open source **Intrusion Detection** for containers

Shane Lawrence
Shopify

Kris Nóva
Sysdig

Welcome to Our Story

Shane Lawrence
Kris Nóva

shopify

$61.1
billion
volume

300
million
shoppers

$1.5
million
sales per minute
(peak)

shopify

50+

clusters

10+

thousand

services

170+

thousand

requests per second (peak)

Commerce at scale

over 1,000,000 shops.

some may look familiar.

we enable their brand.

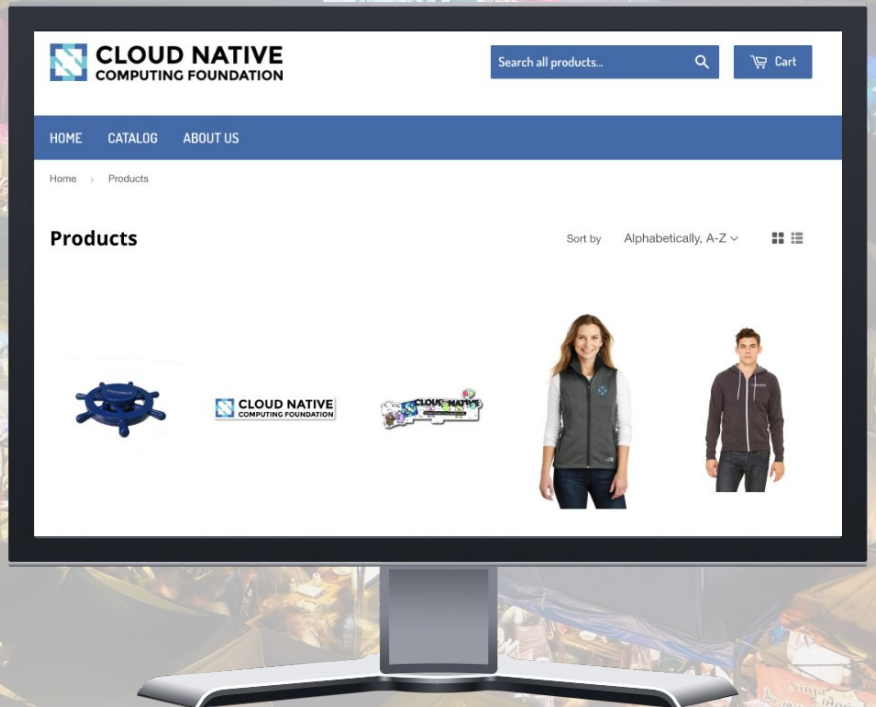

shopify

Security is our business

defence in depth

external audits

security bug bounty:
