



# Beyond the Buzzword: BPF's unexpected role in Kubernetes

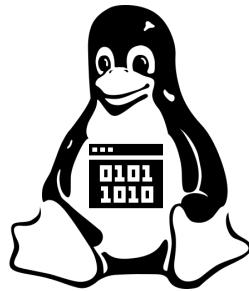
Andrew Randall  
Alban Crequy

November 19, 2020

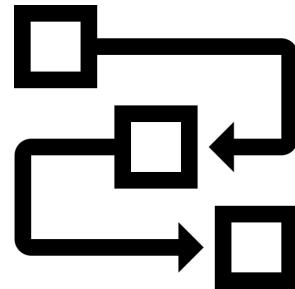


# What is (e)BPF?

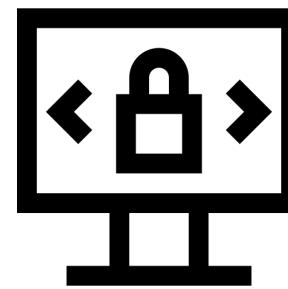
(extended) Berkeley Packet Filter



custom programs  
that run in the  
Linux kernel



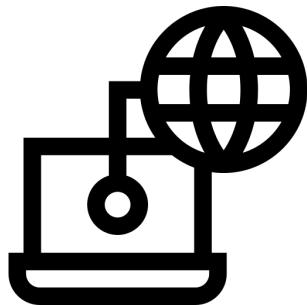
hooks and data  
structures (maps)



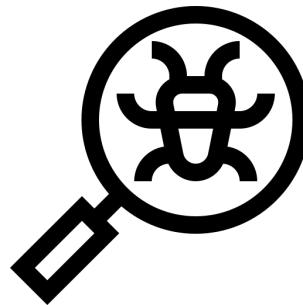
restricted virtual  
machine sandbox +  
code verifier



