ISOVALENT

Cilium – Cloud Native Networking & Security

Cloud Native Days Stockholm

Thomas Graf

CTO & Co-Founde, Isovalent

cilium

Created by ISOVALENT

eBPF-based:

- Networking
- Security
- Observability
- Service Mesh & Ingress

Foundation

CLOUD NATIVE COMPUTING FOUNDATION

Technology

eBPF envoy

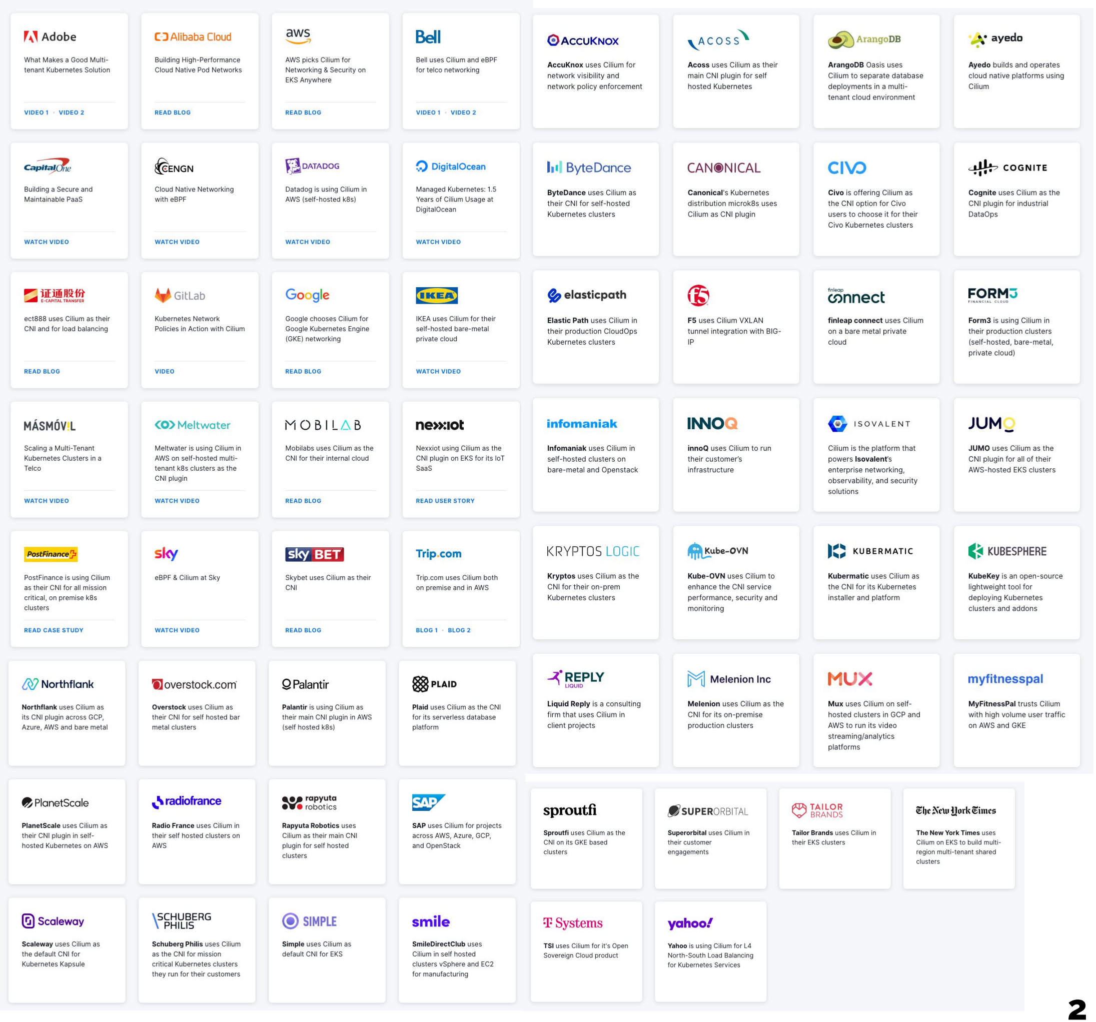

Adobe

What Makes a Good Multi-tenant Kubernetes Solution

VIDEO 1 · VIDEO 2

Alibaba Cloud

Building High-Performance Cloud Native Pod Networks

READ BLOG

aws

AWS picks Cilium for Networking & Security on EKS Anywhere

READ BLOG

Bell

Bell uses Cilium and eBPF for telco networking

VIDEO 1 · VIDEO 2

AccuKnox

AccuKnox uses Cilium for network visibility and network policy enforcement

ACOSS

Acoss uses Cilium as their main CNI plugin for self hosted Kubernetes

ArangoDB

ArangoDB Oasis uses Cilium to separate database deployments in a multi-tenant cloud environment

ayedo

Ayedo builds and operates cloud native platforms using Cilium

Capital One

Building a Secure and Maintainable PaaS

WATCH VIDEO

CENGN

Cloud Native Networking with eBPF

WATCH VIDEO

DATADOG

Datadog is using Cilium in AWS (self-hosted k8s)

WATCH VIDEO

DigitalOcean

Managed Kubernetes: 1.5 Years of Cilium Usage at DigitalOcean

WATCH VIDEO

ByteDance

ByteDance uses Cilium as their CNI for self-hosted Kubernetes clusters

CANONICAL

Canonical's Kubernetes distribution microk8s uses Cilium as CNI plugin

CIVO

Civo is offering Cilium as the CNI option for Civo users to choose it for their Civo Kubernetes clusters

COGNITE

Cognite uses Cilium as the CNI plugin for industrial DataOps

证通股份 E-CAPITAL TRANSFER

ect888 uses Cilium as their CNI and for load balancing

READ BLOG

GitLab

Kubernetes Network Policies in Action with Cilium

VIDEO

Google

Google chooses Cilium for Google Kubernetes Engine (GKE) networking

READ BLOG

IKEA

IKEA uses Cilium for their self-hosted bare-metal private cloud

WATCH VIDEO

elasticpath

Elastic Path uses Cilium in their production CloudOps Kubernetes clusters

f5

F5 uses Cilium VXLAN tunnel integration with BIG-IP

finleap connect

finleap connect uses Cilium on a bare metal private cloud

FORM3 FINANCIAL CLOUD

Form3 is using Cilium in their production clusters (self-hosted, bare-metal, private cloud)

MÁSMÓV!L

Scaling a Multi-Tenant Kubernetes Clusters in a Telco

WATCH VIDEO

Meltwater

Meltwater is using Cilium in AWS on self-hosted multi-tenant k8s clusters as the CNI plugin

WATCH VIDEO

MOBILAB

Mobilabs uses Cilium as the CNI for their internal cloud

READ BLOG

nexxiot

Nexxiot using Cilium as the CNI plugin on EKS for its IoT SaaS

READ USER STORY

infomaniak

Infomaniak uses Cilium in self-hosted clusters on bare-metal and Openstack

INNOQ

innoQ uses Cilium to run their customer's infrastructure

ISOVALENT

Cilium is the platform that powers **Isovalent**'s enterprise networking, observability, and security solutions

JUMO

JUMO uses Cilium as the CNI plugin for all of their AWS-hosted EKS clusters

PostFinance

PostFinance is using Cilium as their CNI for all mission critical, on premise k8s clusters

READ CASE STUDY

sky

eBPF & Cilium at Sky

WATCH VIDEO

sky BET

Skybet uses Cilium as their CNI

READ BLOG

Trip.com

Trip.com uses Cilium both on premise and in AWS

BLOG 1 · BLOG 2

KRYPTOS LOGIC

Kryptos uses Cilium as the CNI for their on-prem Kubernetes clusters

Kube-OVN

Kube-OVN uses Cilium to enhance the CNI service performance, security and monitoring

KUBERMATIC

Kubermatic uses Cilium as the CNI for its Kubernetes installer and platform

KUBESPHERE

KubeKey is an open-source lightweight tool for deploying Kubernetes clusters and addons

Northflank

Northflank uses Cilium as its CNI plugin across GCP, Azure, AWS and bare metal

overstock.com

Overstock uses Cilium as their CNI for self hosted bar metal clusters

Palantir

Palantir is using Cilium as their main CNI plugin in AWS (self hosted k8s)

PLAID

Plaid uses Cilium as the CNI for its serverless database platform

REPLY LIQUID

Liquid Reply is a consulting firm that uses Cilium in client projects

Melenion Inc

Melenion uses Cilium as the CNI for its on-premise production clusters

MUX

Mux uses Cilium on self-hosted clusters in GCP and AWS to run its video streaming/analytics platforms

myfitnesspal

MyFitnessPal trusts Cilium with high volume user traffic on AWS and GKE

PlanetScale

PlanetScale uses Cilium as their CNI plugin in self-hosted Kubernetes on AWS

radiofrance

Radio France uses Cilium in their self hosted clusters on AWS

rapyuta robotics

Rapyuta Robotics uses Cilium as their main CNI plugin for self hosted clusters

SAP

SAP uses Cilium for projects across AWS, Azure, GCP, and OpenStack

sproutfi

Sproutfi uses Cilium as the CNI on its GKE based clusters

SUPERORBITAL

Superorbital uses Cilium in their customer engagements

TAILOR BRANDS

Tailor Brands uses Cilium in their EKS clusters

The New York Times

The New York Times uses Cilium on EKS to build multi-region multi-tenant shared clusters

Scaleway

Scaleway uses Cilium as the default CNI for Kubernetes Kapsule

SCHUBERG PHILIS

Schuberg Philis uses Cilium as the CNI for mission critical Kubernetes clusters they run for their customers

SIMPLE

Simple uses Cilium as default CNI for EKS

smile

SmileDirectClub uses Cilium in self hosted clusters vSphere and EC2 for manufacturing

T Systems

TSI uses Cilium for it's Open Sovereign Cloud product

yahoo!

Yahoo is using Cilium for L4 North-South Load Balancing for Kubernetes Services

eBPF

eBPF

Makes the Linux kernel programmable in a secure and efficient way.

“What JavaScript is to the browser, eBPF is to the Linux Kernel”

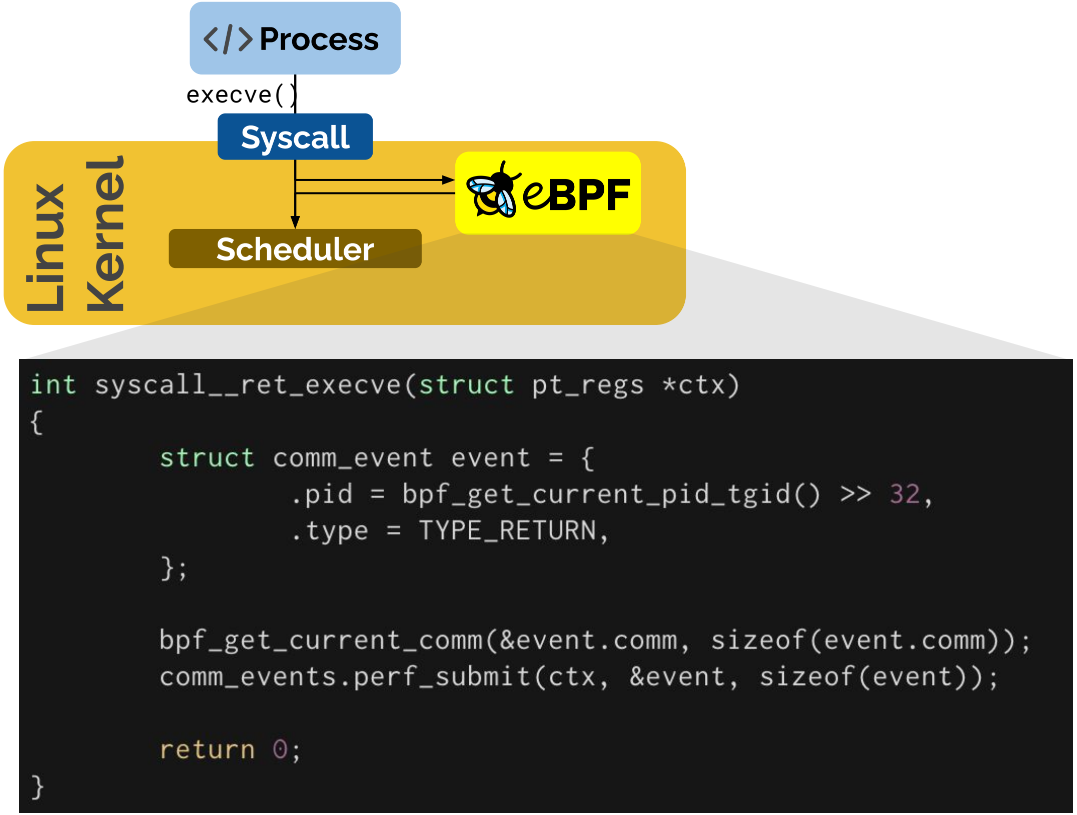

```
int syscall__ret_execve(struct pt_regs *ctx)
{
        struct comm_event event = {
                .pid = bpf_get_current_pid_tgid() >> 32,
                .type = TYPE_RETURN,
        };

        bpf_get_current_comm(&event.comm, sizeof(event.comm));
        comm_events.perf_submit(ctx, &event, sizeof(event));

        return 0;
}
```

