From Managed Kubernetes to App Platform: 1.5 Years of Cilium Usage at DigitalOcean

Timo Reimann, DigitalOcean

October 28, 2020

History / Context

- *DigitalOcean Kubernetes Service* aka DOKS: our managed Kubernetes offering
- Started out using Flannel but decided to move to Cilium in late 2018 for a couple of reasons:
  - support for `NetworkPolicies`
  - feature-rich CNI implementation
  - actively maintained project
  - healthy, supportive community
- Today, all O(10000) DOKS clusters run on Cilium

Cilium in the DOKS Architecture

- cilium-agent managed as DaemonSet on each worker node
- cilium-operator managed as Deployment (2 replicas / HA mode in latest releases) on workers
- cilium-agent running on control plane to enable control/data plane connectivity
- Cilium state-keeping in shared cluster etcd

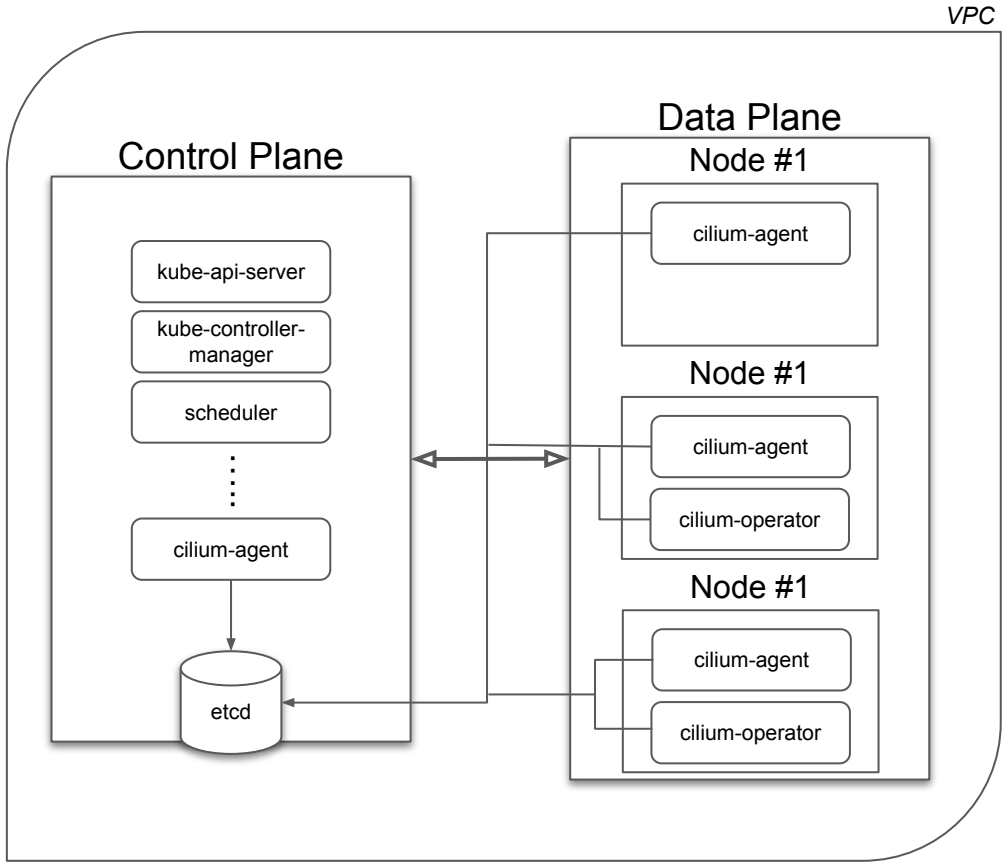

How's Cilium been working for us?

- Good experience overall
  - reasonably easy to maintain
  - ~3.5% of DOKS clusters are using `{Cilium}NetworkPolicies` directly
- Upgrades have been pretty smooth
  - moved from Cilium 1.4 initially to 1.8 today
  - retain old RBAC rules across certain cluster upgrades to avoid disruptions
- (Health checking) tooling really helpful in troubleshooting issues
  - readiness probes, `cilium status`, and friends

Planned adoptions

- Move from VXLAN to direct routing (leveraging the next generation of our internal VPC)
  - still a challenge to provide a disruption-free migration path for clusters
- Move to using eBPF kube-proxy replacement
- Look more into Hubble

App Platform

- App Platform: push-to-deploy PaaS offering by DigitalOcean
- Built on top of DOKS with multi-tenancy
  - various measures applied to guarantee isolation between tenants
- Makes extensive use of `CiliumNetworkPolicies` to:
  - restrict connectivity between apps (allow for same customer, deny otherwise)
  - restrict connectivity for ingress (Envoy) and egress (public Internet with exceptions, e.g., SMTP)
  - allow connectivity to needed infrastructure (DNS) and between our own management services
- Feature wish list for Cilium: 🙂
  - specifying port ranges
  - traffic shaping (resorting to iptables for now)
  - tracking connections in specific TCP states (SYN/FIN only) (cilium/cilium#12827)

Thanks for listening 😀

Timo Reimann
treimann@digitalocean.com