



Contribution ID: 46

Type: **not specified**

Managing Memory Bandwidth Antagonism at Scale

Wednesday, 14 November 2018 16:00 (45 minutes)

Providing a consistent and predictable performance experience for applications is an important goal for cloud providers. Creating isolated job domains in a multi-tenant shared environment can be extremely challenging. At Google, performance isolation challenges due to memory bandwidth has been on the rise with newer workloads. This talk covers our attempt to understand and mitigate isolation issues caused by memory bandwidth saturation.

The recent Intel RDT support in Linux helps us both monitor and manage memory bandwidth use on newer platforms. However, it still leaves a large chunk of our fleet at risk of memory bandwidth issues. The talk covers three aspects of our isolation attempts:

1. At Google and Borg, we run all application in containers. Our first attempt was to estimate memory bandwidth utilization for each container on all supported platform by using existing performance counters. The talk will cover details on our approximation methodology and issues we identified in monitoring as well as some usage trends across different workloads.
2. The second part of our effort was focussed on building actuators and policies for memory bandwidth control. We will cover multiple iterations of our enforcement efforts at node and cluster level with production use-cases and lessons learnt.
3. For newer platforms, we attempted to use Intel RDT support via the resctrl interface. We ran into issues on both the monitoring and isolation side. We'll discuss the fixes and workarounds we used and changes we proposed for resource-control support in Linux.

We believe the problems and trends we have observed are universally applicable. We hope to inform and initiate discussion around common solutions across the community.

I agree to abide by the anti-harassment policy

Yes

Primary authors: Mr JNAGAL, Rohit (Google Inc); Mr LO, David (Google Inc); Mr SBIRLEA, Dragos (Google Inc)

Presenters: Mr JNAGAL, Rohit (Google Inc); Mr LO, David (Google Inc); Mr SBIRLEA, Dragos (Google Inc)

Session Classification: LPC Main Track

Track Classification: Refereed talk