Calico Networking with eBPF

Shaun Crampton, Core Developer for Project Calico
Chris Hoge, Developer Advocate for Project Calico

What prompted the team to add another dataplane to Calico?

Calico's Pluggable Dataplane

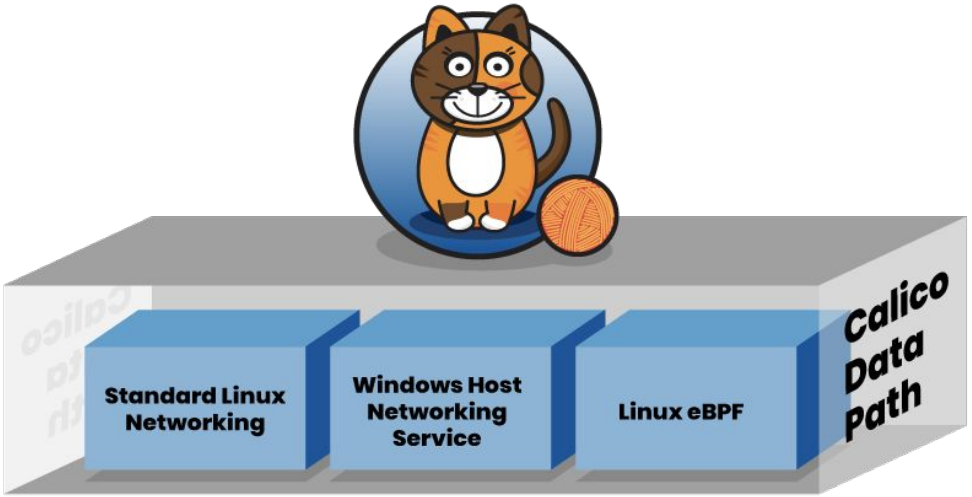

What is eBPF?

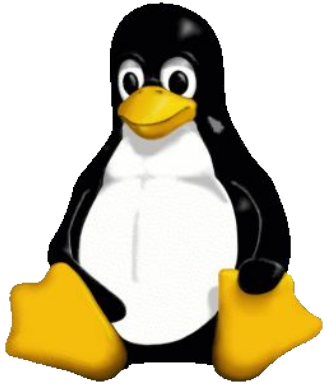

(extended) Berkeley Packet Filter

- An in-kernel virtual machine that “gives super-powers to Linux”
- Allows you to attach mini-programs to low-level hooks in the kernel
- Programs verified to ensure they are “safe”
  - e.g. can’t crash the system, access invalid memory addresses, will terminate
- Programs can only interact with the rest of the kernel through helper functions (there’s a limit to super powers!)
- The `clang` compiler can be used to build eBPF programs or you can write them directly in byte-code

What sort of things can you do with eBPF?

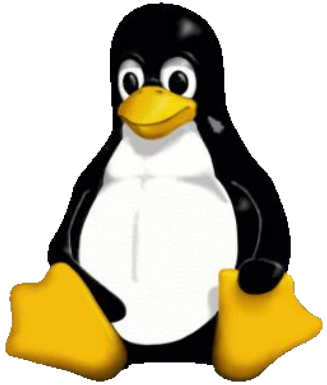

eBPF Features and Uses

- Security!
  - A *seccomp* filter mode allows users to write a program to determine if a system call is allowed.
- Logging and Tracing!
  - Gather information directly from the kernel about what calls are being run and how much time is being spent in them.
- Network Routing and Packet Filtering!
  - It's right there in the name. There are many different networking hooks - with varying performance and richness in capabilities.

How did you figure out what to build?

What was your design and development process?

How is this different from the current implementation?

What improvements does eBPF bring to Calico?

