EKS Runtime Security

 Falco Security를 중심으로

HELLO! I'm...

IT 인프라 / 보안에 관심이 많아요.

지금은 아이디어스에서
기술 보안을 담당하고 있어요.

Lee Jiyoung

TABLE OF CONTENTS.

TABLE OF CONTENTS.

01 Basic of Container Security

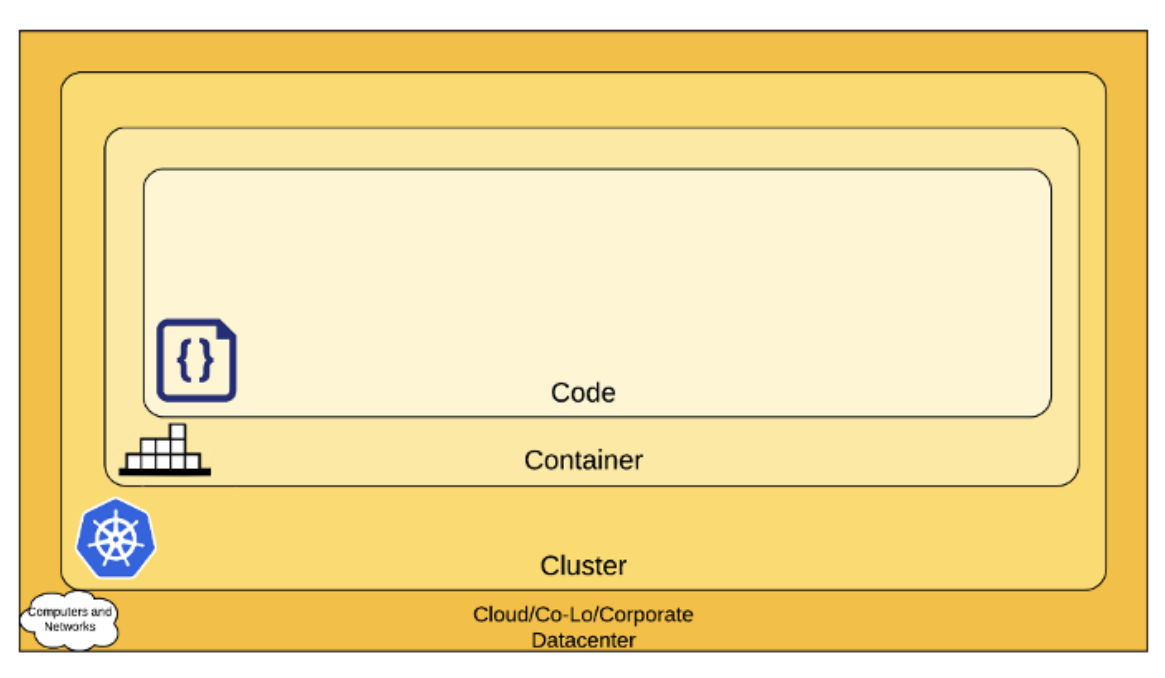

클라우드 네이티브 보안의 4C

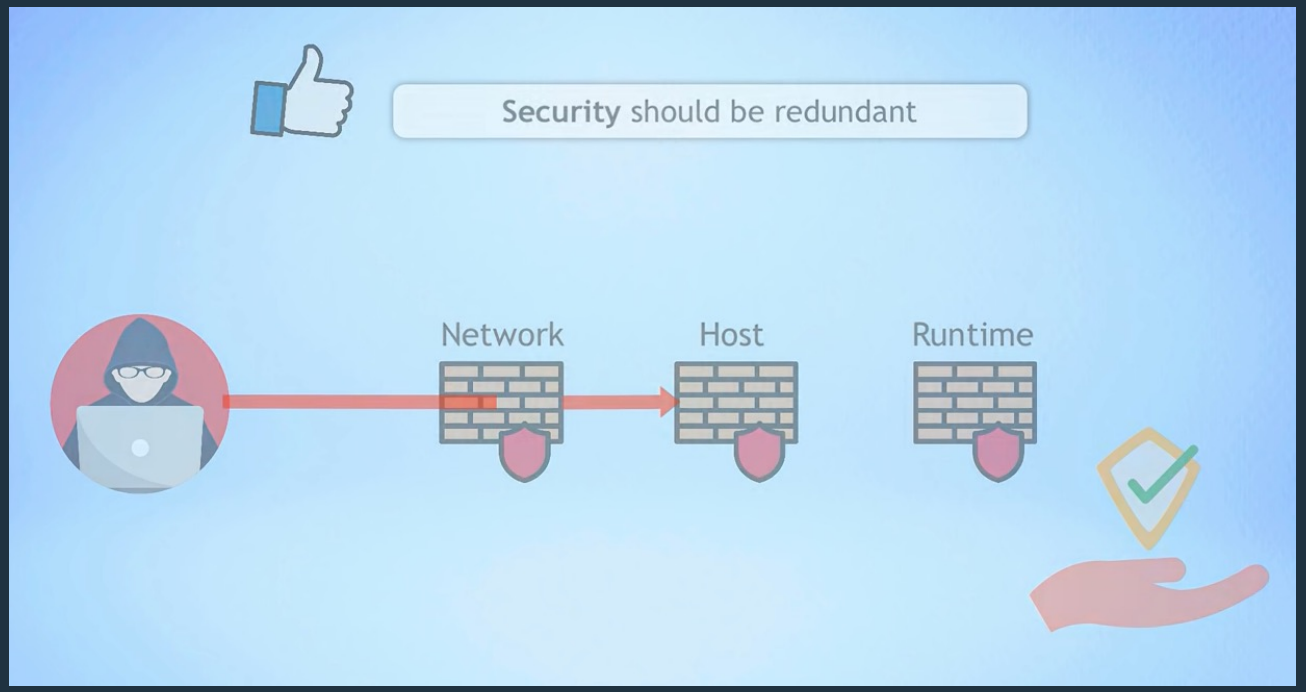

보안은 공격이 일어날 수 있는 모든 레이어에서 이루어져야 함.

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Impact
Using Cloud credentials	Exec into container	Backdoor container	Privileged container	Clear container logs	List K8S secrets	Access the K8S API server	Access cloud resources	Data Destruction
Compromised images in registry	bash/cmd inside container	Writable hostPath mount	Cluster-admin binding	Delete K8S events	Mount service principal	Access Kubelet API	Container service account	Resource Hijacking
Kubeconfig file	New container	Kubernetes CronJob	hostPath mount	Pod / container name similarity	Access container service account	Network mapping	Cluster internal networking	Denial of service
Application vulnerability	Application exploit (RCE)		Access cloud resources	Connect from Proxy server	Applications credentials in configuration files	Access Kubernetes dashboard	Applications credentials in configuration files	
Exposed Dashboard	SSH server running inside container					Instance Metadata API	Writable volume mounts on the host	
							Access Kubernetes dashboard	
							Access tiller endpoint	

https://www.microsoft.com/en-us/security/blog/2020/04/02/attack-matrix-kubernetes/

보안은 공격이 일어날 수 있는 모든 레이어에서 이루어져야 함.

