



Contribution ID: 140

Type: **not specified**

TCP memory isolation on multi-tenant servers

Tuesday, 13 September 2022 15:45 (45 minutes)

On Linux, `tcp_mem` `sysctl` is used to limit the amount of memory consumed by active TCP connections. However that limit is shared between all the jobs running on the system. Potentially a low priority job can hog all the available TCP memory and starve the high priority jobs collocated with it. Indeed we have seen production incidences of low priority jobs negatively impacting the network performance of collocated high priority jobs.

Through `cgroups`, Linux does provide TCP memory accounting and isolation for the jobs running on the system but that comes with its own set of challenges which can be categorized into two buckets:

1. New and unexpected semantics of memory pressure and OOM for `cgroup` based TCP memory accounting.
2. Logistical challenges related to resource and quota management for large infrastructures running millions of jobs.

This is an ongoing work and new challenges keep popping up as we expand `cgroup` based TCP memory in our infrastructure. In this presentation we want to share our experience in tackling these challenges and would love to hear how others in the community have approached the problem of TCP memory isolation on their infrastructure.

I agree to abide by the anti-harassment policy

Yes

Primary authors: WARLOE, Christian (Google); BUTT, Shakeel (Google); WANG, Wei (Google)

Presenters: WARLOE, Christian (Google); BUTT, Shakeel (Google); WANG, Wei (Google)

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track