READ ME

This deck has been successfully presented at conferences and webinars.
Make a copy of this material prior to making any changes. Thank you.

sysdig

Intro to Kubernetes Runtime Security

Speaker Name | **Contact info, LinkedIn or Github repo**

Title, Sysdig

Shift Left & Shield Right

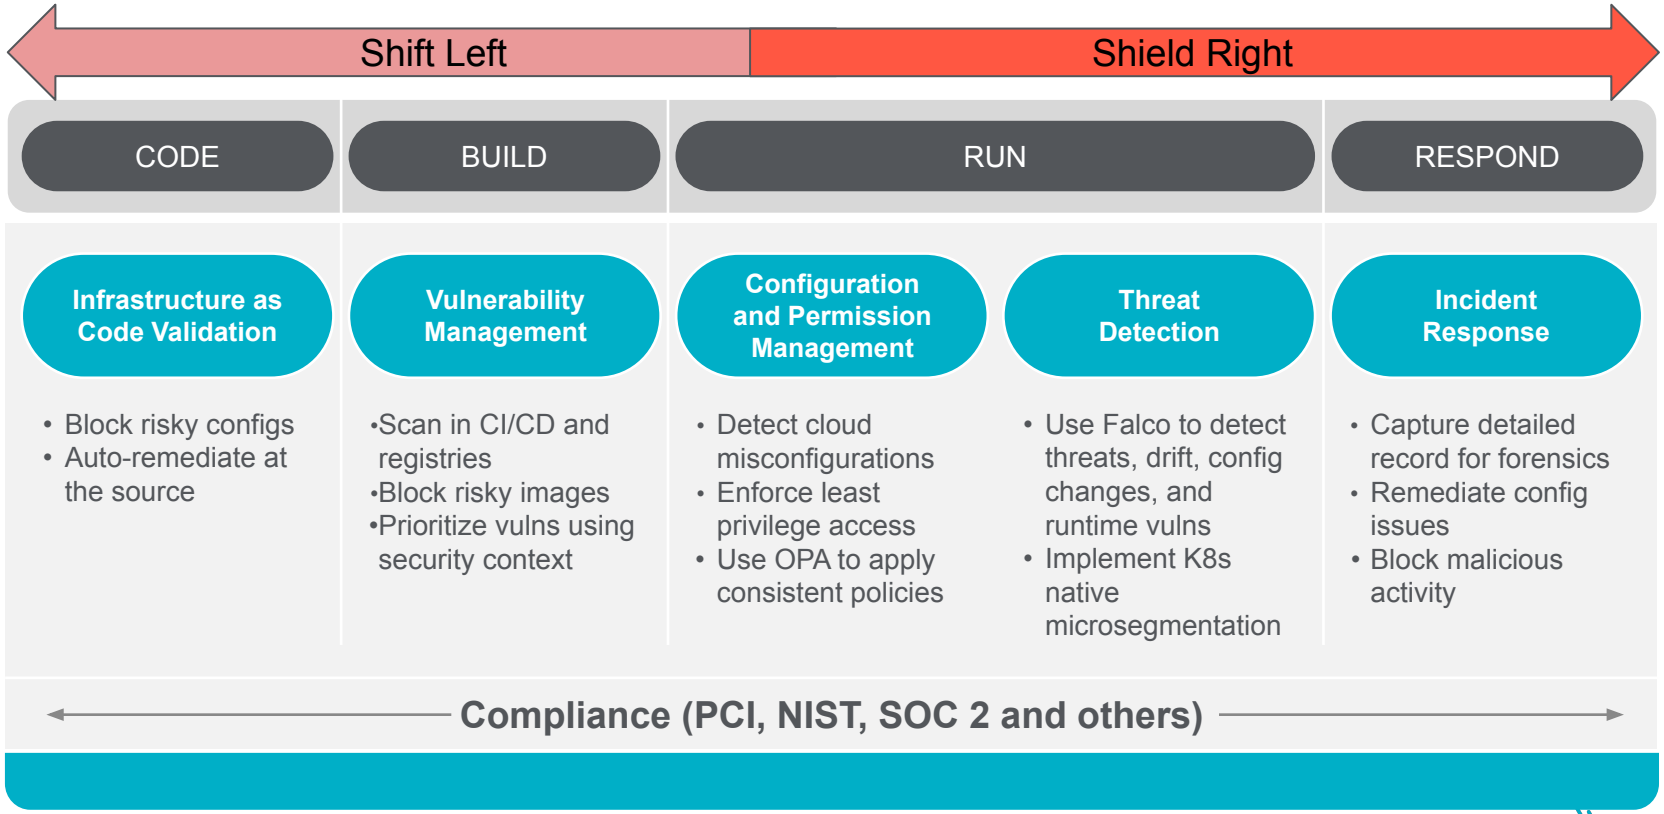

sysdig

OWASP Kubernetes Top 10

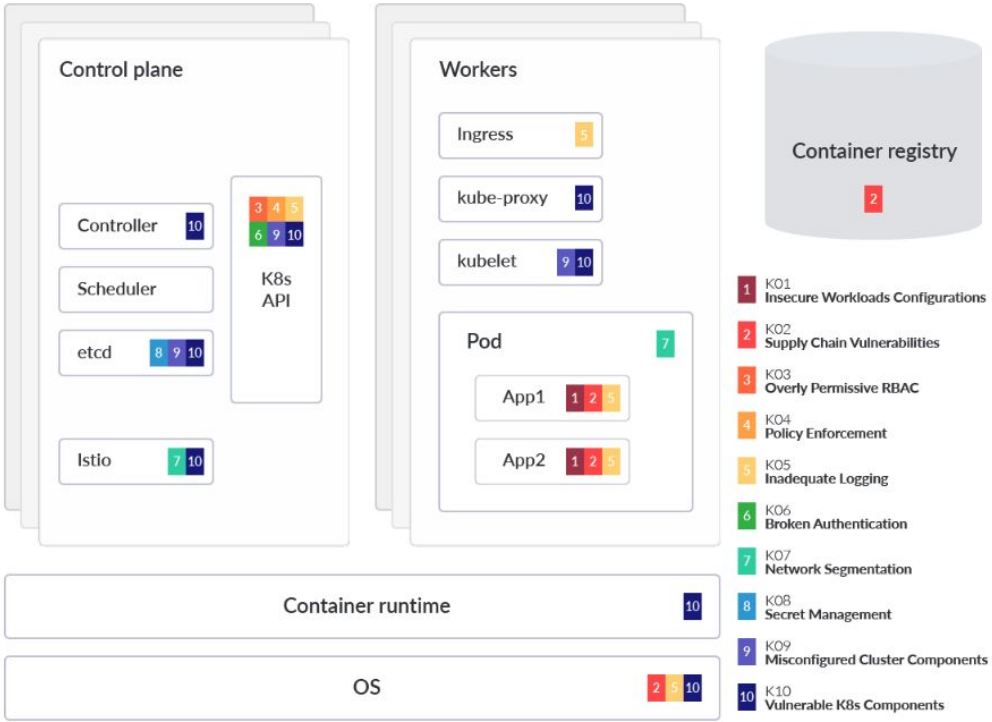

K01 - Insecure Workload Configurations

K02 - Supply Chain Vulnerabilities

K03 - Overly Permissive RBAC

K04 - Policy Enforcement

K05 - Inadequate Logging

K06 - Broken Authentication

K07 - Network Segmentation

K08 - Secret Management

K09 - Misconfigured Cluster Components

K10 - Vulnerable K8s Components

https://owasp.org/www-project-kubernetes-top-ten/
https://sysdig.com/blog/top-owasp-kubernetes/

sysdig

Containers are NOT…

Containers ARE:

- A group of processes running on the same host, sharing the same host kernel
- Isolation of these processes are facilitated by native Linux capabilities
  - cgroups
  - namespaces
- Attackers use many of the same methodologies to create container escape and access other resources on the host