02 What is Runtime Security Enforcement

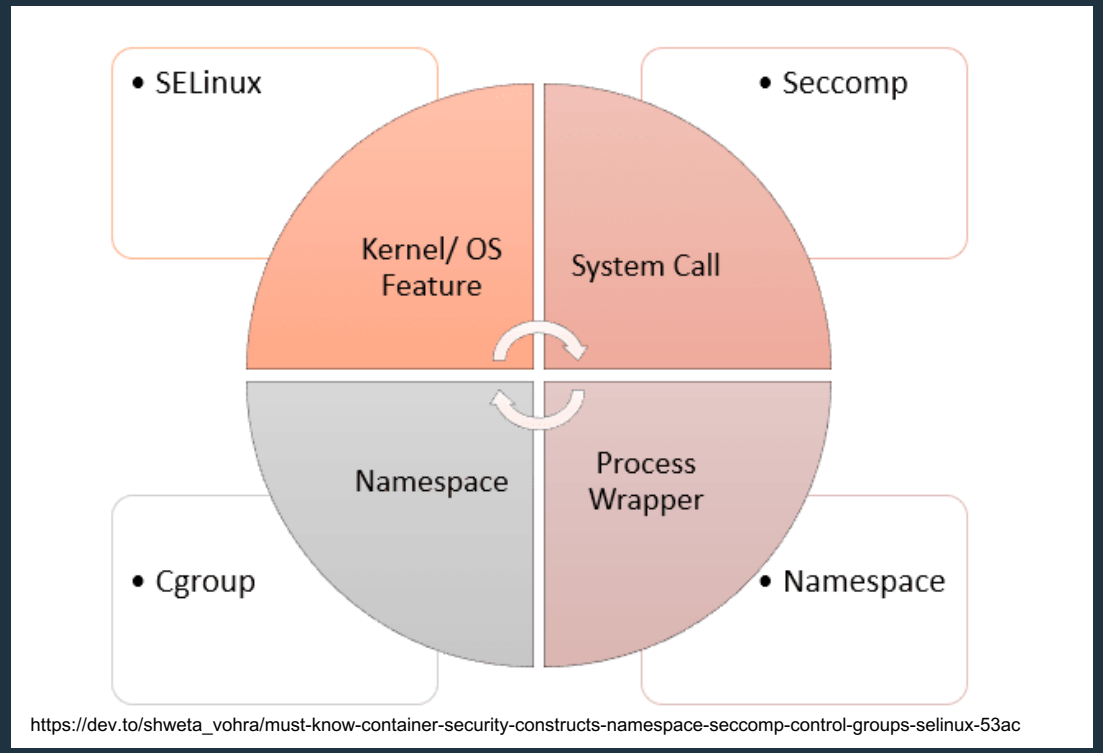

https://dev.to/shweta_vohra/must-know-container-security-constructs-namespace-seccomp-control-groups-selinux-53ac

SELINUX / SECCOMP 좋지만 실제로 적용은 ?

03 What is Runtime Security Audit

컨테이너 / 워크로드에서 실행 시간 동안 일어나는 모든 위협 행위에 대한 보안

04 What Can Falco Detect

- 인가되지 않은 컨테이너의 배포
- 설정 오류로 의도하지 않은 Secret 등의 중요 정보 노출
- 컨테이너의 실행 권한 상승을 통한 취약점 공격 등

- 안전한 권한을 벗어나는 대부분의 설정 탐지

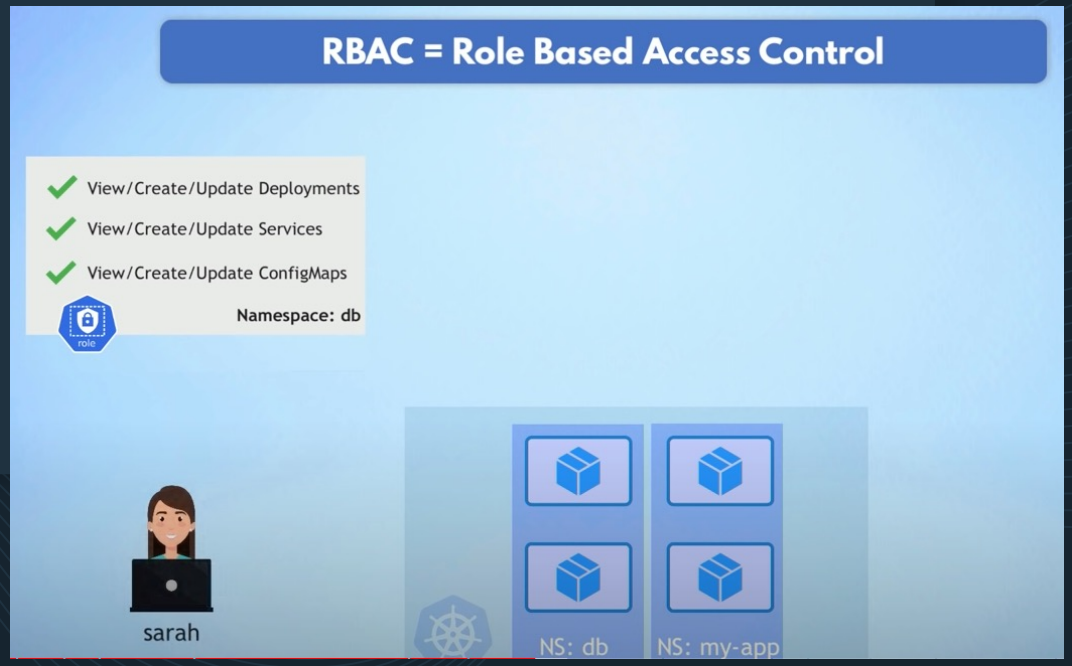

- 누군가 Root(Privileged) 권한으로 컨테이너를 실행한다면 ?
- 누군가 /etc/passwd를 변경한다면?
- 누군가 디바이스가 아닌 파일을 /dev에 파일을 생성한다면 ?
- 누군가 컨테이너에서 쉘을 실행한다면?