cilium

Cilium

CNI

Scalable, Secure,
High Performance
CNI Plugin

eBPF

Cilium CNI

Scalable, Secure, High Performance CNI Plugin

Cilium Service Mesh

Sidecar-free Mesh & Ingress

eBPF

Cilium CNI

Scalable, Secure, High Performance CNI Plugin

Cilium Service Mesh

Sidecar-free Mesh & Ingress

Hubble

Network Observability

eBPF

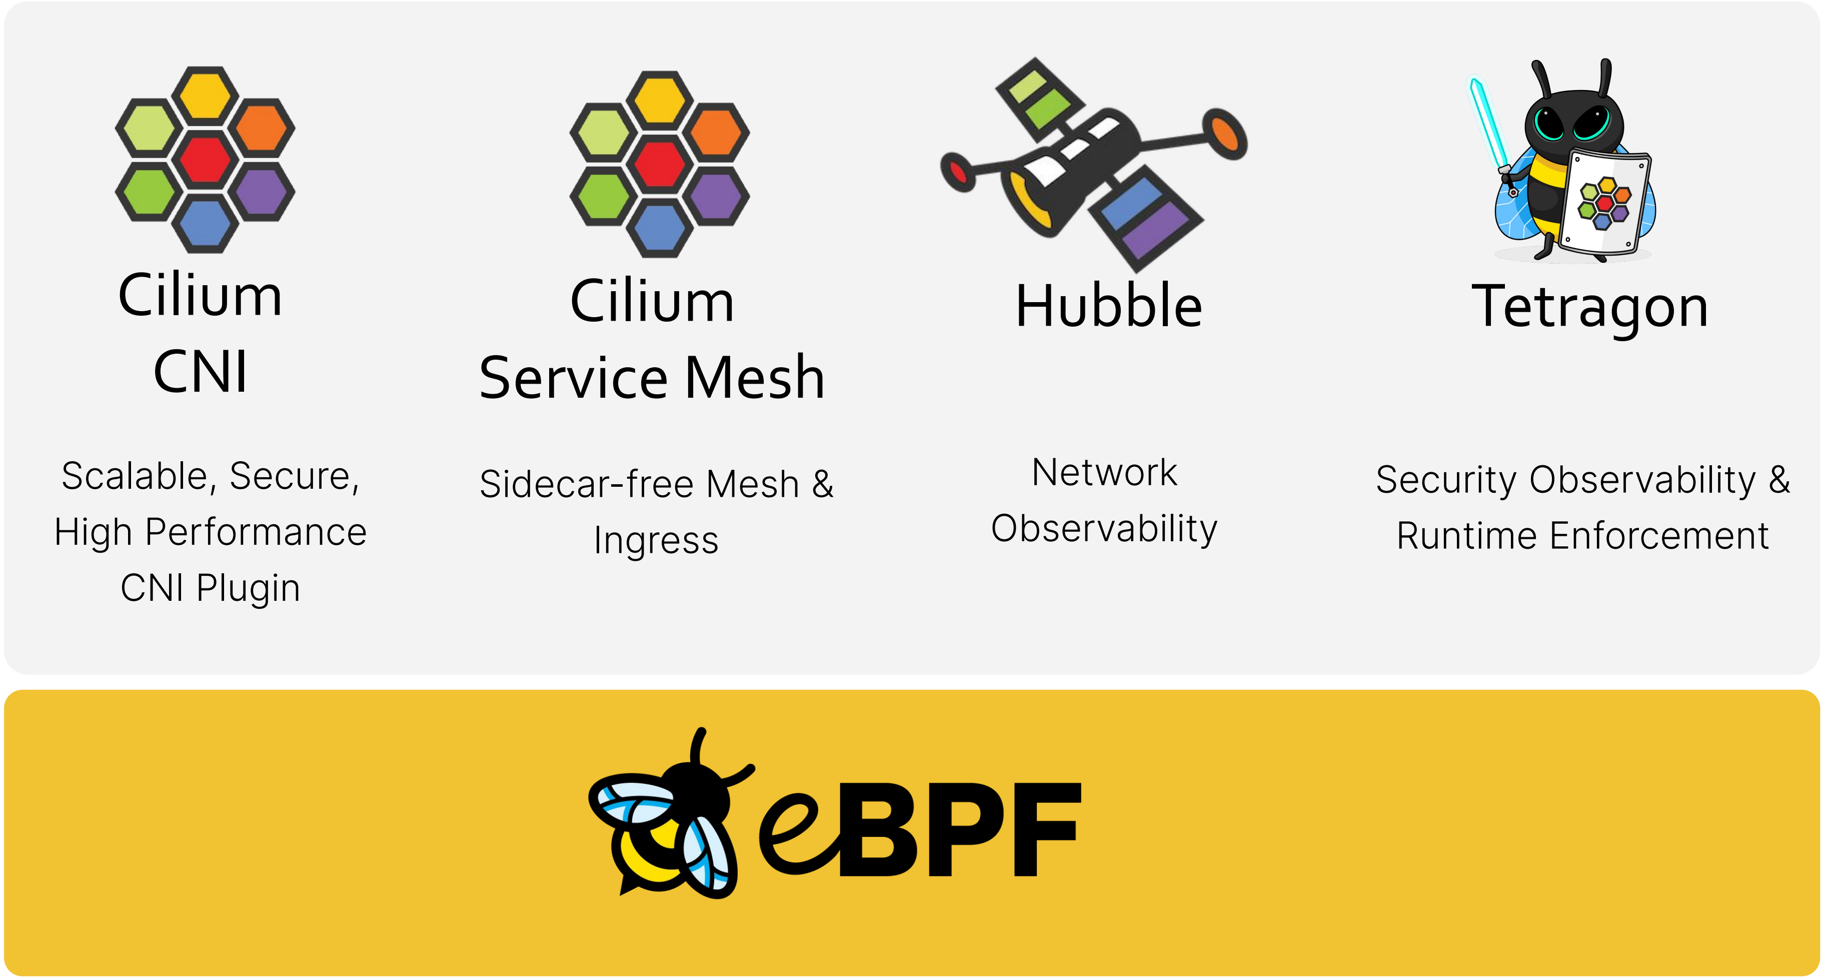

Cilium CNI

Scalable, Secure, High Performance CNI Plugin

Cilium Service Mesh

Sidecar-free Mesh & Ingress

Hubble

Network Observability

Tetragon

Security Observability & Runtime Enforcement

eBPF

cilium

The Origins....

Then

Cilium Design Summit,
Diavolezza, 2016

Now

Isovalent Team,
Punt Muragl, 2022

First Conference Talk

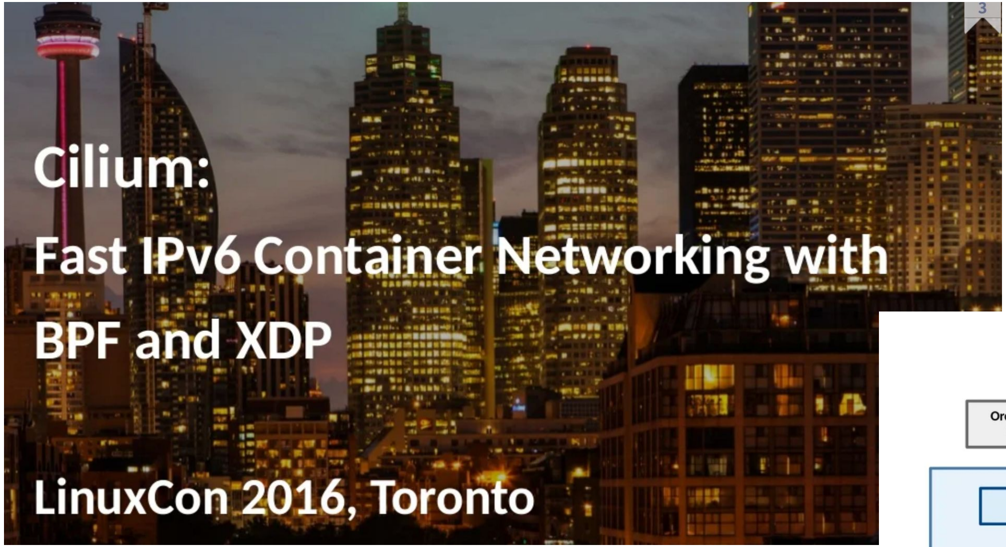

3

Cilium:

Fast IPv6 Container Networking with BPF and XDP

LinuxCon 2016, Toronto

BPF – Berkley Packet Filter

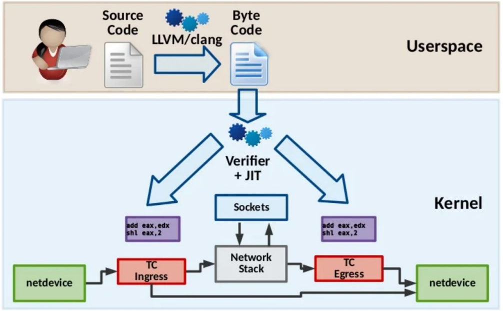

Cilium Architecture

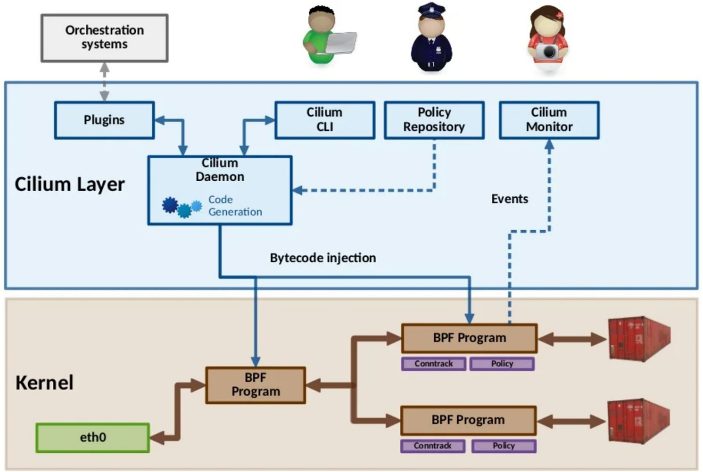

- IPv6 Only
- No concept of networks
- Policy decoupled from addressing
- Giant flat L3

DockerCon 2017
When everything changed

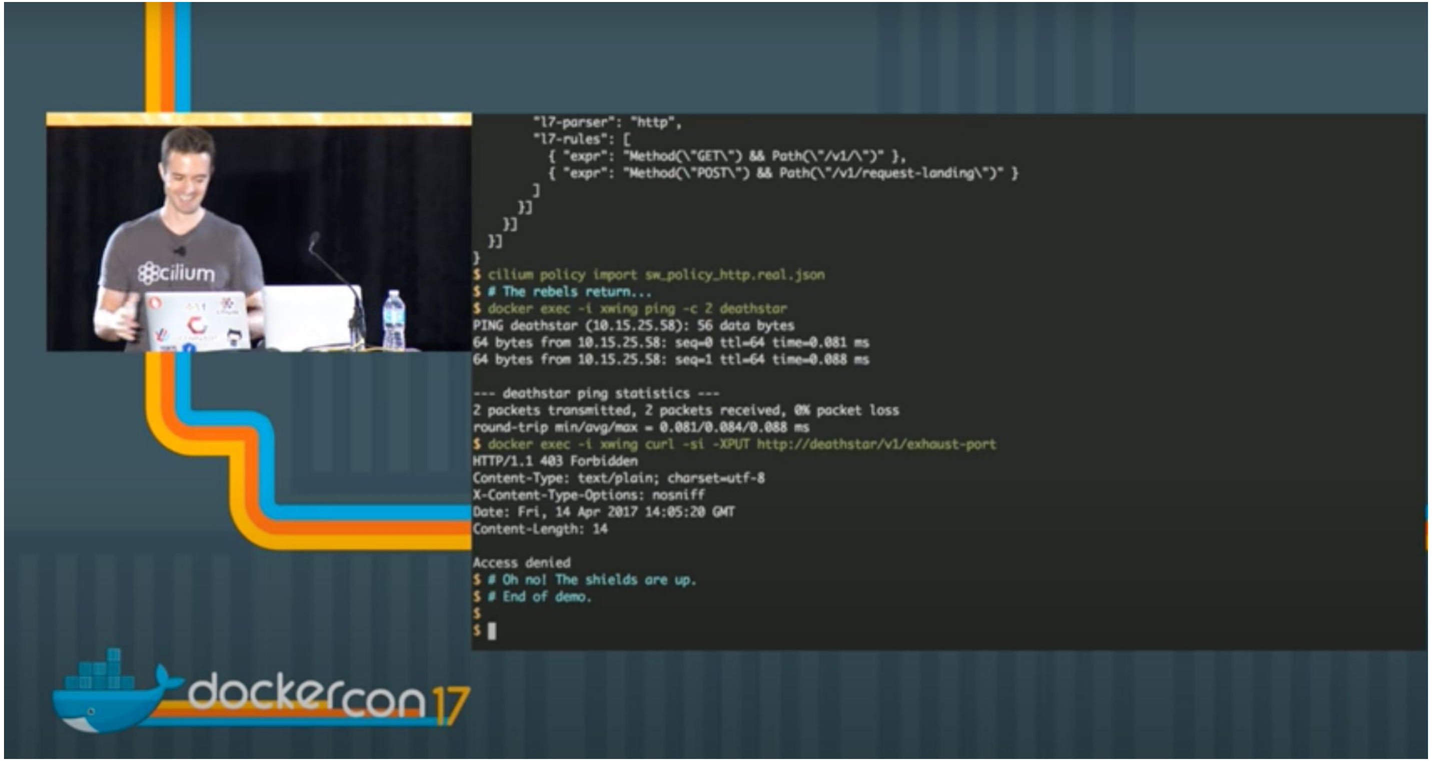

```
        "l7-parser": "http",
        "l7-rules": [
          { "expr": "Method(\"GET\") && Path(\"/v1/\")" },
          { "expr": "Method(\"POST\") && Path(\"/v1/request-landing\")" }
        ]
      }]
    }]
  }]
}
$ cilium policy import sw_policy_http.real.json
$ # The rebels return...
$ docker exec -i xwing ping -c 2 deathstar
PING deathstar (10.15.25.58): 56 data bytes
64 bytes from 10.15.25.58: seq=0 ttl=64 time=0.081 ms
64 bytes from 10.15.25.58: seq=1 ttl=64 time=0.088 ms

--- deathstar ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0.081/0.084/0.088 ms
$ docker exec -i xwing curl -si -XPUT http://deathstar/v1/exhaust-port
HTTP/1.1 403 Forbidden
Content-Type: text/plain; charset=utf-8
X-Content-Type-Options: nosniff
Date: Fri, 14 Apr 2017 14:05:20 GMT
Content-Length: 14

Access denied
$ # Oh no! The shields are up.
$ # End of demo.
$
$ 
```

