

seccomp feature development

Wednesday, 26 August 2020 09:00 (45 minutes)

As outlined in <https://lore.kernel.org/lkml/202005181120.971232B7B@keescook/> the topics include:

- fd passing
- deep argument inspection
- changing structure sizes
- syscall bitmasks

Specifically, seccomp needs to grow the ability to inspect Extensible Argument syscalls, which requires that it inspect userspace memory without Time-of-Check/Time-of-Use races and without double-copying. Additionally, since the structures can grow and be nested, there needs to be a way to deal with flattening the arguments into a linear buffer that can be examined by seccomp's BPF dialect. All of this also needs to be handled by the USER_NOTIF implementation. Finally, fd passing needs to be finished, and there needs to be an exploration of syscall bitmasks to augment the existing filters to gain back some performance.

I agree to abide by the anti-harassment policy

I agree

Primary author: COOK, Kees (Google)

Presenter: COOK, Kees (Google)

Session Classification: Kernel Summit

Track Classification: Kernel Summit