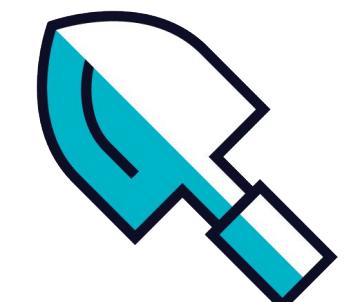


Sysdig

eBPF-powered distributed Kubernetes performance analysis

Lorenzo Fontana.

Open Source Software Engineer, Sysdig.



eBPF

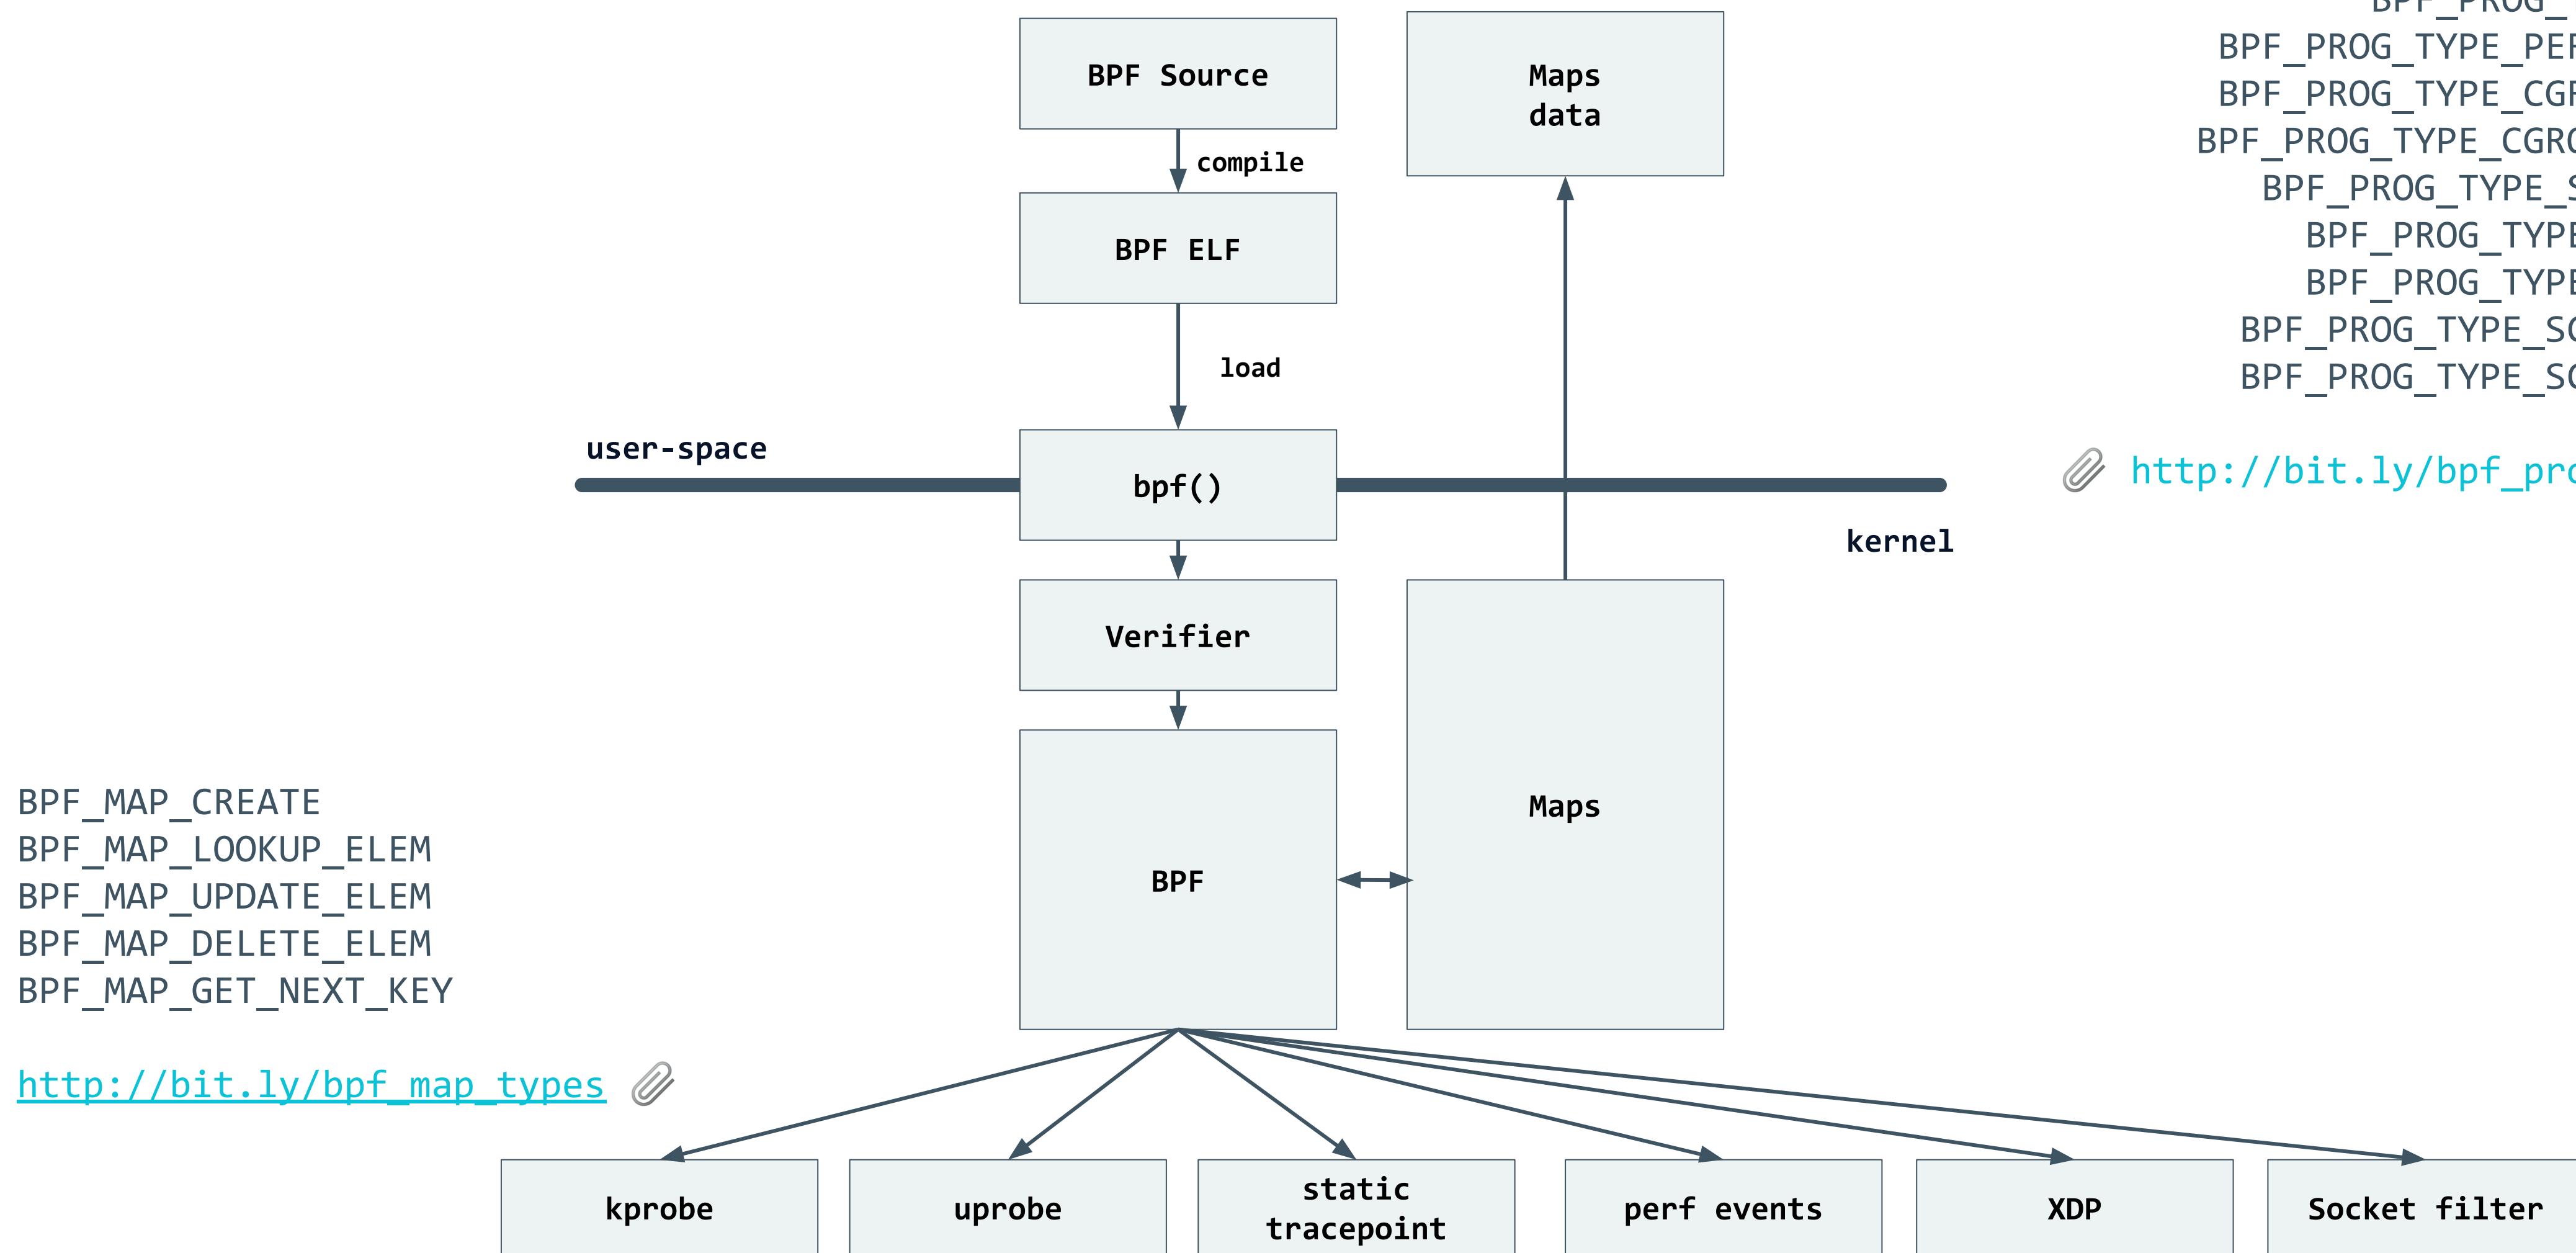
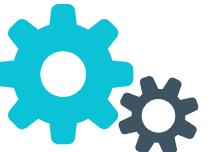
extended BPF