hackerone.com/Shopify

shopify

Journey to the cloud

Kubernetes
scales workloads to meet demand

Cloud Provider
provides resources on-demand

Data Centers?
not for us

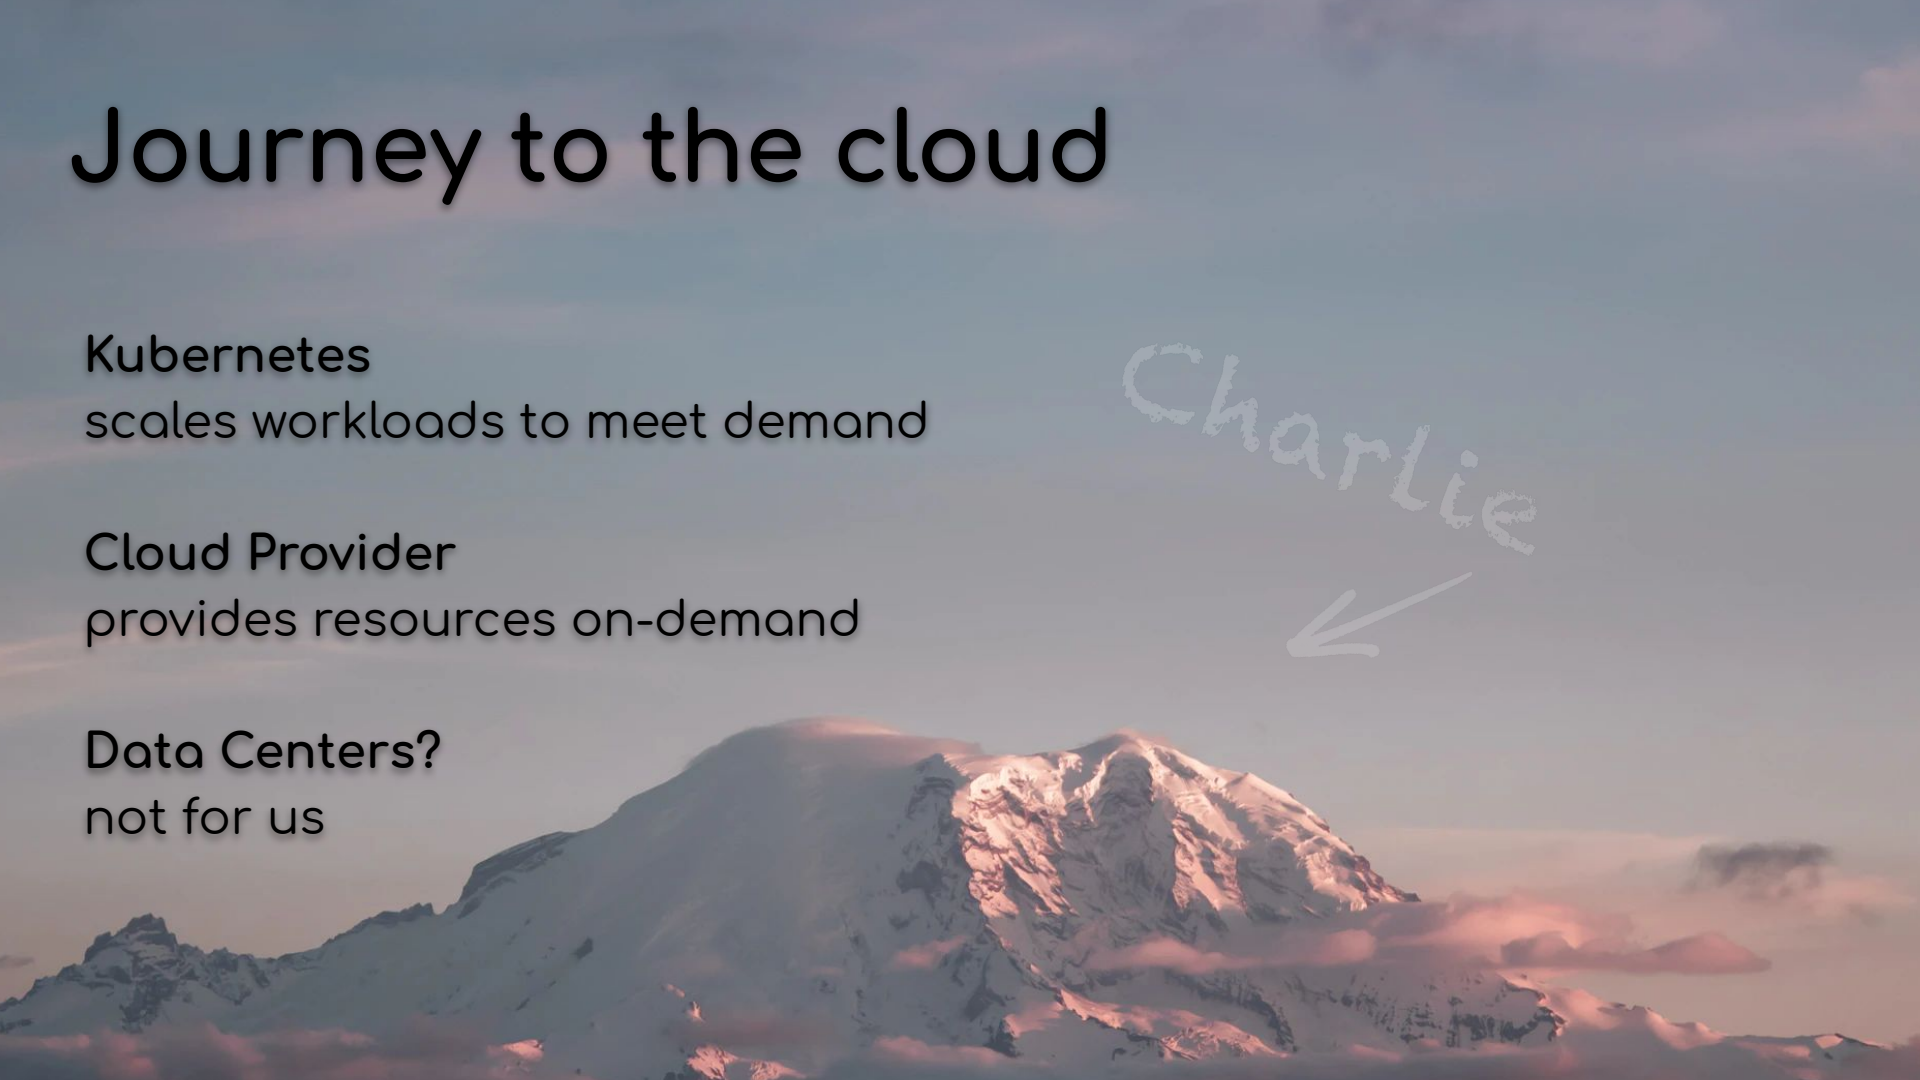

Network Security Boundaries

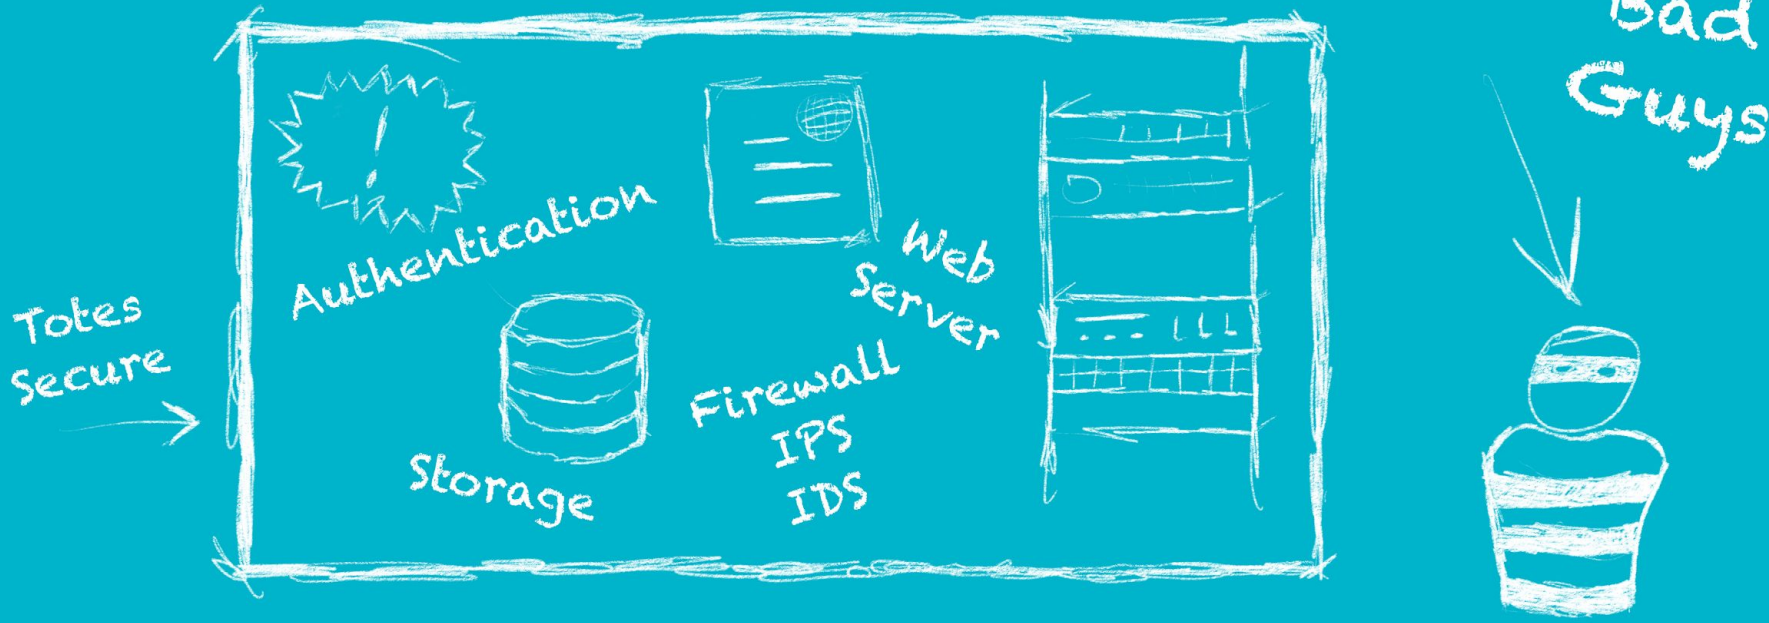

Home Net: 10.0.0.0/8

@krishnova

Network Security in the Cloud

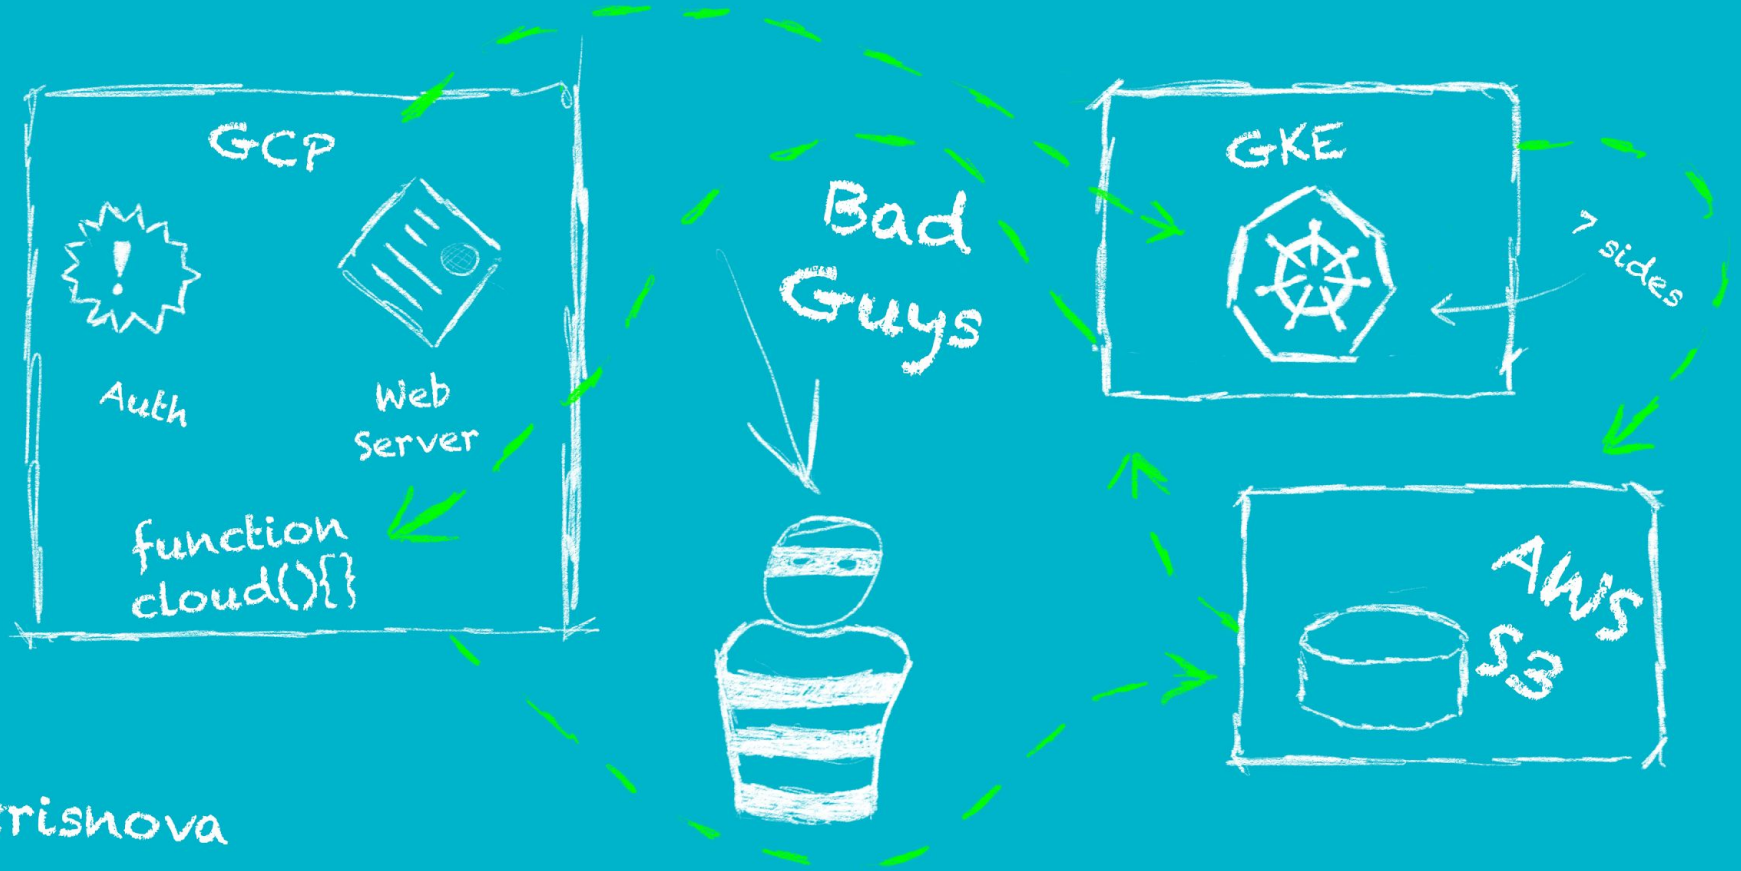

@krishnova

shopify

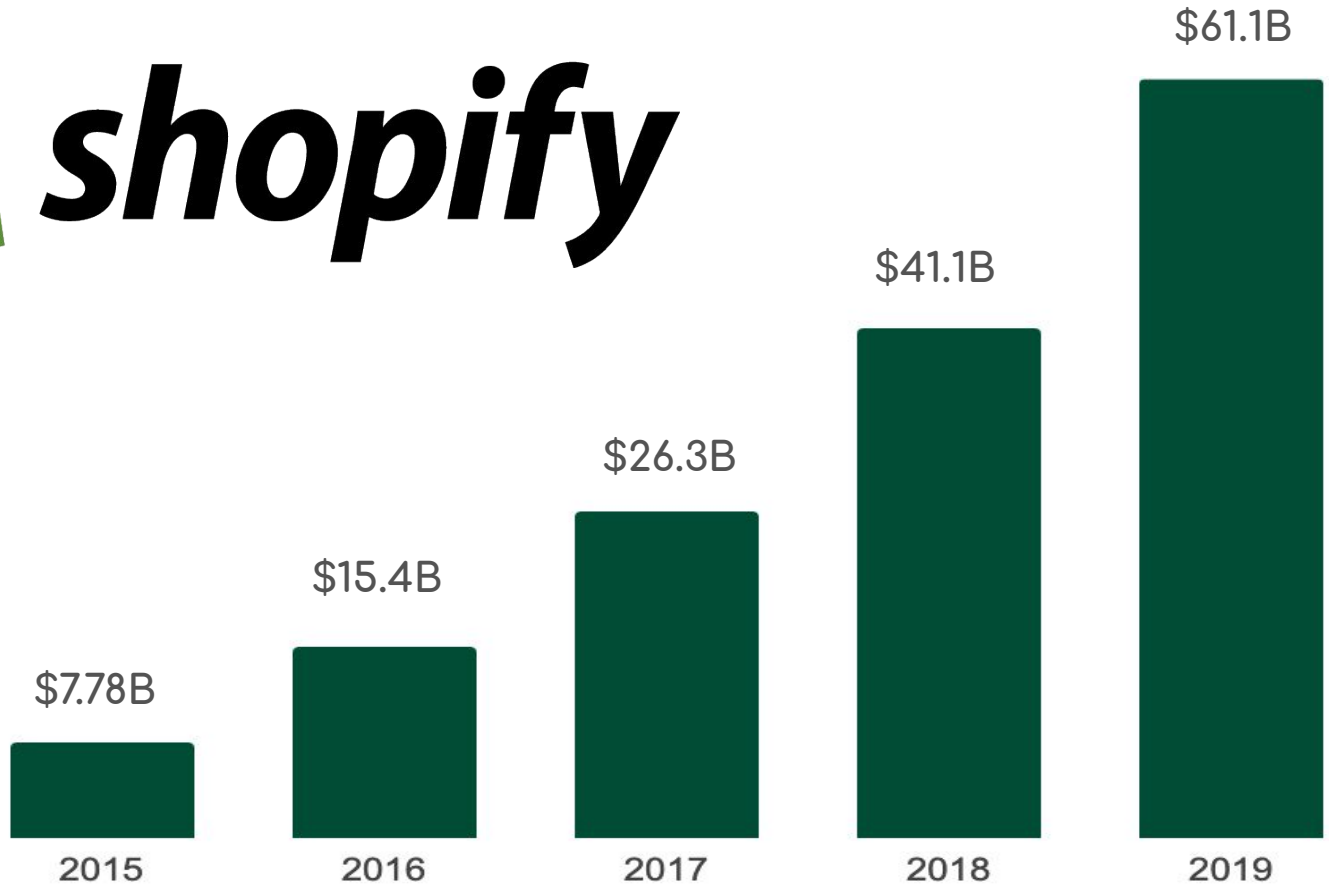

Gross Merchandise Volume

Greenfield Kubernetes

Kubernetes on GKE in early 2017:

- ABAC only
- dashboards enabled with admin access
- open computeMetadata
- no cloud audit logs

PCI Compliance

auditors vs. early adopters

don't take our word for it.

Security Approach: Prevention

Kubernetes

- disable old APIs, unused features
- metadata proxy
- kubelet bootstrap
- Role Based Access Control
- seccomp & apparmor profiles
- network policies

In-house

- github.com/Shopify/kubeaudit
- security-auditors

Security Approach: Detection

Misconfiguration

- insecure API endpoints
- overprivileged roles
- weak security context
- assumptions about safety/identity

Software supply chain

- deliberate backdoors
- bugs in dependencies
- typo, tag hijacking

Unmitigated Vulnerabilities

- **Heartbleed**
  CVE-2014-0160
- **Spectre v1**
  CVE-2017-5753
- **Spectre v2**
  CVE-2017-5715
- **Meltdown**
  CVE-2017-5754
- **CVE-2019-5736**
  runc Container Escape
- **CVE-2018-15664**
  Symlink directory traversal
- **CVE-2018-1002105**
  Unauthenticated Remote Privilege Escalation
- **CVE-2020-8558**
  Localhost Boundary Bypass

Falco
Cloud-Native Runtime Security

Kris Nóva

What is Falco?

