Contribution ID: **250** Type: **not specified**

Improving *at(2) to make more secure container runtimes

Tuesday, 13 November 2018 11:35 (15 minutes)

Currently, container runtimes are faced with a large attack surface when it comes to a malicious container guest. This most obvious attack surface is the filesystem, and the wide variety of filesystem races and other such tricks that can be used to trick a container runtime into accessing files they shouldn't. To tackle this, most container runtimes have come up with necessary userspace hacks to work around these issues – but many of the improvements are inherently flawed as they are not done from kernel-space.

In this session, a discussion of the various kernel APIs that could benefit container runtime security will be opened. Topics on the agenda would be the use of AT_EMPTY_PATH with openat(2), whether there are any more blockers for the AT_NO_JUMPS patchset, and a proposal of AT_THIS_ROOT which would allow for much more secure interaction with
container filesystems.

I agree to abide by the anti-harassment policy

Presenter: BRAUNER, Christian (Canonical)

Session Classification: Containers MC