



Contribution ID: 197

Type: **not specified**

Unifying and improving test regression reporting and tracking

Monday, 13 November 2023 12:25 (35 minutes)

The current CI systems for the kernel offer basic and low-level regression detection and handling capabilities based on test results, but they do that in their own specific way. We wonder if we can find more common ways of tackling the problem through post-processing the data provided by the different CI systems. We could then extract additional “hidden” information, look at failure trends, analyze and compare logs, as well as understand the impact of the test infra setup in a failure. This would allow us to address the need to track the reported regressions status, creating a standard workflow across the different tools.

Regzbot is proving itself as part of this standard workflow, but reporting the regressions from CI systems there is hard as the volume could be high, the relevance of some regressions could be low, and false positives do exist (a lot - specially for hardware testing). Better post-processing of the data and more standard workflows can help more reports to get to the community without destroying the quality of the data provided by regzbot.

We’d like to briefly talk about the current status and limitations of the testing ecosystem regarding regressions, exhibit the current efforts and strategies being developed, and start a discussion about contributions and plans for the future.

Primary authors: PADOVAN, Gustavo (Collabora); CAÑUELO, Ricardo

Presenters: PADOVAN, Gustavo (Collabora); CAÑUELO, Ricardo

Session Classification: Kernel Testing & Dependability MC

Track Classification: LPC Microconference: Linux Kernel Debugging MC