Parses system calls at runtime

Rebuilds system state

Enriches with Kubernetes data

Event triggered during anomaly

Enabling Prevention

Falco

Unexpected change in system

Runtime based detection

Utility for uncovering future security policy

CVE response

Exploit response

Surveillance

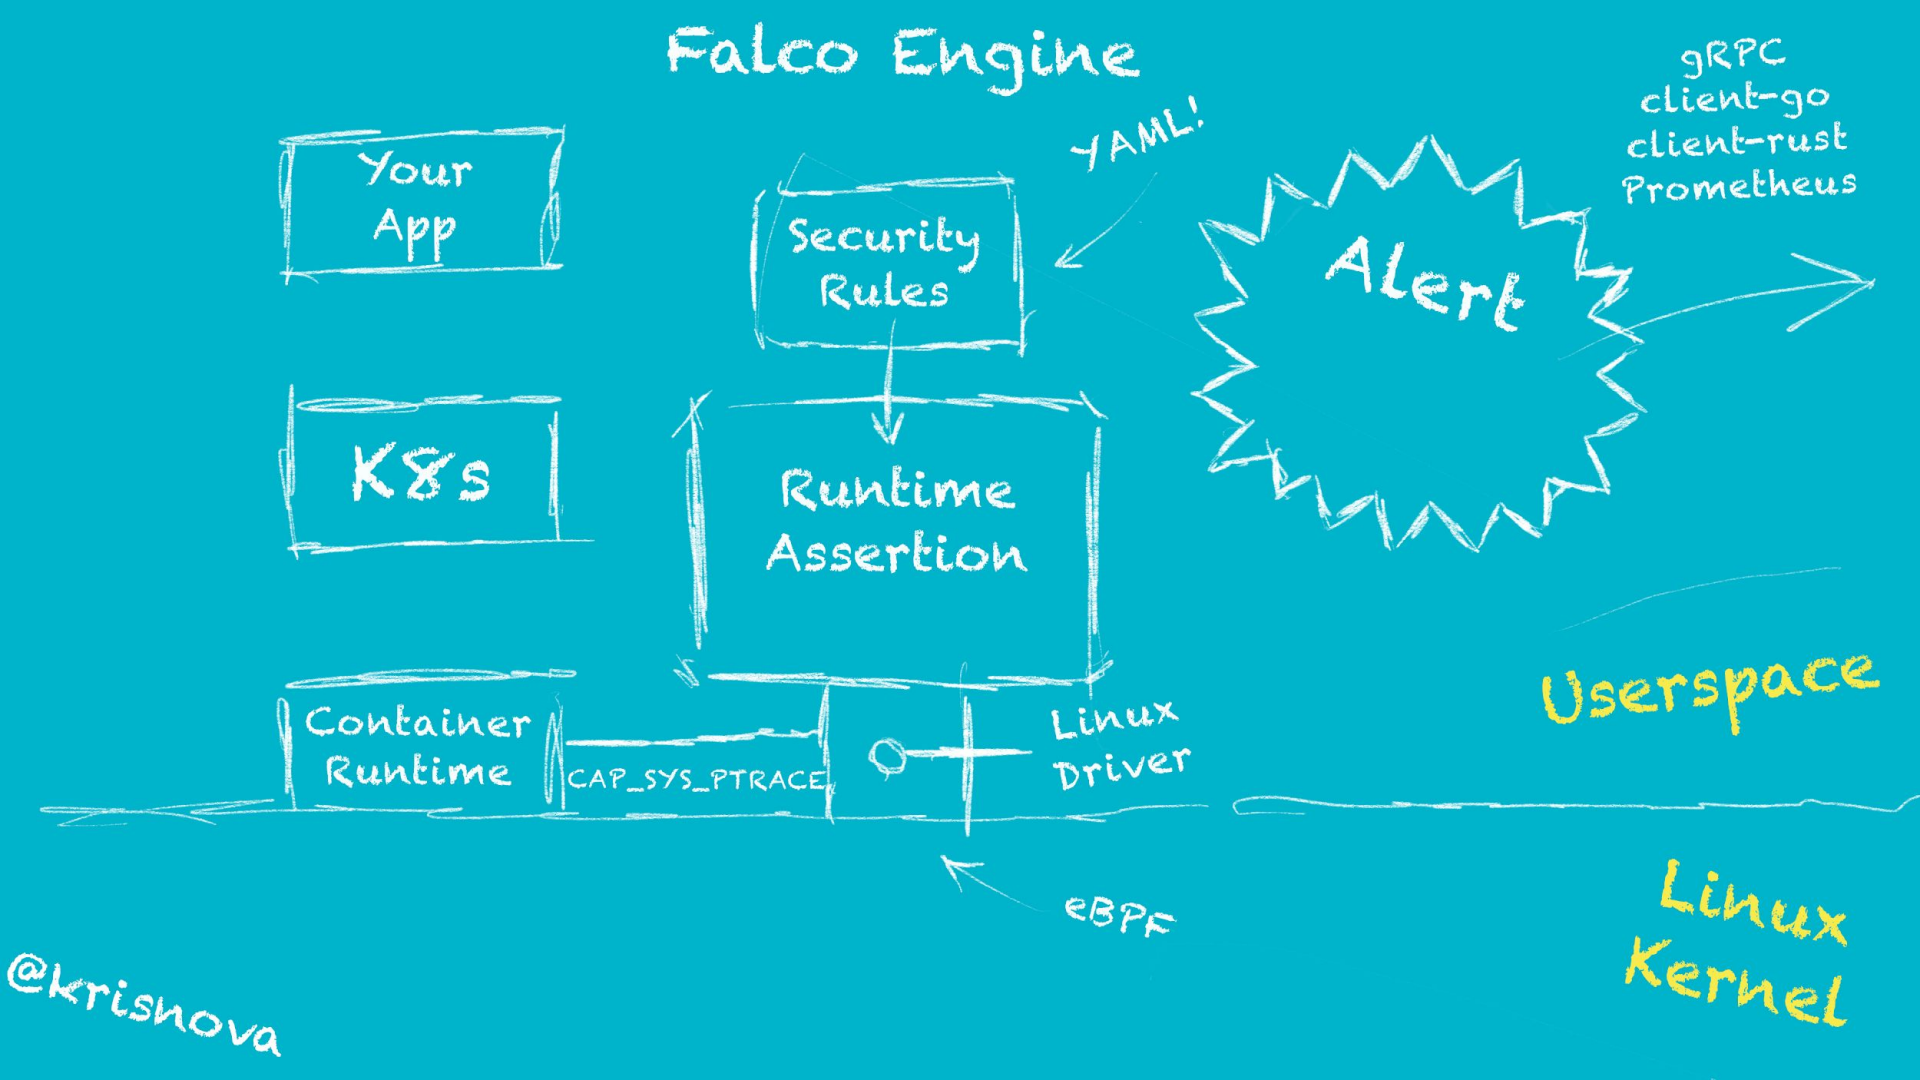
Falco Engine
Your App
K8s
Container Runtime
CAP_SYS_PTRACE
Security Rules
YAML!
Runtime Assertion
Linux Driver
eBPF
Alert
gRPC
client-go
client-rust
Prometheus
Userspace
Linux Kernel
@krisnova

