Linux Plumbers Conference 2023



Contribution ID: 297

Type: not specified

Evolution of Direct Server Return (DSR) implementation for containerized applications

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

The industry extensively relies on direct server response, DSR, and Meta has a long history of employing this technology for L4 load balancing. At the same time, our fleet went through an evolution of being an isolated subset of machines per team, to a more efficient model with a single shared pool that provides multi-tenant capacity. Moving services to network namespace becomes necessary to implement stackable workload solutions allowing multiple services running on the same ports to be able to schedule on the same host. Our approach for DSR has transformed together with those changes.

Initially, network service running in datacenters used a "rootlet", system wide XDP program array to jump sequentially from one XDP program to another. It's an in-house built solution to attach XDP program, but has its own limitations making it not work well with shared hosts. To migrate traffic services to shared hosts, we built XDP Chainer, an in-house solution to attach multiple XDP programs on a single interface, which also addresses some of the shortcomings of the "rootlet"-based solution.

However, the introduction of multi-tenancy and network namespaces has brought new challenges, requiring reevaluation of how we facilitate decapsulation support for backends. This evolution involves the migration of the decapsulation data path to the TC-BPF solution. This approach is tightly integrated with internal container orchestration at Meta's and requires no configuration and less privileges from users.

During our presentation, we will address the challenges encountered, discuss the alternatives considered, describe wins that we achieved with a new approach and reflect on the lessons learned throughout the duration of this project

Alternative Titles:

- * Direct Server Return (DSR) and Multi-tenancy: obstacles and solutions
- * Adaptation of BPF implementation for Direct Server Return in response to containers evolution

Primary authors: GUPTA, Lalit (Meta); DUBOVITSKY, Pavel; SHUKHAU, Raman
Presenters: GUPTA, Lalit (Meta); DUBOVITSKY, Pavel; SHUKHAU, Raman
Session Classification: eBPF & Networking

Track Classification: eBPF & Networking Track