

System Introspection for Micro-Services Pentests in K8S

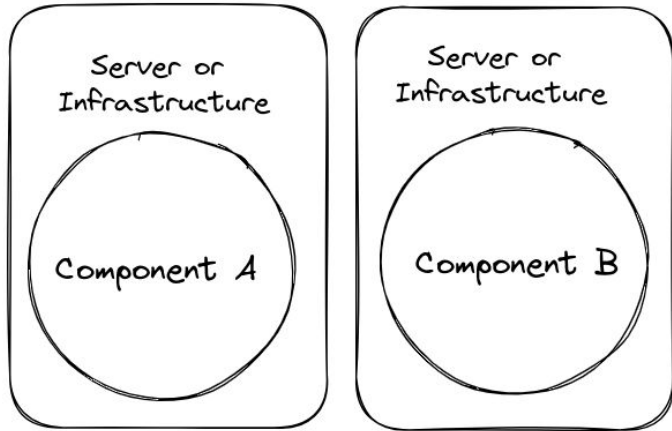


1. Introduction to the “in-between” world
2. Capturing Component System Interactions
3. Micro-Services Pentesting in K8S 101
4. How-To Falco & My Favorite SRE
5. Pros, Cons, & WhatNot
6. Conclusion & Kudos

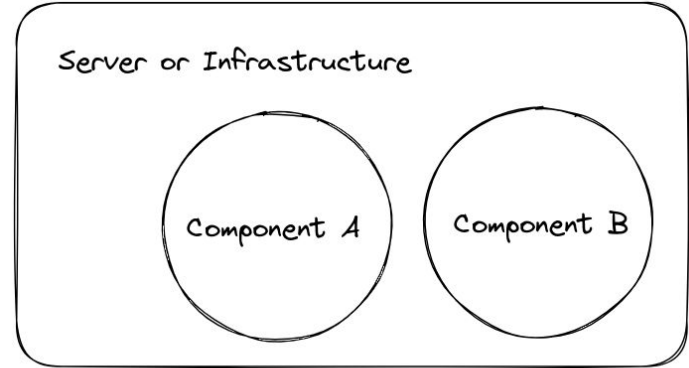


Introduction to the “in-between” world

The root of evil: “You & I are safe, so we are safe”



SAFE ;)



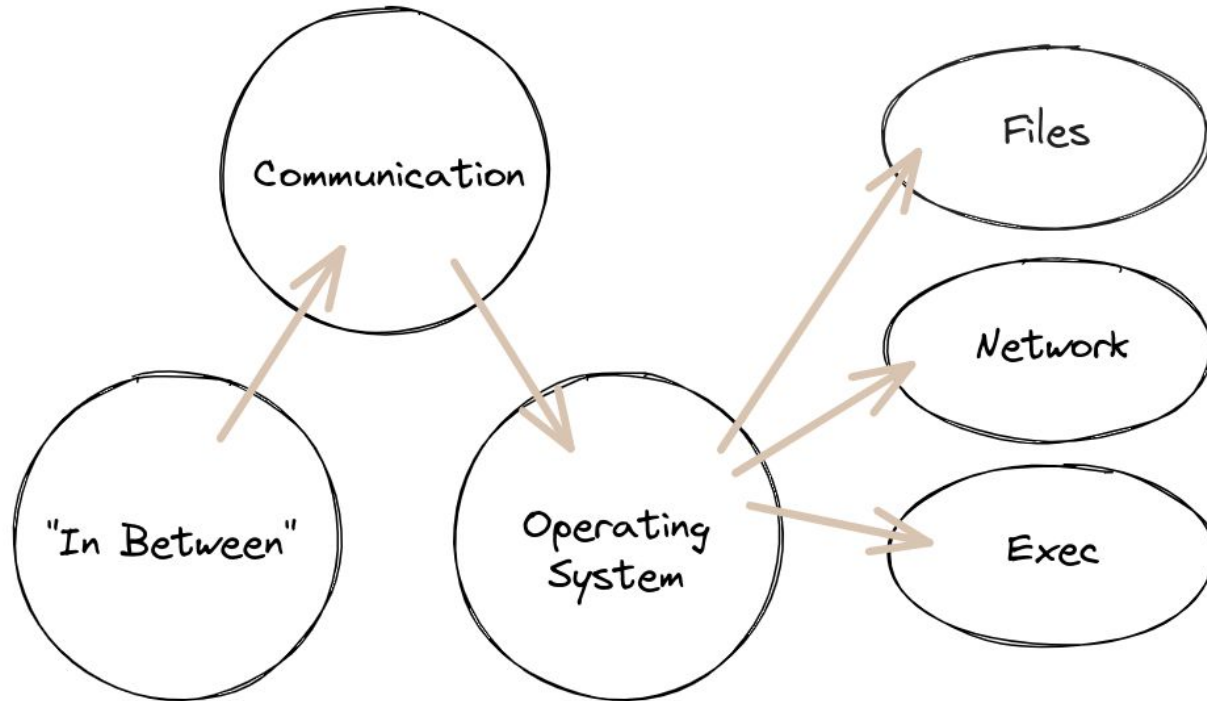
BOOM

The “in between”: Examples


- File upload + SSH server with key
- FTP service + PHP server
- Fetch by URL + Cloud Infra
- URL parser front != URL parser back
- Admin iface on loopback + Proxy service
- Customizable username + PDF export
- PNG file upload + PHAR format



The “in between”: Primitives & Detection



The “in between”: Primitives & Detection



The “in between”: Pick your BOX, but stay blind.