# Why do you care?



fast,  
customizable  
networking



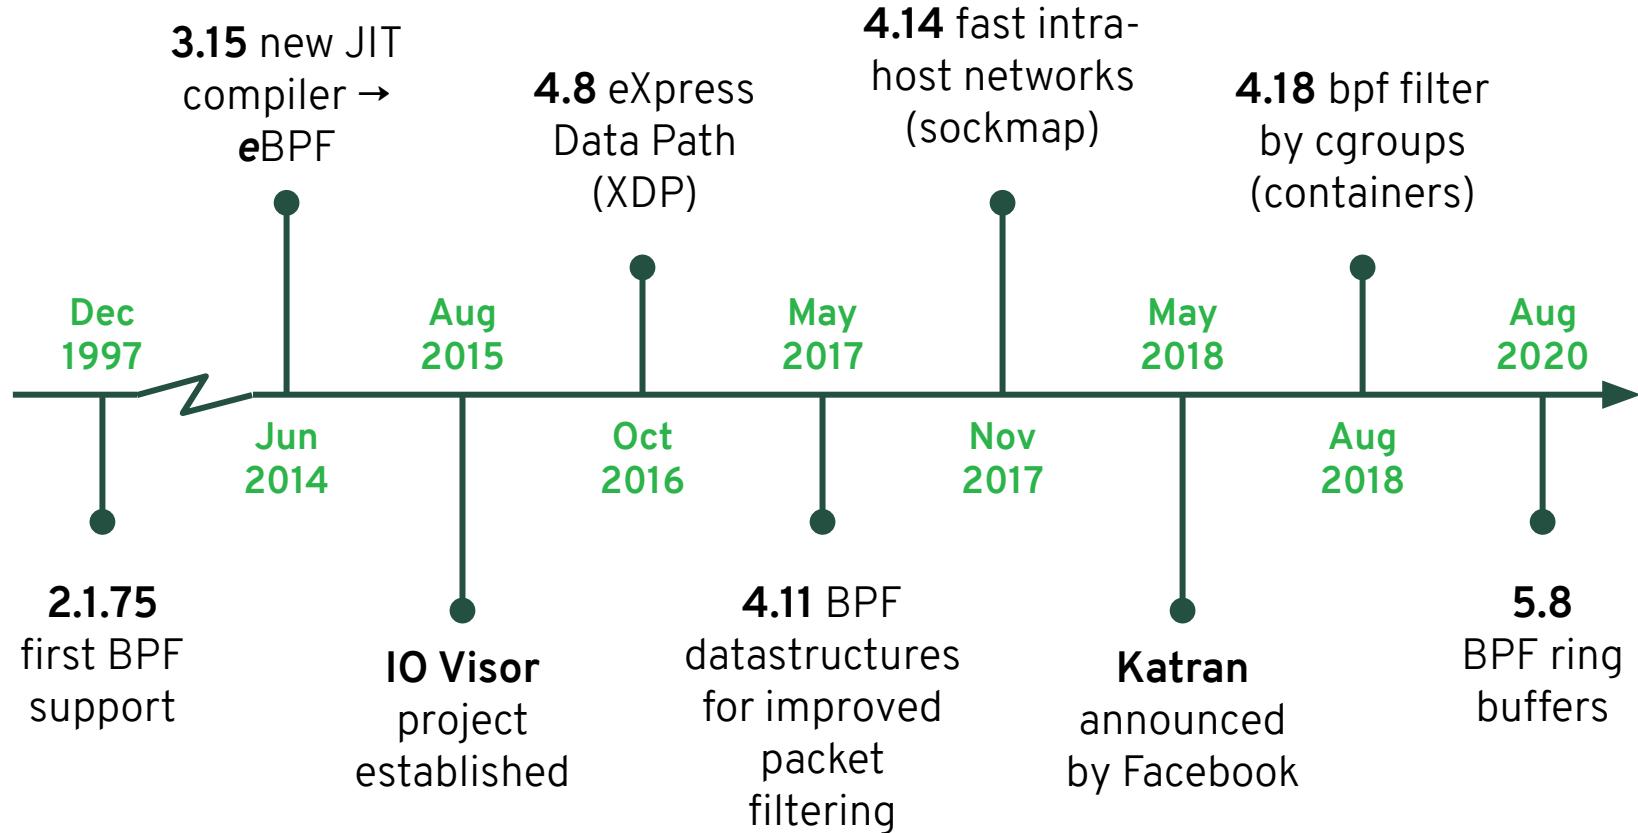
debugging /  
performance  
analysis



application  
monitoring &  
security



# Evolution of (e)BPF



# An eBPF OSS Landscape

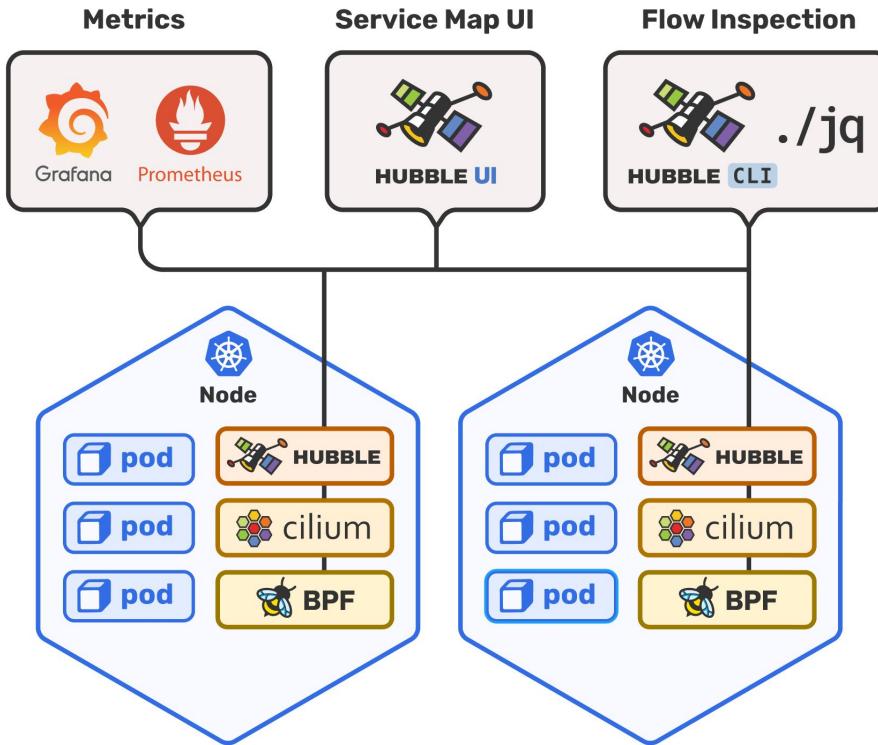
Other

kinvolk / netcost  
nirmata / kube-netc  
cloudflare / ebfpxpporter

Visibility						
Security & Networking						
API Libraries						
Low-level tools						
	bcc	bpftace	llvm	kernel tools e.g. bpftool	tcptracer-bpf	