sysdig

Unique Problems - Real World Insights

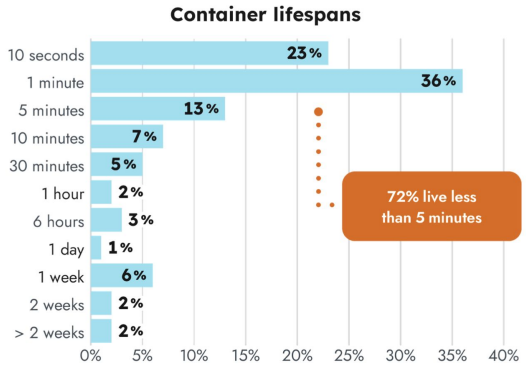

Containers are ephemeral by nature

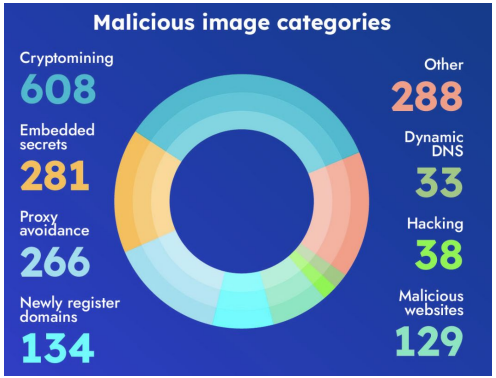

Supply Chain injection is target rich

Companies are having trouble keeping pace with the vulnerability landscape

Path of least resistance is still a real thing!

https://sysdig.com/2023-cloud-native-security-and-usage-report/

Once, There was a Perimeter

You had a perimeter
guarded by a firewall

Detecting intrusions was
your breach indicator

sysdig

Now, There is No Perimeter in the Cloud

- Cloud providers own external connections
- Cloud is exposed to the outside world
- You need to control access to services your team uses
- You need to detect unusual activity

sysdig

Without a Perimeter, a Security Camera is More Important than a Good Lock

- Watch for changes that create security gaps
- Identify intruders and suspicious insider behavior
- Send an alert and take immediate action

sysdig

created by Sysdig

CLOUD NATIVE
COMPUTING FOUNDATION

CNCF INCUBATED PROJECT

The Security Camera for Modern Apps

What is Falco?

- Runtime security engine
- Observability for endpoints and cloud infrastructure
- Built on eBPF
- Integrated with Kubernetes