ZERO KNOWLEDGE



SOME KNOWLEDGE



FULL KNOWLEDGE

Capturing Component System Interactions



Bienvenue sur wikiHow, le site de tutoriels le plus fiable d'internet.

Qu'allez-vous découvrir sur wikiHow aujourd'hui ?





Hook within the binary in User-Land?

- 🕒 28 351 vues
- 🕒 Actualisé il y a 7 ans



Hook with a Kernel Module?

- 🕒 21 579 vues
- 🕒 Actualisé il y a 2 ans



Rely on eBPF probes?

- 🕒 159 989 vues
- 🕒 Actualisé il y a 2 ans



Comment survivre dans 7 Days **to** Die

- 28 351 vues
- Actualisé il y a 7 ans




Comment cultiver dans 7 Days **to** Die

- 21 579 vues
- Actualisé il y a 2 ans



Comment éviter d'avoir un camel toe (l'orteil de chameau)

- 159 989 vues
- Actualisé il y a 2 ans



Getting Started [↗](#)

Run Sysdig in a container:

```
sudo docker run --rm -i -t --privileged --net=host \  
-v /var/run/docker.sock:/host/var/run/docker.sock \  
-v /dev:/host/dev \  
-v /proc:/host/proc:ro \  
-v /boot:/host/boot:ro \  
-v /src:/src \  
-v /lib/modules:/host/lib/modules:ro \  
-v /usr:/host/usr:ro \  
-v /etc:/host/etc:ro \  
docker.io/sysdig/sysdig
```



And then run the `sysdig` or `csysdig` tool from the container shell!

Or install the [latest release](#) with a `deb` or `rpm` package for your distribution.

```
root@ML-PF2W5DKR: /# whoami
```

```
root
```