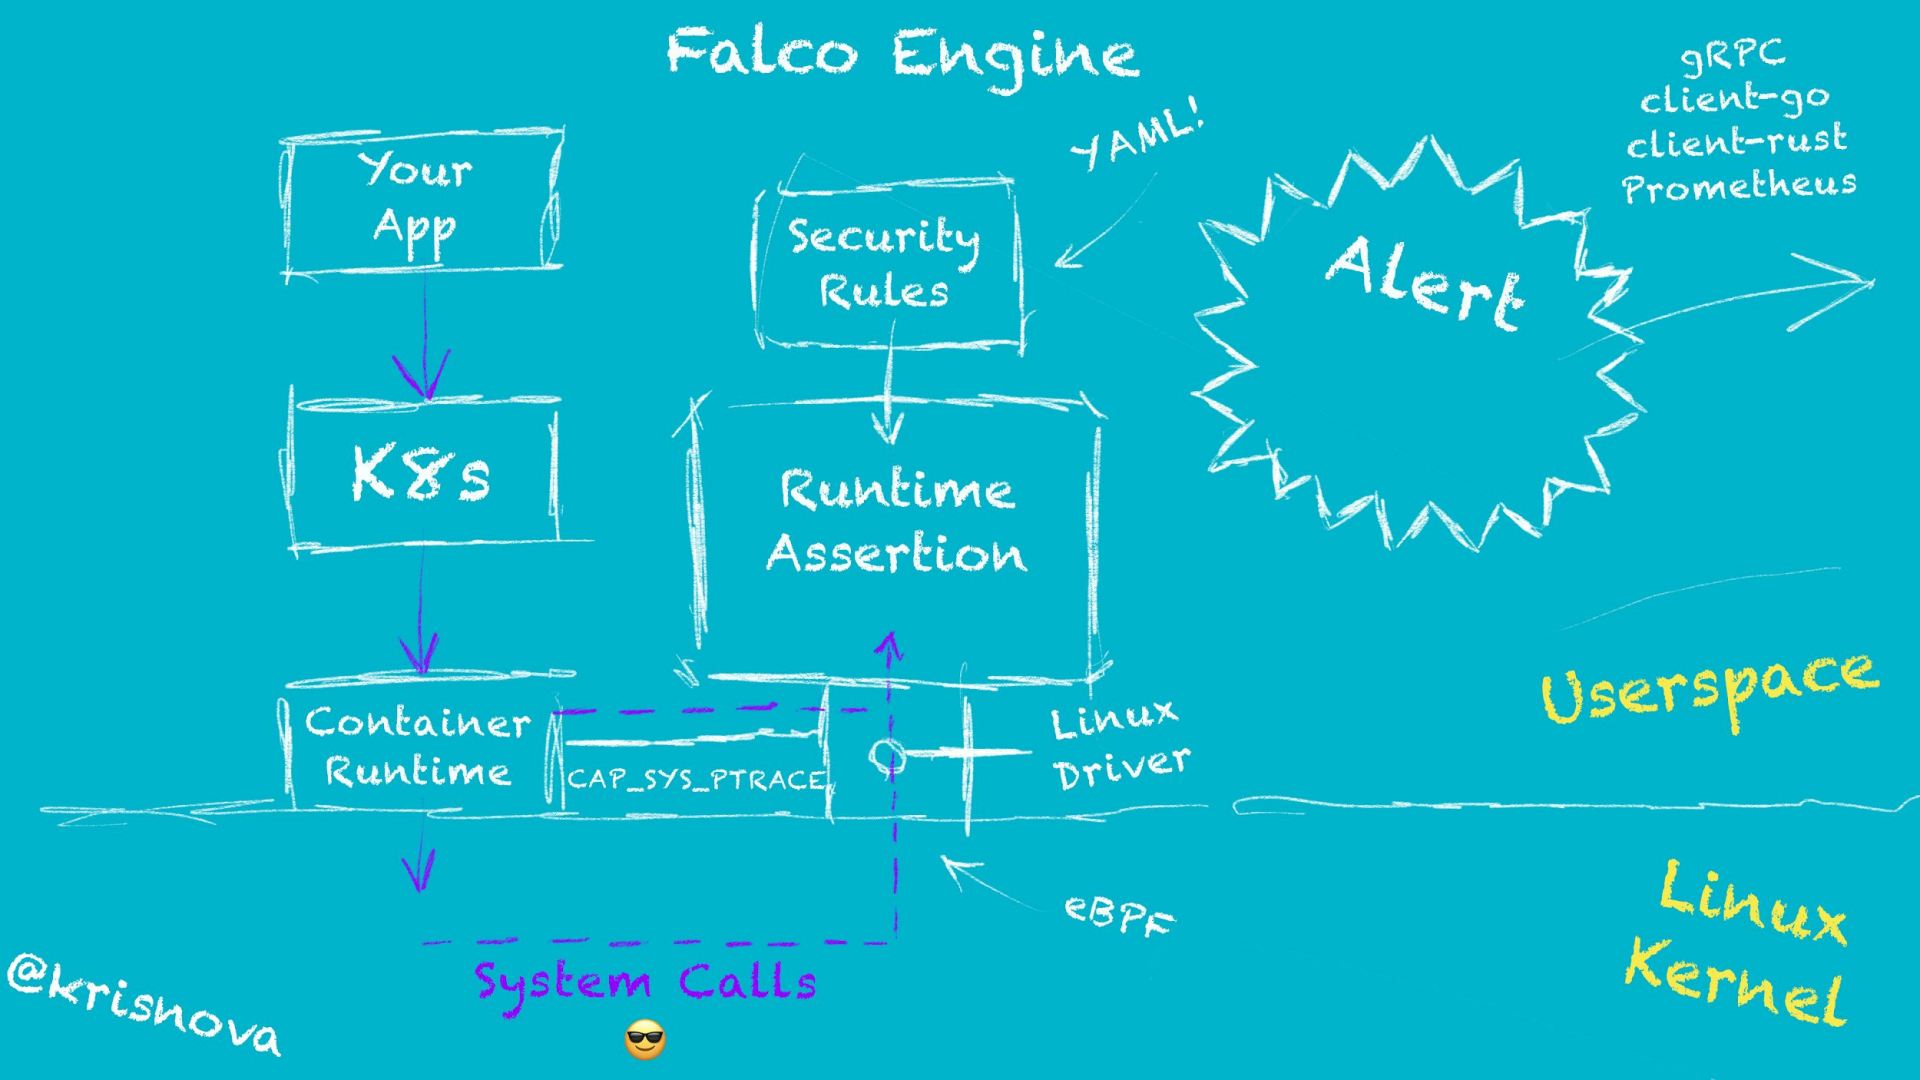
Falco Engine
Your App
K8s
Container Runtime
Security Rules
YAML!
Runtime Assertion
CAP_SYS_PTRACE
Linux Driver
eBPF
System Calls
Alert
gRPC
client-go
client-rust
Prometheus
Userspace
Linux Kernel
@krishnova

