ISOVALENT

Cloud Native Superpowers with **eBPF**

Liz Rice | @lizrice
Chief Open Source Officer, Isovalent
Emeritus Chair, CNCF Technical Oversight Committee | CNCF & OpenUK boards

ISOVALENT

What is  eBPF ?

Makes the kernel programmable

Run custom code in the kernel

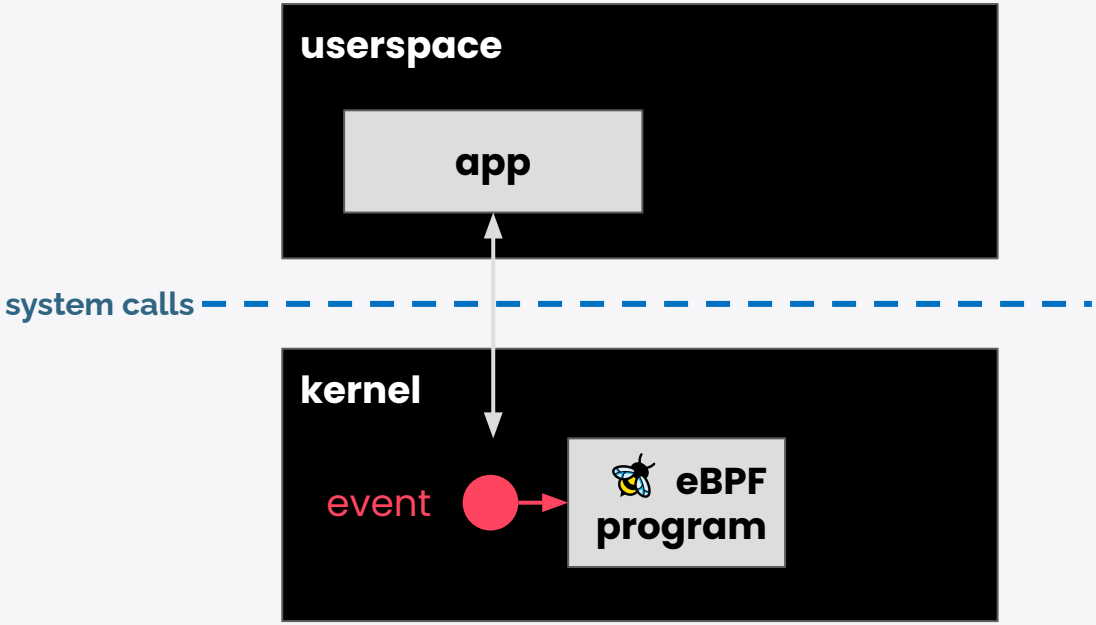

You, 14 months ago | 1 author (You)

```python
#!/usr/bin/python3
from bcc import BPF

program = r"""
int hello(void *ctx) {
    bpf_trace_printk("Hello World!");
    return 0;
}
"""

b = BPF(text=program)
syscall = b.get_syscall_fnname("execve")
b.attach_kprobe(event=syscall, fn_name="hello")

b.trace_print()
```

```
root@lima-learning-ebpf >
```

TERMINAL PROBLEMS 3 OUTPUT DEBUG CONSOLE PORTS COMMENTS

```
root@lima-learning-ebpf >
```

hello
ping

ISOVALENT

eBPF Hello World

```c
SEC("kprobe/sys_execve")
int hello(void *ctx)
{
    bpf_printk("Hello!");
    return 0;
}
```

\+ userspace code to load eBPF program

Info about process that called execve syscall

```
$ sudo ./hello
  bash-20241   [004] d... 84210.752785: 0: I'm alive!
  bash-20242   [004] d... 84216.321993: 0: I'm alive!
  bash-20243   [004] d... 84225.858880: 0: I'm alive!
```

ISOVALENT

Dynamically change kernel behaviour

@lizrice

Application Developer:

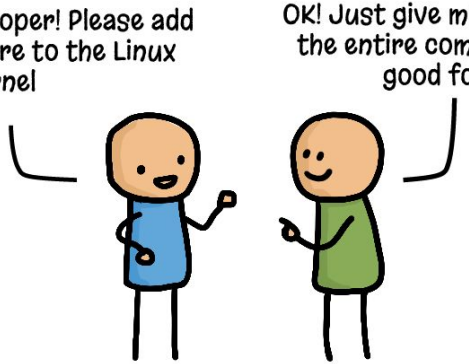

1 year later...

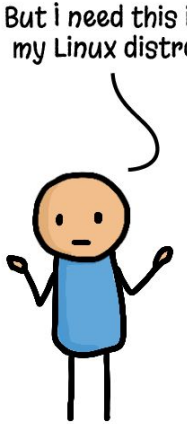

5 years later...

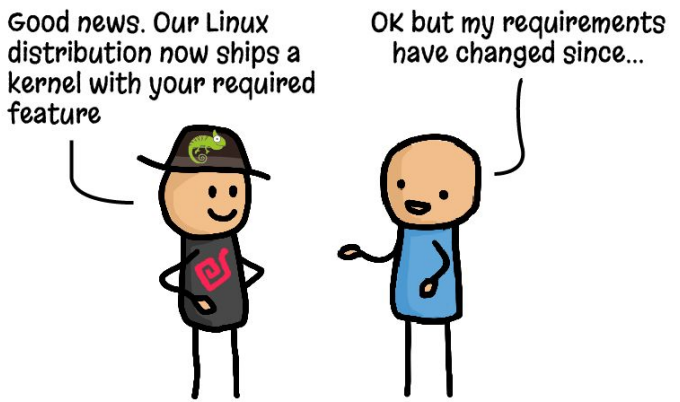

Application Developer:

i want this new feature to observe my app

eBPF Developer:

OK! The kernel can't do this so let me quickly solve this with eBPF.

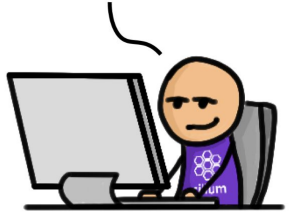

A couple of days later...

Here is a release of our eBPF project that has this feature now. BTW, you don't have to reboot your machine.

