



Contribution ID: 308

Type: **not specified**

Percpu hashtable traversal measurement study

Tuesday, 13 September 2022 18:00 (30 minutes)

As platforms grow in cpu count (200+ cpu), using per cpu data structures is becoming more and more expensive. Copying the percpu data from the bpf hashtable map to userspace buffers can take up to 22 us per entry on a platform with 256 cores.

This talk presents a detailed measurement study of the cost of percpu hashtable traversal, covering various methods and systems with core counts.

We will discuss how the current implementation of this data structure makes it hard to amortize cache misses, and solicit proposal for possible enhancements.

I agree to abide by the anti-harassment policy

Yes

Primary author: VAZQUEZ, Brian (Google)

Co-authors: RIZZO, Luigi (Google); FOMICHEV, Stanislav (Google); SHARMA, Sagarika (Columbia University)

Presenter: VAZQUEZ, Brian (Google)

Session Classification: eBPF & Networking

Track Classification: eBPF & Networking Track