

Notes About These Tests[#]

Each of these test programs basically uses a kernel call hundreds or thousands of times, and records the average time. Even with this, there is some margin of error - usually about 1%, if the tests are run repeatedly. The tests for interrupt/scheduling latency seem to be particularly inconsistent, and shouldn't be interpreted as absolute truth.

For the x86 sheet, all the percentages up to the "Improvement" columns are against the stock Watcom kernel build from the Trinity release.

Wiki syntax is lacking compared to html, so the actual chart is attached below.