Pod-to-pod throughput and CPU

40 Gbps network, running qperf in single pod

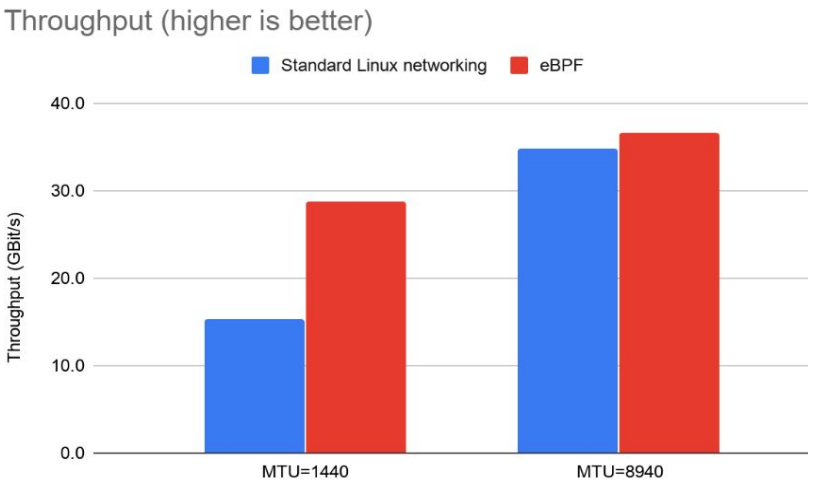

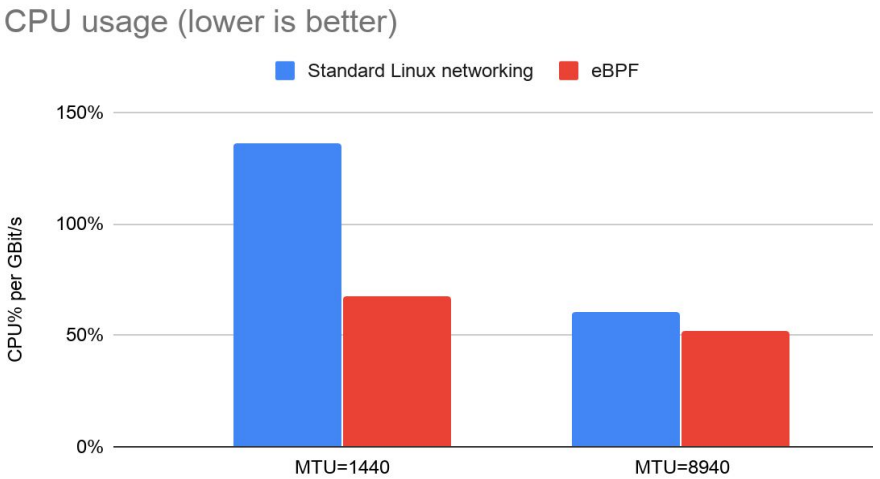

Native handling of Services: First packet latency

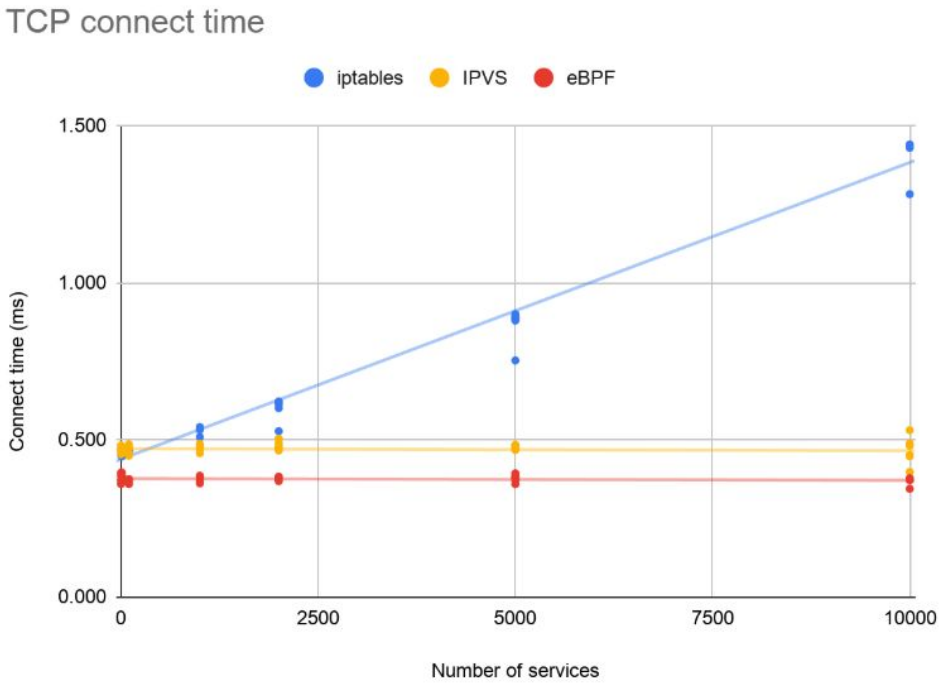

Native handling of Services: More efficient updates

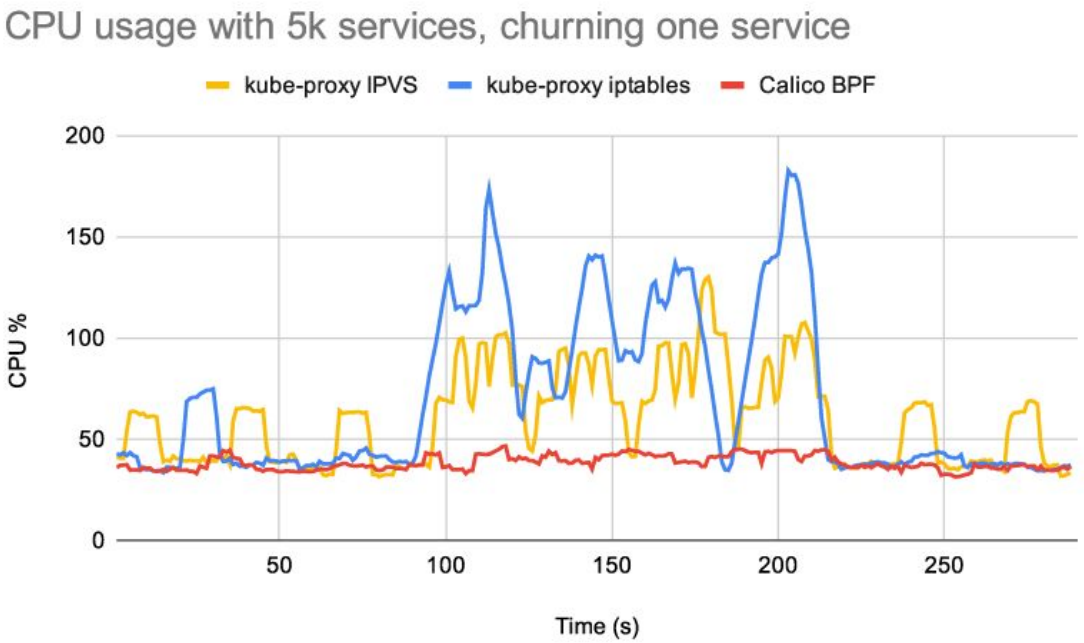

Native handling of Services: Direct Server Return

Kube-proxy packet path

