Linux Plumbers Conference 2022

Contribution ID: **275**

Type: **not specified**

A Clash of Things: Name Ambiguity, Tracing and kallsyms

Monday, 12 September 2022 10:00 (1h 30m)

A brief look at sort -k 3 /proc/kallsyms makes it clear that the kernel symbol table is stuffed full of duplicate names with distinct addresses. These are all symbols in distinct translation units, so it would presumably be nice for tracers to be able to distinguish between them (e.g. to trace only specific functions known to be called from TUs of interest). Right now, for symbols in the core kernel, there is no way to tell these things apart: you can either trace all of them, or none, or guess and hope. Any tracing tool clearly shows this by not being able to disambiguate.

A few years ago a kallmodsyms "idea/feature" was proposed upstream which does most of what is needed to solve this ambiguity problem at minimal space and time cost. With additional not-very-large extensions it could solve it completely, disambiguating every function symbol in /proc/kallsyms from every other. A few iterations of the patch were made [1]. This additional symbol information would also be available through the BPF ksym iterator[2] that was recently added. But more issues remain to be resolved.

This BOF will show how kallmodsyms works and what it does and how it might be useful, accompanied by a few examples to illustrate the problems that it is trying to solve. Discussion will follow to figure out what else is necessary for acceptance upstream.

I agree to abide by the anti-harassment policy

Yes

Primary author: ALCOCK, Nick (Oracle Corporation)

Presenter: ALCOCK, Nick (Oracle Corporation)

Session Classification: BOFs Session

Track Classification: Birds of a Feather (BoF)