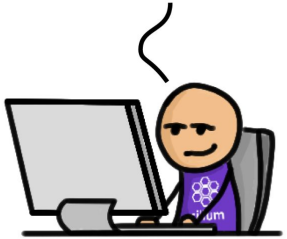

ISOVALENT

Packet of Death mitigation

@lizrice

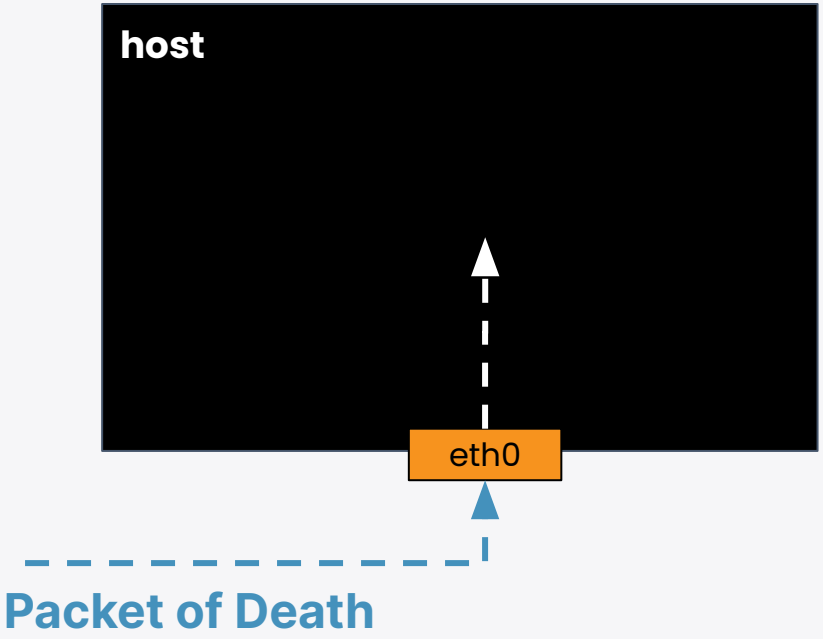
host
eth0
Packet of Death

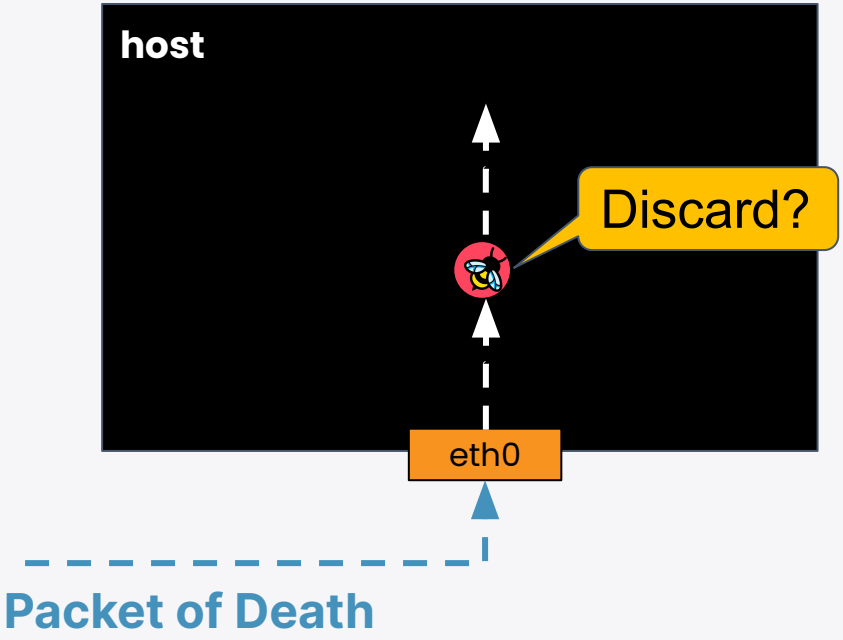
host
Discard?
eth0
Packet of Death

ISOVALENT

eBPF Packet Drop

```
SEC("xdp/bye")
int goodbye_ping(struct xdp_md *ctx)
{
    ...
    if (iph->protocol == IPPROTO_ICMP)
       return XDP_DROP;

    return XDP_PASS;
}
```

```
34
35  int xdp(struct xdp_md *ctx) {
36    void *data = (void *)(long)ctx->data;
37    void *data_end = (void *)(long)ctx->data_end;
38
39    if (is_icmp_ping_request(data, data_end)) {
40      struct iphdr *iph = data + sizeof(struct ethhdr);
41      struct icmphdr *icmp = data + sizeof(struct ethhdr) + sizeof(struct iphdr);
42      bpf_trace_printk("[xdp] ICMP request for %x type %x", iph->daddr,
43                       icmp->type);
44      return XDP_PASS;
45    }
46
47    return XDP_PASS;
48  }
49
```

```
root@lima-learning-ebpf >
```

TERMINAL PROBLEMS 3 OUTPUT DEBUG CONSOLE PORTS COMMENTS

```
root@lima-learning-ebpf >
```