05 Why Falco Security

CNCF 유일의
Runtime Security Project,
It's Free

06 How to Install Falco Security

Easy to Deploy

Helm 을 통한 방법을 추천

커맨드 한번에 아래 내용을 한번에 !

- 슬랙 연동을 포함한 Alert 모듈
- UI
- Falcoctl

```
kubectl create namespace falco
helm repo add falcosecurity
https://falcosecurity.github.io/charts
helm install falco falcosecurity/falco ₩
--set falcosidekick.enabled=true ₩
--set falcosidekick.webui.enabled=true ₩
 --set falcoctl.artifact.install.enabled=false ₩
 --set falcoctl.artifact.follow.enabled=false ₩
--set
falcosidekick.config.slack.webhookurl="https://hooks.sl
ack.com/services/T4XXXXXXXXXGsd" ₩
-n falco
```

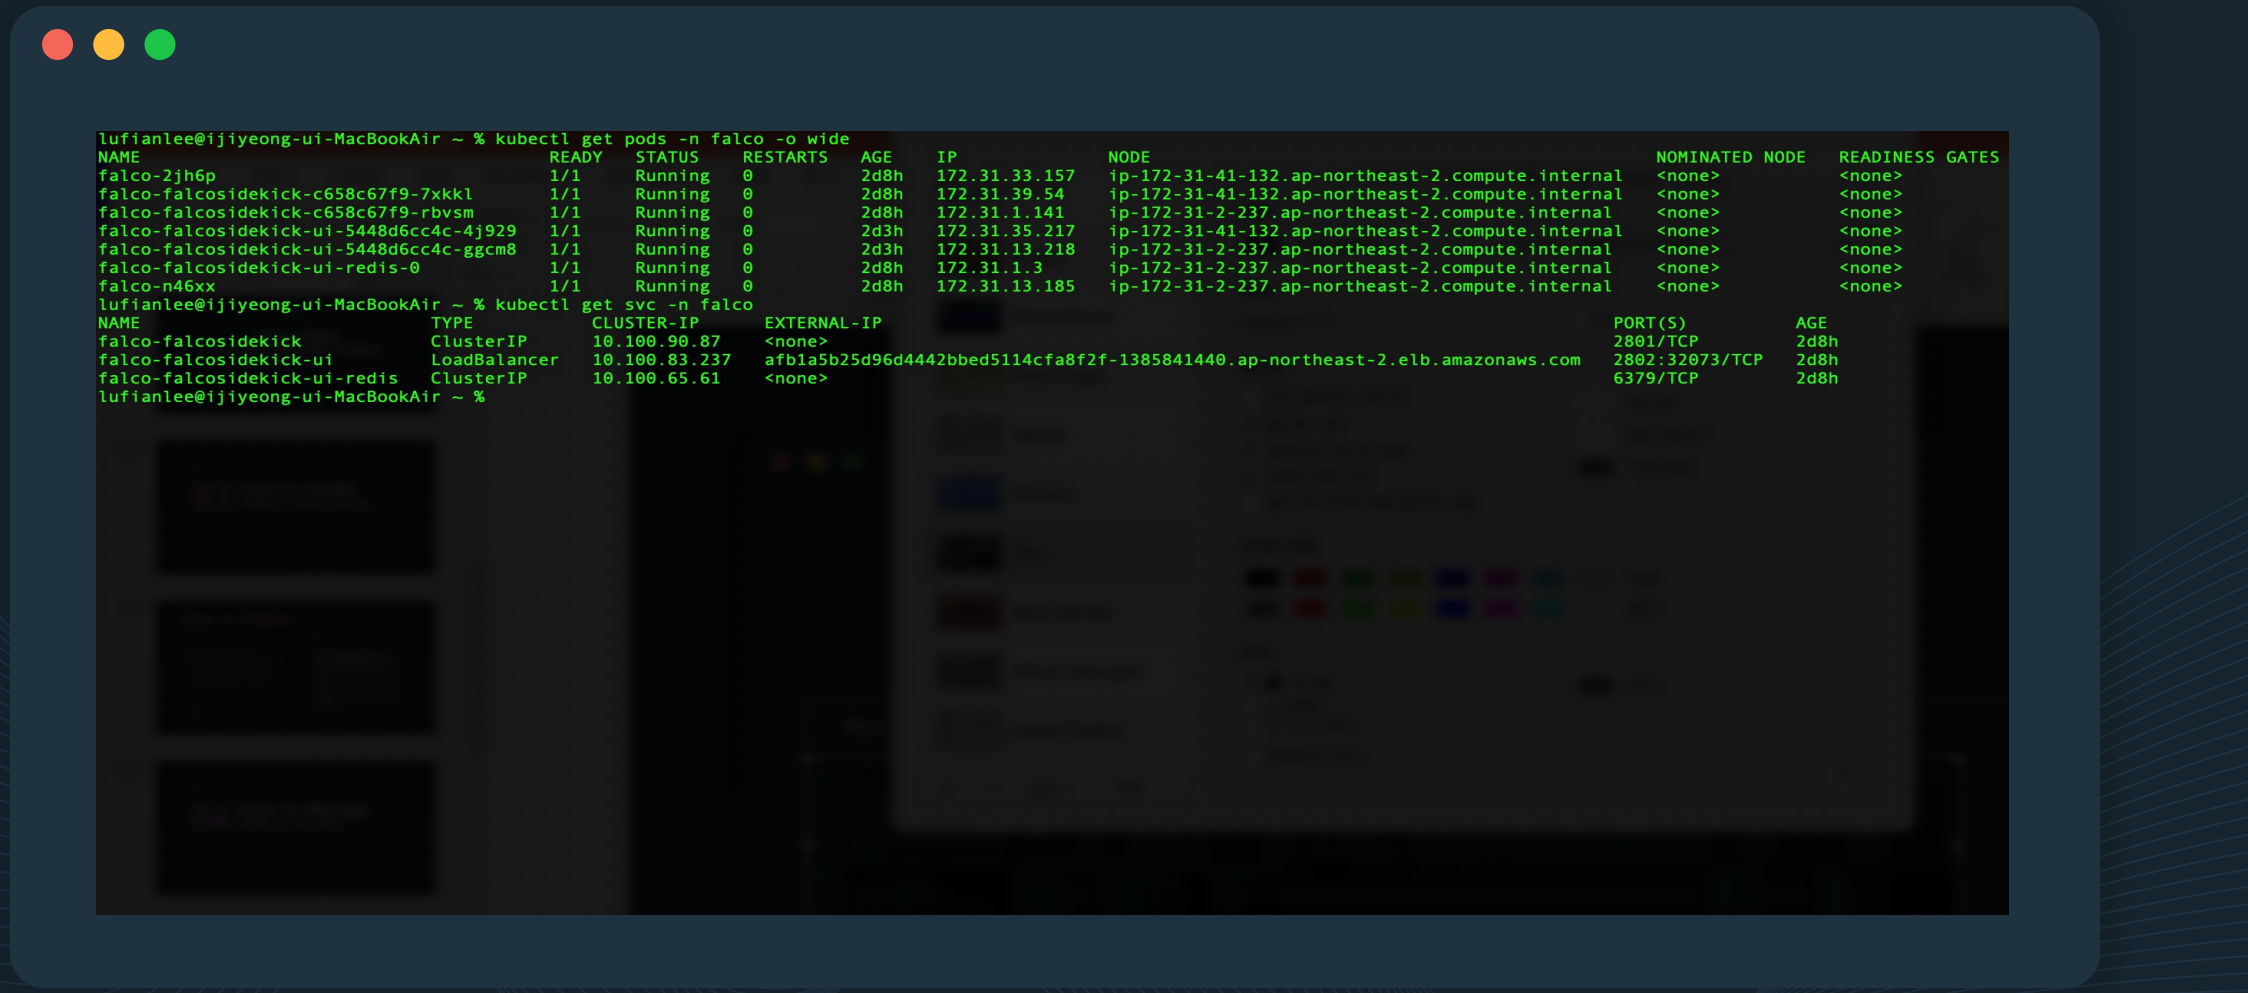
lufianlee@ijiyeong-ui-MacBookAir ~ % kubectl get pods -n falco -o wide