2017 - First Office

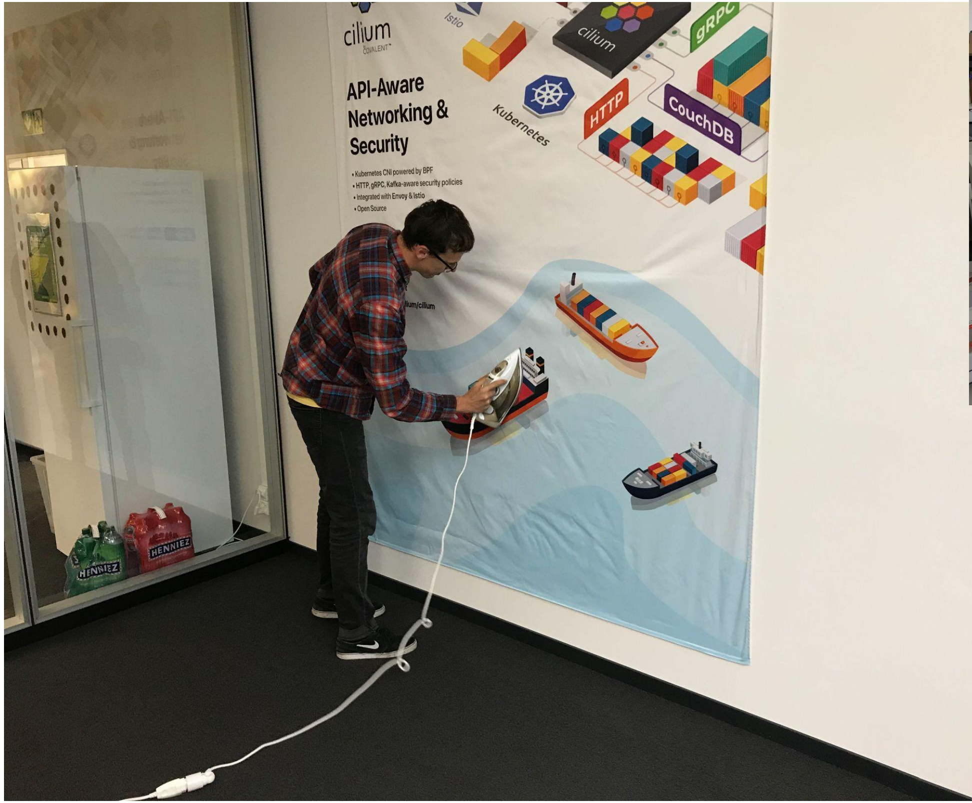

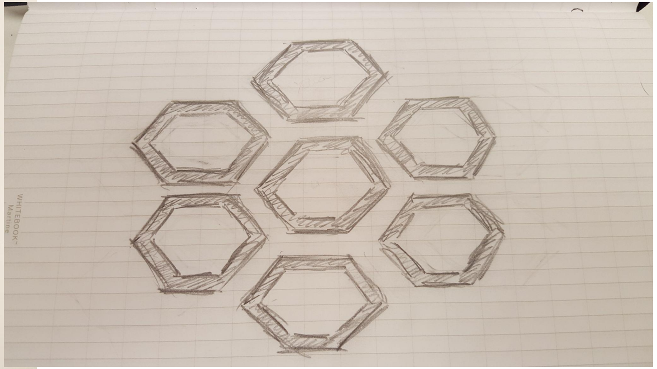

2018 Starting to Grow

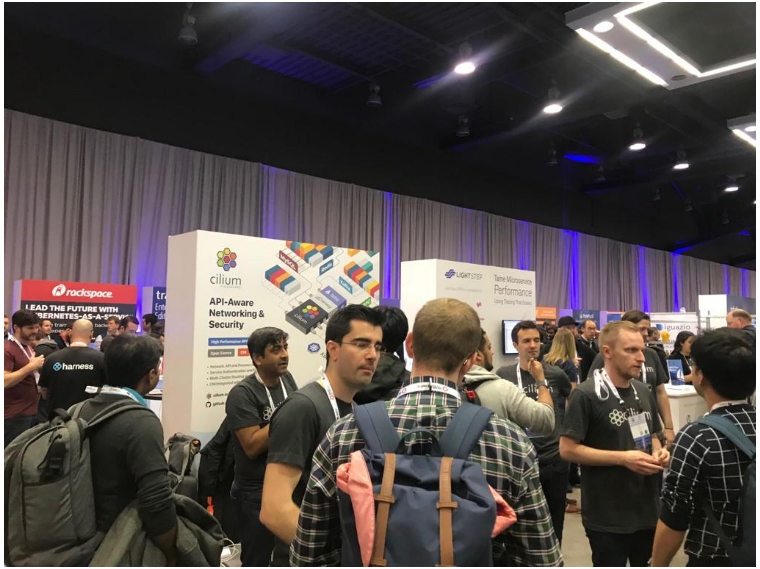
cilium
API-Aware Networking & Security
rackspace
LEAD THE FUTURE WITH KUBERNETES-AS-A-SERVICE
LIGHTSTEP
Tame Microservice Performance
harness
iguazio

2019
Team Building

The Snow Chains Incident, Julier Pass

2019
Swiss Culture

Fondue,
AirBnB,
Palo Alto

2021- Things get crazy…

Cilium joins the CNCF

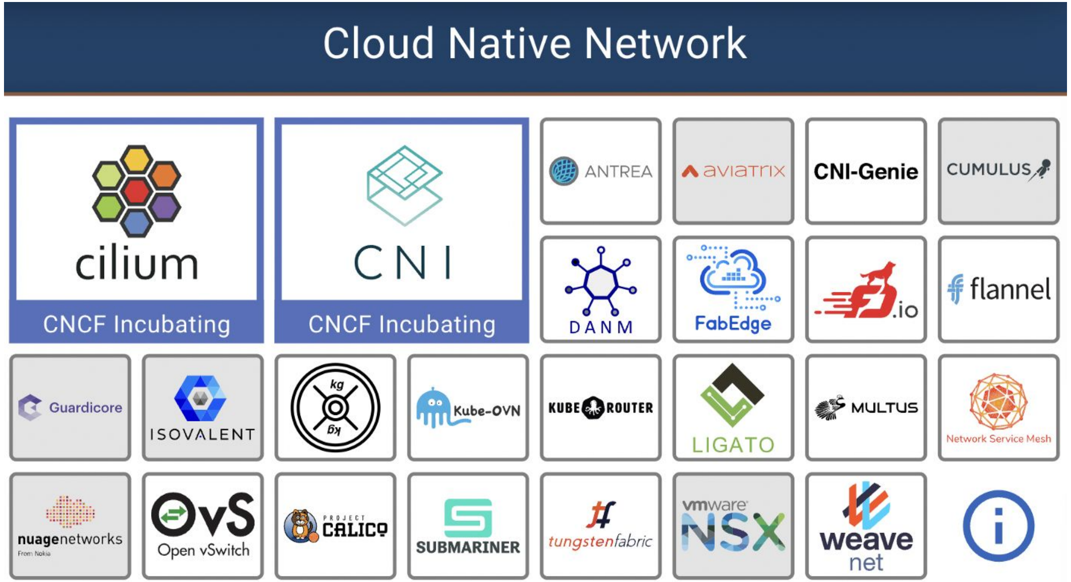

AWS picks Cilium

Google picks Cilium

2022 - Crazy^2

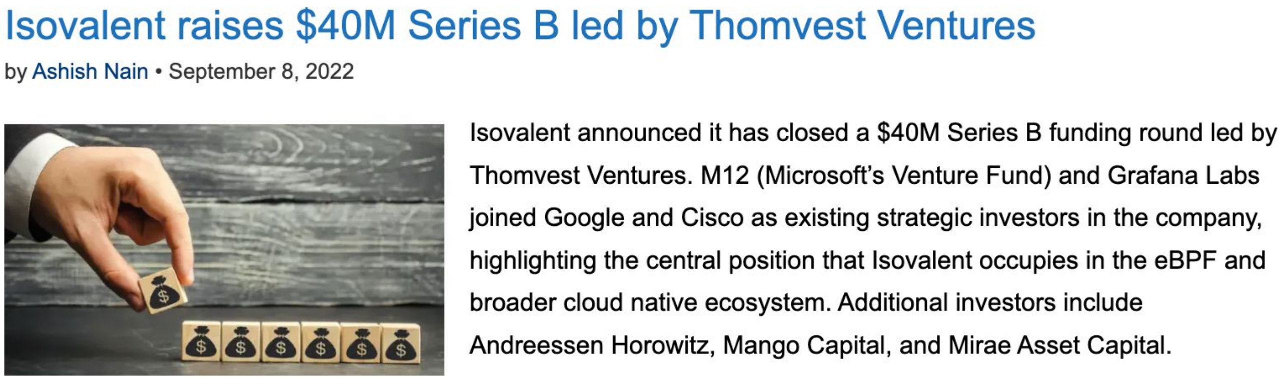

Isovalent raises $40M Series B led by Thomvest Ventures

by Ashish Nain • September 8, 2022

Isovalent announced it has closed a $40M Series B funding round led by Thomvest Ventures. M12 (Microsoft's Venture Fund) and Grafana Labs joined Google and Cisco as existing strategic investors in the company, highlighting the central position that Isovalent occupies in the eBPF and broader cloud native ecosystem. Additional investors include Andreessen Horowitz, Mango Capital, and Mirae Asset Capital.

THOMVEST Grafana Microsoft

a16z ANDREESSEN HOROWITZ Google cisco

cilium

CNI

What is Cilium CNI?

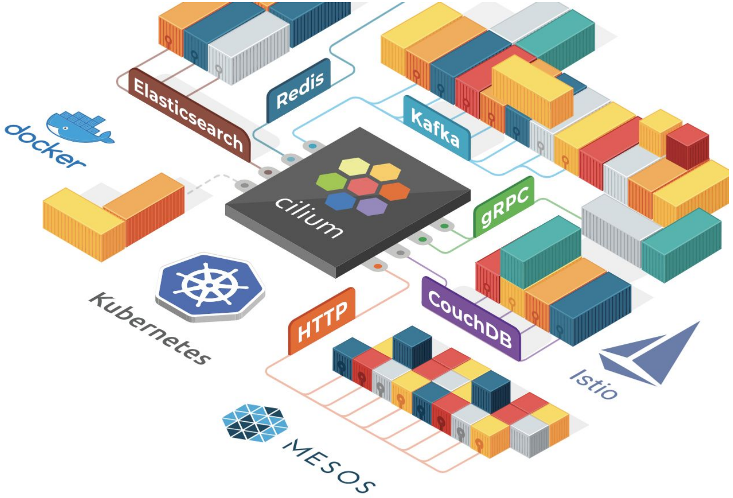

Efficient and Scalable Kubernetes CNI

- IPv4, IPv6, NAT46, SRv6, ...
- Overlays, BGP, Cloud Provider SDNs

High-performance Load-Balancing

- Kubernetes Services
- North-South load-balancer
- Kubernetes Ingress

Network Policies & Encryption

- Kubernetes Network Policy
- Cilium Network Policy (FQDN, L7, ...)
- Transparent Encryption

Multi-Cluster & External Workloads

- Global Services, Service Discovery, Network Policy
- Integration of Metal & VMs
- Egress Gateway