extended because it's not just packets anymore

Berkley Packet Filter



How does eBPF work?



Aggregate events at kernel level
and deal with just a few
instead of thousands of them



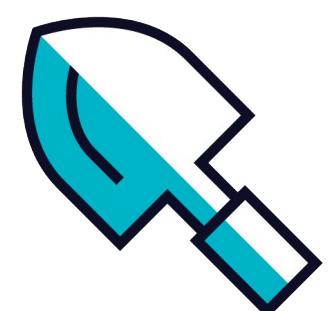
What is performance analysis about?

Performance analysis is a **quantitative** and **systematic** approach to **identify** performance issues in a software by doing:

- Measurement of time
- Measurement of space
- Measurement of complexity
- Profiling
- Code Instrumentation



What about Kubernetes ?



Just use a container

```
● ● ●  
apiVersion: v1  
kind: Pod  
metadata:  
  name: happy-ebpf  
spec:  
  shareProcessNamespace: true  
  containers:  
    - name: execsnoop  
      image: calavera/execsnoop # <-- the actual image containing the eBPF program  
      securityContext:  
        - privileged: true  
      volumeMounts:  
        - name: sys # mount the debug filesystem  
          mountPath: /sys  
          readOnly: true  
        - name: headers # mount the kernel headers required by bcc  
          mountPath: /usr/src  
          readOnly: true  
        - name: modules # mount the kernel modules required by bcc  
          mountPath: /lib/modules  
          readOnly: true  
    - name: container doing random work  
      image: yourcompany/yourapp # <-- your actual application  
  ...
```

- 👉 A sidecar container sharing the process namespace
- 👉 You just provide an image with an eBPF loader and program in it
- 👉 Not extremely generic but does the job!
- 👉 A very flexible approach!





@fntlnz

Want
something
more generic?