NAME READY STATUS RESTARTS AGE IP NODE NOMINATED NODE READINESS GATES
falco-2jh6p 1/1 Running 0 2d8h 172.31.33.157 ip-172-31-41-132.ap-northeast-2.compute.internal <none> <none>
falco-falcosidekick-c658c67f9-7xkkl 1/1 Running 0 2d8h 172.31.39.54 ip-172-31-41-132.ap-northeast-2.compute.internal <none> <none>
falco-falcosidekick-c658c67f9-rbvsm 1/1 Running 0 2d8h 172.31.1.141 ip-172-31-2-237.ap-northeast-2.compute.internal <none> <none>
falco-falcosidekick-ui-5448d6cc4c-4j929 1/1 Running 0 2d3h 172.31.35.217 ip-172-31-41-132.ap-northeast-2.compute.internal <none> <none>
falco-falcosidekick-ui-5448d6cc4c-ggcm8 1/1 Running 0 2d3h 172.31.13.218 ip-172-31-2-237.ap-northeast-2.compute.internal <none> <none>
falco-falcosidekick-ui-redis-0 1/1 Running 0 2d8h 172.31.1.3 ip-172-31-2-237.ap-northeast-2.compute.internal <none> <none>
falco-n46xx 1/1 Running 0 2d8h 172.31.13.185 ip-172-31-2-237.ap-northeast-2.compute.internal <none> <none>
lufianlee@ijiyeong-ui-MacBookAir ~ % kubectl get svc -n falco
NAME TYPE CLUSTER-IP EXTERNAL-IP PORT(S) AGE
falco-falcosidekick ClusterIP 10.100.90.87 <none> 2801/TCP 2d8h
falco-falcosidekick-ui LoadBalancer 10.100.83.237 afb1a5b25d96d4442bbed5114cfa8f2f-1385841440.ap-northeast-2.elb.amazonaws.com 2802:32073/TCP 2d8h
falco-falcosidekick-ui-redis ClusterIP 10.100.65.61 <none> 6379/TCP 2d8h
lufianlee@ijiyeong-ui-MacBookAir ~ %