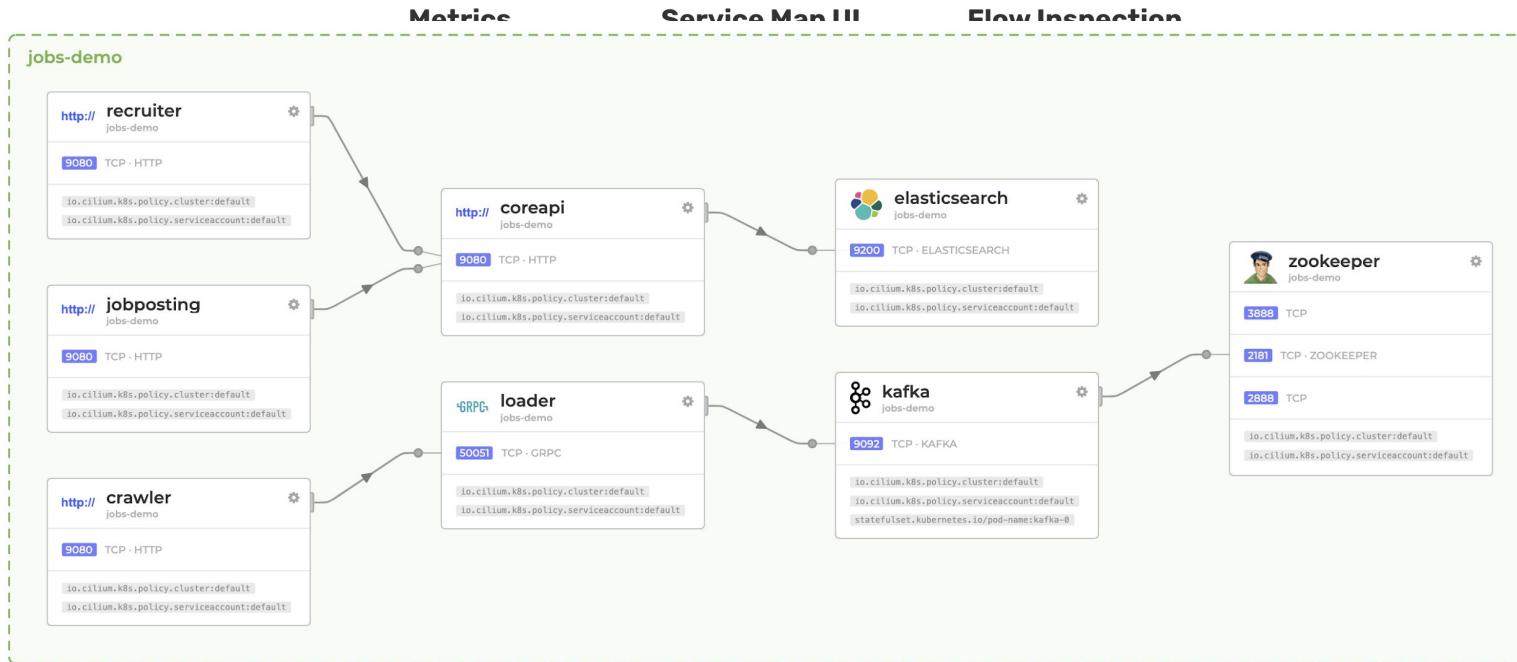


# Hubble





# Hubble





# Weave Scope

# weave scope

PROCESSES  
BY NAME

11 NODES

CPU Memory ↗

System containers Application containers Both

Stopped containers Running containers Both

## tns\_app2\_1

peterbourgon/tns-app:lat... dlite