- 👏 Here's an experiment I've been working with
@leodido

👏 It loads eBPF ELF objects using a CRD

👏 Same as the container example but you don't have to write the loader

👏 Exposes a Prometheus endpoint

It's called kube-bpf
<https://github.com/bpftools/kube-bpf>



<https://github.com/bpftools/kube-bpf/blob/master/examples/pkts.c>



```

struct bpf_map_def SEC("maps/packets") countmap = {
    .type = BPF_MAP_TYPE_HASH,
    .key_size = sizeof(int),
    .value_size = sizeof(int),
    .max_entries = 256,
};

SEC("socket/prog")
int socket_prog(struct __sk_buff *skb) {
    int proto = load_byte(skb, ETH_HLEN + offsetof(struct iphdr, protocol));
    int one = 1;
    int *el = bpf_map_lookup_elem(&countmap, &proto);
    if (el) {
        (*el)++;
    } else {
        el = &one;
    }
    bpf_map_update_elem(&countmap, &proto, el, BPF_ANY);
    return 0;
}

char _license[] SEC("license") = "GPL";

unsigned int _version SEC("version") = 0xFFFFFFFF; // this tells to the ELF loader to set the current running
kernel version

```

pkts.c

- 👏 Counts all the packets
- 👏 Uses a map to keep a counter
- 👏 It's an HASH map so that it can assign the counter to a packet type