Hubble Observability

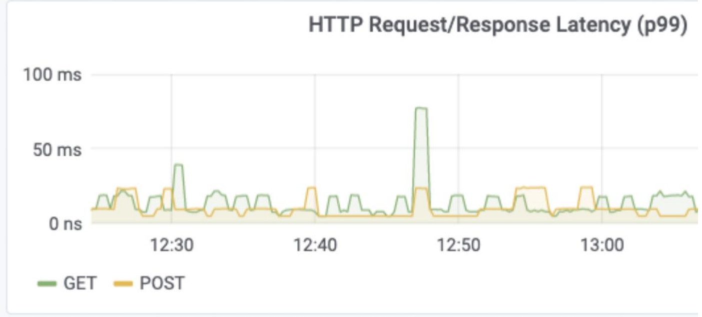

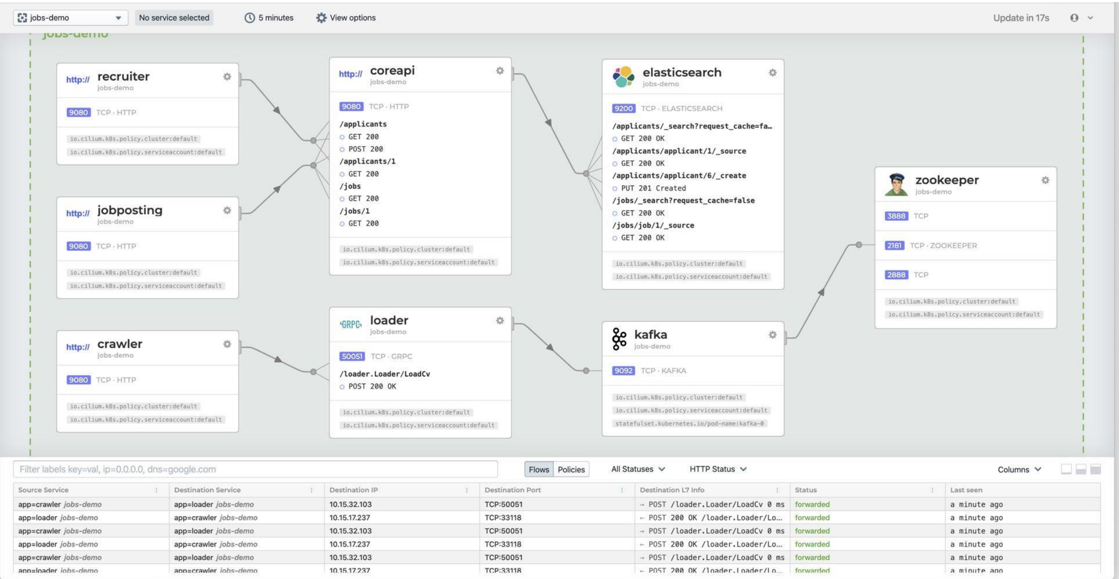
jobs-demo
No service selected
5 minutes
View options
Update in 17s
jobs-demo
http:// recruiter jobs-demo
9080 TCP · HTTP
io.cilium.k8s.policy.cluster:default
io.cilium.k8s.policy.serviceaccount:default
http:// jobposting jobs-demo
9080 TCP · HTTP
io.cilium.k8s.policy.cluster:default
io.cilium.k8s.policy.serviceaccount:default
http:// crawler jobs-demo
9080 TCP · HTTP
io.cilium.k8s.policy.cluster:default
io.cilium.k8s.policy.serviceaccount:default
http:// coreapi jobs-demo
9080 TCP · HTTP
/applicants
GET 200
POST 200
/applicants/1
GET 200
/jobs
GET 200
/jobs/1
GET 200
io.cilium.k8s.policy.cluster:default
io.cilium.k8s.policy.serviceaccount:default
GRPC loader jobs-demo
50051 TCP · GRPC
/loader.Loader/LoadCv
POST 200 OK
io.cilium.k8s.policy.cluster:default
io.cilium.k8s.policy.serviceaccount:default
elasticsearch jobs-demo
9200 TCP · ELASTICSEARCH
/applicants/_search?request_cache=fa…
GET 200 OK
/applicants/applicant/1/_source
GET 200 OK
/applicants/applicant/6/_create
PUT 201 Created
/jobs/_search?request_cache=false
GET 200 OK
/jobs/job/1/_source
GET 200 OK
io.cilium.k8s.policy.cluster:default
io.cilium.k8s.policy.serviceaccount:default
kafka jobs-demo
9092 TCP · KAFKA
io.cilium.k8s.policy.cluster:default
io.cilium.k8s.policy.serviceaccount:default
statefulset.kubernetes.io/pod-name:kafka-0
zookeeper jobs-demo
3888 TCP
2181 TCP · ZOOKEEPER
2888 TCP
io.cilium.k8s.policy.cluster:default
io.cilium.k8s.policy.serviceaccount:default
Filter labels key=val, ip=0.0.0.0, dns=google.com
Flows
Policies
All Statuses
HTTP Status
Columns

Source Service	Destination Service	Destination IP	Destination Port	Destination L7 Info	Status	Last seen
app=crawler *jobs-demo*	app=loader *jobs-demo*	10.15.32.103	TCP:50051	→ POST /loader.Loader/LoadCv 0 ms	forwarded	a minute ago
app=loader *jobs-demo*	app=crawler *jobs-demo*	10.15.17.237	TCP:33118	← POST 200 OK /loader.Loader/Lo…	forwarded	a minute ago
app=crawler *jobs-demo*	app=loader *jobs-demo*	10.15.32.103	TCP:50051	→ POST /loader.Loader/LoadCv 0 ms	forwarded	a minute ago
app=loader *jobs-demo*	app=crawler *jobs-demo*	10.15.17.237	TCP:33118	← POST 200 OK /loader.Loader/Lo…	forwarded	a minute ago
app=crawler *jobs-demo*	app=loader *jobs-demo*	10.15.32.103	TCP:50051	→ POST /loader.Loader/LoadCv 0 ms	forwarded	a minute ago
app=loader *jobs-demo*	app=crawler *jobs-demo*	10.15.17.237	TCP:33118	← POST 200 OK /loader.Loader/Lo…	forwarded	a minute ago

Metrics, Logs, & Service Ma

- L3/L4
- L7 (HTTP, DNS, Kafka, ...)
- Network Policy
- ...

fluentd JSON Grafana OpenTelemetry

cilium

Service Mesh

Observable, secure, resilient cloud native connectivity

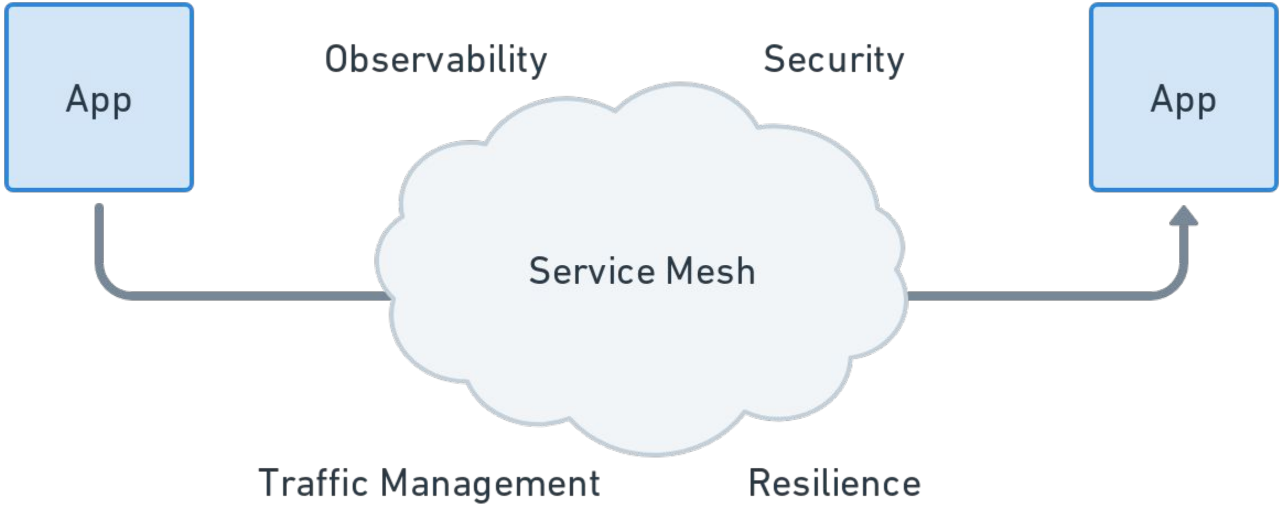

Traditional networking is falling short

tcpdump, NetFlow, sFlow, ...

SSH, IPsec, Wireguard, VPN, ..

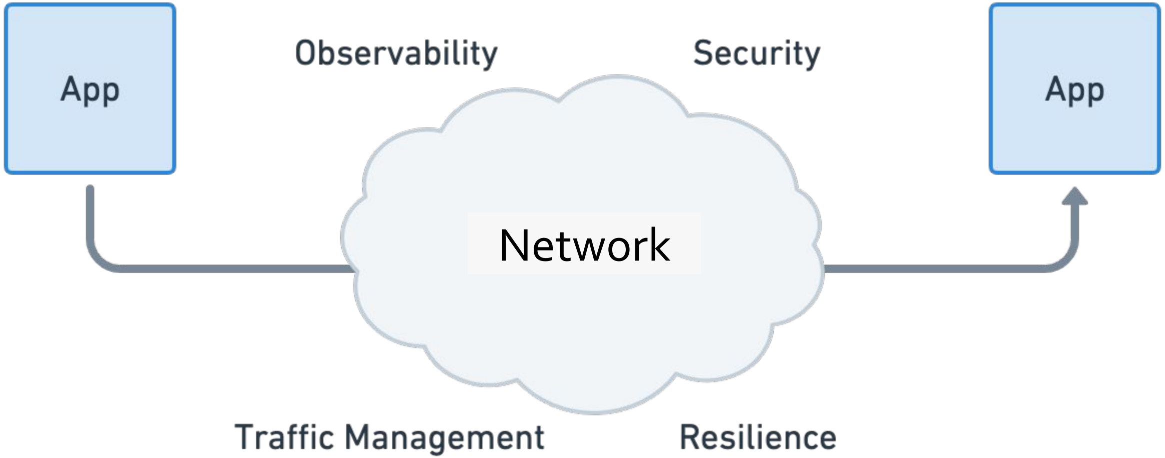

L4 Load-Balancing (IPVS, iptables, ...)

TCP retransmissions

Service Mesh Origins

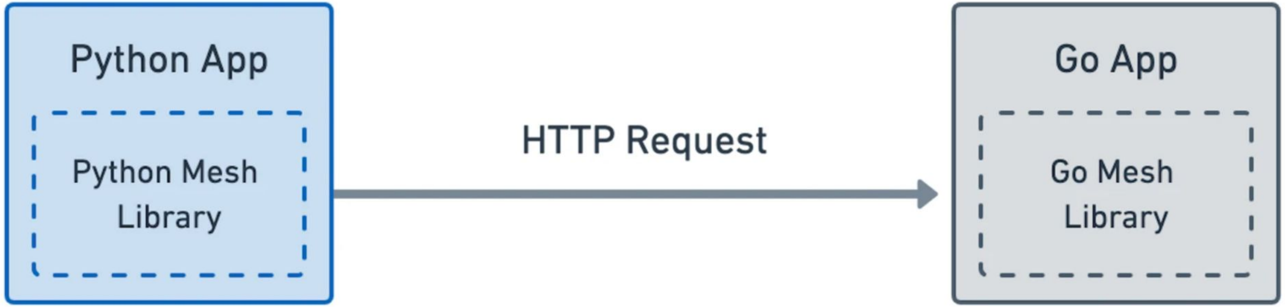

Each application requires a service mesh library written in the language framework of the application.

Service Mesh with Sidecars

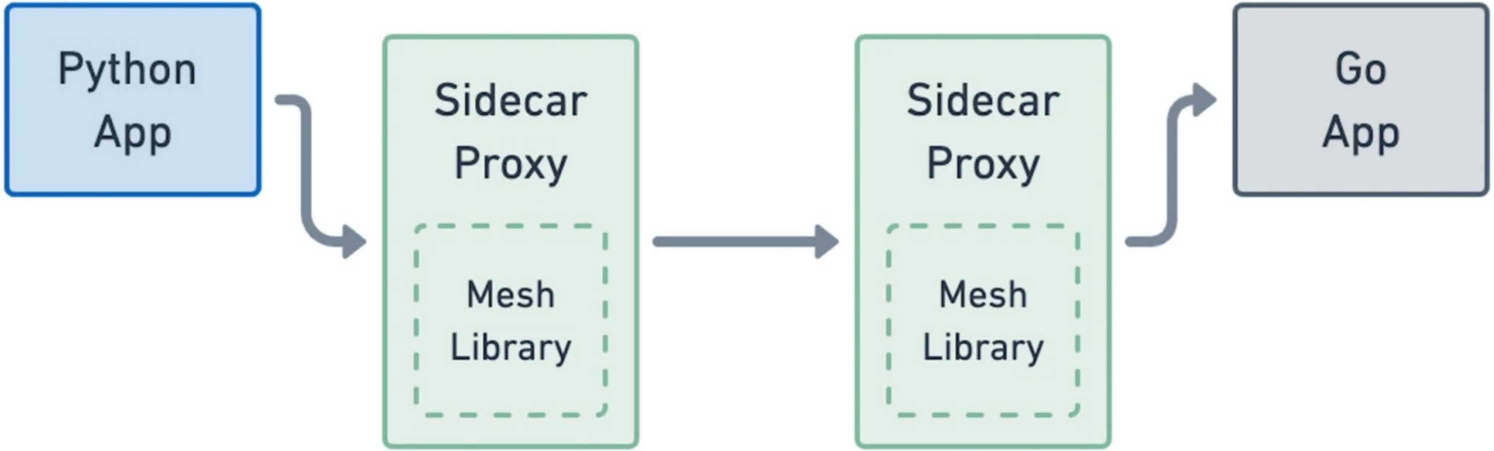

Service mesh is is embedded in a proxy running outside of the application.

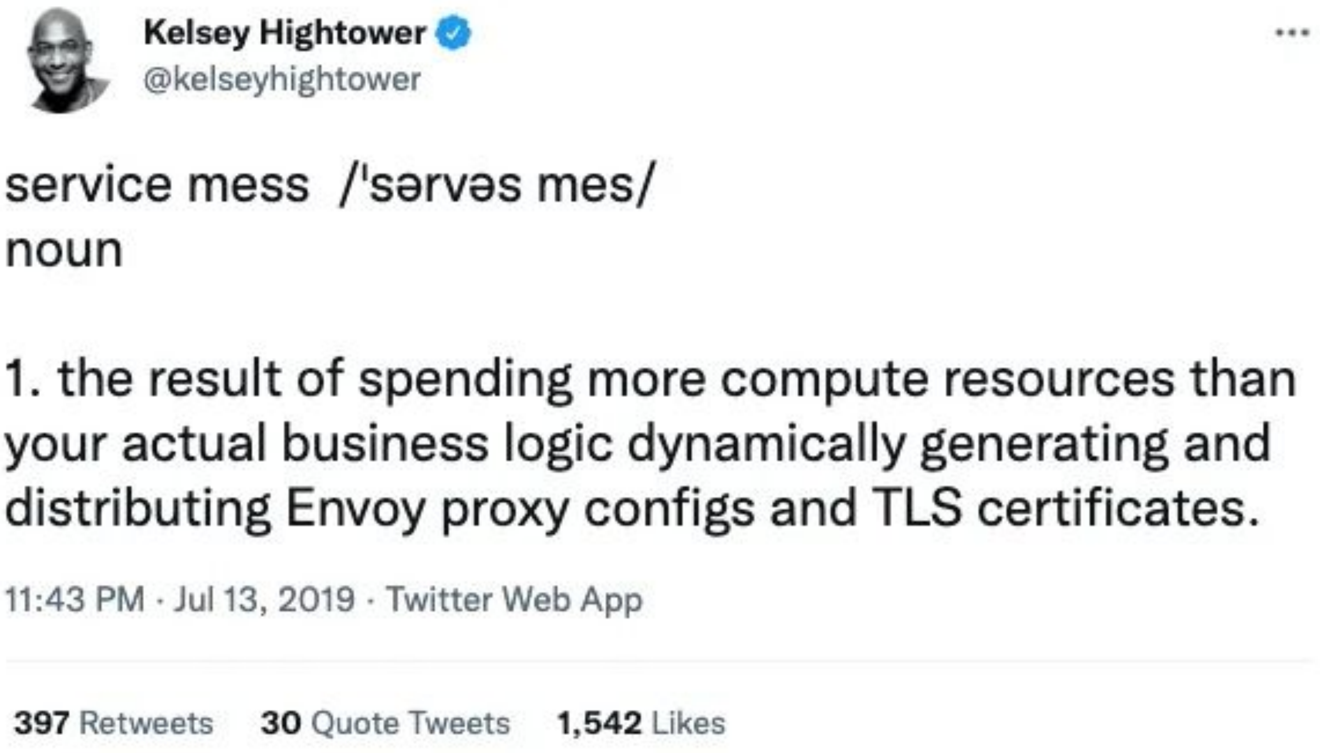

Kelsey Hightower
@kelseyhightower

service mess /ˈsərvəs mes/
noun

1. the result of spending more compute resources than your actual business logic dynamically generating and distributing Envoy proxy configs and TLS certificates.