STATUS

0.00 % CPU    6.1 MB MEMORY

INFO

ID: 0562f3bcfa8a  
STATE: Up About an hour  
COMMAND: /app db1 db2 db3

+6 ▾

INBOUND	PORT	COUNT
tns_lb2_1	8080	1
tns_lb1_1	8080	1

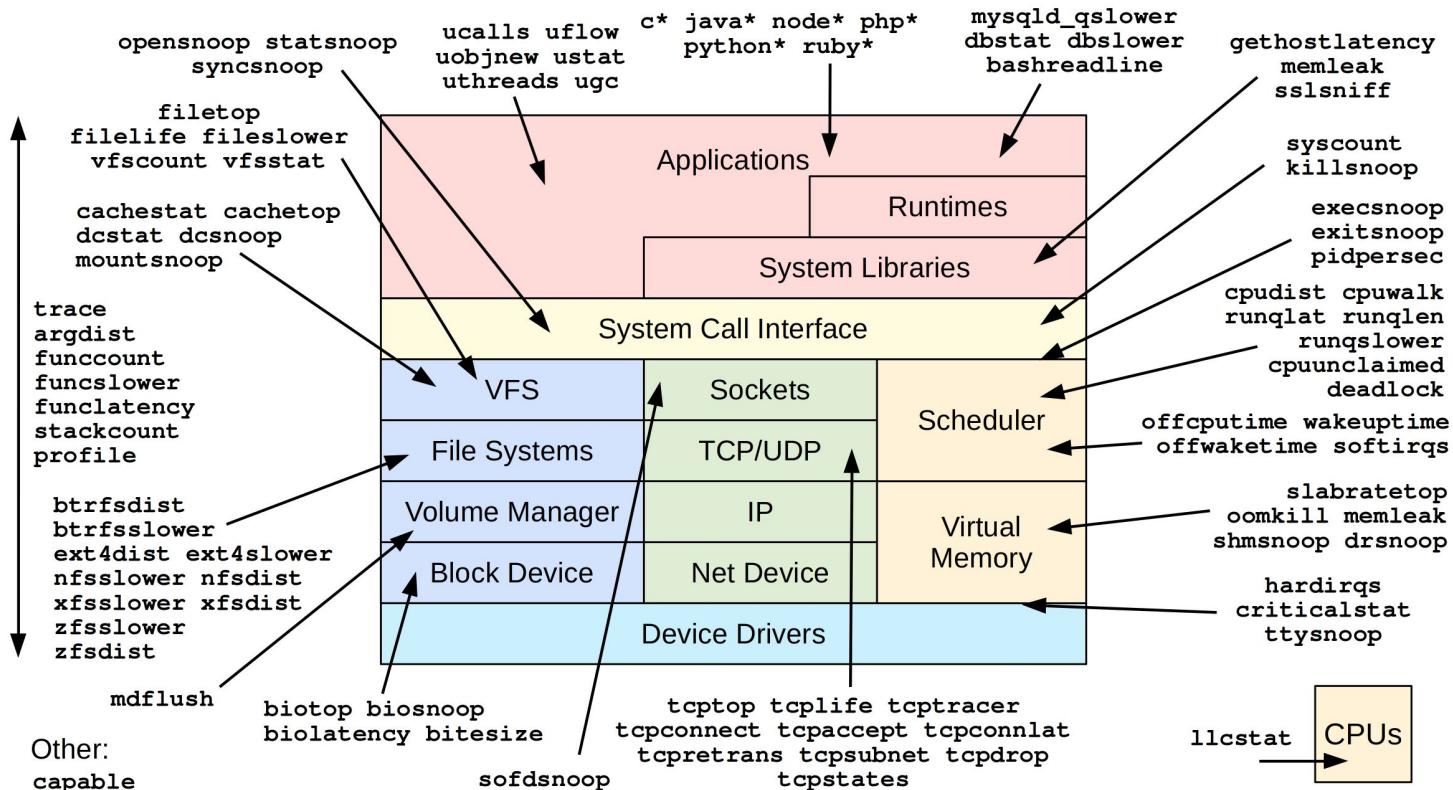
OUTBOUND	PORT	COUNT
tns_db3_1	9000	1

VERSION b9d93fb ON dlite PLUGINS: iowait

