

Generic functionality for system call and trap entry and exit

Tuesday, 25 August 2020 09:00 (25 minutes)

The system call entry and exit code is needlessly duplicated and different in all architectures. The work carried after the real low level ASM bits should not be different across architectures as well as the code that handles the pending work before returning from a system call to user space. Likewise, the interrupt and exception handling has to establish the state for various kernel subsystems like lockdep, RCU and tracing and there is no good reason to have twenty-some similar and pointlessly different implementations.

A common infrastructure for kernel entry handling was merged in v5.9 release cycle and for now it is only used by x86.

Let's discuss how this infrastructure is adopted by other architectures.

I agree to abide by the anti-harassment policy

I agree

Primary author: GLEIXNER, Thomas

Presenter: GLEIXNER, Thomas

Session Classification: linux/arch/* MC

Track Classification: linux/arch/* MC