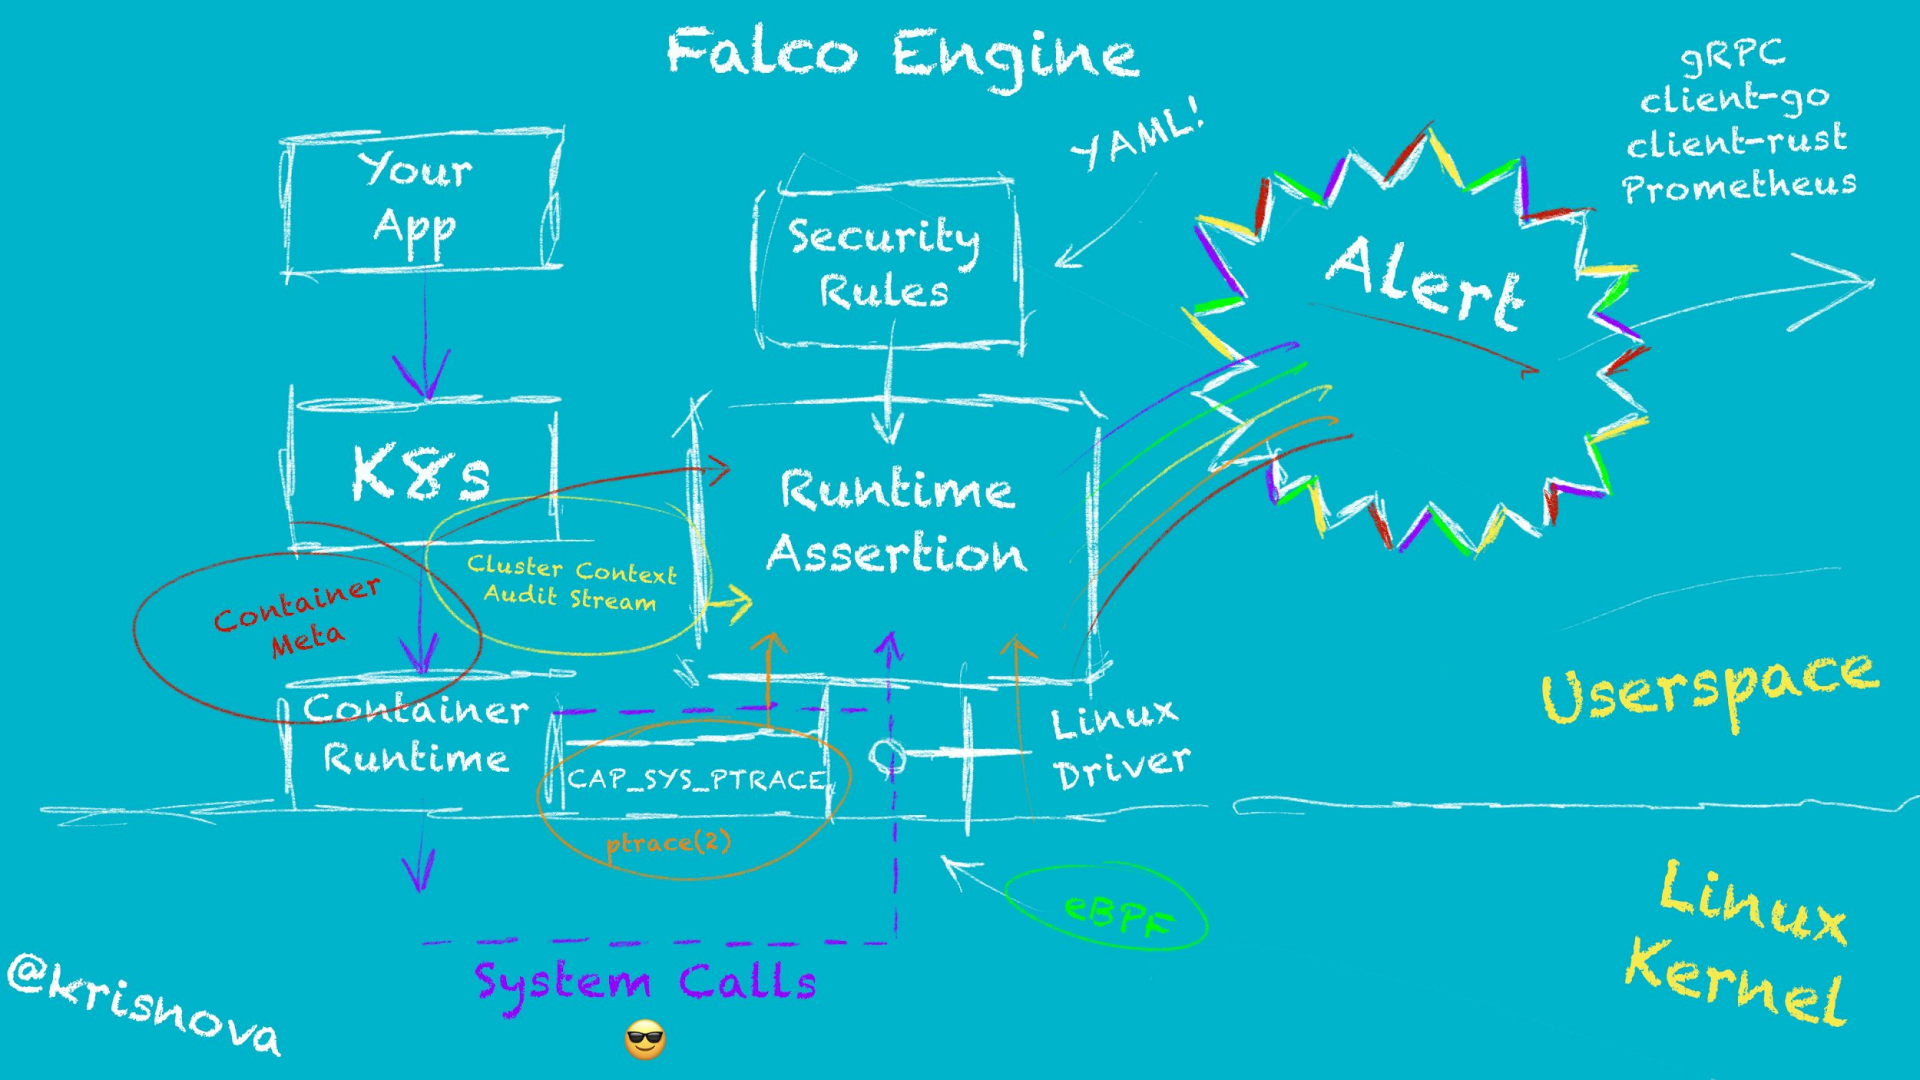
Falco Engine
Your App
K8s
Container Runtime
Security Rules
YAML!
Runtime Assertion
Cluster Context
Audit Stream
Container Meta
CAP_SYS_PTRACE
ptrace(2)
Linux Driver
eBPF
System Calls
Alert
gRPC
client-go
client-rust
Prometheus
Userspace
Linux Kernel
@krishnova

Building With Falco

Access
system calls
kubernetes

Assert
state of the system
runtime with security rules

Action
dynamically respond
infrastructure

Falco in Production

Falco

Data egress, coin mining

- unexpected outbound connection

Metadata server abuse

- contact cloud metadata service from container

Remote code execution

- spawned_process

Clients, SDKs, gRPC

- see falco.org

```
- rule: Contact cloud metadata service from container
  desc: Detect attempts to contact the Cloud Instance Metadata Service from a container
  condition: outbound and fd.sip="169.254.169.254" and container and consider_metadata_access and not user_known_metadata_access
  output: Outbound connection to cloud instance metadata service (command=%proc.cmdline connection=%fd.name %container.info image=%container.image.repository:%container.image.tag)
  priority: NOTICE
  tags: [network, container, mitre_discovery]
```

Falco APP 12:02 AM

k8s.pod.name
test-hostnetwork

proc.cmdline
wget -qO- --header Metadata-Flavor: Google 169.254.169.254/computeMetadata/v1/instance/service-accounts/default/email