☰

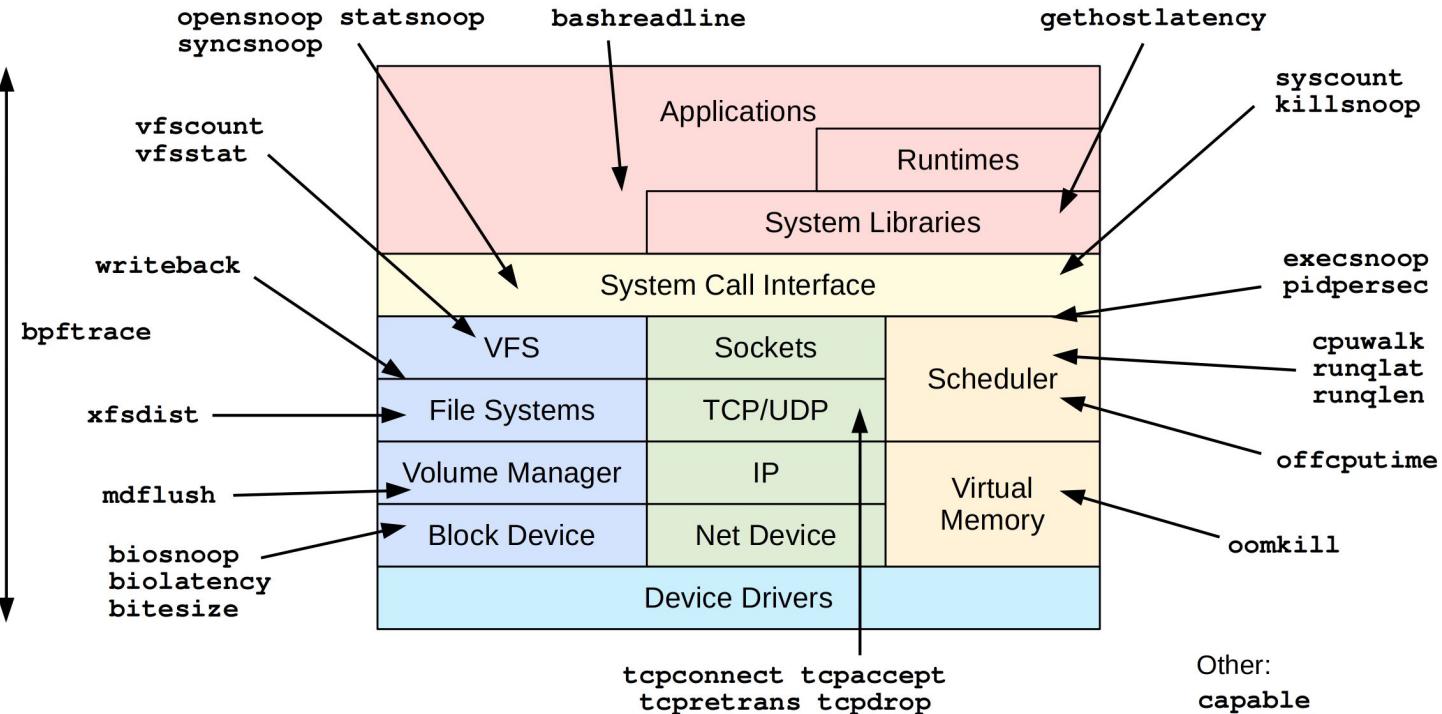


# IOvisor BPF Compiler Collection (bcc)





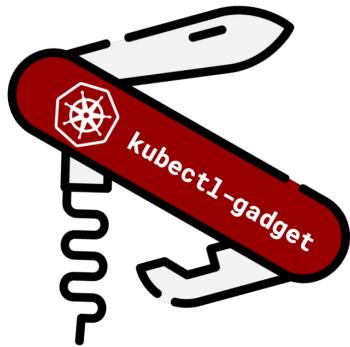
# bpftrace



# SO MANY POWERFUL BPF TOOLS



# Enter: Inspektor Gadget



a “swiss army knife”  
collection of various  
bpf tools (gadgets)