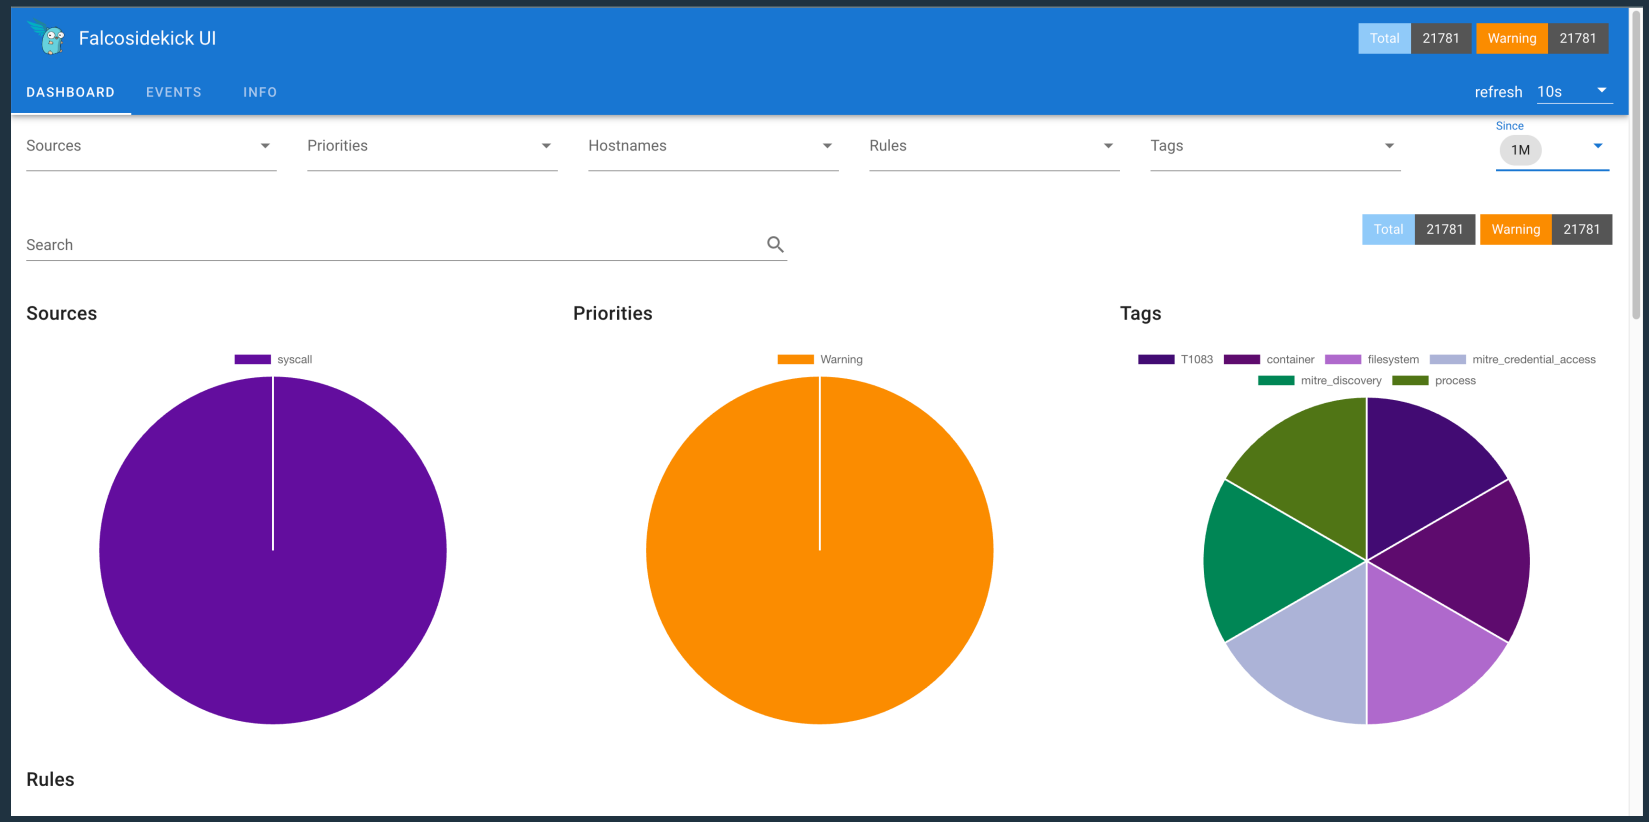
Falcosidekick UI
Total 21781
Warning 21781
DASHBOARD
EVENTS
INFO
refresh 10s
Sources
Priorities
Hostnames
Rules
Tags
Since
1M
Search
Total 21781
Warning 21781
Sources
syscall
Priorities
Warning
Tags
T1083
container
filesystem
mitre_credential_access
mitre_discovery
process
Rules