11:43 PM · Jul 13, 2019 · Twitter Web App

397 Retweets 30 Quote Tweets 1,542 Likes

The network cost of sidecar proxies

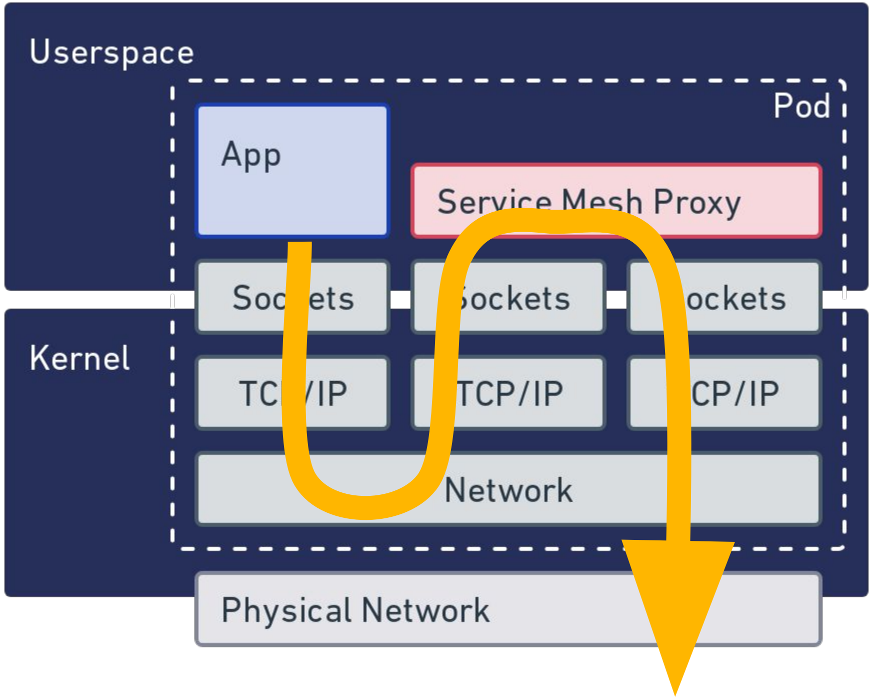

@lizrice

Removing sidecars from Service Mesh

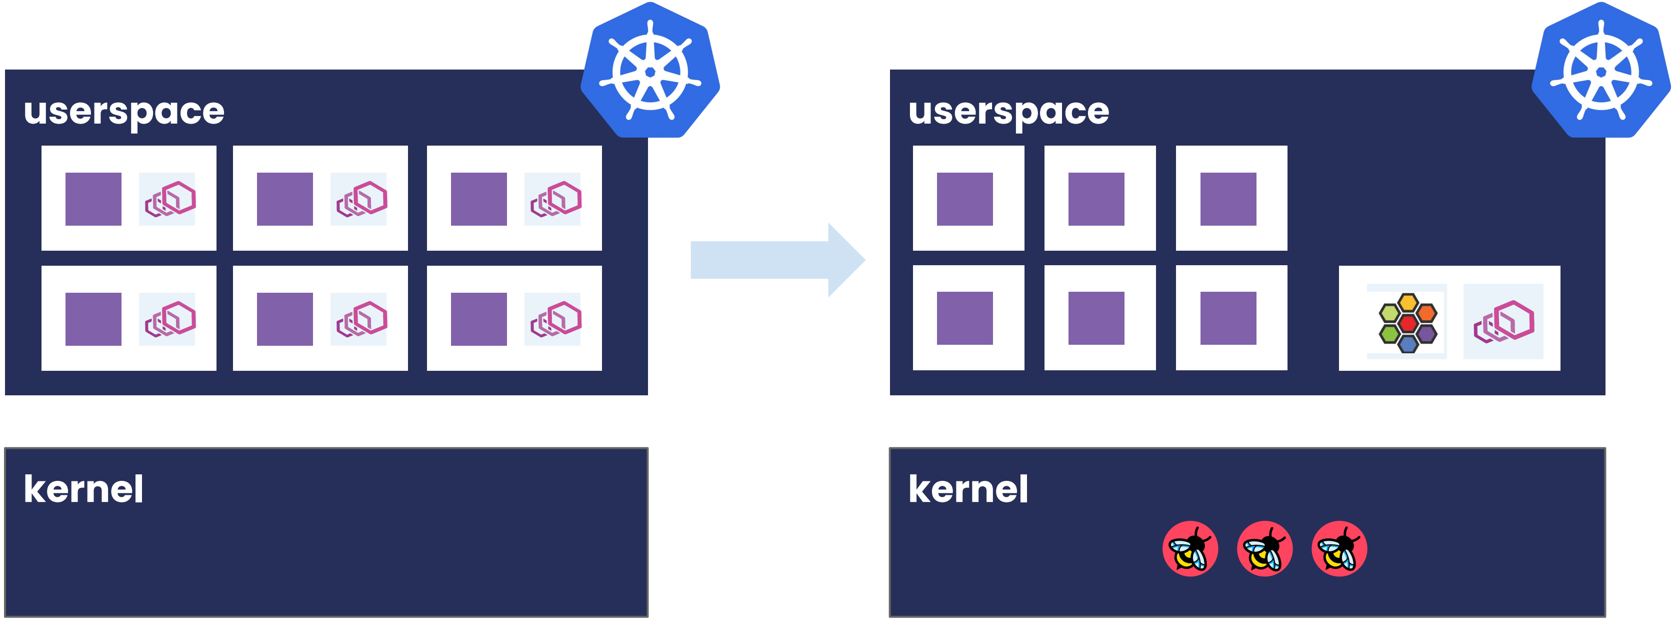

Cilium Service Mesh

Option 1:
Sidecar-free

eBPF + envoy

Option 2:
Istio Integration

eBPF + Istio + envoy

Control plane of your choice

Istio

Ingress / Services

Gateway API

SPIFFE

Observability Integrations

fluentd

JSON

Grafana

OpenTelemetry

eBPF Native
(no proxy)

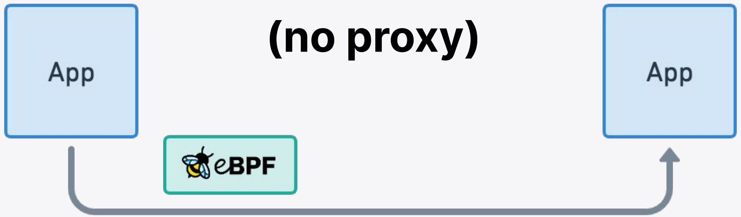

Whenever possible

Traffic Management

- L3/L4 forwarding & Load-balancing
- Canary, Topology Aware Routing
- Multi-cluster

Security

- Network Policy
- mTLS

Observability

- Tracing, OpenTelemetry, & Metrics
- HTTP, TLS, DNS, TCP, UDP, ...

Proxy

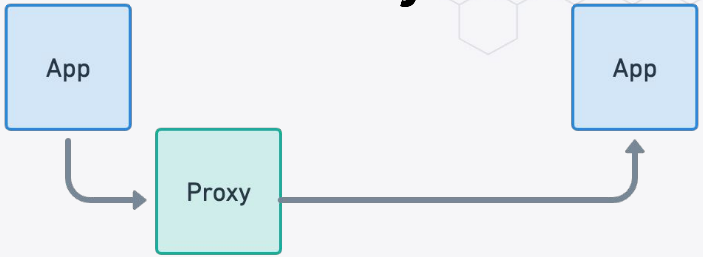

When eBPF cannot do it

Traffic Management

- L7 Load-balancing & Ingress

Resilience

- Retries, L7 Rate Limiting

Security

- TLS Termination & Origination
- L7 Network Policy*

ISOVALENT

Tetragon

Security Observability &
Runtime Enforcement

CLOUD NATIVE
COMPUTING FOUNDATION

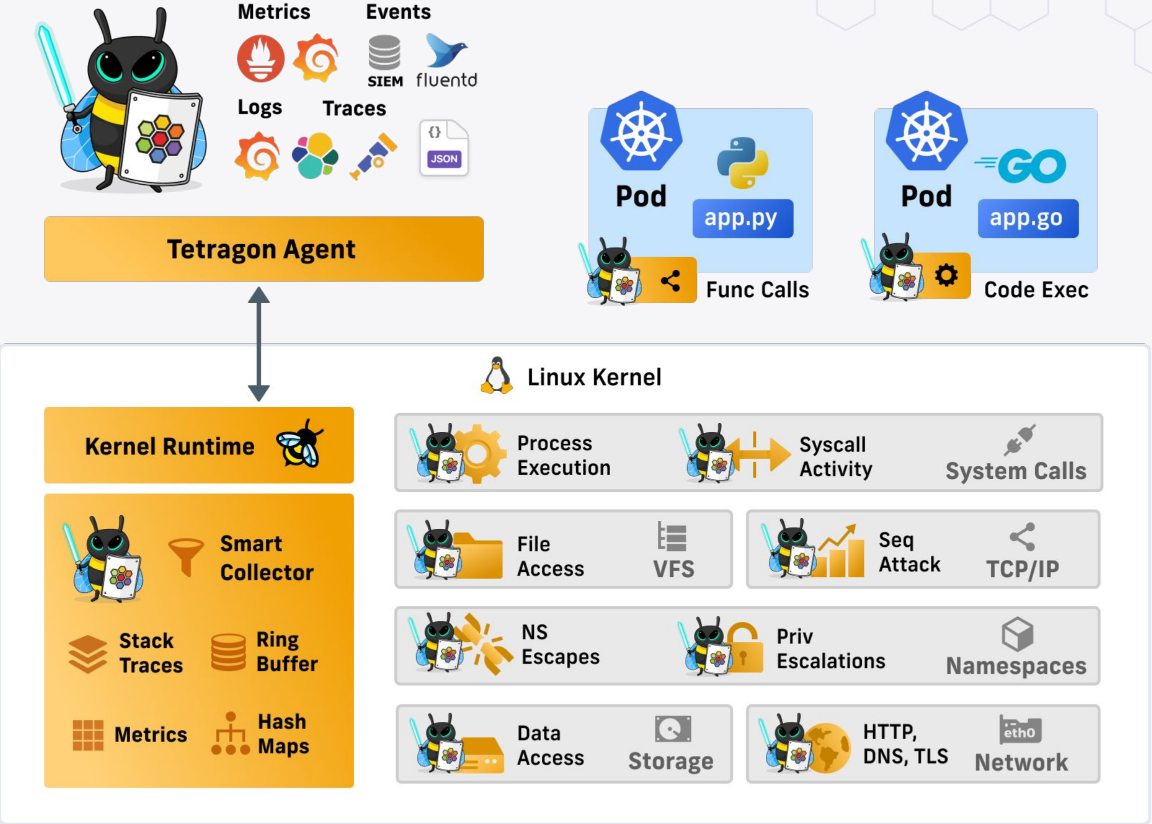

ISOVALENT

LD_PRELOAD

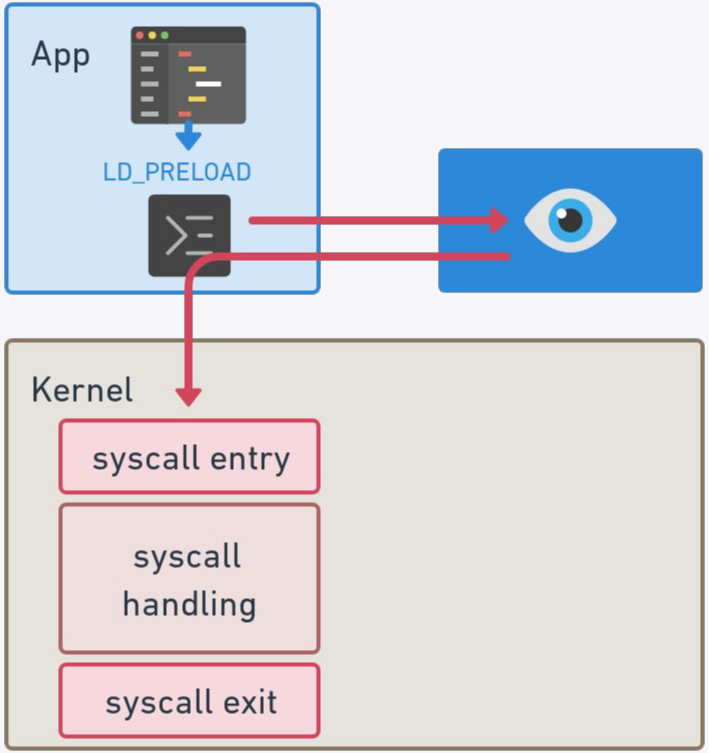

- Standard C library, dynamically linked
- System call API
- Replace the “standard” library

