



Contribution ID: 291

Type: **not specified**

## connect() - why you so slow?!

*Wednesday, 15 November 2023 14:30 (30 minutes)*

What happens when your application opens upwards of 50k connections to a single destination? Short answer - connect() syscall becomes slow. Cloudflare found out the hard way.

Through this talk we would like to share our story of what we have learned about connect() implementation for TCP in Linux, both its strong and weak sides. How connect() latency changes under pressure, and how to open connection so that the syscall latency is deterministic and time-bound.

In this talk we would like to cover:

- Why Cloudflare services sometimes experience pressure, where we need to open lots of connections to just one destination.
- How we have been avoiding the connect() latency pitfall so far, and why it is no longer a viable option.
- Our efforts to benchmark connect() syscall and characterize its latency as the the number of open connections increases.
- Existing difficulties in tracing and monitoring connect() performance at scale in a production environment.
- A look at how connect() is implemented in Linux for TCP; its evolution and previous attempts dealing with high-latency under pressure.
- How to control how connect() takes with existing Linux APIs - recipes for how to open TCP connections with predictable syscall latency.

**Primary author:** LAWLER, Frederick (Cloudflare)

**Presenter:** LAWLER, Frederick (Cloudflare)

**Session Classification:** eBPF & Networking

**Track Classification:** eBPF & Networking Track