Contribution ID: **196** Type: **not specified**

seccomp filtering for struct pointers

Thursday 19 September 2024 16:10 (20 minutes)

With the introduction of extensible-struct syscalls such as openat2 and clone3, the inability to usefully filter syscalls with pointer arguments makes it harder for various programs to make use of newer kernel features because of both default container and self-hardening seccomp profiles. The inability for systemd and other system utilities to use RESOLVE_IN_ROOT and related openat2 features is a particular issue.

This talk will describe a proposal for an extension to seccomp to allow for the filtering of extensible-struct syscalls on an opt-in basis, as well as some of the potential issues with creating forward-compatible filters due to the restrictions of cBPF and some possible solutions.

Primary author: SARAI, Aleksa (SUSE LLC)

Presenter: SARAI, Aleksa (SUSE LLC)

Session Classification: Containers and checkpoint/restore MC

Track Classification: Containers and checkpoint/restore MC