint/Restore MC