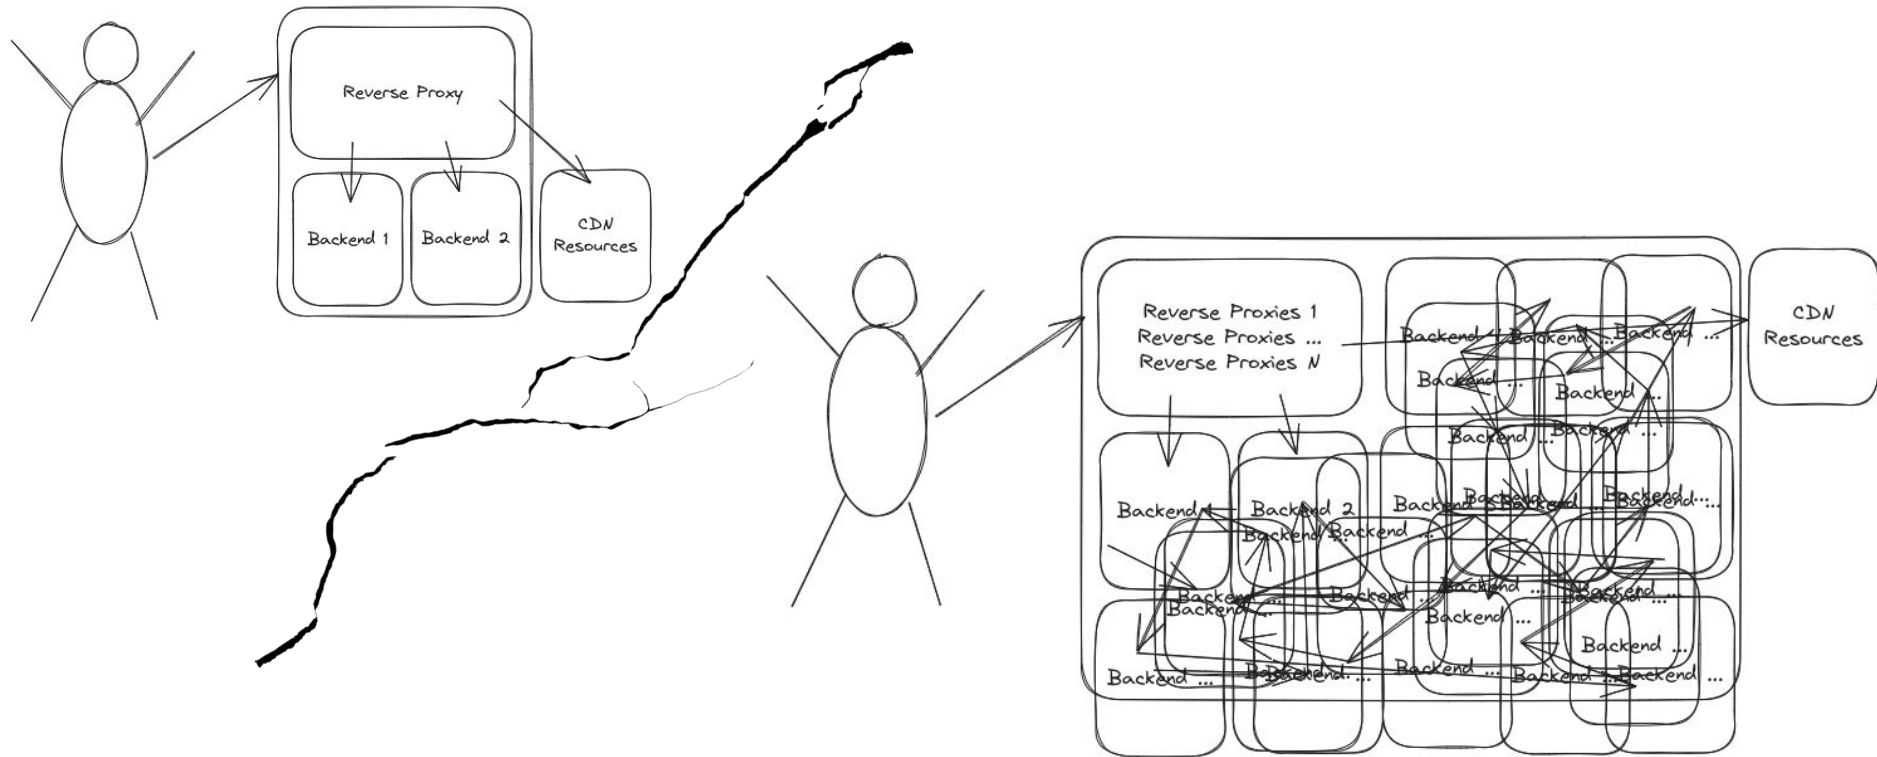
```
root@ML-PF2W5DKR: /#
```

```
[root@ML-PF2W5DKR /]# sysdig -A container.name=ubuntu | grep -P 'exec|open|write|read'
```

```
6712 16:48:20.882391241 5 bash (106256) > read fd=0(<f>/dev/pts/0) size=1
6714 16:48:20.882400157 5 bash (106256) < read res=1 data=
6716 16:48:20.882406794 5 bash (106256) > write fd=2(<f>/dev/pts/0) size=1
6718 16:48:20.882413531 5 bash (106256) < write res=1 data=
6719 16:48:20.882421681 5 bash (106256) > write fd=2(<f>/dev/pts/0) size=9
6722 16:48:20.882424145 5 bash (106256) < write res=9 data=
7021 16:48:20.883043224 4 bash (108244) > read fd=3(<p>) size=1
7023 16:48:20.883046538 4 bash (108244) < read res=0 data=NULL
7100 16:48:20.883149165 4 bash (108244) > execve filename=/usr/bin/whoami
7198 16:48:20.883552137 4 whoami (108244) < execve res=0 exe=whoami args=NULL tid=108244(whoami) pid=108244(whoami) ptid=106256(bash) cwd= fdlimit=1048576
m_rss=4 vm_swap=0 comm=whoami cgroups=
7219 16:48:20.883742774 4 whoami (108244) > openat dirfd=-100(AT_FDCWD) name=/etc/ld.so.cache flags=4097(O_RDONLY|O_CLOEXEC) mode=0
7228 16:48:20.883760911 4 whoami (108244) < openat fd=3(<f>/etc/ld.so.cache) dirfd=-100(AT_FDCWD) name=/etc/ld.so.cache flags=4097(O_RDONLY|O_CLOEXEC) mode=0
7250 16:48:20.883791145 4 whoami (108244) > openat dirfd=-100(AT_FDCWD) name=/lib/x86_64-linux-gnu/libc.so.6 flags=4097(O_RDONLY|O_CLOEXEC) mode=0
7254 16:48:20.883806160 4 whoami (108244) < openat fd=3(<f>/lib/x86_64-linux-gnu/libc.so.6) dirfd=-100(AT_FDCWD) name=/lib/x86_64-linux-gnu/libc.so.6 flags=4097(O_RDONLY|O_CLOEXEC) mode=0
ino=10003473
7255 16:48:20.883808525 4 whoami (108244) > read fd=3(<f>/lib/x86_64-linux-gnu/libc.so.6) size=832
7261 16:48:20.883813228 4 whoami (108244) < read res=832 data=
7267 16:48:20.883820820 4 whoami (108244) > pread fd=3(<f>/lib/x86_64-linux-gnu/libc.so.6) size=784 pos=64
7269 16:48:20.883823209 4 whoami (108244) < pread res=784 data=
7270 16:48:20.883824525 4 whoami (108244) > pread fd=3(<f>/lib/x86_64-linux-gnu/libc.so.6) size=48 pos=848
7272 16:48:20.883825936 4 whoami (108244) < pread res=48 data=
7273 16:48:20.883827018 4 whoami (108244) > pread fd=3(<f>/lib/x86_64-linux-gnu/libc.so.6) size=68 pos=896
7274 16:48:20.883828432 4 whoami (108244) < pread res=68 data=
7277 16:48:20.883837008 4 whoami (108244) > pread fd=3(<f>/lib/x86_64-linux-gnu/libc.so.6) size=784 pos=64
7278 16:48:20.883838433 4 whoami (108244) < pread res=784 data=
7427 16:48:20.884627998 4 whoami (108244) > openat dirfd=-100(AT_FDCWD) name=/etc/nsswitch.conf flags=4097(O_RDONLY|O_CLOEXEC) mode=0
7431 16:48:20.884642651 4 whoami (108244) < openat fd=3(<f>/etc/nsswitch.conf) dirfd=-100(AT_FDCWD) name=/etc/nsswitch.conf flags=4097(O_RDONLY|O_CLOEXEC) mode=0
7435 16:48:20.884665915 4 whoami (108244) > read fd=3(<f>/etc/nsswitch.conf) size=4096
7436 16:48:20.884671829 4 whoami (108244) < read res=494 data=
7437 16:48:20.884692309 4 whoami (108244) > read fd=3(<f>/etc/nsswitch.conf) size=4096
7438 16:48:20.884693817 4 whoami (108244) < read res=0 data=NULL
7444 16:48:20.884719309 4 whoami (108244) > openat dirfd=-100(AT_FDCWD) name=/etc/passwd flags=4097(O_RDONLY|O_CLOEXEC) mode=0
7448 16:48:20.884731758 4 whoami (108244) < openat fd=3(<f>/etc/passwd) dirfd=-100(AT_FDCWD) name=/etc/passwd flags=4097(O_RDONLY|O_CLOEXEC) mode=0 dev=
```