hello
ping

eBPF code has to be safe

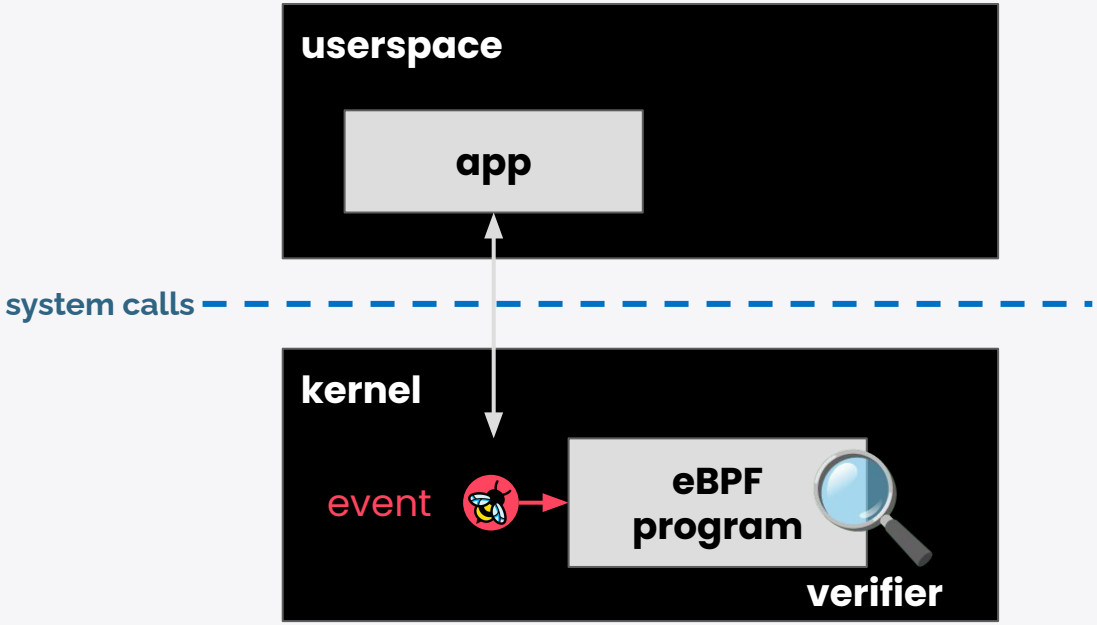

ISOVALENT

Programmable kernel in Kubernetes

@lizrice

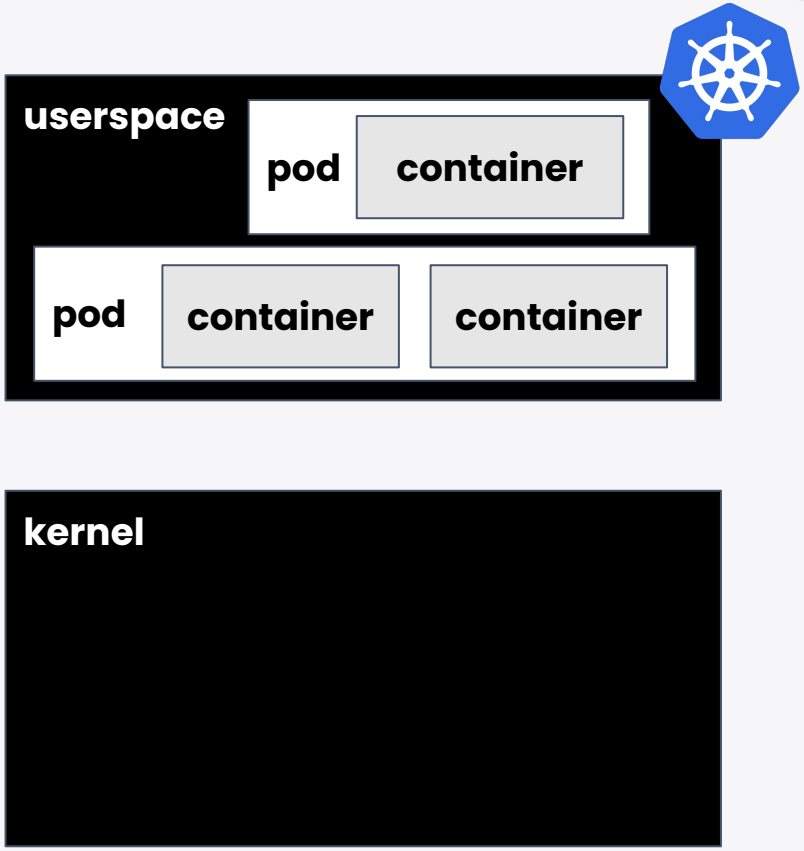

One kernel per host

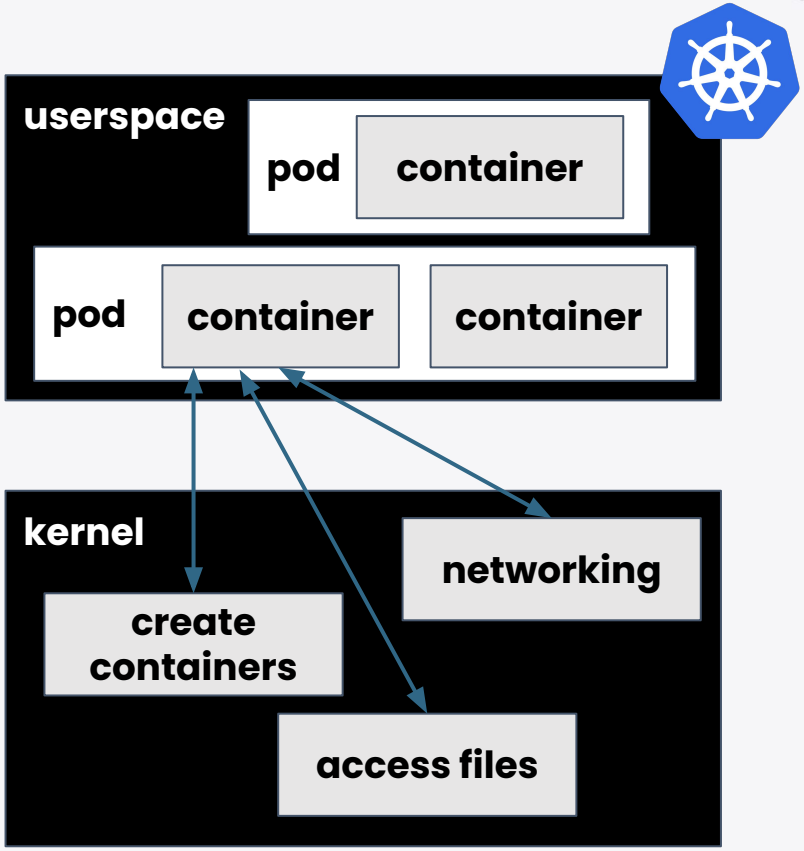

One kernel per host

ISOVALENT

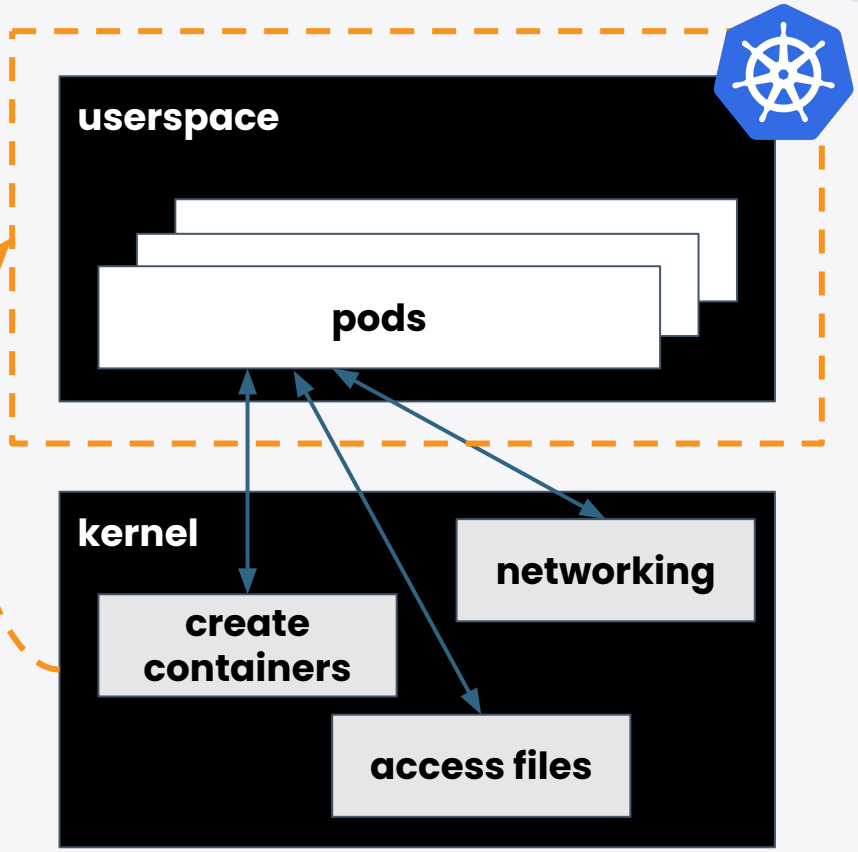

Kernel aware of everything on the host

ISOVALENT

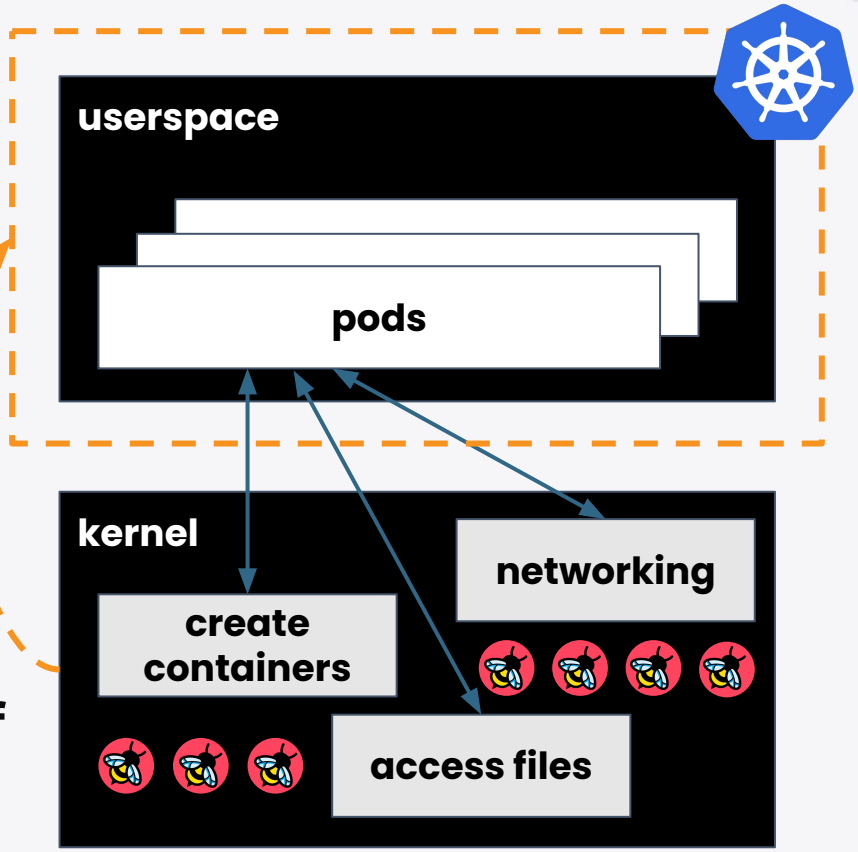

eBPF programs can be aware of everything

ISOVALENT

eBPF apps have a view across the entire node enabling deep observability

@lizrice

ISOVALENT

Cilium Hubble observability

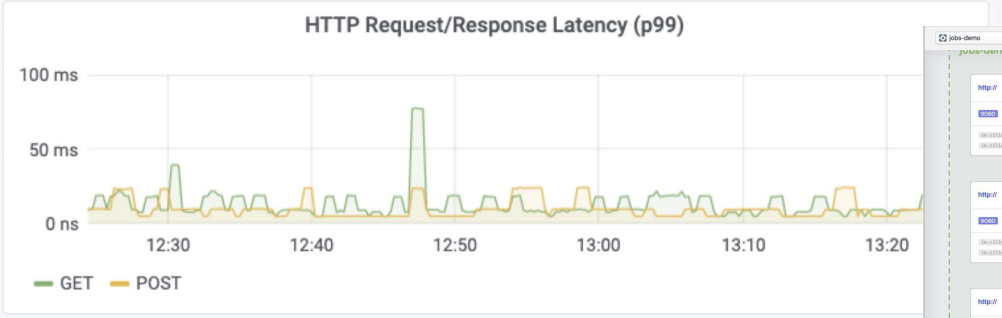

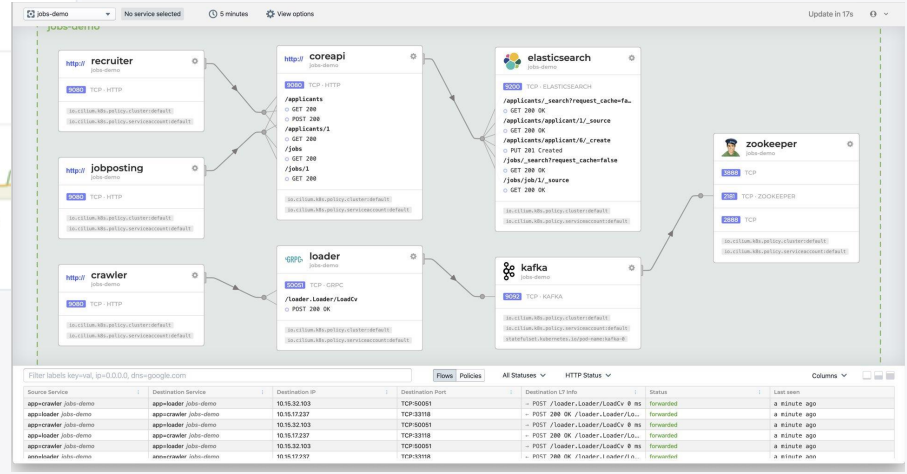

- Network flow logs
- Prometheus metrics
- Service map
- L3/4 & L7 (HTTP, DNS, Kafka, ...)
- Aware of Kubernetes identities

