



Contribution ID: 172

Type: **not specified**

## #snapsafe: restoring uniqueness in Virtual Machine clones

*Monday 12 September 2022 15:25 (40 minutes)*

### short version

When a virtual machine gets cloned, it still contains old data that believes are unique - random number generation seeds, UUIDs, etc. Linux recently included support for VMGenID to reseed its in-kernel PRNG, but all other RNGs and UUIDs are still identical after a clone.

In this session, we will discuss approaches to solve this and reveal experiments on which we worked on, such as creating a user space readable system generation counter and going through a systemd inhibitor list for pre-snapshot/post-snapshot phases.

### long(er) version

Linux recently added support for the Virtual Machine Generation ID (VMGenID) feature, an emulated device that informs the guest kernel about VM restore events by exposing a 128-bits UUID which changes every time a VM is restored from a snapshot. The kernel uses the UUID to reseed its PRNG, thus de-duplicating the PRNG state across VMs.

Although, VMGenID definitely works towards the correct direction, it does not provide a mechanism for notifying user-space applications of VM restore events. In this presentation, we introduce Virtual Machine Generation Counter, an extension to vmgenid which provides a low-latency and race-free mechanism for communicating restore events to user-space. Moreover, we will speak about why VM Generation Counter is not enough for ensuring across-the-stack snapshot safety. We will present an effort which builds on top of Systemd inhibitor locks to make snapshot-restore cycle a first-class citizen in the life-cycle of a system, achieving end-to-end snapshot safety

## I agree to abide by the anti-harassment policy

Yes

**Primary author:** CHALIOS, Babis (Amazon Web Services)

**Co-author:** GRAF, Alexander

**Presenter:** CHALIOS, Babis (Amazon Web Services)

**Session Classification:** Service Management and systemd MC

**Track Classification:** LPC Microconference: Service Management and systemd MC