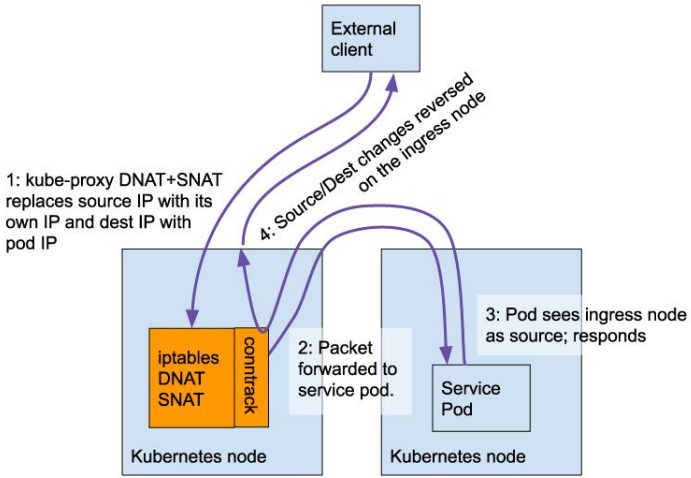

Calico eBPF

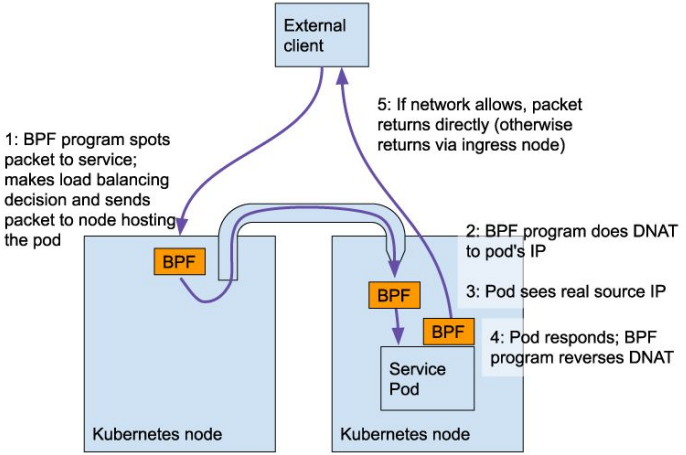

Native handling of Services: Direct Server Return

40 Gbps network, 1k services

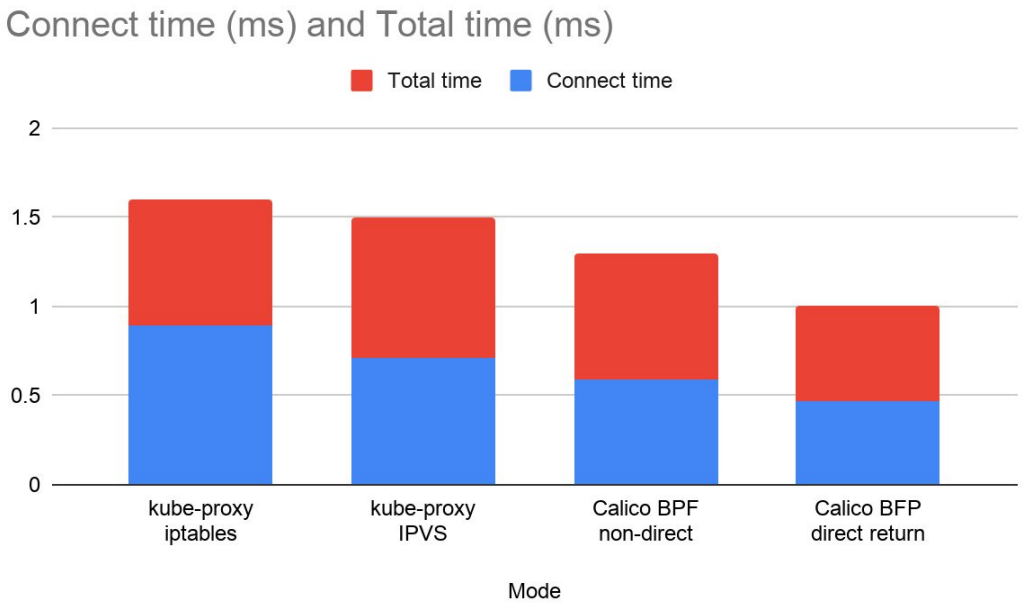

How can I try it out?

How to try it out!

- This is a tech preview, which means it’s not ready for production… yet!

- https://docs.projectcalico.org/getting-started/kubernetes/trying-ebpf

What's next?

Thank you!

Questions?

PROJECT CALICO

https://projectcalico.org

@projectcalico

https://github.com/projectcalico/community

https://slack.projectcalico.org

https://discuss.projectcalico.org

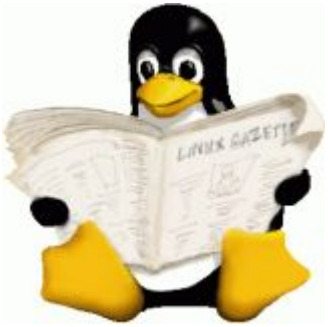

References

- Introducing the Calico eBPF Dataplane (projectcalico)
- A Thorough Introduction to eBPF (lwn)
- A seccomp overview (lwn)
- eBPF Tracing Tools