@lizrice

ISOVALENT

PIXIE flame graph

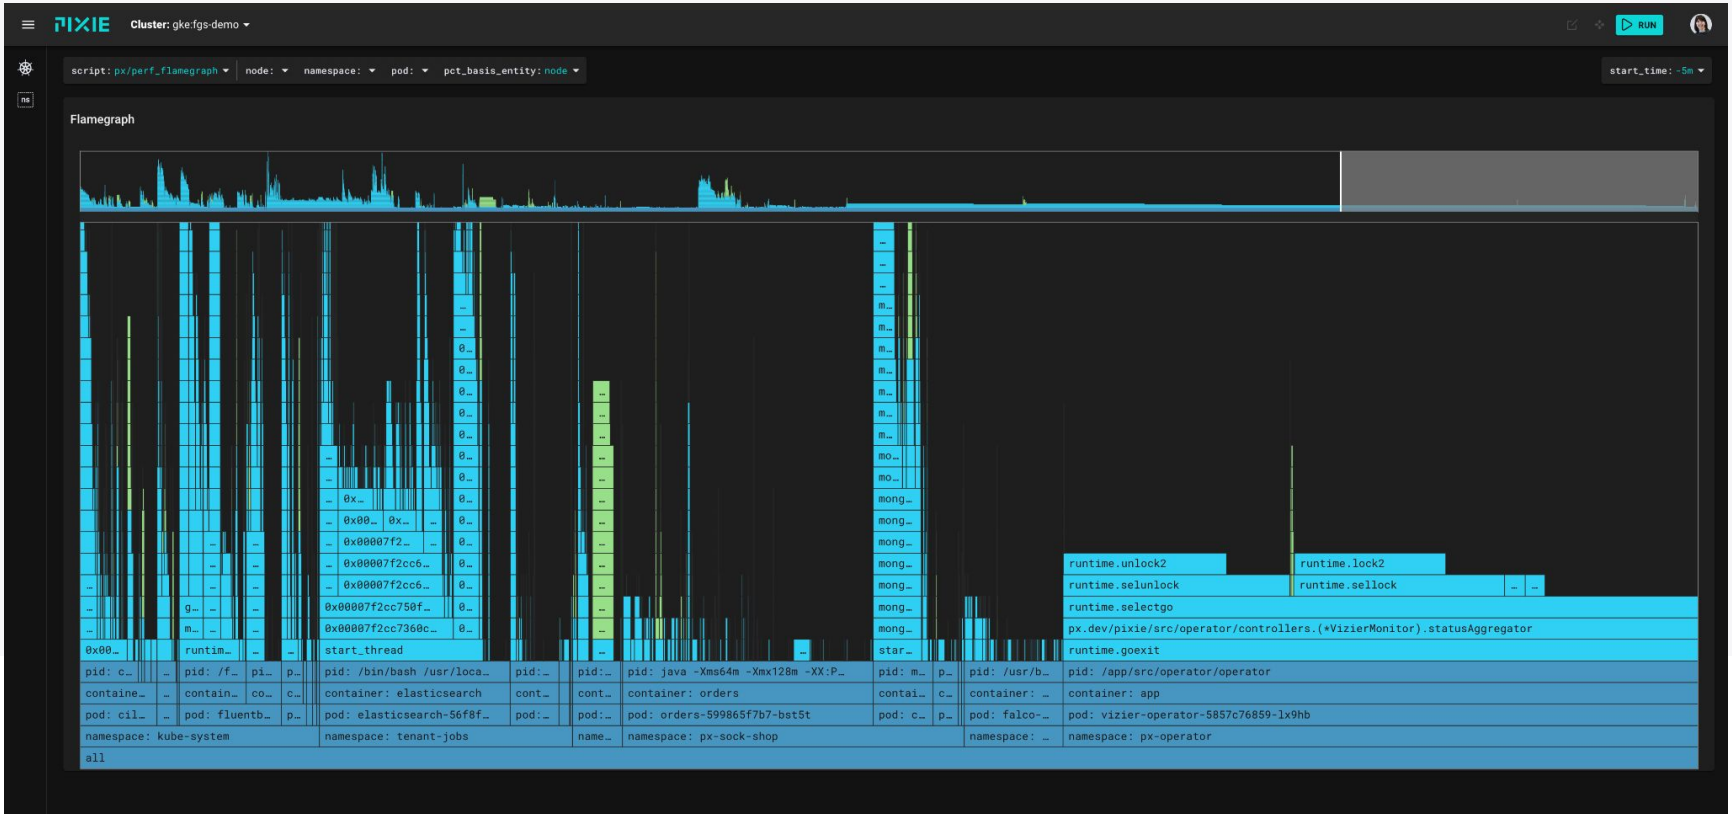

ISOVALENT

eBPF tools have a view across the entire node enabling network efficiency

@lizrice

ISOVALENT

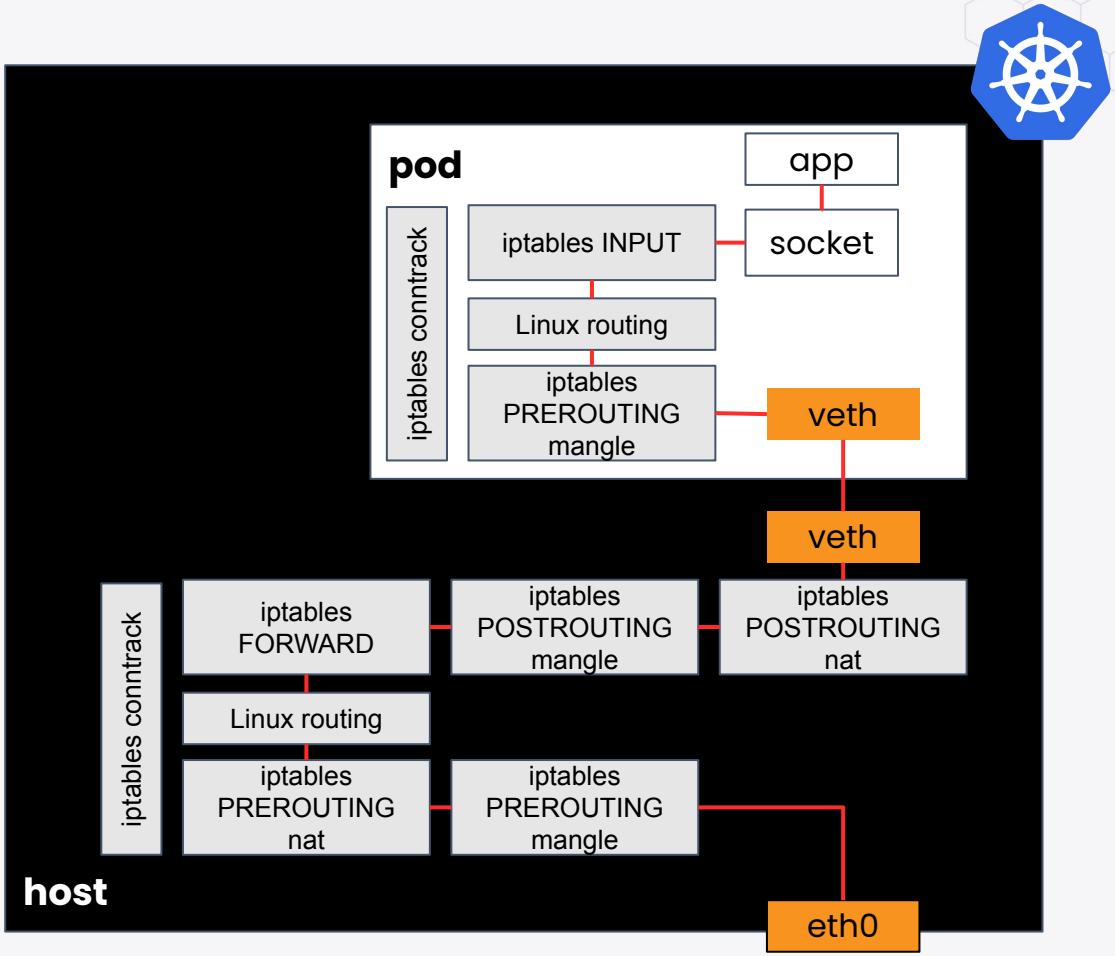

@lizrice

ISOVALENT

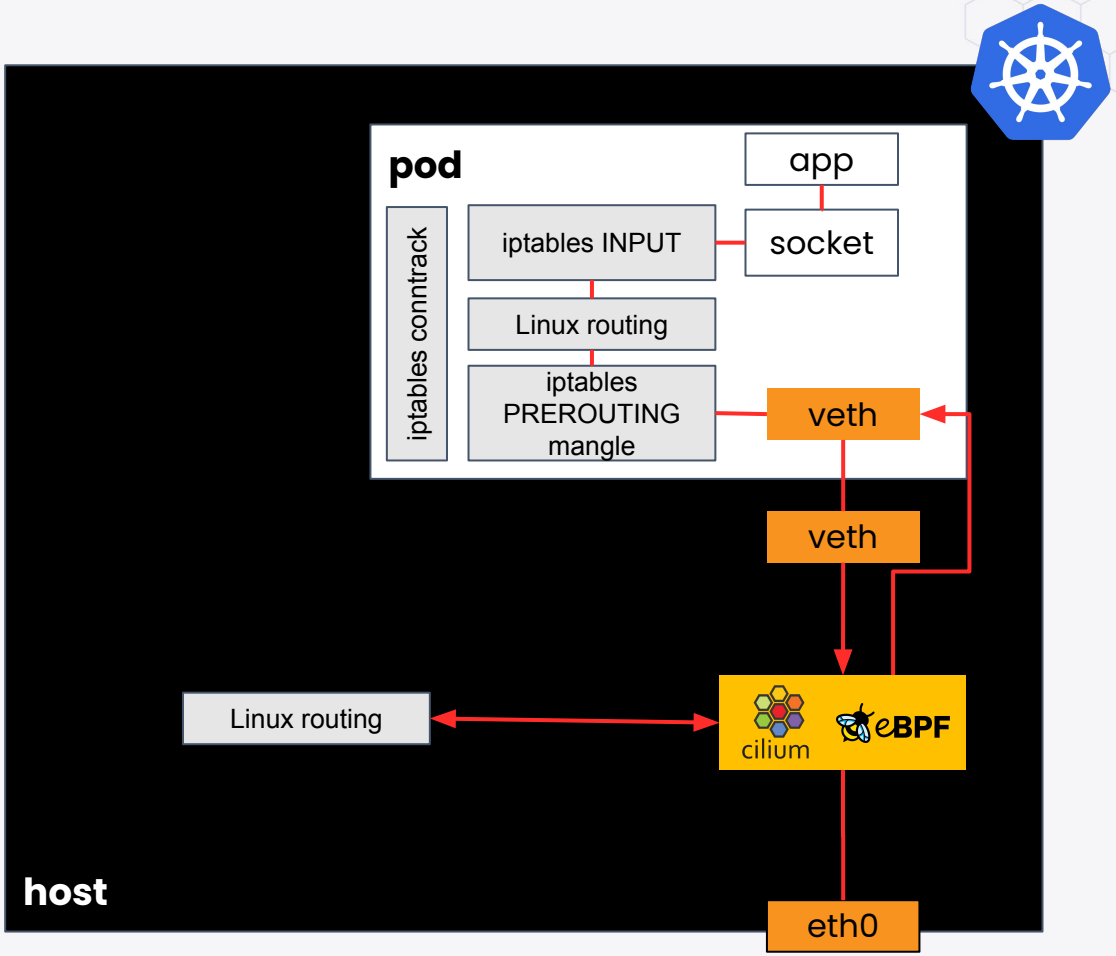

@lizrice

ISOVALENT

TCP RR (higher is better)