Micro-Services Pentesting in K8S 101

The pentest I Know Vs The one I Fear




Totally stolen yet accurate graphs

More Complexity → More Bugs



More Complexity → WTF's happening ?



Huge structures have the more -security- bugs

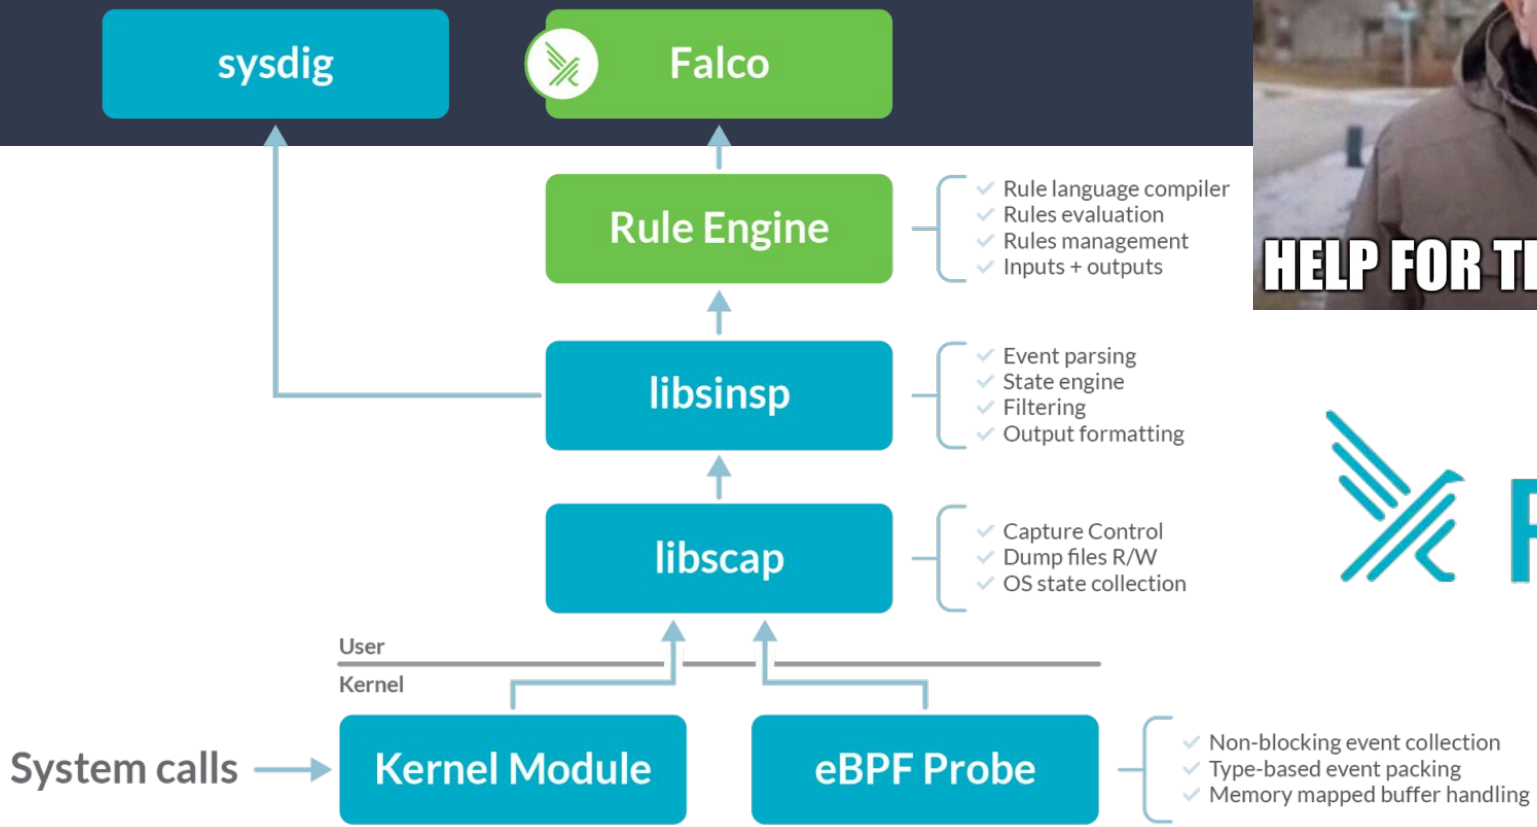


Huge structures offer no “easy assumptions”



How can we bring our magic-glasses to K8S?





