



Contribution ID: 281

Type: **not specified**

An Evaluation of Host Bandwidth Manager

Wednesday, 11 September 2019 12:45 (45 minutes)

Host Bandwidth Manager (HBM) is a BPF based framework for managing per-cgroupv2 egress and ingress bandwidths in order to provide a better experience to workloads/services coexisting within a host. In particular, HBM allows us to divide a host's egress and ingress bandwidth among workloads residing in different v2 cgroups. Note that although sample BPF programs are included in the BPF patches, one can easily use different algorithms for managing bandwidth.

This talk presents an evaluation of HBM and associated BPF programs. It explores the performance of various approaches to bandwidth management for TCP flows that use Cubic, Cubic with ECN or DCTCP for their congestion control. For evaluating performance, we consider how well flows can utilize the allocated bandwidth, how many packets are dropped by HBM, increases to RTTs due to queueing, RPC size fairness, as well as RPC latencies. This evaluation is done independently for egress and ingress. In addition, we explore the use of HBM for protecting against incast congestion by also using HBM on the root v2 cgroup.

Our testing shows that HBM, with the appropriate BPF program, is very effective at managing egress bandwidths regardless of which TCP congestion control algorithm is used, preventing flows from exceeding the allocated bandwidth while allowing them to use most of their allocation. Not surprisingly, effectively managing ingress bandwidth requires ECN, and preferably DCTCP. Finally, we show that using HBM is very effective at preventing packet losses due to incast congestion, as long as we are willing to sacrifice some ingress bandwidth.

I agree to abide by the anti-harassment policy

Yes

I confirm that I am already registered for LPC 2019

Primary author: BRAKMO, Lawrence (Facebook)

Presenter: BRAKMO, Lawrence (Facebook)

Session Classification: Networking Summit Track