Syscall checks within the kernel

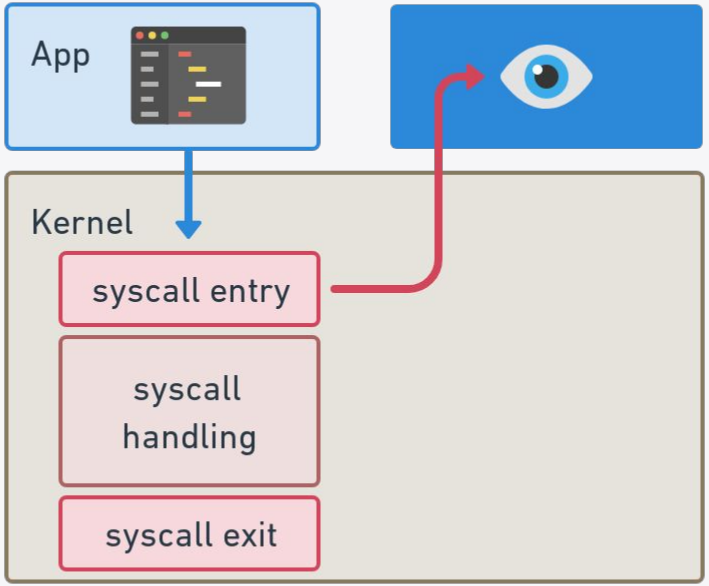

ptrace,

seccomp,

eBPF kprobes on syscall entry

TOCTTOU with syscalls

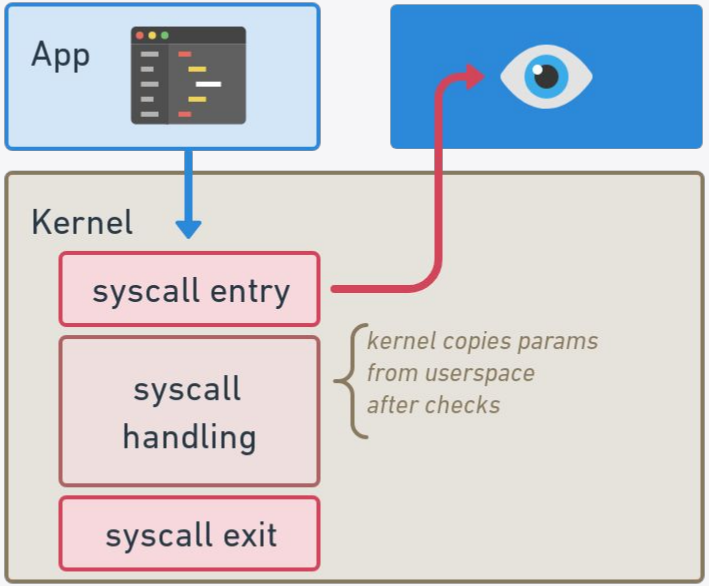

ptrace,

seccomp,

eBPF kprobes on syscall entry

For more details

- Leo Di Donato & KP Singh at CN eBPF Day 2021
- Rex Guo & Junyuan Zeng at DEFCON 29 on Phantom attacks

Need to make the check at the right place

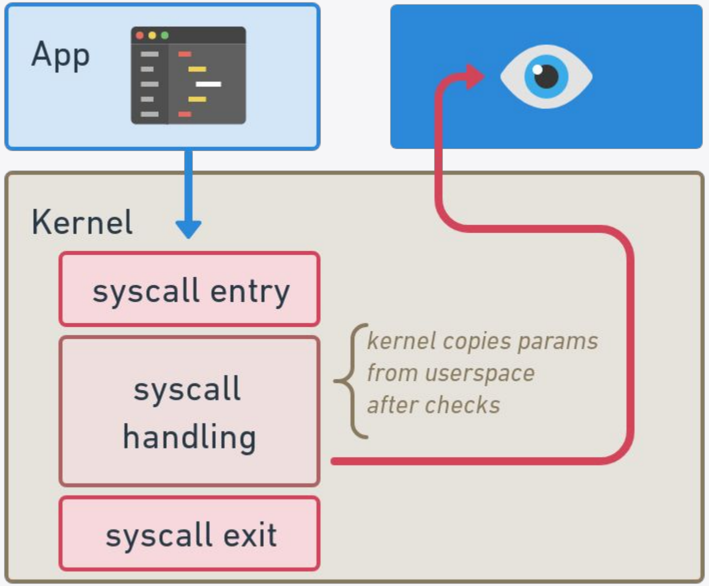

ISOVALENT

Linux Security Modules

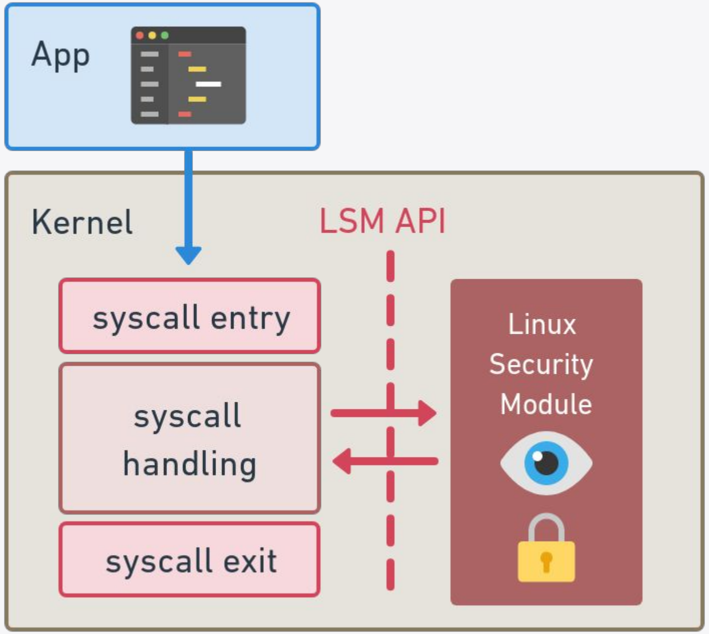

- Stable interface
- Safe places to make checks

BPF LSM

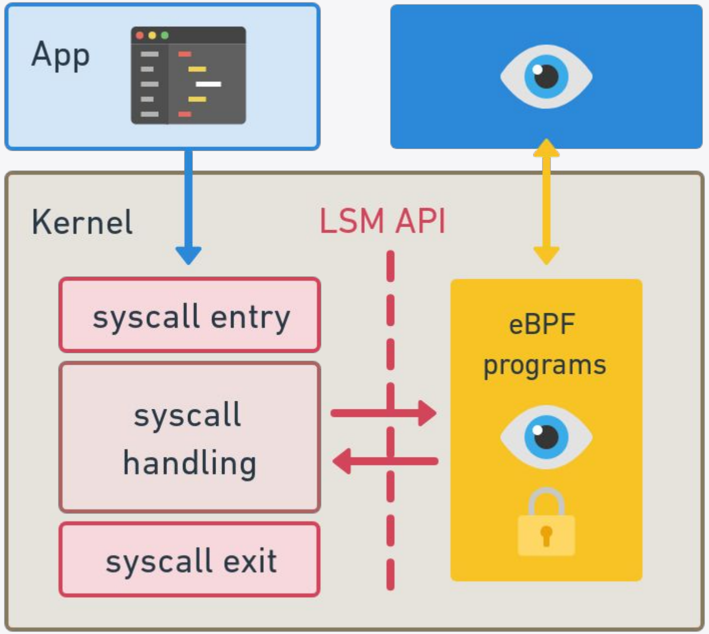

- Stable interface
- Safe places to make checks
- eBPF makes it dynamic
- Protect pre-existing processes

ISOVALENT

BPF LSM

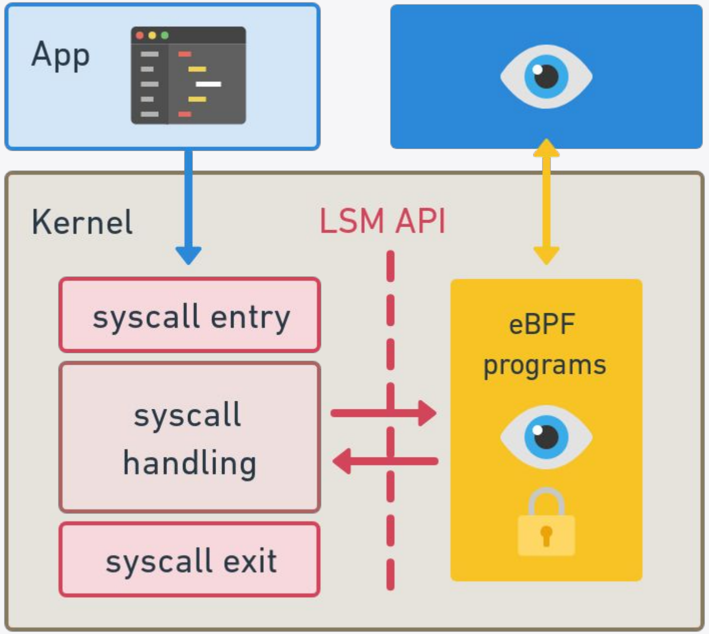

- Stable interface
- Safe places to make checks
- eBPF makes it dynamic
- Protect pre-existing processes
- Needs **kernel 5.7+**

Cilium Tetragon

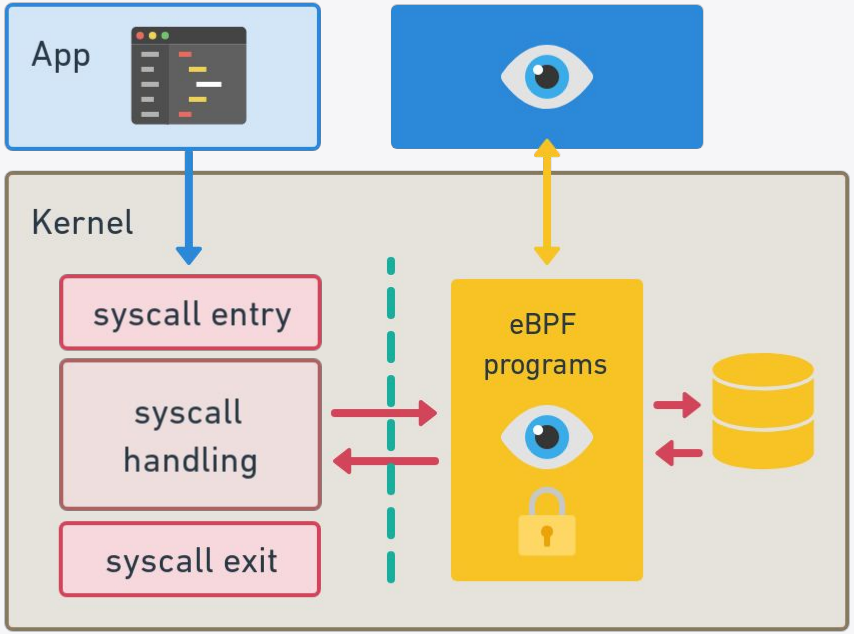

- eBPF makes it dynamic
- Protect pre-existing processes
- Uses **kernel knowledge** to hook into sufficiently stable functions
- Multiple **co-ordinated** eBPF programs
- In-kernel event **filtering**

