Contribution ID: **103** Type: **not specified**

How we built Magic Transit

Thursday, August 27, 2020 9:45 AM (45 minutes)

In this talk we will present Magic Transit, Cloudflare's layer 3 DDoS protection service, as a case study in building a network product from the standard linux networking stack. Linux provided us with flexibility and isolation that allowed us to stand up this product and on-board more than fifty customers within a year of conceptualization. Cloudflare runs all of our services on every server on our edge, and Magic Transit is not an exception to that rule - one of our biggest design challenges was working a layer 3 product into a networking environment tuned for proxy and server products. We'll cover how we built Magic Transit, what worked really well, and what challenges we encountered along the way.

Magic Transit is largely implemented as a "configurator", that is our software manages the network setup, and lets the kernel do the heavy lifting with network namespaces, policy routing and netfilter to safely direct and scrub IP traffic for our customers. This design allows drop-in integration with our DDoS protection systems, and our proxying and L7 products, and in a way that our operations team was familiar with. These benefits do not come without their caveats; specifically route placement/reporting inconsistencies, quirks revolving around icmp packets being generated from within a namespace when fragmentation occurs, problems stemming from conntrack and a mystery around offload···Finally we'll touch on our future plans to migrate our web of namespaces to a Rust service that makes use of ebpf/xdp.

I agree to abide by the anti-harassment policy

I agree

Primary authors: HEINE, Erich (Cloudflare); JONES, Connor (Cloudflare)

Presenters: HEINE, Erich (Cloudflare); JONES, Connor (Cloudflare)

Session Classification: Networking and BPF Summit

Track Classification: Networking & BPF Summit