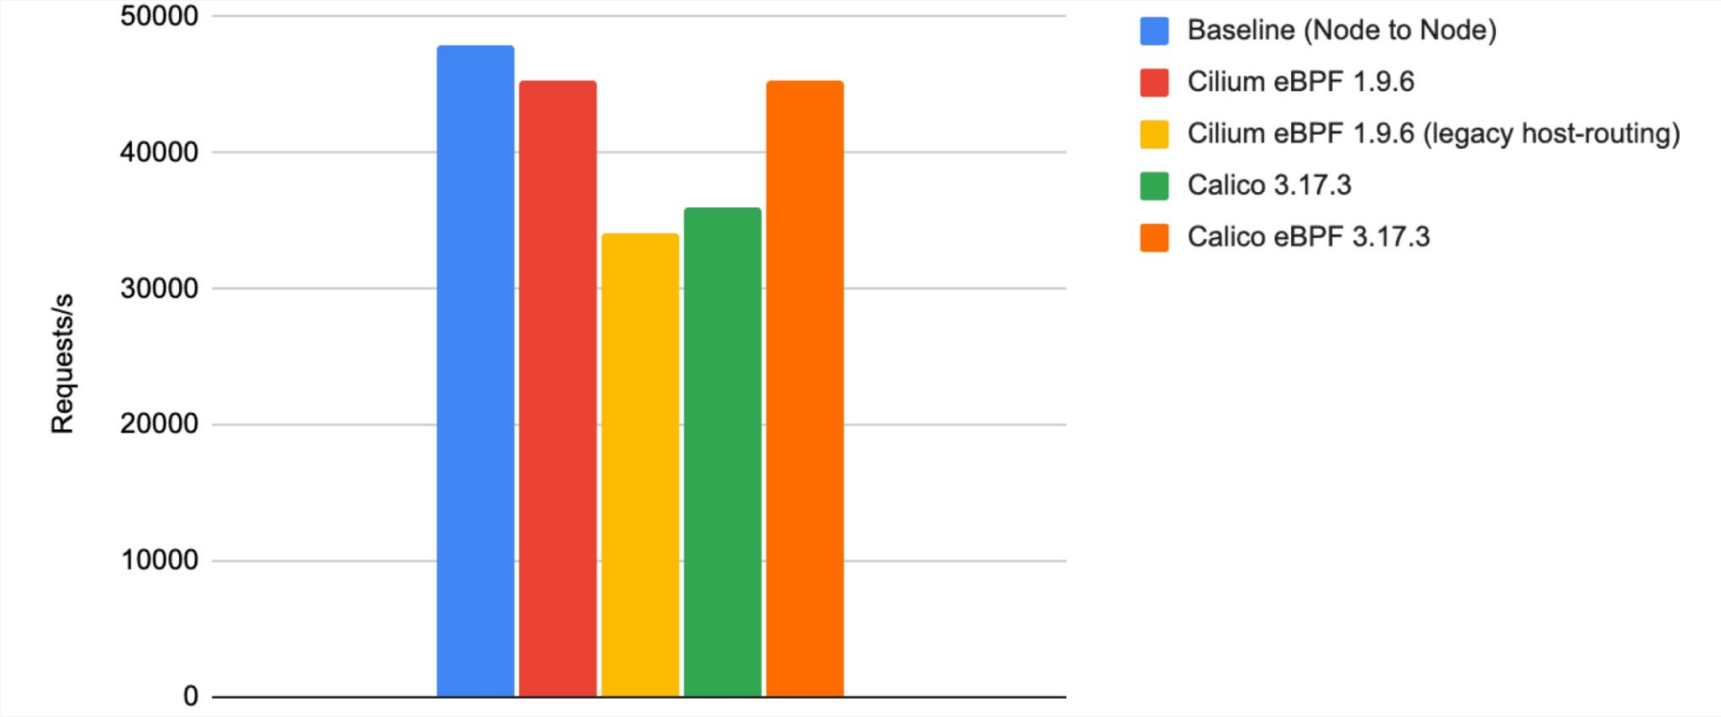

https://cilium.io/blog/2021/05/11/**cni-benchmark**

@lizrice

ISOVALENT

eBPF tools have a view across the entire node without any app or config changes

@lizrice

My other sidecar is a **kernel**

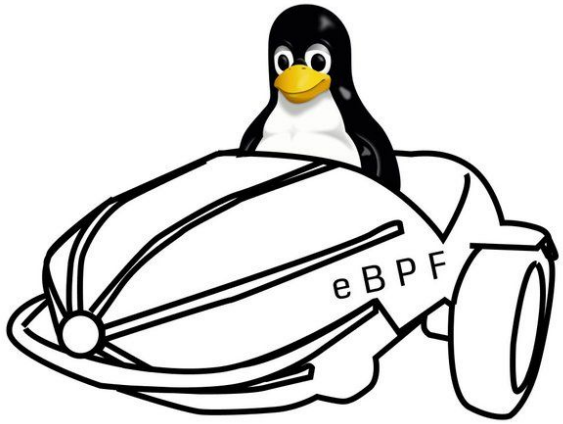

"Get in loser. We're going tracing"

- Nathan LeClaire @dotpem

ISOVALENT

A sidecar has a view across one pod

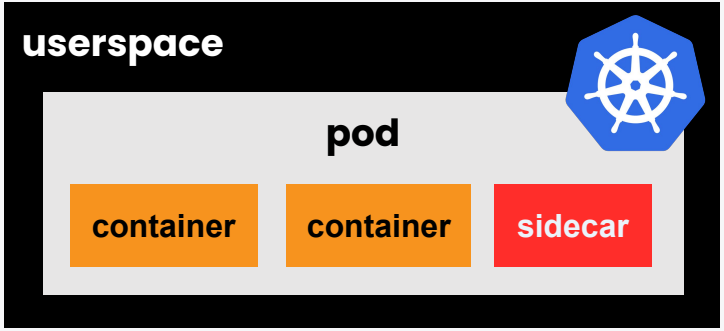

@lizrice

ISOVALENT

Sidecars need YAML

my-app.yaml

```
containers:
  -   name: my-app
      ...
  -   name: my-app-init
      …
  -   name: my-sidecar
      ...
```

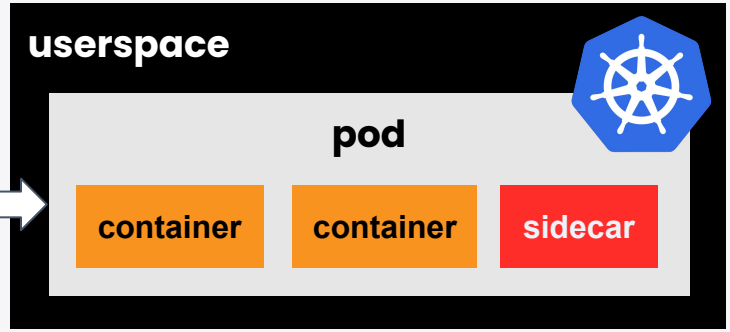

ISOVALENT

eBPF does not need any app changes

my-app.yaml

```
containers:
  -   name: my-app
      ...
  -   name: my-app-init
      …
```

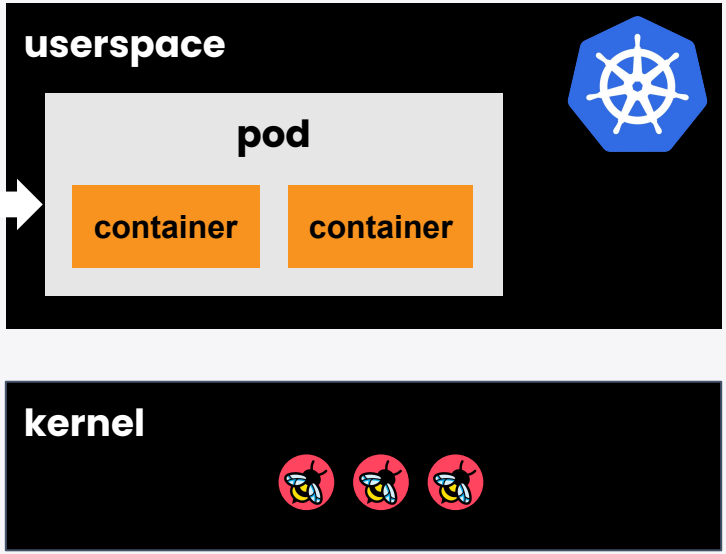

ISOVALENT

eBPF can see ALL activity on the node

my-app.yaml

```
containers:
  -   name: my-app
      ...
  -   name: my-app-init
      …
```