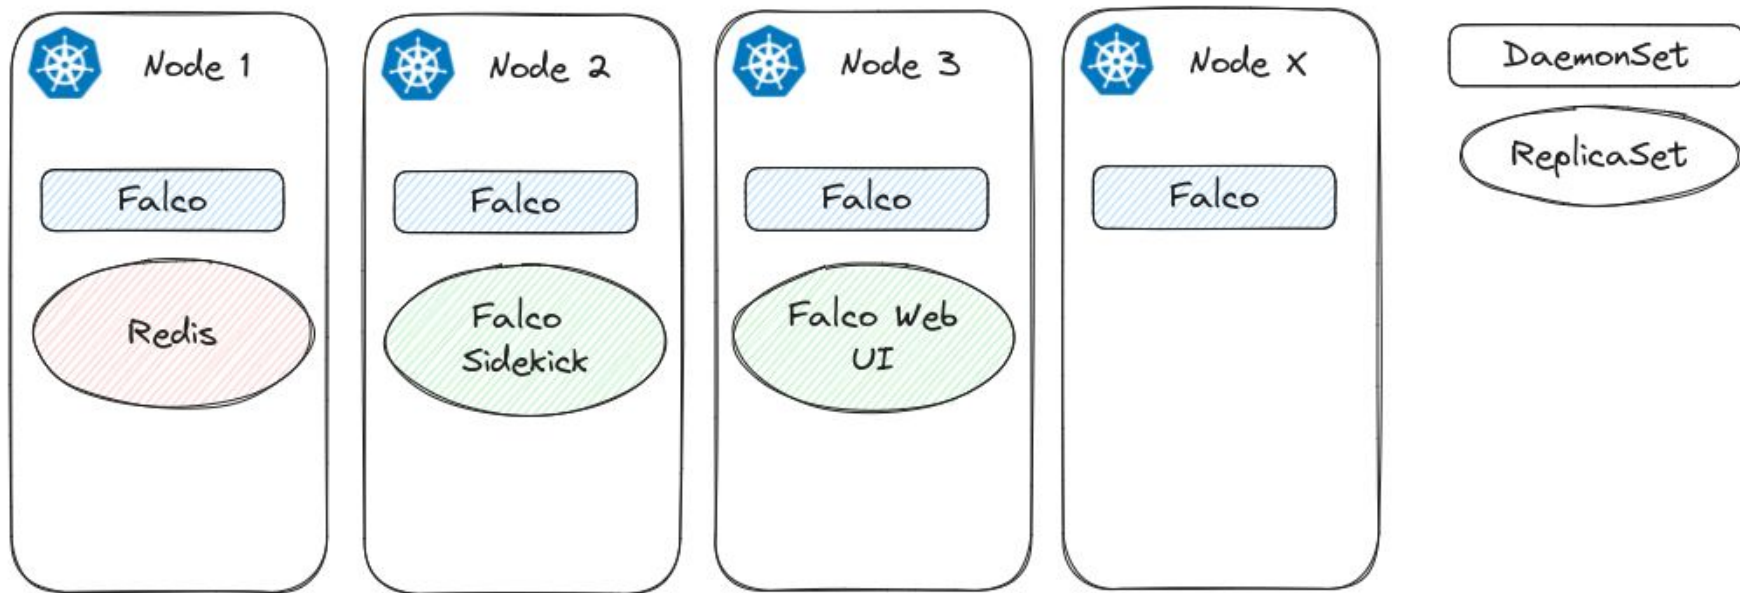
How-To Falco & My Favorite SRE

Falco containers services

- Falco agent
 - Monitor kernel events base on custom rules
- Falco sidekick
 - Take Falco events and forward to different output (Slack, Prometheus, Kafka ...)
- Falco Web UI (Optional)
 - Simple web ui to displaying latest event
- Redis (Optional)
 - Store Falco events for Falco Web UI



Falco on Kubernetes



Hello Helm chart !

```
~$ kubectl -n falco get deployment
```

NAME	READY	UP-TO-DATE	AVAILABLE
infra-falco-sidekick-falcosidekick	1/1	1	1
infra-falco-sidekick-falcosidekick-ui	1/1	1	1


```
~$ kubectl -n falco get daemonset
```

NAME	DESIRED	CURRENT	READY	UP-TO-DATE	AVAILABLE
infra-falco-controller	28	28	28	28	28

```
—
image:
  repository: "public/falcosecurity/falco-no-driver"

resources:
  requests:
    cpu: 200m
    memory: 512Mi
  limits:
    cpu: 1000m
    memory: 1024Mi

affinity:
  nodeAffinity:
    requiredDuringSchedulingIgnoredDuringExecution:
      nodeSelectorTerms:
        - matchExpressions:
            - key: eks.amazonaws.com/compute-type
              operator: NotIn
              values:
                - fargate
            - key: karpenter.k8s.aws/instance-gpu-count
              operator: DoesNotExist
```



```
—
{{- if eq .Values.controller.kind "daemonset" }}
apiVersion: apps/v1
kind: DaemonSet
metadata:
  name: {{ include "falco.fullname" . }}
  namespace: {{ include "falco.namespace" . }}
  labels:
    {{- include "falco.labels" . | nindent 4 }}
  {{- if .Values.controller.annotations }}
  annotations:
    {{ toYaml .Values.controller.annotations | nindent 4 }}
  {{- end }}
spec:
  selector:
    matchLabels:
      {{- include "falco.selectorLabels" . | nindent 6 }}
  template:
    {{- include "falco.podTemplate" . | nindent 4 }}
    {{- with .Values.controller.daemonset.updateStrategy }}
  updateStrategy:
    {{- toYaml . | nindent 4 }}
    {{- end }}
  {{- end }}
```


A little bit of Magic !

