

Secrets in cloned snapshots

Monday, 20 September 2021 08:35 (25 minutes)

Starting things is slow. Even if only 1 second slow, saving 1s on a million container restores means we can save 11 days of useless work that every container will perform identically.

That's where snapshots come in. Snapshots in theory allow us to save an initialized container once, but then restore it a million times at less overhead than cold starting it takes.

Unfortunately, Linux applications (and the kernel in VM based container setups) expect that during their lifetime they don't get cloned from the outside. Applications create user space PRNGs (Pseudo Random Number Generators) which would generate identical random numbers after a clone. They create UUIDs that would no longer be unique. They generate unique temporary key material that is no longer unique.

Eventually, if we want to enable cloning properly, user space applications will need to learn that they have to adapt to clone events. For that they need notifications.

This session will discuss the requirements such a notification mechanism has as well as possible paths forward to implement it and drive adoption.

References:

- <https://github.com/systemd/systemd/issues/19269>
- <https://lkml.org/lkml/2021/3/8/677>

I agree to abide by the anti-harassment policy

I agree

Primary authors: GRAF, Alexander; Mr CATANGIU, Adrian

Presenters: GRAF, Alexander; Mr CATANGIU, Adrian

Session Classification: Containers and Checkpoint/Restore MC

Track Classification: Containers and Checkpoint/Restore MC