CLOUD NATIVE COMPUTING FOUNDATION

CNCF INCUBATED PROJECT

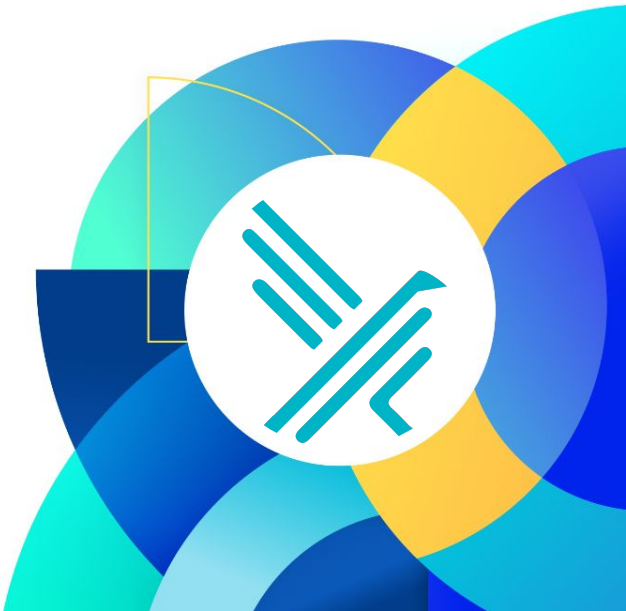

Introducing Falco

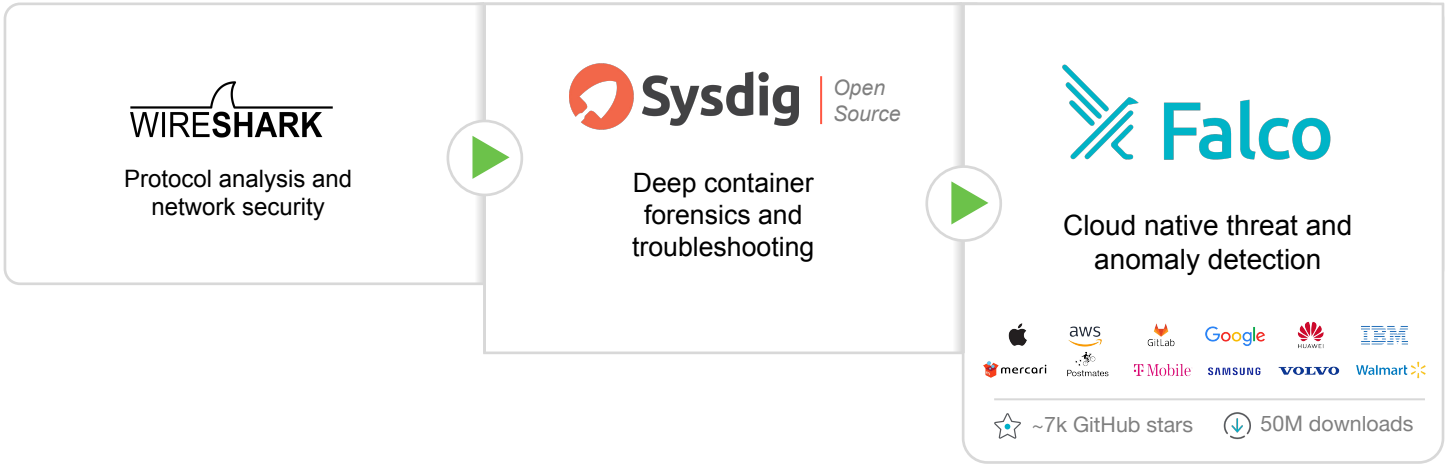

Falco rule engine

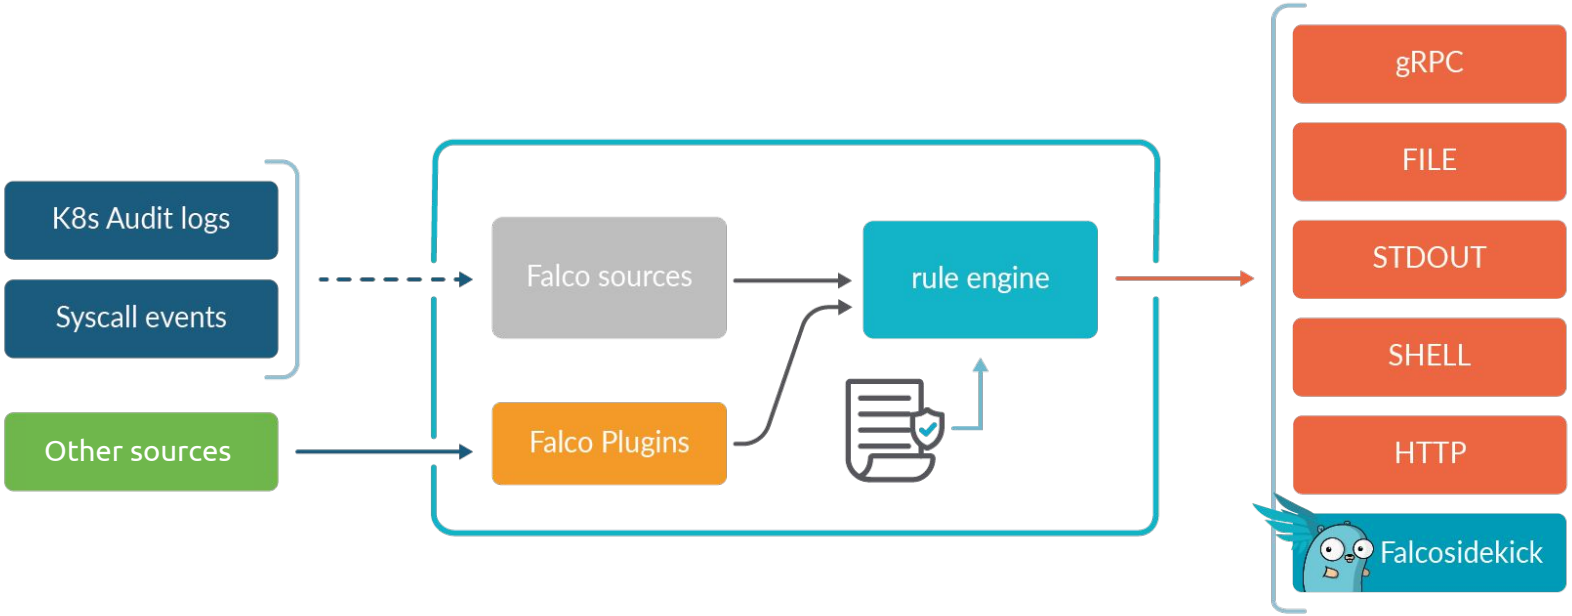

sysdig

Falco rule engine

```
- macro: create_symlink
  condition: (evt.type in (symlink, symlinkat) and
evt.dir=<)
```

```
Feb 21 13:04:32 ubuntu-2004 falco: 13:04:32.460103947: Warning Symlinks created over sensitive
files (user=root user_loginuid=-1 command= ln -sf /etc/shadow /tmp/marcel pid=1950
target=/etc/shadow linkpath=/tmp/marcel parent_process=create_symlink_)
```

```
- rule: Create Symlink Over Sensitive Files
  desc: Detect symlink created over sensitive files
  condition: >
    create_symlink and
    (evt.arg.target in (sensitive_file_names) or evt.arg.target in (sensitive_directory_names))
  output: >
    Symlinks created over sensitive files (user=%user.name user_loginuid=%user.loginuid command=%proc.cmdline
pid=%proc.pid target=%evt.arg.target linkpath=%evt.arg.linkpath parent_process=%proc.pname)
  priority: WARNING
  tags: [host, container, filesystem, mitre_exfiltration, mitre_credential_access, T1020, T1083, T1212, T1552,
T1555]
```

```
- list: sensitive_file_names
  items: [/etc/shadow, /etc/sudoers, /etc/pam.conf,
/etc/security/pwquality.conf]
```

sysdig

Common Examples

A shell is run in a container	container.id != host and proc.name = bash