some from bcc

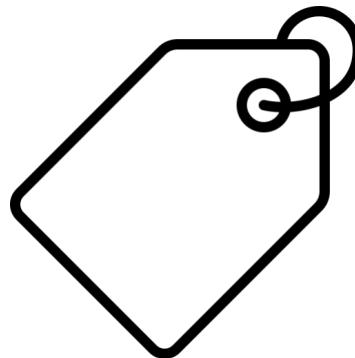


+ some new ones  
developed by  
kinvolk

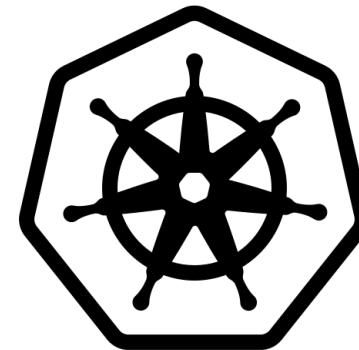
# What do we need for Kubernetes?



granularity:  
“pod, not pid”



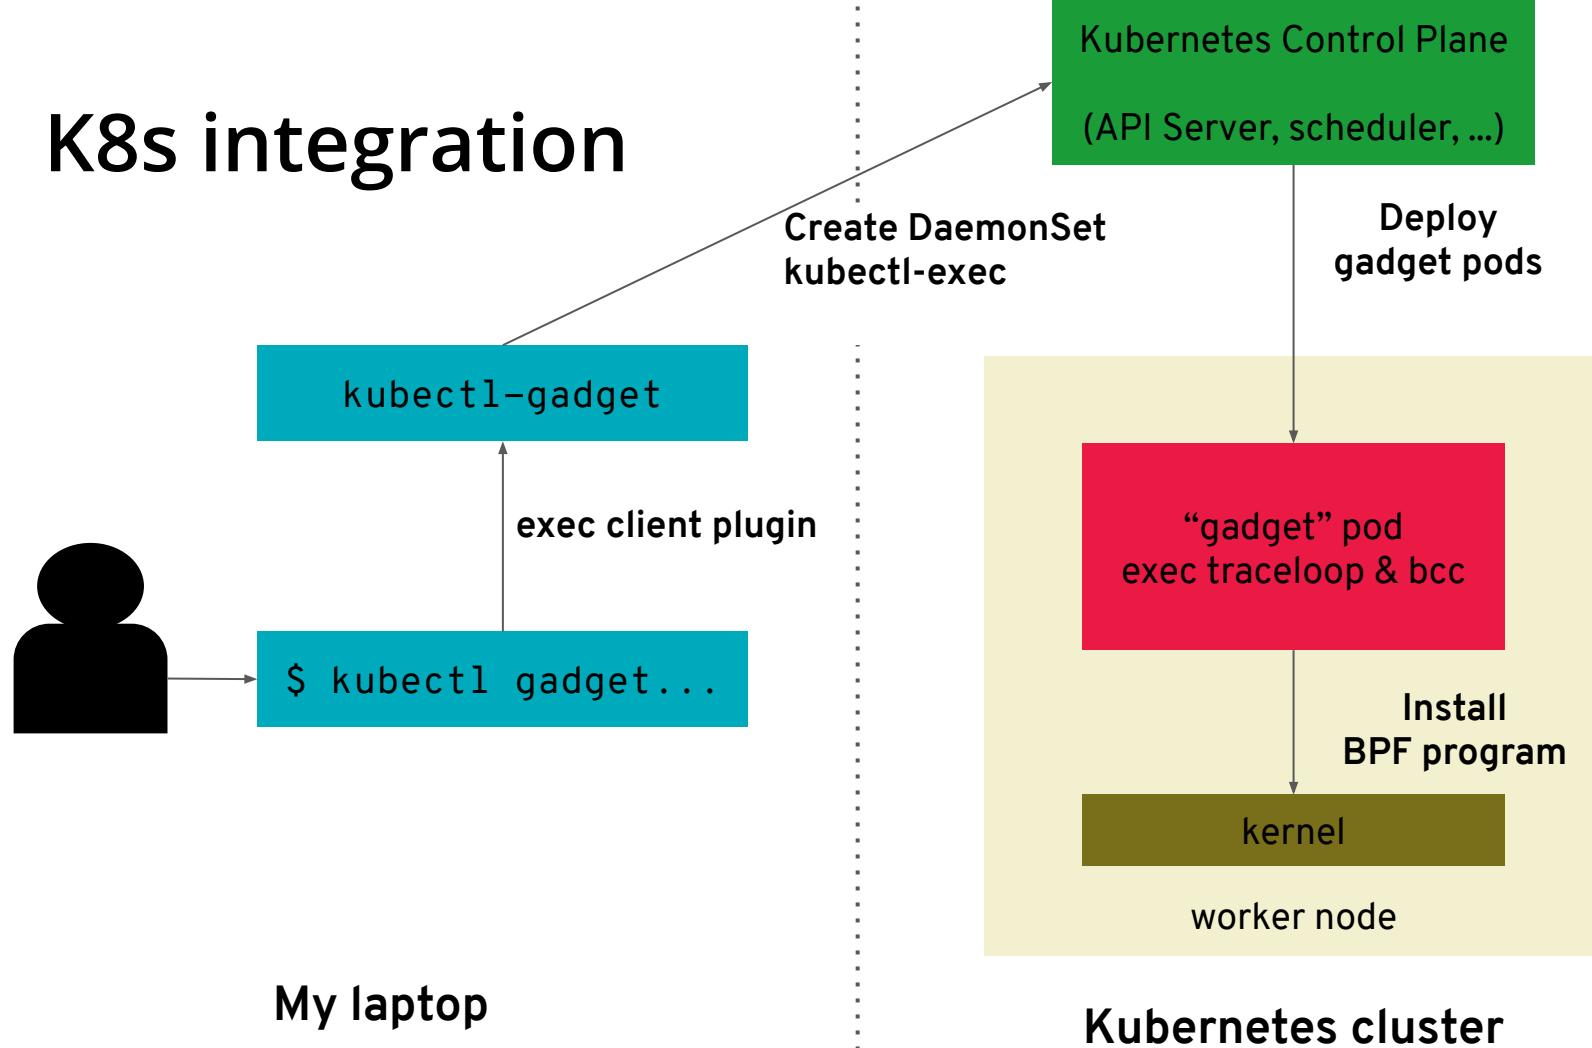
