

How do we kick our RT habit?

Monday, August 24, 2020 7:55 AM (25 minutes)

Inside our large database application setup, we have a few critical processes. Some of the functions include, heartbeat (for the cluster), monitoring what was happening (to debug in case a cluster does go down) amongst others.

Elaborating on a single example, if the heartbeat process doesn't run when it should, the cluster could remove the node, and then the node would have to shutdown, which would then need the monitoring process to do more work to identify why we failed.

Clearly a database consumes a lot of CPU, and so these critical processes became RT and have been RT for a long time. With containers coming in, and RT cgroups being sub-optimal, maybe it is time to revisit this decision. We have some observations. Are these RT processes? Maybe not in the strict academic RT sense, but these are critical, time sensitive processes (with a deterministic function). Or does SCHED_OTHER need to be fixed for a clearly SCHED_OTHER problem?

Helmets advised for this discussion!

I agree to abide by the anti-harassment policy

I agree

Primary authors: GIANI, Dhaval (Oracle); SANGAPPA, Prakash (Oracle)

Presenters: GIANI, Dhaval (Oracle); SANGAPPA, Prakash (Oracle)

Session Classification: Real-time MC

Track Classification: Real-time MC