

Ongoing Development in rteval to measure realtime latency

John Kacur



What is rteval?

- rteval is a python program that runs a measurement module such a cyclicttest to measurement real time latency while running various load modules to stress a system.

What are load modules

- Load modules try to stress a system and see if that will interfere with realtime performance.
- There are built-in load modules such as kcompile that builds the linux kernel with allmodconfig
- There are load modules that run external programs such as stressng that runs stress-ng
 - stressng is an exclusive load-module meaning no other load modules run when stressng is running

What are measurement modules?

- We used to have only one measurement module, cyclicttest, which ran cyclicttest and parsed the histogram output saving the results in an xml file
- By default load and measurement modules are run on all cpus
- The user has the ability to specify which cpus run load modules and which cpus run measurement modules

New rteval measurement module - timerlat

- Recently there is a new measurement module – rta timerlat
- The output to timerlat is similar to cyclictst
 - Timerlat breaks down the latency into IRQ, kernel-thread and user-thread
- Tracing capabilities of timerlat are better than cyclictst

rteval uses

- Red Hat uses rteval to certify customer's hardware for realtime purposes
- Customer's typically want to divide load modules and measurement modules amongst the various cpus
- You can think of load modules as representing non-rt programs, and measurement modules as rt-programs
- Why not use rteval to determine how to distribute loads and still achieve adequate rt latency?

Recent innovations - isolcpus

- `--measurement-run-on-isolcpus`
Include isolated CPUs in default cpulist
- Added by Tomas Glozar

cpuset

- Cpusets are being added as an alternative to isolcpus

Power Saving

- Power saving. Instead of writing a zero to `/dev/cpu_dma_latency/dev/cpu_dma_latency`, can we achieve good latency results with power saving.
- Multiple ongoing efforts
- `--idle-set IDLESTATE`
 - Idle state depth to set on cpus running measurement modules (Anubhav Shelat)
- Tomax Glozar looking at using `cpupower` in `rtla timerlat`
- John Wyatt looking at adding python bindings to `cpupower`

Partitioning machines

- Cgroups
- Containers (Chris White is working on this)