From Falco's yaml rules-file

```
- macro: isLogLoopSafe
  condition: >
    ((k8s.pod.name exists)
    and (not k8s.pod.name icontains "falco")
    and (not k8s.pod.name icontains "fluent")
    and (not k8s.pod.name icontains "guardduty")
    and (not k8s.pod.name icontains "vector-forward"))

- rule: Privilege Escalation Script
  desc: A shell script has been used to attempt privilege escalation and enumeratoin.
  condition: >
    spawned_process and (
    proc.cmdline icontains LinEnum or
    proc.cmdline icontains lse.sh or
    proc.cmdline icontains smartenum or
    proc.cmdline icontains exploit or
    proc.cmdline icontains privcheck or
    proc.cmdline icontains linpeas)
    and isLogLoopSafe
  output: >
    A shell script has been used to attempt privilege escalation and enumeratoin. (evt.type=%evt.type
    user=%user.name user_loginuid=%user.loginuid shell=%proc.name parent=%proc.pname cmdline=%proc.cmdline
    pid=%proc.pid terminal=%proc.tty)
  priority: CRITICAL
  tags: [manomano, shell, privesc]
```

Kustomize

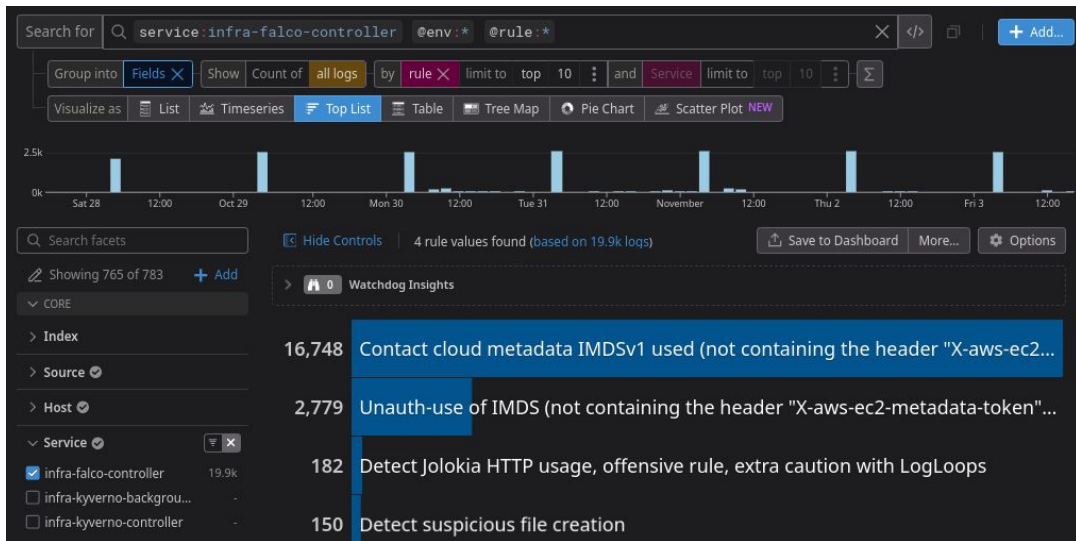


... to Kubernetes ConfigMap

```
apiVersion: v1
data:
  falco_rules.local.yaml: "# MM Custom rules\n\n macro: isLogLoopSafe\n condition:
  >\n ((k8s.pod.name exists)\n and (not k8s.pod.name icontains \"falco\")\n
  \ and (not k8s.pod.name icontains \"fluent\")\n and (not k8s.pod.name icontains
  \"guardduty\")\n and (not k8s.pod.name icontains \"vector-forward\"))\n\n-
  list: mysql-db-fqdns\n items: [int-app-common-db.manomano.tech, int-app-common-db-
  ro.manomano.tech,
  sig-app-common-db.manomano.tech, sig-app-common-db-ro.manomano.tech, sig-app-common-db-
  bi.manomano.tech,
  prd-app-common-db.manomano.tech, prd-app-common-db-ro.manomano.tech, prd-app-common-db-
  bi.manomano.tech]\n\n-
  list: allowed-files\n items: [dd_init.php, App_KernelProdContainer.preload.php,
  flatted.php]\n\n- list: monitored-dangerous-sockets\n items: [docker.sock, containerd.sock,
  docker-containerd.sock, criu.sock]\n\n- rule: Privilege Escalation Script\n desc:
  A shell script has been used to attempt privilege escalation and enumeratoin.\n
  \ condition: >\n spawned_process and (\n proc.cmdline icontains LinEnum
  or\n proc.cmdline icontains lse.sh or\n proc.cmdline icontains smartenum
  or\n proc.cmdline icontains exploit or\n proc.cmdline icontains privcheck
  or\n proc.cmdline icontains linpeas)\n and isLogLoopSafe\n output: >\n
  \ A shell script has been used to attempt privilege escalation and enumeratoin.
```


Houston we have a problem !

