



Contribution ID: 165

Type: **not specified**

Beyond per-CPU atomics and rseq syscall: subset of eBPF bytecode for the do_on_cpu syscall

Wednesday, 11 September 2019 17:40 (20 minutes)

The Restartable Sequences system call [1,2,3,4] introduced in Linux 4.18 has limitations which can be solved by introducing a bytecode interpreter running in inter-processor interrupt context which accesses user-space data.

This discussion is about the subset of the eBPF bytecode and context needed by this interpreter, and extensions of that bytecode to cover load-acquire and store-conditional memory accesses, as well as memory barrier instructions. The fact that the interpreter needs to allow loading data from userspace (tainted data), which can then be used as address for loads and stores, as well as conditional branches source register, will also be discussed.

[1] "PerCpu Atomics" <http://www.linuxplumbersconf.org/2013/ocw/system/presentations/1695/original/LPC%20-%20PerCpu%20Atomics.pdf>

[2] "Restartable sequences" <https://lwn.net/Articles/650333/>

[3] "Restartable sequences restarted" <https://lwn.net/Articles/697979/>

[4] "Restartable sequences and ops vectors" <https://lwn.net/Articles/737662/>

I agree to abide by the anti-harassment policy

Yes

Primary author: DESNOYERS, Mathieu (EfficiOS Inc.)

Presenter: DESNOYERS, Mathieu (EfficiOS Inc.)

Session Classification: BPF MC