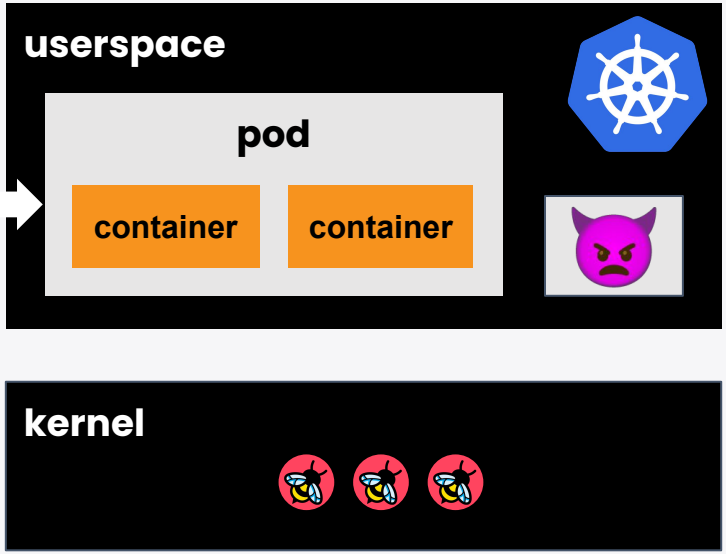

@lizrice

ISOVALENT

eBPF enables sidecarless Service Mesh

@lizrice

ISOVALENT

Reduce resource usage

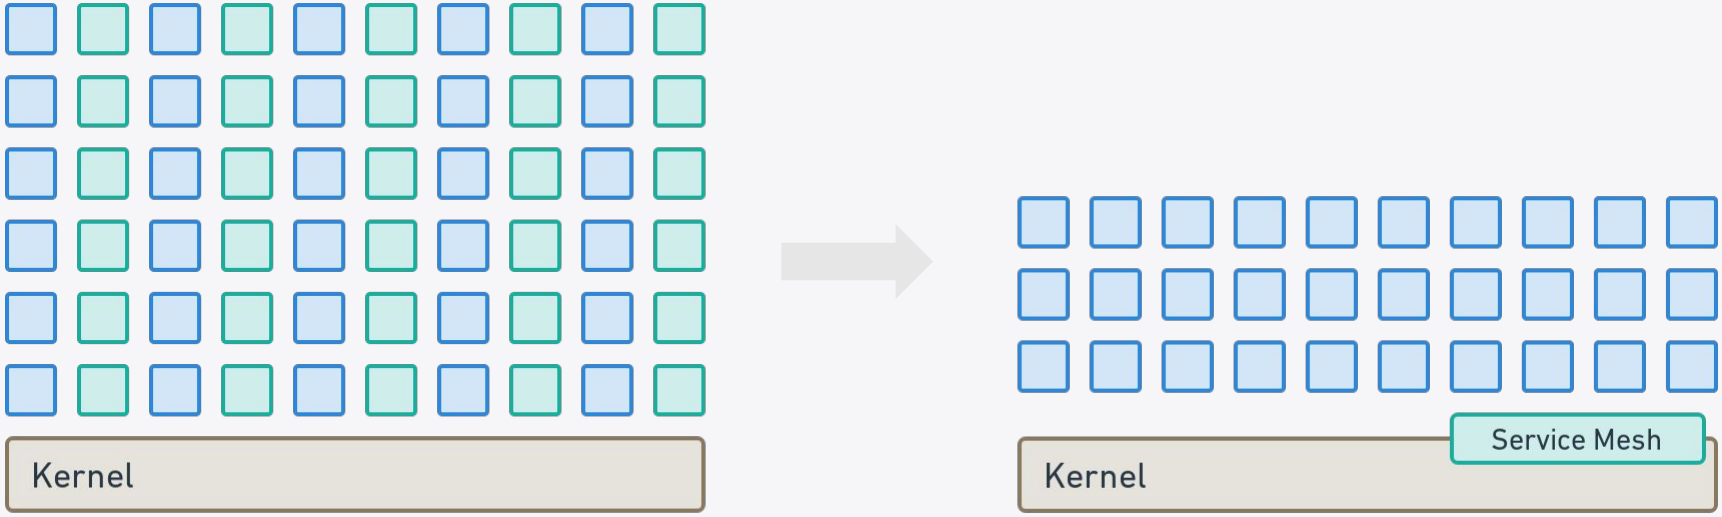

ISOVALENT

Network path with sidecar

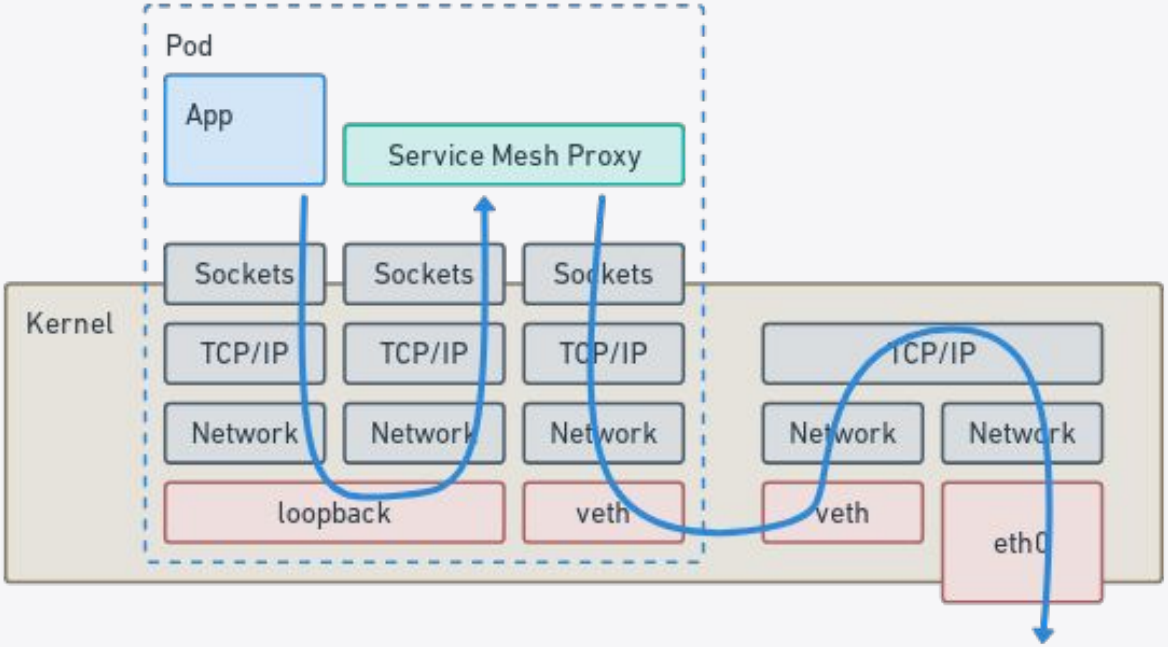

ISOVALENT

Network path for L3/4 traffic

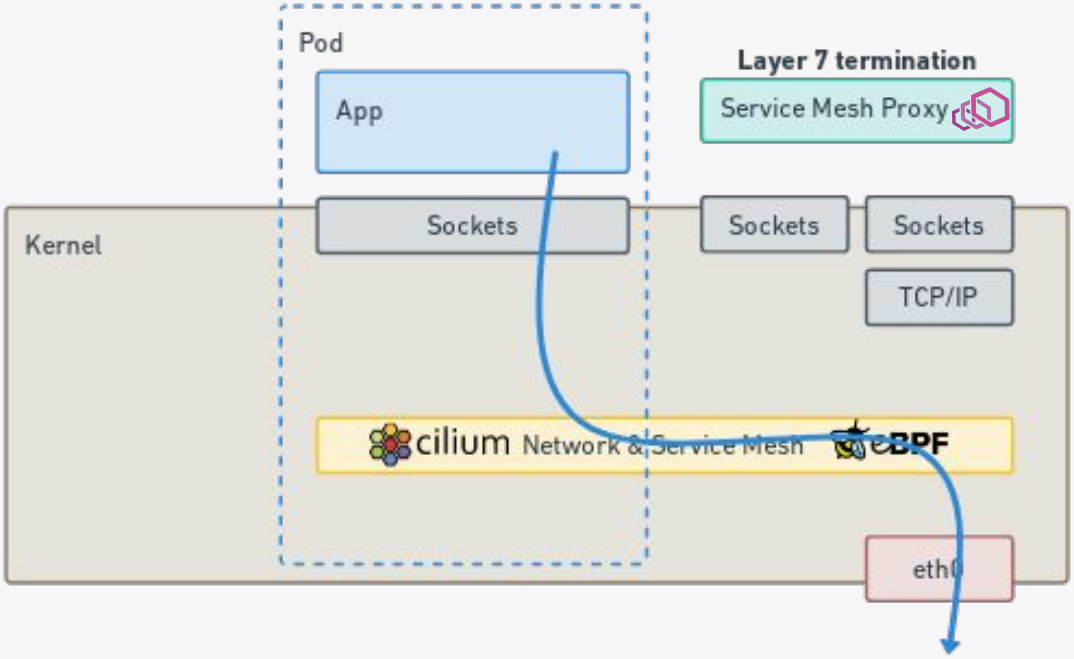

ISOVALENT

Envoy for Layer 7 terminations when needed

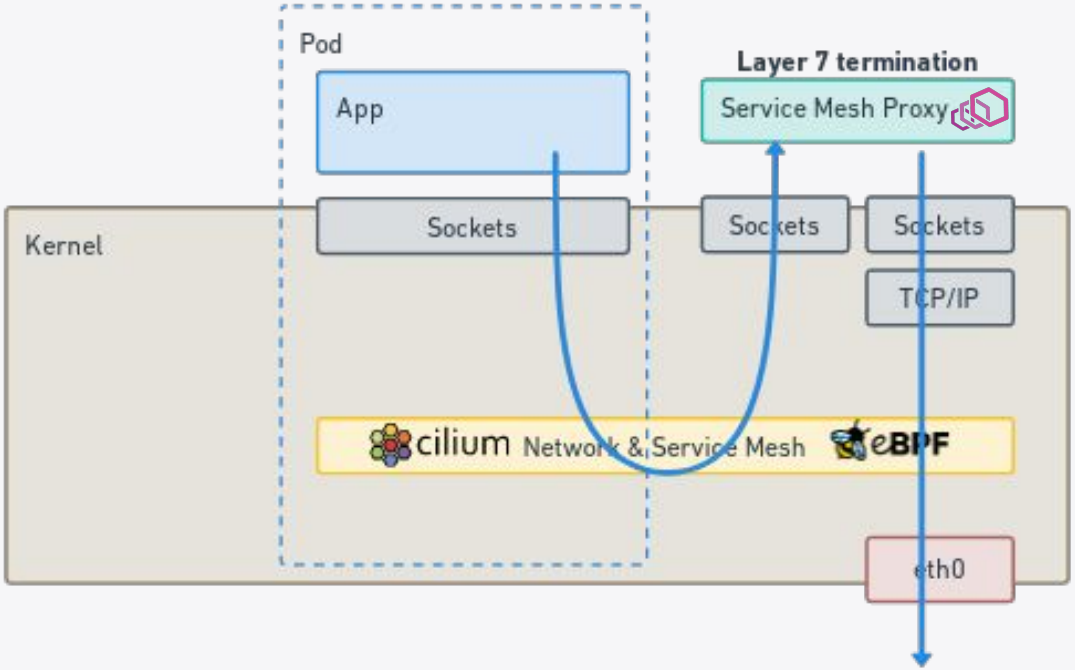

ISOVALENT

eBPF tools have a view across the entire node enabling powerful security

@lizrice

Cilium transparent encryption

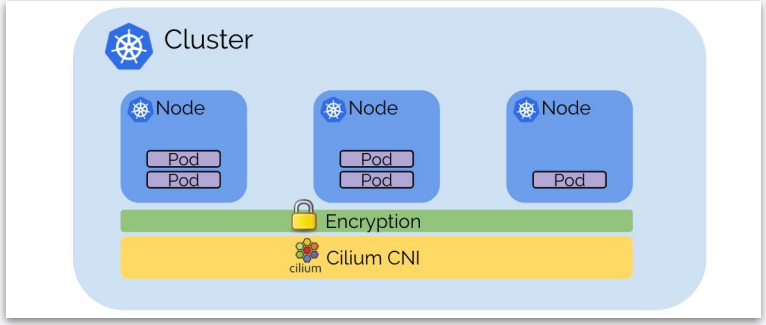

- IPSec / WireGuard
- Traffic encrypted in the kernel
- No app changes needed

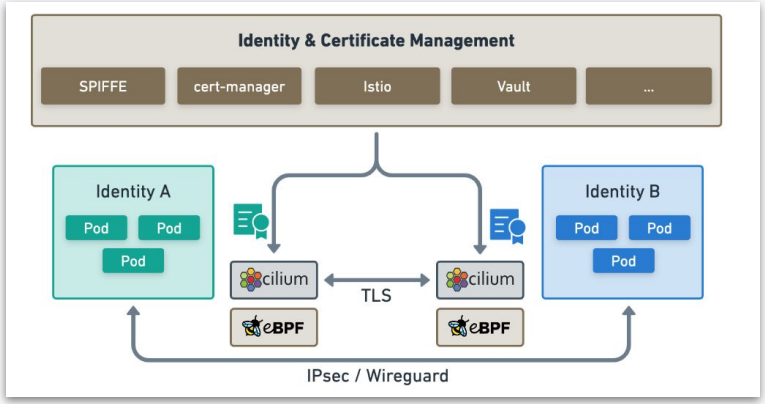

Next-generation mutual auth

- App identity authenticated with TLS handshake

ISOVALENT

Cilium network policy → eBPF programs drop packets

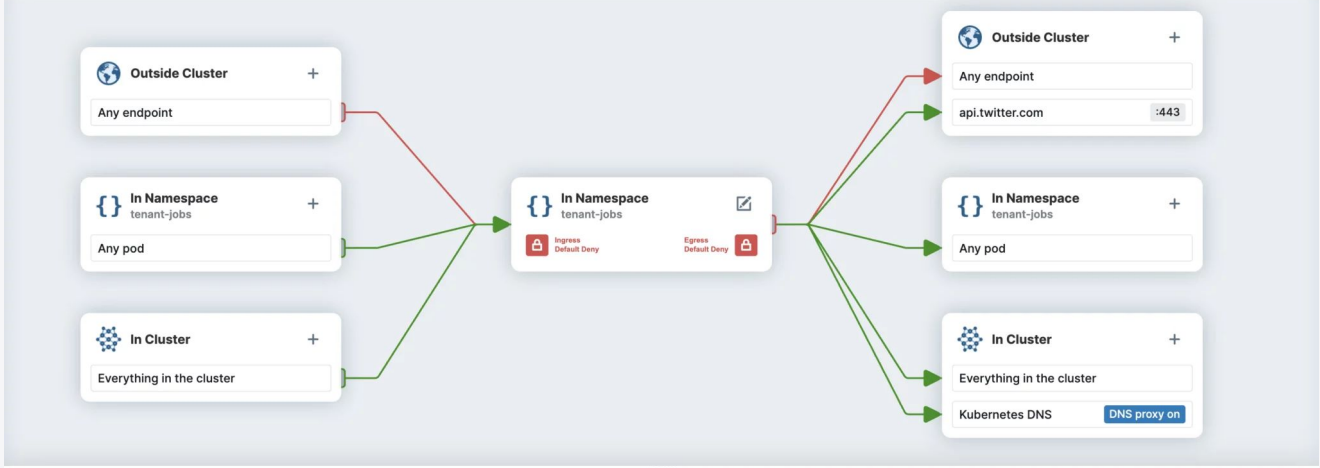

Kubernetes | Cilium | Download

```
21        - ports:
22            - port: "53"
