Linux Plumbers Conference 2019

Contribution ID: **322** Type: **not specified**

Scaling container policy management with kernel features

Wednesday 11 September 2019 10:00 (45 minutes)

Cilium is an open source project which implements the Container Network Interface (CNI) to provide networking and security functions in modern application environments. The primary focus of the Cilium community recently has been on scaling these functions to support thousands of nodes and hundreds of thousands of containers. Such environments impose a high rate of churn as containers and nodes appear and leave the cluster. For each change, the networking plugin needs to handle the incoming events and ensure that policy is in sync with network configuration state. This creates a strong incentive to efficiently interpret and map down cluster events into the required Linux networking configuration to minimize the window during which there are discrepancies between the desired and realized state in the cluster—something that is made possible through eBPF and other kernel features.

Cilium realizes these policy and container events through the use of many aspects of the networking stack, from rules to routes, tc to socket hooks, skb->mark to the skb->cb. Modelling the changes to datapath state involves a non-trivial amount of work in the userspace daemon to structure the desired state from external entities and allow incremental adjustments to be made, keeping the amount of work required to handle an event proportional to its impact on the kernel configuration. Some aspects of datapath configuration such as the implementation of L7 policy have gone through multiple iterations, which provides a window for us to explore the past, present and future of transparent proxies.

This talk will discuss the container policy model used by Cilium to apply whitelist filtering of requests at layers 3, 4 and 7; memoization techniques used to cache intermediate policy computation artifacts; and impacts on dataplane design and kernel features when considering large container based deployments with high rates of change in cluster state.

I agree to abide by the anti-harassment policy

Yes

I confirm that I am already registered for LPC 2019

Primary author: STRINGER, Joe (Cilium.io)

Presenter: STRINGER, Joe (Cilium.io)

Session Classification: Networking Summit Track