Overwrite system binaries	fd.directory in (/bin, /sbin, /usr/bin, /usr/sbin) and write
Container namespace change	evt.type = setns and not proc.name in (docker, sysdig)
Non-device files written in /dev	(evt.type = create or evt.arg.flags contains O_CREAT) and proc.name != blkid and fd.directory = /dev and fd.name != /dev/null
Process tries to access camera	evt.type = open and fd.name = /dev/video0 and not proc.name in (skype, webex)

Popular Falco Rules

Best practices

Update packages
Modify /bin /usr
Write below /etc
Read sensitive file
DB spawned proc
Change namespace
Privileged container
Sensitive mount
Privileged shell
Terminal shell

Compliance

FIM (File Integrity)
Privileged pod
ConfigMap creds
kubectl exec/attach
Role changes audit
PCI
NIST

Vulnerabilities

CVE-2019-11246
kubectl cp

CVE-2019-5736
runc breakout

CVE-2020-14386
container escape

Cloud Native Stack

K8s control plane
Nginx
Elasticsearch
Redis
HAproxy
Rook
MongoDB
PostgreSQL

```
2022-04-07T12:51:08: Notice A shell was spawned in a container with an attached terminal (user=root
user_loginuid=-1 elastic_borg (id=a10bd3b1b2a8) shell=bash parent=<NA> cmdline=bash terminal=34816
container_id=a10bd3b1b2a8 image=ubuntu)
2022-04-07T12:51:41: Warning Netcat runs inside container that allows remote code execution
(user=root user_loginuid=-1 command=nc -e container_id=a10bd3b1b2a8 container_name=elastic_borg
image=ubuntu:latest)
```

