Contribution ID: **45**   Type: **not specified**

Adding benchmark results support to KTAP/kselftest

Friday 20 September 2024 10:00 (30 minutes)

Benchmark test results are difficult to interpret in an automated fashion. They often require human interpretation to detect regressions because they depend on a number of variables, including configuration, cpu count, processor speed, storage speed, memory size, and other factors. Tim proposes a new system for managing benchmark data and interpretation in kselftest. It consist of 3 parts: 1) adding syntax to KTAP to support a consistent format for benchmark values in KTAP/kselftest test output, 2) the use of a set of criteria, external to the test itself, for interpreting benchmark result values, and 3) an automated tool to determine and set appropriate reference values to use in the test result criteria. A prototype system will be demonstrated, that supports converting benchmark values into KTAP "ok" and "not ok" results, consumable by humans and automated tools (such as CI systems). This system is intended to enable the detection of regressions in benchmark outputs, using appropriate threshold values that are customizable (in an automated fashion) by a tester for their own configuration and hardware.

Primary author: BIRD, Tim (Sony)

Presenter: BIRD, Tim (Sony)

Session Classification: Kernel Testing & Dependability MC

Track Classification: Kernel Testing & Dependability MC