



Contribution ID: 201

Type: **not specified**

C-state latency measurement infrastructure

Tuesday, 10 September 2019 17:50 (25 minutes)

We in Intel developed instrumentation for measuring C-state wake latency. The instrumentation, which we call “waltr” (WAKE up Latency Tracer) consists of user-space and kernel modules parts.

In principle, waltr works by scheduling delayed interrupts and measuring the wake latency close to the x86 ‘mwait’ x86 instruction. This requires an external device equipped with high precision clock and capable of delayed interrupts. We have been

using the Intel i210 Ethernet adapter for these purposes. But theoretically this could be a completely different device, e.g., a GFX card.

The C-state latency measurement instrumentation should be very useful for the open-source community and we would like to upstream the kernel parts of it. We are seeking for feedback on how to properly modify the kernel in a maintainable and reusable way, to benefit everyone.

Here are few examples for the dilemmas have.

- * How do we design a framework for compliant devices like the i210 adapter?
- * What would be the right user-space API for the delayed interrupts provider?
- * How do we take snapshots of C-state counters and deliver them to user-space?

I am asking for a 20-30 minutes time-slot. And I am hoping to talk to people more about this in hallway discussions.

I agree to abide by the anti-harassment policy

Yes

I confirm that I am already registered for LPC 2019

Primary author: BITYUTSKIY, Artem

Presenter: BITYUTSKIY, Artem

Session Classification: Power Management and Thermal Control MC