Datadog TopList & Logs




Slack

The screenshot shows a Slack message from the 'Blue team' channel, dated 5 months ago. The message contains a critical alert about a shell script attempt to escalate privileges. The alert details include the time, rule name, source, tags, environment, and container information.

```
13:16:19.105451711: Critical A shell script has been used to attempt privilege escalation and enumeratoin. (user=root user_loginuid=-1 shell=touch parent=bash cmdline=touch /tmp/lse.sh.christophe pid=7375 terminal=34816) k8s.ns=data-science k8s.pod=ml-data-science-devspace-devspace-worker-devspace-7dd48989h54f7 container=af8bd279ecff
rule                               priority
Privilege Escalation Script        Critical
source                               hostname
syscall                             infra-falco-controller-6hw7v
tags                                 container.id
manomano, privesc, shell            af8bd279ecff
environment                          k8s.ns.name
int                                  data-science
k8s.pod.name                         ml-data-science-devspace-devspace-worker-devspace-7dd48989h54f7
proc.cmdline                         proc.name
touch /tmp/lse.sh.christophe         touch
proc.pname                           user.name
bash                                  root
time                                  2023-05-31 13:16:19.105451711 +0000 UTC
https://github.com/falcosecurity/falcosidekick
```



- Go to...
- Watchdog
- Service Mgmt
- Dashboards
- Infrastructure
- Monitors
- Metrics
- Integrations
- APM
- CI
- Notebooks
- Logs
- Security
- UX Monitoring



Our Kubernetes dashboard gives you broad visibility into the scale, status, and resource usage of your cluster and its containers.

Further reading for Kubernetes monitoring

- [Autoscale Kubernetes workloads with any Datadog metric](#)
- [How to monitor Kubernetes + Kubernetes Pods](#)

More

Overview

Pods Running

Pods Runnin...

falco, infra-falco-sideki...
falco, infra-falco-sideki...
falco, N/A

Pods in Bad P...

4 pendin...

Pods running by Namespace

Tags in Po...	Avg	Max	Value
kube_clus...	34.1	45	31.0

Containers

Containers States

Ready Running Terminated Waiting

Containers OOM Killed (by Pod)

Containers in CrashloopBackOff (by Pod)

Tags Metric Avg Max Value

Container Restarts by Pod

pod_name in R...	Avg	Min	Max	Sum	Value
infra-falco-cont...	0	0	0	0	—
infra-falco-cont...	0	0	0	0	—
infra-falco-cont...	0	0	0	0	0

Pods

CPU-intensive...

129.2	infra-falco...
123.8	infra-falco...
108.4	infra-falco...
105.1	infra-falco...
99.7	infra-falco...
97.2	infra-falco...
85.4	infra-falco...
83.9	infra-falco...
83.5	infra-falco...
81.2	infra-falco...

CPU Usage by Pod

pod...	Avg	Max	Value
in...	34.0 mc...	57.2 mc...	—
in...	22.5 mc...	28.1 mc...	—
in...	57.2 mc...	114.0 mc...	—

Memory-inte...

469	infra-falco...
229	infra-falco...
224	infra-falco...
211	infra-falco...
211	infra-falco...
194	infra-falco...
163	infra-falco...
159	infra-falco...
158	infra-falco...
155	infra-falco...

Memory Usage by Pod

pod...	Avg	Max	Value
infr...	24.5 MiB	25.9 MiB	—
infr...	25.1 MiB	30.6 MiB	25.5 MiB
infr...	50.5 MiB	67.5 MiB	—

CPU Usage by Container

Tags in CPU Usage	Avg	Max	Value
container_id:0a158a9f43e...	80.3 mcores	185.3 mcores	—
container_id:0bd6cbbc3...	16.2 mcores	18.8 mcores	—
container_id:0d5a8ef04f...	40.9 mcores	68.9 mcores	39.4 mcores

Memory Usage by Container

Tags in Memory Usage	Avg	Max	Value
container_id:0a158a9f43e440af0f...	105.2 MiB	165.4 MiB	—
container_id:0bd6cbbc3b38003c...	25.7 MiB	25.9 MiB	—
container_id:0d5a8ef04f2194c1c...	100.1 MiB	118.7 MiB	118.5 MiB

Network Rate by Pod

pod_name	Metric	Avg	Max	Value
infra-falco-contr...	Received ...	45.3 B/s	52 B/s	—
infra-falco-contr...	Received ...	68.6 B/s	104 B/s	69.9 B/s
infra-falco-contr...	Received ...	159.8 B/s	240 B/s	—

Network Errors by Pod

pod_name	Metric	Avg	Max	Value
infra-falco-controller...	Received Err...	0 errs/s	0 errs/s	—
infra-falco-controller...	Received Err...	0 errs/s	0 errs/s	0 errs/s
infra-falco-controller...	Received Err...	0 errs/s	0 errs/s	—

Pods Running...

45 int-infra...

Pods in Ready...

3 int-infra-eks-clu...
2 int-infra-eks-clu...
2 int-infra-eks-clu...
2 int-infra-eks-clu...
2 int-infra-eks-clu...
2 int-infra-eks-clu...
2 int-infra-eks-clu...
2 int-infra-eks-clu...
2 int-infra-eks-clu...