aggregation by label  
selectors



kubectl-like  
experience



# K8s integration



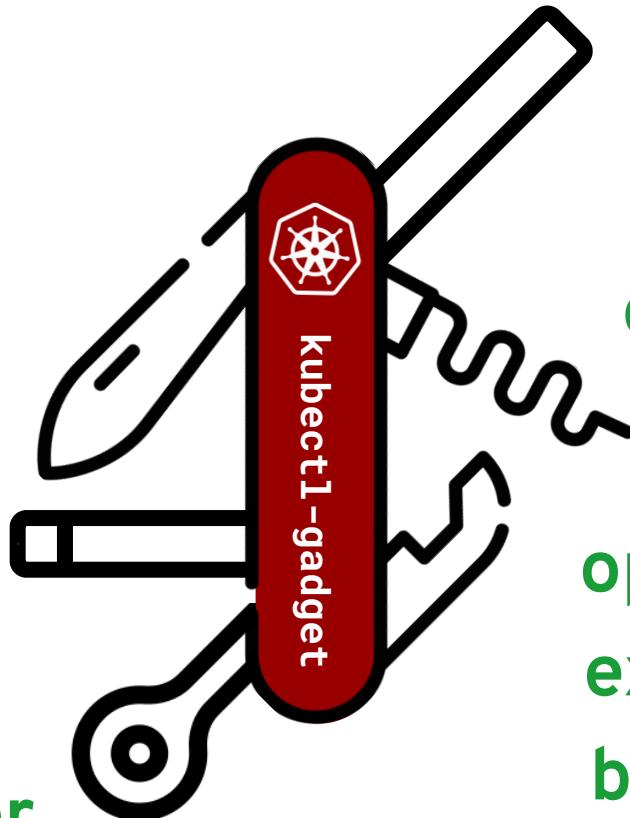
# Gadgets available today

traceloop

profile

tcptop

tcptracer



network policy  
advisor

capabilities

opensnoop

execsnoop

bindsnoop