container.id
d9b4c3cddd6d

container.image.repository
alpine

container.image.tag
latest

fd.name
10.104.204.46:56316->169.254.169.254:80

k8s.ns.name
test-app

rule
Contact cloud metadata service from container

priority
Notice

time
2020-07-24 04:02:18.323169867 +0000 UTC

https://github.com/falcosecurity/falcosidekick

Mastering Falco

Rules

- signal to noise ratio
- false positives

Alerts

- severity
- immediate actions

Infrastructure

- configuration
- upgrades

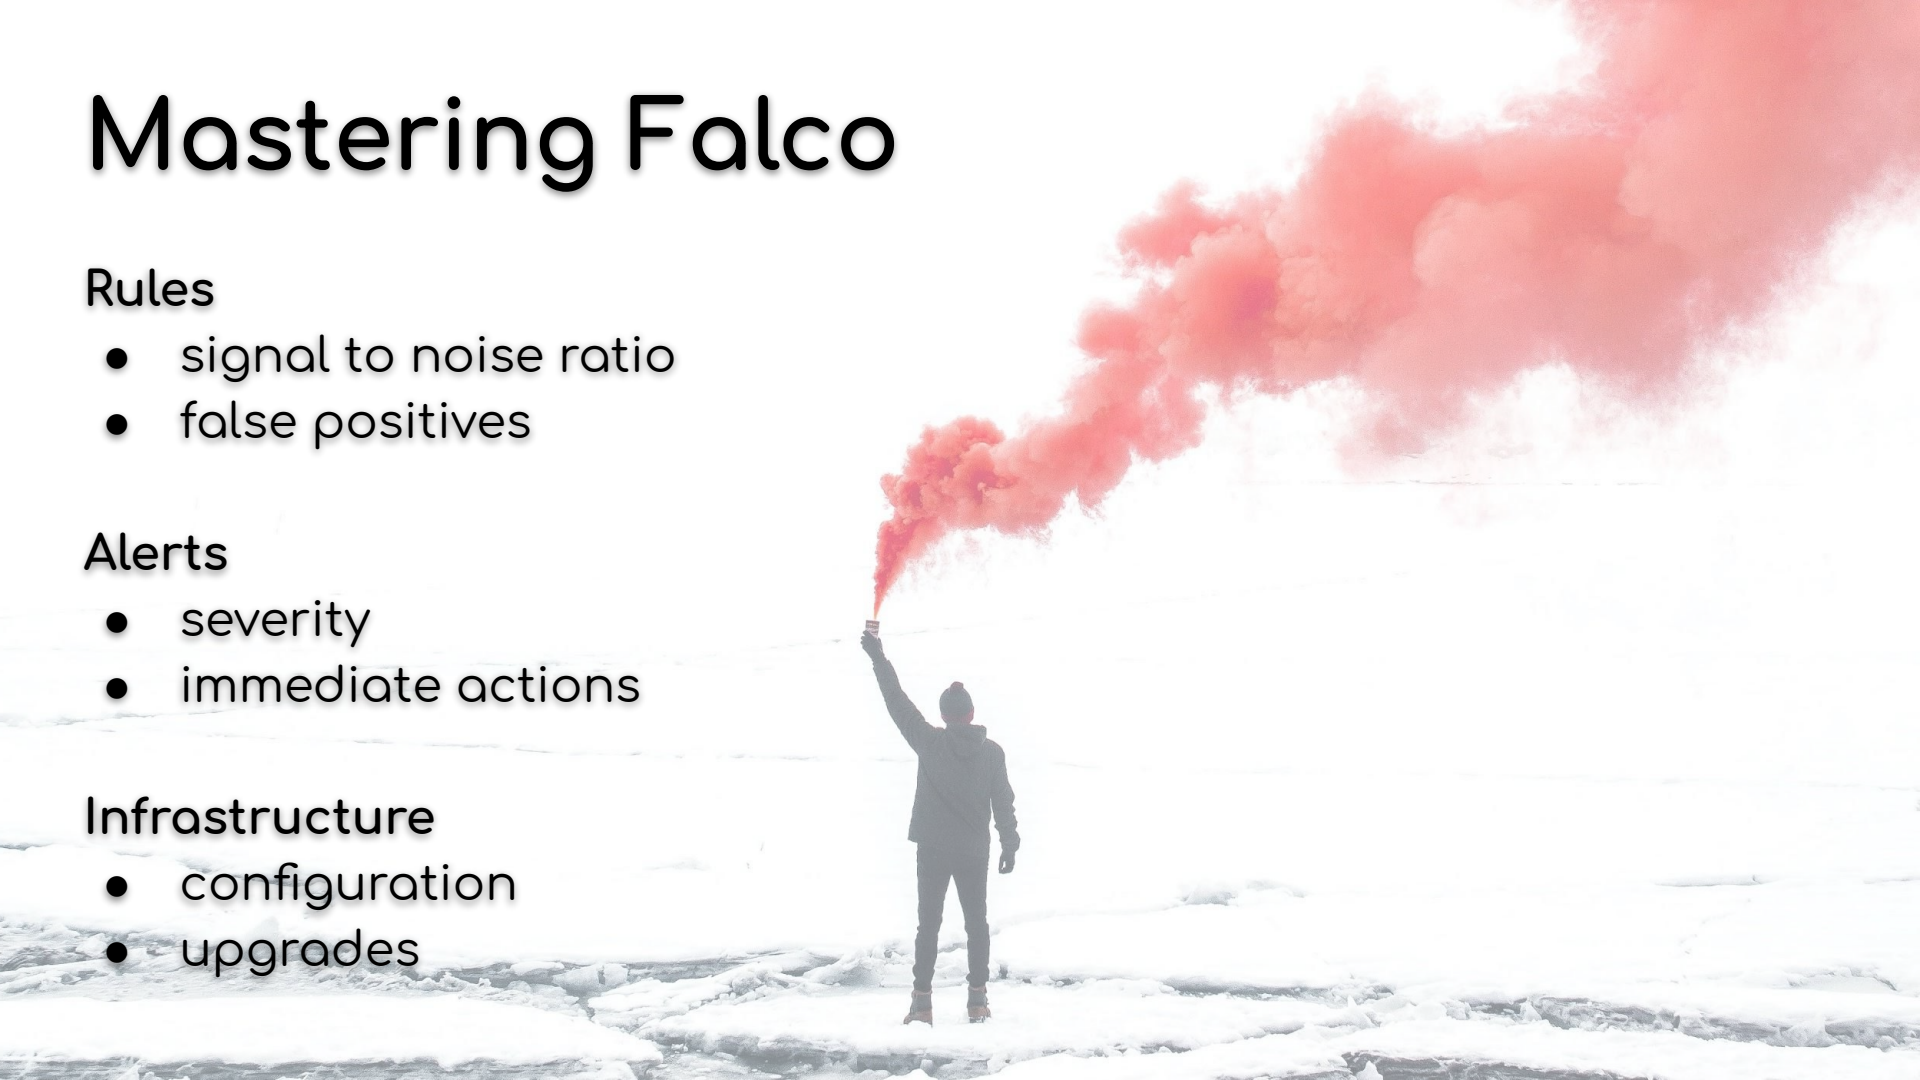

Security with Falco

Prevention

- first priority
- improves over time
- never guaranteed

Detection

- allows human reasoning
- enables manual intervention
- informs future prevention

Come Join the Party!

shopify.com/careers

sysdig.com/jobs

hackerone.com/shopify

falco.org

github.com/Shopify/kubeaudit

github.com/falcosecurity

engineering.shopify.com

sysdig.com/opensource

shopify

CLOUD NATIVE COMPUTING FOUNDATION

sysdig

Falco