

Linux Plumbers Conference 2024



Contribution ID: 264

Type: **not specified**

Improving guest latency & throughput by improving RCU in KVM

Thursday, 19 September 2024 16:28 (8 minutes)

In the mission of reducing latency in KVM guests, we have seen a lot of missed deadlines caused by RCU core invocation, often causing guest exit only to have a timer interrupt invoking `rcu_core()` on host and causing a task switch.

While looking to improve that, it was noticed that no RCU lock is held in guest context, and thus it's possible to report a quiescent state in guest exit, avoiding the whole `rcu_core()` invocation.

In this presentation, there will be more details on how it was possible to improve both latency and throughput in the guest, presented with the numbers we got from standard tests.

Part of this work also caused the creation of a new RCU boot parameter that can reduce latency even in RT workloads ran in host context, which will also be discussed in the presentation.

Primary author: SOARES PASSOS, Leonardo Bras (Red Hat)

Presenter: SOARES PASSOS, Leonardo Bras (Red Hat)

Session Classification: Real-time MC

Track Classification: Real-time MC