ISOVALENT

Reactive actions from user space

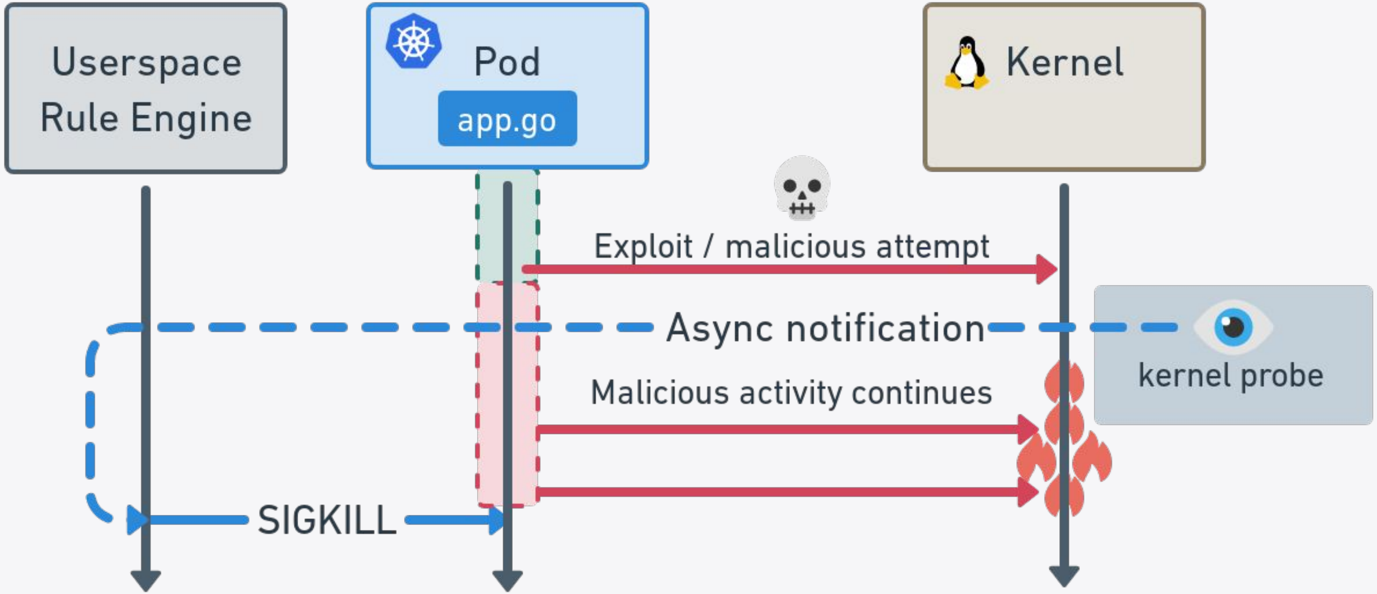

Preventative actions from kernel

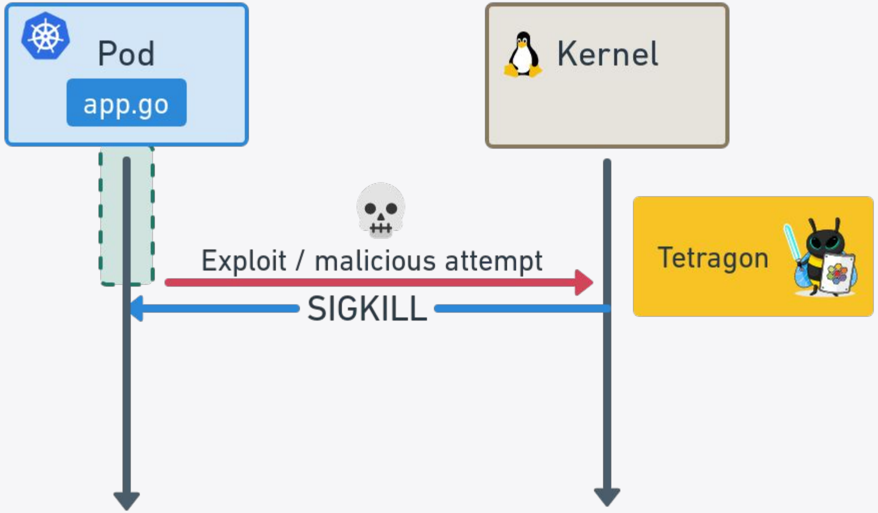

ISOVALENT

Observability

Tetragon

- **Deep Visibility**
  - System, network, protocols, filesystem, applications, ...
- **Transparent**
  - App agonistic
  - No changes to applications
- **Low-Overhead**
  - Minimal overhead
  - Extensive filtering & aggregation
- **Integrations**
  - Prometheus, Grafana, SIEM, fluentd, OpenTelemetry, elasticsearch

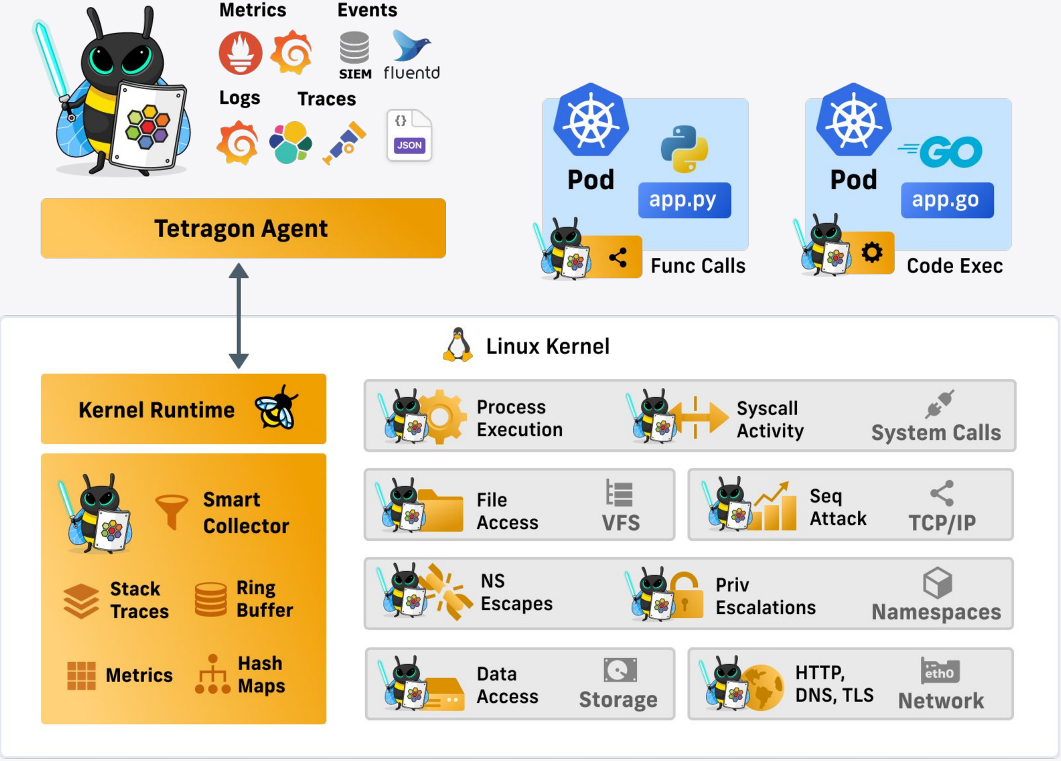

ISOVALENT

Combined Network & Runtime Visibility

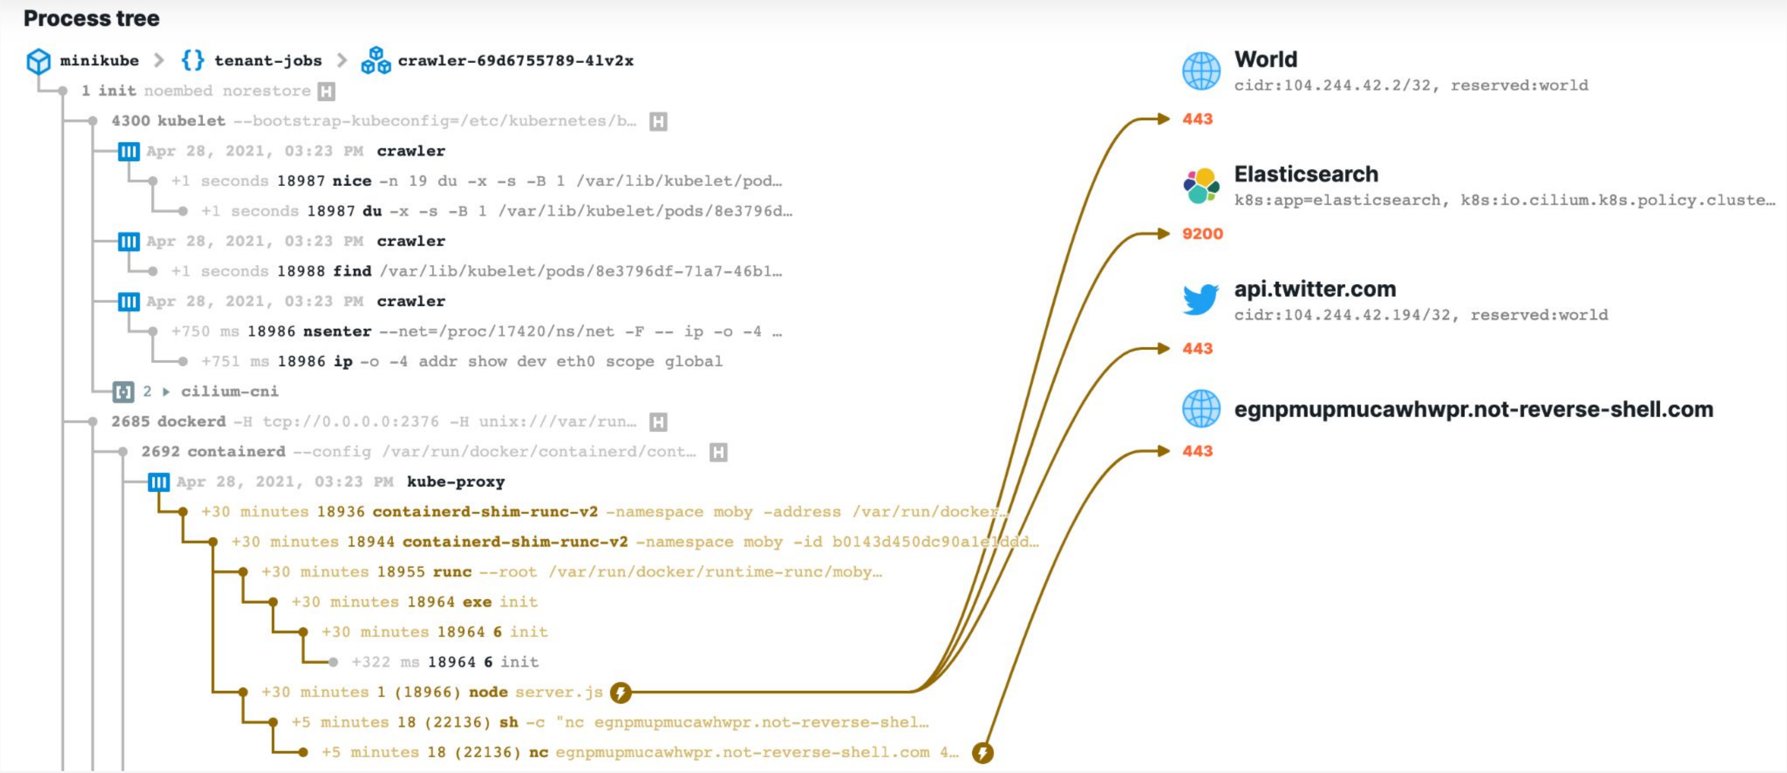

Process tree

```
minikube > {} tenant-jobs > crawler-69d6755789-4lv2x
 1 init noembed norestore H
  4300 kubelet --bootstrap-kubeconfig=/etc/kubernetes/b… H
   Apr 28, 2021, 03:23 PM crawler
    +1 seconds 18987 nice -n 19 du -x -s -B 1 /var/lib/kubelet/pod…
     +1 seconds 18987 du -x -s -B 1 /var/lib/kubelet/pods/8e3796d…
   Apr 28, 2021, 03:23 PM crawler
    +1 seconds 18988 find /var/lib/kubelet/pods/8e3796df-71a7-46b1…
   Apr 28, 2021, 03:23 PM crawler
    +750 ms 18986 nsenter --net=/proc/17420/ns/net -F -- ip -o -4 …
     +751 ms 18986 ip -o -4 addr show dev eth0 scope global
   2 ▸ cilium-cni
  2685 dockerd -H tcp://0.0.0.0:2376 -H unix:///var/run… H
   2692 containerd --config /var/run/docker/containerd/cont… H
    Apr 28, 2021, 03:23 PM kube-proxy
     +30 minutes 18936 containerd-shim-runc-v2 -namespace moby -address /var/run/docker…
      +30 minutes 18944 containerd-shim-runc-v2 -namespace moby -id b0143d450dc90a1e1ddd…
       +30 minutes 18955 runc --root /var/run/docker/runtime-runc/moby…
        +30 minutes 18964 exe init
         +30 minutes 18964 6 init
          +322 ms 18964 6 init
       +30 minutes 1 (18966) node server.js
        +5 minutes 18 (22136) sh -c "nc egnpmupmucawhwpr.not-reverse-shel…
         +5 minutes 18 (22136) nc egnpmupmucawhwpr.not-reverse-shell.com 4…
```

World
cidr:104.244.42.2/32, reserved:world
443

Elasticsearch
k8s:app=elasticsearch, k8s:io.cilium.k8s.policy.cluste…
9200

api.twitter.com
cidr:104.244.42.194/32, reserved:world
443

egnpmupmucawhwpr.not-reverse-shell.com
443

ISOVALENT

TLS/SSL Visibility

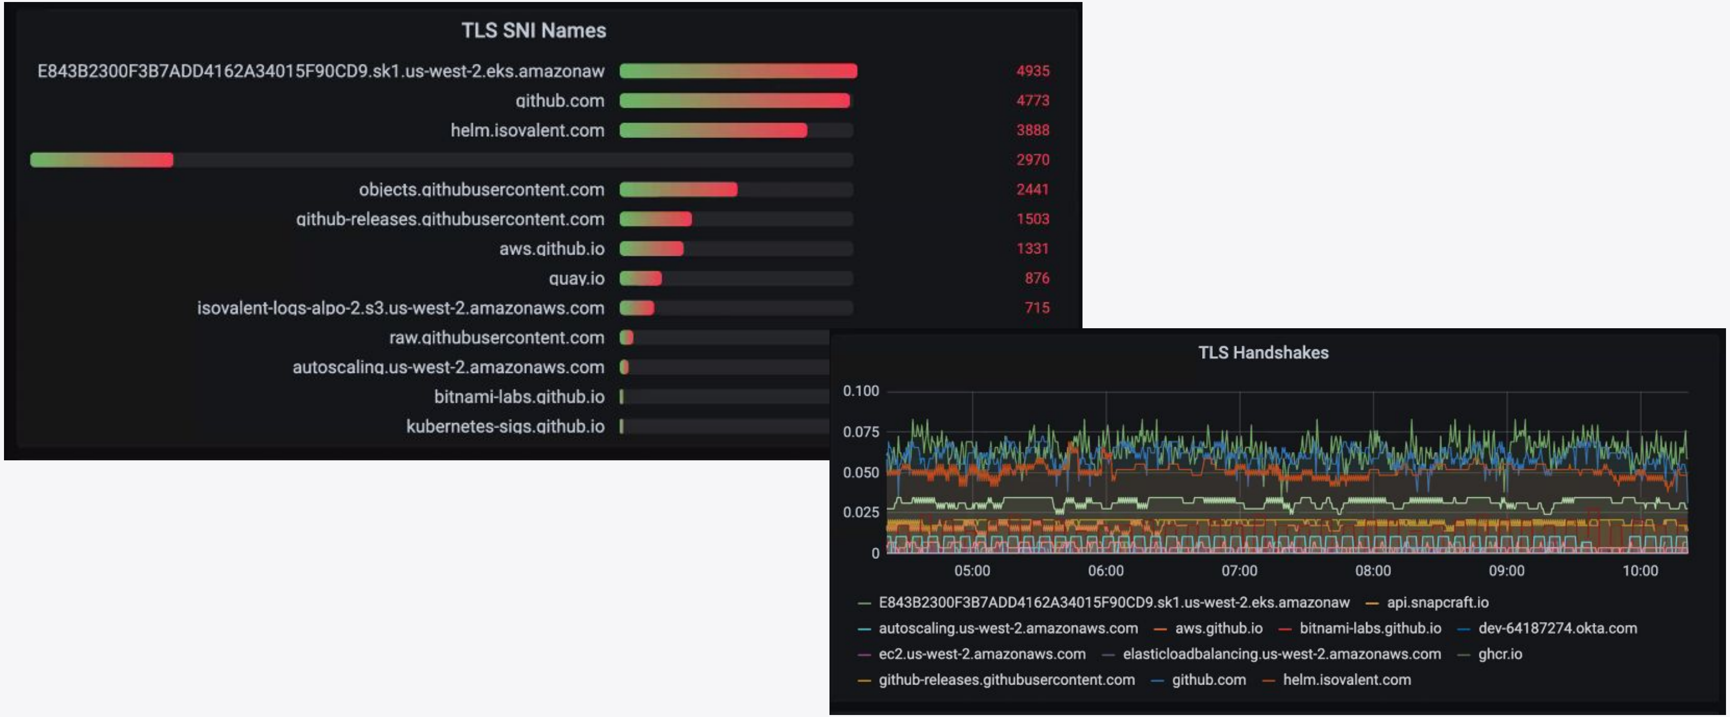

Observing DNS, HTTP, TCP, ...