Pods in Bad S...

2 falco

Pods Status Phase

- Live Chat
- Help NEW STUFF
- Invite Users
- louka.jacques...
- ManoMano

DEMO

```
$ remote-pod-inspection.sh M x ! rules-remote-pod-inspection.yaml
```

```
$ remote-pod-inspection.sh
```

```
1 #!/bin/bash
2 set -eo pipefail # -u
3
4 ## Doc
5 # MMDEBUG=1 ./remote-pod-inspection.sh
6 # Events: https://falco.org/docs/reference/rules/supported-events/
7 # Fields: https://falco.org/docs/reference/rules/supported-fields/
8
9 # Enable debug
10 if [ "$MMDEBUG" == "1" ]; then
11 |   set -x
12 | fi
13
14 # Check prereqs
15 for command in fzf kubectl kubectx tee docker grep jq tmux zsh; do
16 |   if ! [ -x "$(command -v $command)" ]; then
17 |     echo "Command $command could not be found, please install: fzf stern kubectl kubectx"
18 |     exit 1
19 |   fi
20 | done
21
22 # Check tmux presence
23 if [[ -z "${TMUX}" ]]; then
24 |   echo "Must be run in tmux"
25 |   exit 42
26 | fi
```

Ln 10, Col 31 Spaces: 4 UTF-8 LF Shell Script Prettier

```
/opt/manogits/falco-rules (master*) »
```

130 ↵

Pros, Cons, & WhatNot

Costs

- **Resource overhead:** Use the Kubernetes cluster resources (CPU/Memory/Network)
- **Storage:** Need to store the logs and events generated by Falco
- **Deployment and management:** Keeping Falco up to date, setup alerts, rules
- **Training and skill development:** Create rules, interpret alerts, setup Falco
- **Integration:** Monitoring and other security solutions



Limitations



- **Kernel or eBPF features**
- **Limited scope**
- **Complexity**
- **Customization**
- **False positives**
- **Performance impact**
- **Community and support**

Results

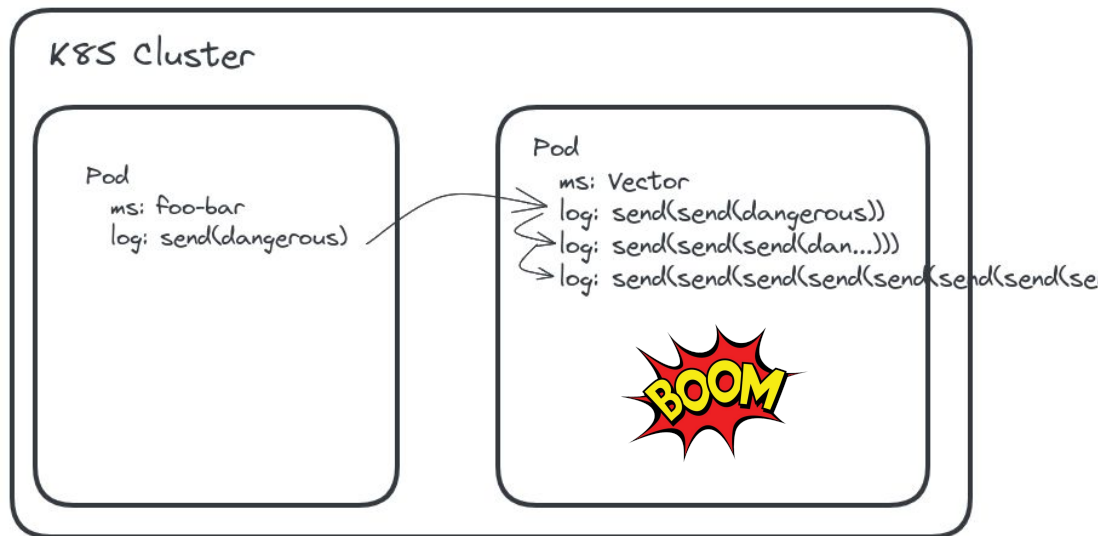
- **Default rules:**
 - Crystal Clear cluster vision
- **Custom rules:**
 - Migration from IMDSv1 to IMDSv2
 - Containerd Socket Observability
 - Red-Team “Blind” Fuzzing
 - WebShells Trusted Alerting
 - Fast detection time for *any event*



Instabilities



- High memory usage of some Falco features
- Log tempest
- Load external kernel module
- Kernel < 5.8 need privileged access
- More complexity



On-Call & Incident Switch-Army-Knife

- Have an historic of weird events during an incident
 - Shell opened in a pod
 - Http request received/sent
 - Files used
- Reproduce and analyse
- Alerts for anomalous activity to on-call Devops
- Monitor service response during security fuzzing



Conclusion & Kudos

Increase Observability & Find More Bugs!

- If you are a pentester you will
 - Start loving breaking things again
 - Benefit from a huge time gain
 - Coder a wider attack surface
- If you are a company you will have
 - A small yet very powerful SoC
 - A powerful way to diagnose anything system-related
 - More findings if you provide this during pentests ! 🎁



System Introspection for Micro-Services Pentests in K8S



@TheLaluka
thinkloveshare.com
offenskill.com



Thank you!