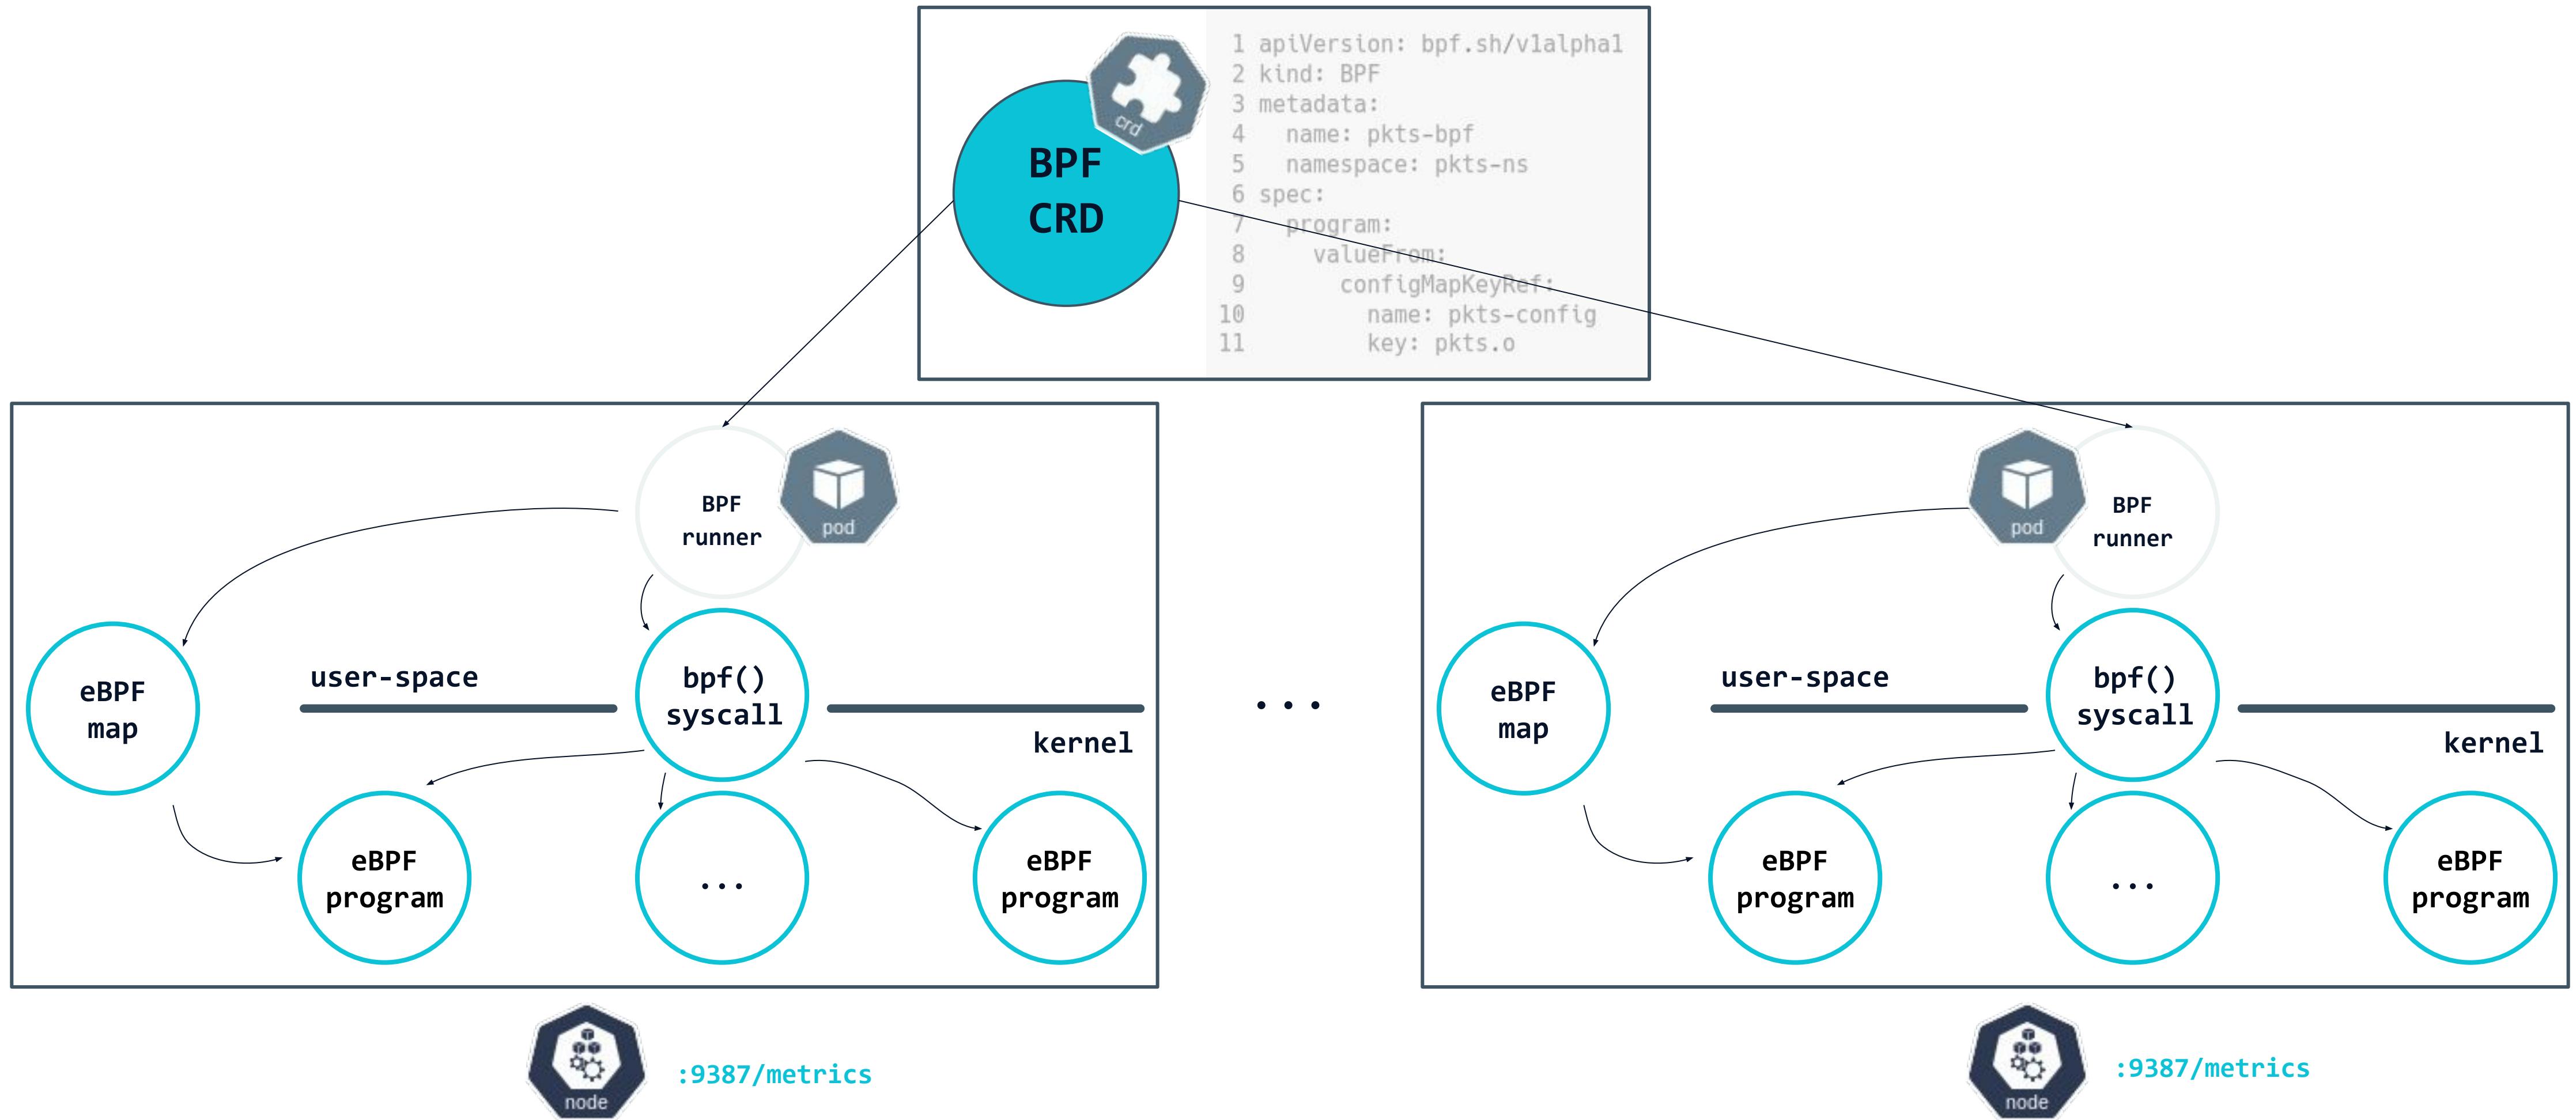


ip-10-12-0-136.ec2.internal:9387/metrics

```
# HELP test_packets No. of packets per protocol (key), node
# TYPE test_packets counter
test_packets{key="00001",node="127.0.0.1"} 8                      # <- ICMP
test_packets{key="00002",node="127.0.0.1"} 1                      # <- IGMP
test_packets{key="00006",node="127.0.0.1"} 551                   # <- TCP
test_packets{key="00008",node="127.0.0.1"} 1                      # <- EGP
test_packets{key="00017",node="127.0.0.1"} 15930                 # <- UDP
test_packets{key="00089",node="127.0.0.1"} 9                      # <- OSPF
test_packets{key="00233",node="127.0.0.1"} 1                      # <- ?
# EOF
```



Here's the evil plan



Get the code!

github.com/bpftools/kube-bpf





demo

<https://github.com/fntlnz/oscon-19-demo/tree/master/kubectl-trace>

eBPF tracing in the kubectl!

<https://github.com/iovisor/kubectl-trace>

kubectl-trace

Run bpftrace program (from file)

```
1 kubectl trace run 127.0.0.1 -f read.bt -a  
2 trace 9df7388a-f0b4-11e8-ae05-8c164500a77e create
```

Ctrl-C tells the program to plot the results using hist()

3 °C ←
4
5 @start[12509]: 4991487155626
6 @start[12856]: 4991483355976
7 @start[12865]: 4991484775952
8 @start[12866]: 4991484856394
9 @start[12867]: 4991487276493