DEMO

- Falco In Flight
- Falco Sidekick

sysdig

Beyond system calls and containers

Plugins are **dynamic shared libraries**

which allow **Falco** to **collect** and **extract fields**

from **streams of events**

sysdig

Users and builders

Giant Swarm

SAMSUNG

Walmart

VOLVO

Booz | Allen | Hamilton

shopify

Frame.io

Phoenix

GitLab

Falco

CLOUD NATIVE
COMPUTING FOUNDATION

CNCF INCUBATED PROJECT

T-Mobile

Skyscanner

HUAWEI

Postmates

sysdig

Resources

Get started at Falco.org

Check out the Falco project in Github

Get involved in the Falco community

Meet the maintainers on the Falco Slack

Follow @falco_org on

Join a Falco workshop

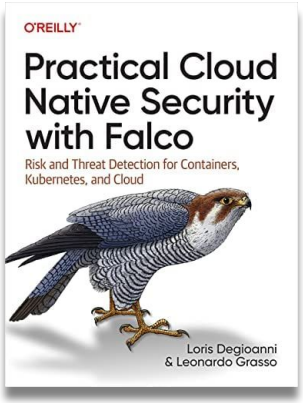

sysdig

Questions

sysdig


sysdig
Seeing is Securing

Open Source Drives Effective Cloud Security

- Collaborative development goes broad and deep
- Standards-based development promotes choice
- Open source is transparent
- Modern platforms are build on open source

The Future of Security is Open

sysdig

The Falco Sensor

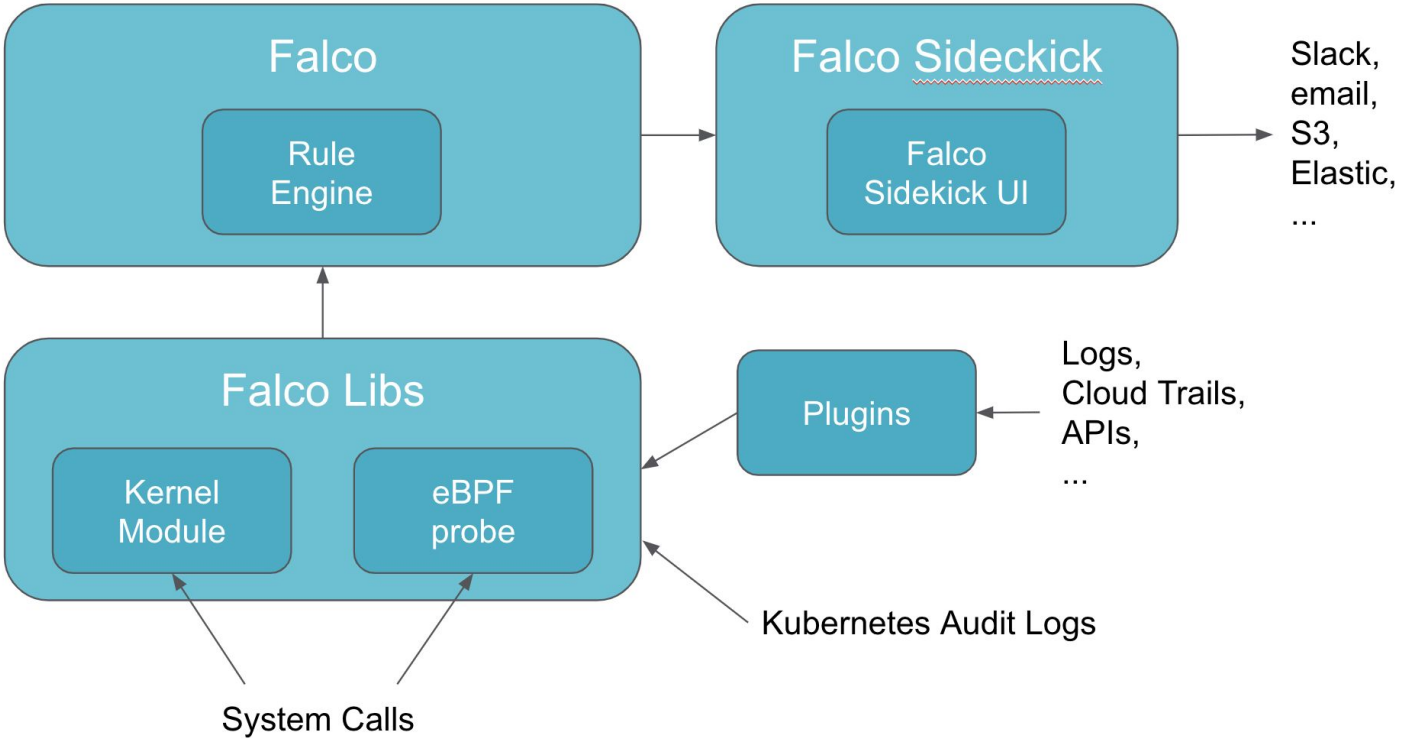

sysdig