23              protocol: ANY
24          rules:
25            dns:
26              - matchPattern: "*"
27      - toFQDNs:
28          - matchName: api.twitter.com
29        toPorts:
30          - ports:
31              - port: "443"
```

Source Identity	Destination Identity	Verdict
crawler *tenant-jobs*	api.twitter.com	forwarded
crawler *tenant-jobs*	elasticsearch *tenant-jobs*	forwarded
crawler *tenant-jobs*	elasticsearch *tenant-jobs*	forwarded
crawler *tenant-jobs*	api.twitter.com	forwarded
loader *tenant-jobs*	kafka *tenant-jobs*	forwarded
coreapi *tenant-jobs*	elasticsearch *tenant-jobs*	forwarded
jobposting *tenant-jobs*	coreapi *tenant-jobs*	forwarded
jobposting *tenant-jobs*	coreapi *tenant-jobs*	forwarded
coreapi *tenant-jobs*	elasticsearch *tenant-jobs*	forwarded
kafka *tenant-jobs*	zookeeper *tenant-jobs*	forwarded
coreapi *tenant-jobs*	elasticsearch *tenant-jobs*	forwarded
recruiter *tenant-jobs*	coreapi *tenant-jobs*	forwarded
recruiter *tenant-jobs*	coreapi *tenant-jobs*	forwarded