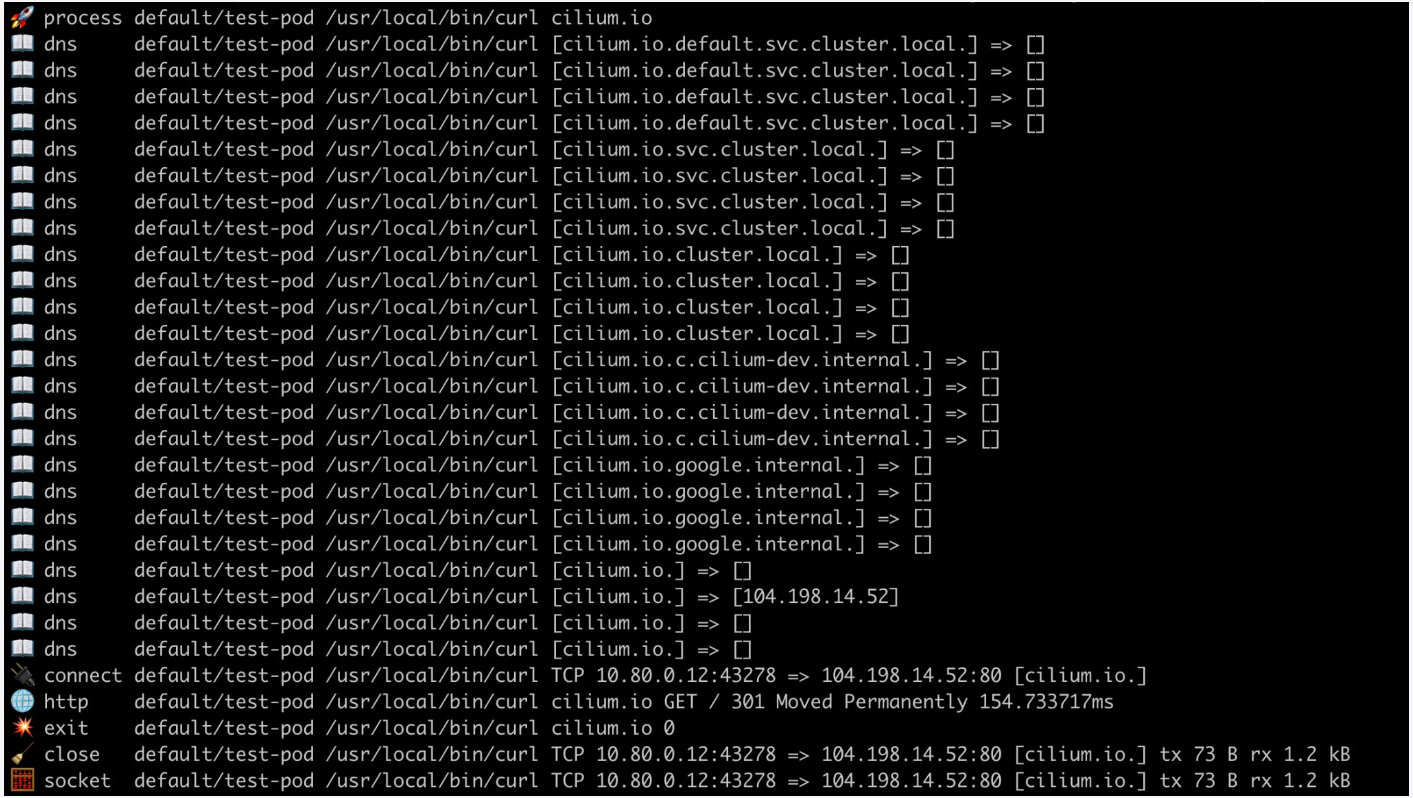

```
🚀 process default/test-pod /usr/local/bin/curl cilium.io
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.default.svc.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.default.svc.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.default.svc.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.default.svc.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.svc.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.svc.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.svc.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.svc.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.cluster.local.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.c.cilium-dev.internal.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.c.cilium-dev.internal.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.c.cilium-dev.internal.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.c.cilium-dev.internal.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.google.internal.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.google.internal.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.google.internal.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.google.internal.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.] => [104.198.14.52]
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.] => []
📖 dns     default/test-pod /usr/local/bin/curl [cilium.io.] => []
🔌 connect default/test-pod /usr/local/bin/curl TCP 10.80.0.12:43278 => 104.198.14.52:80 [cilium.io.]
🌐 http    default/test-pod /usr/local/bin/curl cilium.io GET / 301 Moved Permanently 154.733717ms
💥 exit    default/test-pod /usr/local/bin/curl cilium.io 0
🧹 close   default/test-pod /usr/local/bin/curl TCP 10.80.0.12:43278 => 104.198.14.52:80 [cilium.io.] tx 73 B rx 1.2 kB
🧮 socket  default/test-pod /usr/local/bin/curl TCP 10.80.0.12:43278 => 104.198.14.52:80 [cilium.io.] tx 73 B rx 1.2 kB
```

Monitoring & Preventing Capabilities Abuse

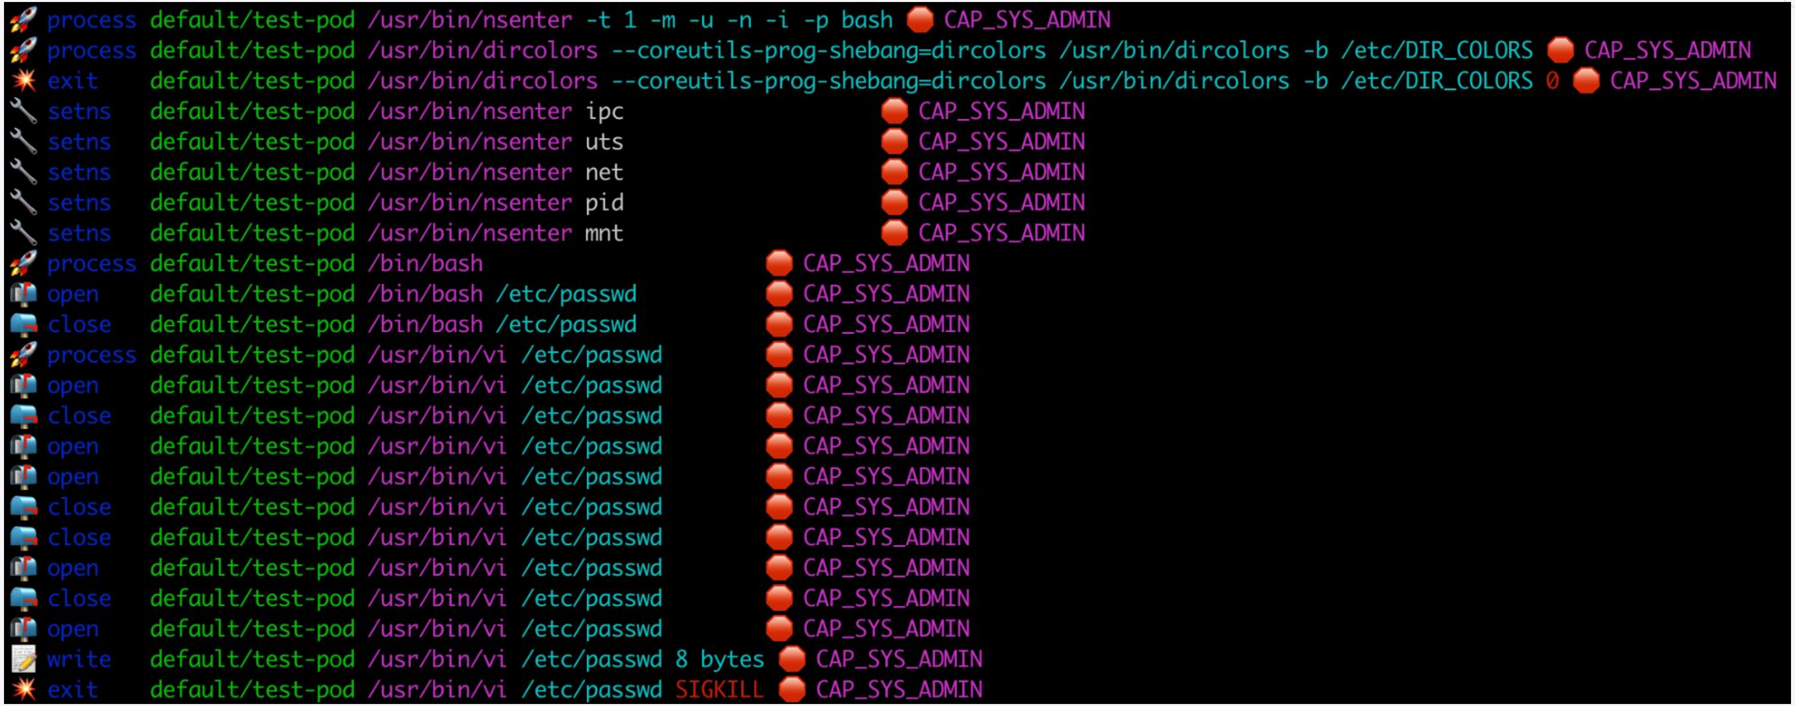

```
🚀 process default/test-pod /usr/bin/nsenter -t 1 -m -u -n -i -p bash 🛑 CAP_SYS_ADMIN
🚀 process default/test-pod /usr/bin/dircolors --coreutils-prog-shebang=dircolors /usr/bin/dircolors -b /etc/DIR_COLORS 🛑 CAP_SYS_ADMIN
💥 exit    default/test-pod /usr/bin/dircolors --coreutils-prog-shebang=dircolors /usr/bin/dircolors -b /etc/DIR_COLORS 0 🛑 CAP_SYS_ADMIN
🔧 setns   default/test-pod /usr/bin/nsenter ipc                    🛑 CAP_SYS_ADMIN
🔧 setns   default/test-pod /usr/bin/nsenter uts                    🛑 CAP_SYS_ADMIN
🔧 setns   default/test-pod /usr/bin/nsenter net                    🛑 CAP_SYS_ADMIN
🔧 setns   default/test-pod /usr/bin/nsenter pid                    🛑 CAP_SYS_ADMIN
🔧 setns   default/test-pod /usr/bin/nsenter mnt                    🛑 CAP_SYS_ADMIN
🚀 process default/test-pod /bin/bash                      🛑 CAP_SYS_ADMIN
📬 open    default/test-pod /bin/bash /etc/passwd          🛑 CAP_SYS_ADMIN
📪 close   default/test-pod /bin/bash /etc/passwd          🛑 CAP_SYS_ADMIN
🚀 process default/test-pod /usr/bin/vi /etc/passwd        🛑 CAP_SYS_ADMIN
📬 open    default/test-pod /usr/bin/vi /etc/passwd        🛑 CAP_SYS_ADMIN
📪 close   default/test-pod /usr/bin/vi /etc/passwd        🛑 CAP_SYS_ADMIN
📬 open    default/test-pod /usr/bin/vi /etc/passwd        🛑 CAP_SYS_ADMIN
📬 open    default/test-pod /usr/bin/vi /etc/passwd        🛑 CAP_SYS_ADMIN
📪 close   default/test-pod /usr/bin/vi /etc/passwd        🛑 CAP_SYS_ADMIN
📪 close   default/test-pod /usr/bin/vi /etc/passwd        🛑 CAP_SYS_ADMIN
📬 open    default/test-pod /usr/bin/vi /etc/passwd        🛑 CAP_SYS_ADMIN
📪 close   default/test-pod /usr/bin/vi /etc/passwd        🛑 CAP_SYS_ADMIN
📬 open    default/test-pod /usr/bin/vi /etc/passwd        🛑 CAP_SYS_ADMIN
📝 write   default/test-pod /usr/bin/vi /etc/passwd 8 bytes 🛑 CAP_SYS_ADMIN
💥 exit    default/test-pod /usr/bin/vi /etc/passwd SIGKILL 🛑 CAP_SYS_ADMIN
```

ISOVALENT

ISOVALENT

Thank you!