```
10  
11  
12 @times:  
13 [512, 1K)           85 | @@@@  
14 [1K, 2K)           767 | @@@@@@  
15 [2K, 4K)           700 | @@@@  
16 [4K, 8K)           920 | @@@@  
17 [8K, 16K)          751 | @@@@
```

18 [16K, 32K) 393 |@@@
19 [32K, 64K) 90 |@@@

20 [64K, 128K) 14 |
21 [128K, 256K) 3 |

21 [128K, 256K]	3	
22 [256K, 512K]	4	

23 [512K, 1M) 2 |
24 [1M, 2M) 2 |

24 [1M, 2M) 2 |
25 [2M, 4M) 2 |

Maps

The output histogram



A cartoon illustration of a cowboy wearing a hat and vest, riding a brown horse through a field of tall, yellow flames. The cowboy is looking back over his shoulder. In the background, there are green hills and a cloudy sky.

demo

<https://github.com/fntlnz/oscon-19-demo/tree/master/kube-bpf>

Wait wait wait wait!

There's a book!

O'REILLY®

Linux Observability with BPF

Advanced Programming for Performance
Analysis and Networking



David Calavera &
Lorenzo Fontana
Foreword by Jessie Frazelle



From me and David Calavera



Almost published



Preorder on Amazon.com, DO IT!



Early Release on O'Reilly Safari



Foreword by Jessie Frazelle



All the acronyms

Computer people loves acronyms

BPF: Berkley Packet Filter

eBPF: Extended Berkley Packet Filter

CRD: Custom Resource Definition (Kubernetes)



Thanks.



Reach me out [@fntlnz](#) on [twitter](#) & [github](#)!