@lizrice

ISOVALENT

Run custom code in the kernel

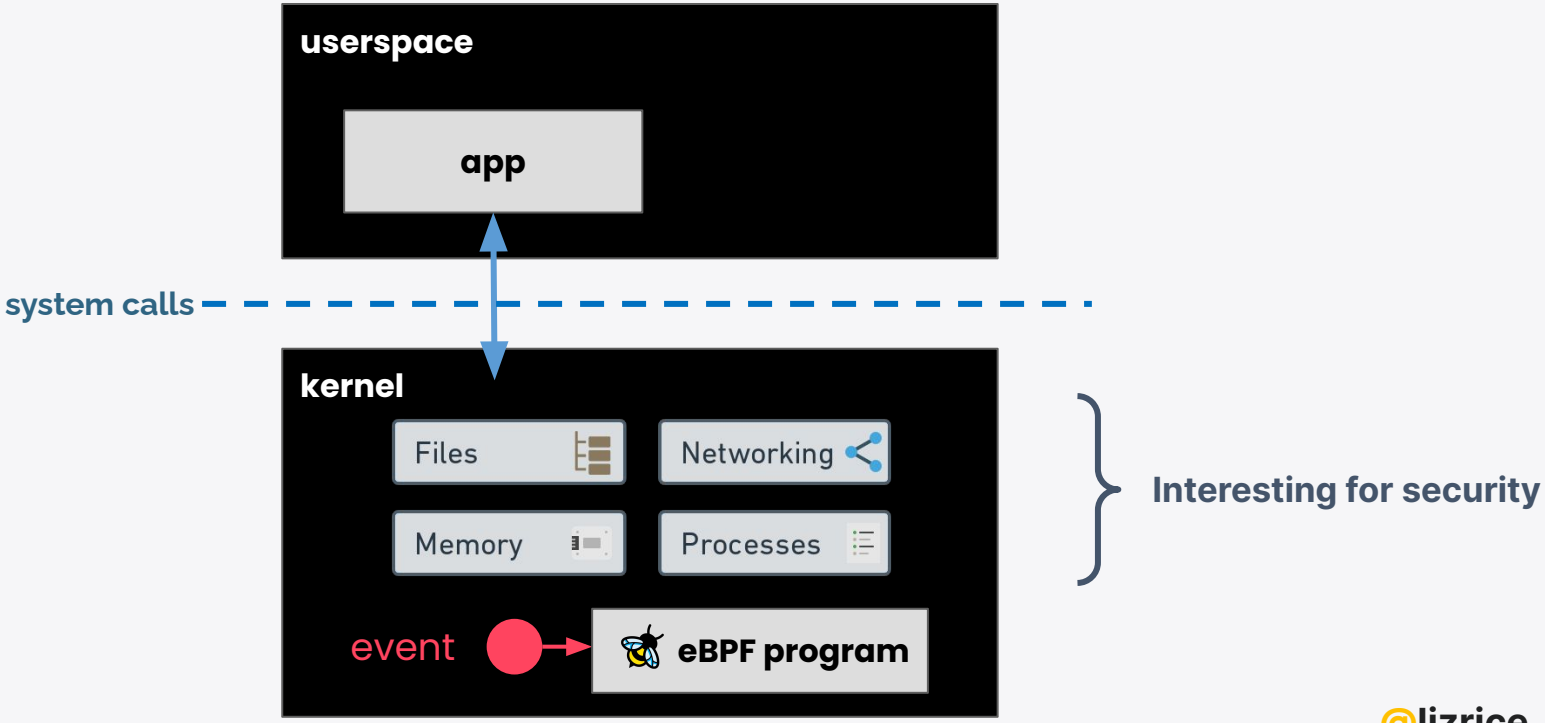

@lizrice

Runtime security - Falco

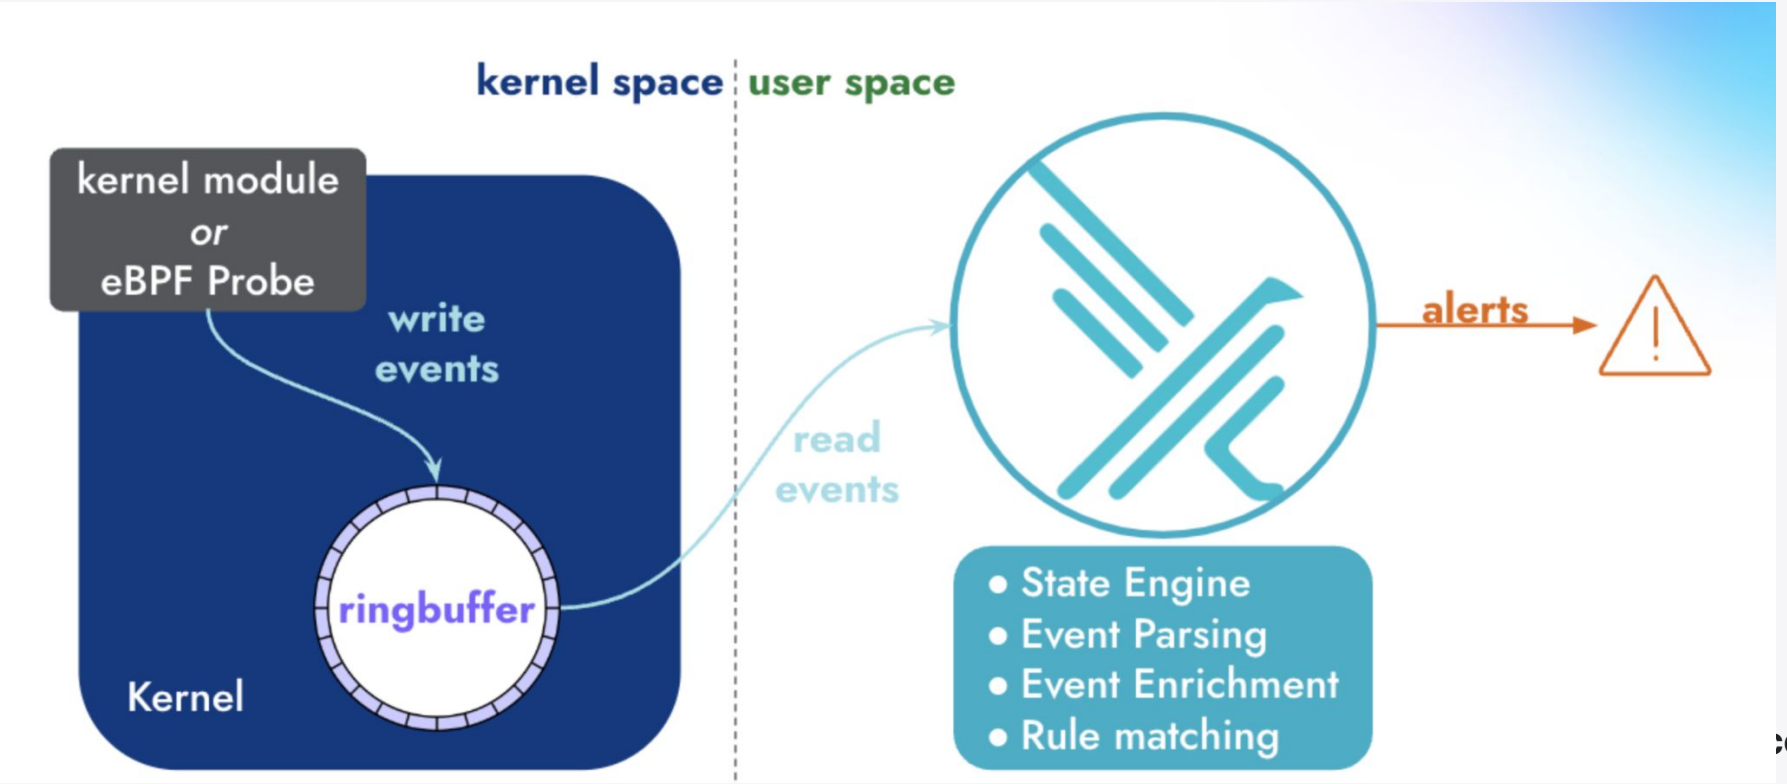

ISOVALENT

Runtime security - Cilium Tetragon

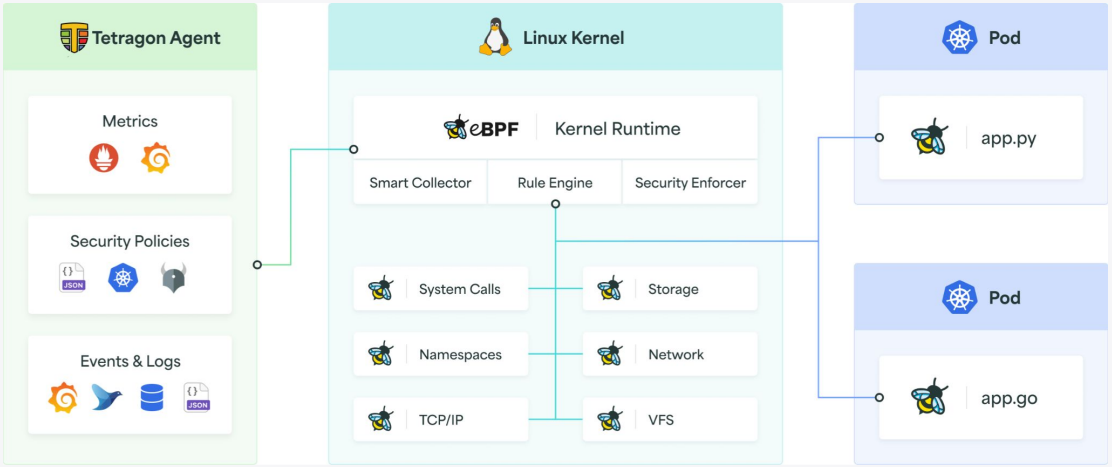

- Observe security events
  - System, network, filesystem, and applications
- Low overhead
  - In-kernel policy filtering
  - <2% overhead
- Synchronous enforcement

@lizrice

ISOVALENT

Cilium Tetragon runtime security

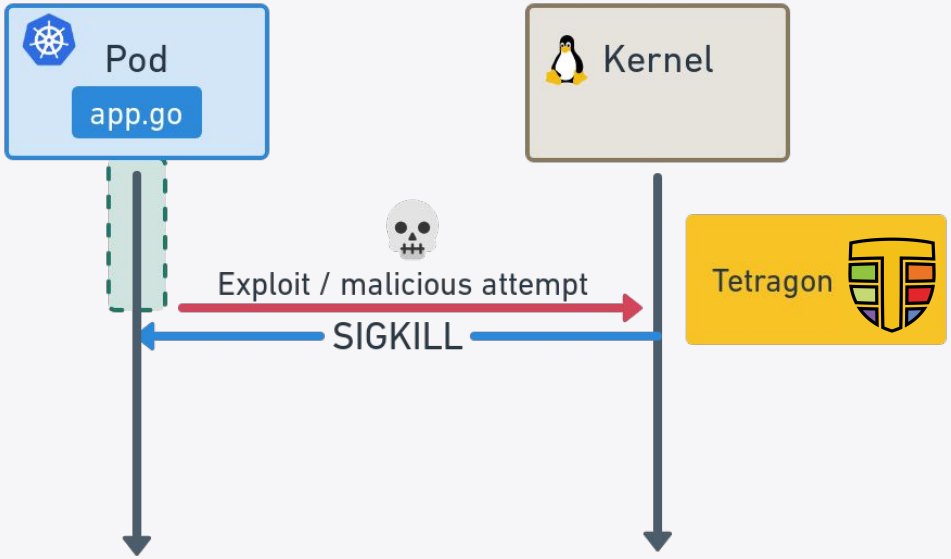

- Observe security events
  - System, network, filesystem, and applications
- Low overhead
  - In-kernel policy filtering
  - <2% overhead
- Synchronous enforcement

ISOVALENT

eBPF makes the Linux kernel programmable enabling a new generation of powerful Cloud Native tools

@lizrice

ISOVALENT

Thank you

- cilium/cilium
- cilium.io
- isovalent.com

@lizrice