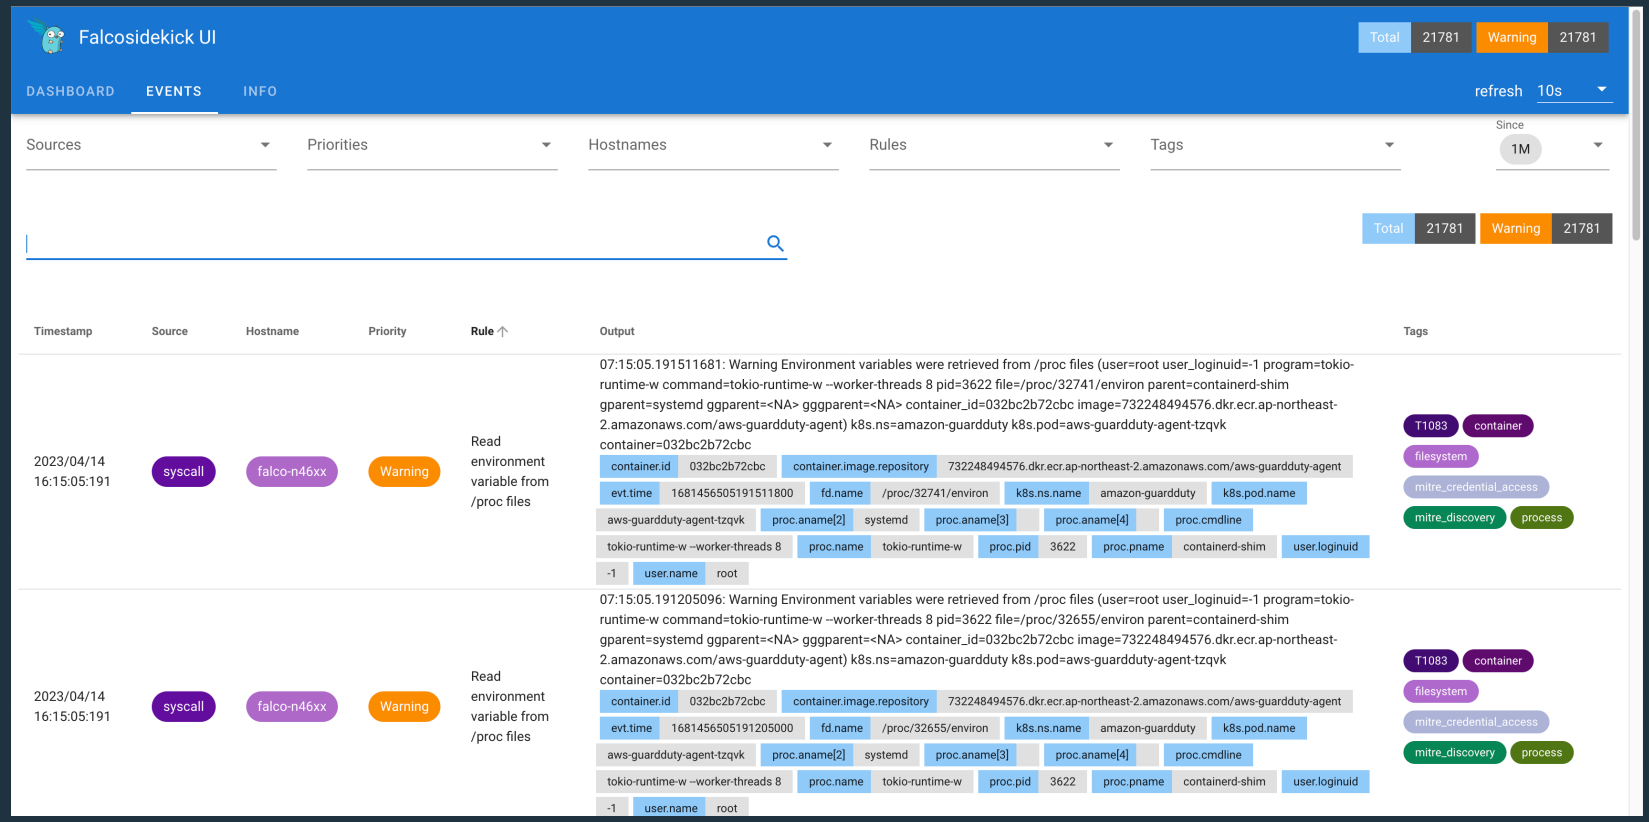
Falcosidekick UI
Total 21781
Warning 21781
DASHBOARD
EVENTS
INFO
refresh 10s
Sources
Priorities
Hostnames
Rules
Tags
Since
1M
Total 21781
Warning 21781
Timestamp
Source
Hostname
Priority
Rule
Output
Tags
2023/04/14 16:15:05:191
syscall
falco-n46xx
Warning
Read environment variable from /proc files
07:15:05.191511681: Warning Environment variables were retrieved from /proc files (user=root user_loginuid=-1 program=tokio-runtime-w command=tokio-runtime-w --worker-threads 8 pid=3622 file=/proc/32741/environ parent=containerd-shim gparent=systemd ggparent=<NA> gggparent=<NA> container_id=032bc2b72cbc image=732248494576.dkr.ecr.ap-northeast-2.amazonaws.com/aws-guardduty-agent) k8s.ns=amazon-guardduty k8s.pod=aws-guardduty-agent-tzqvk container=032bc2b72cbc
container.id 032bc2b72cbc
container.image.repository 732248494576.dkr.ecr.ap-northeast-2.amazonaws.com/aws-guardduty-agent
evt.time 1681456505191511800
fd.name /proc/32741/environ
k8s.ns.name amazon-guardduty
k8s.pod.name aws-guardduty-agent-tzqvk
proc.aname[2] systemd
proc.aname[3]
proc.aname[4]
proc.cmdline tokio-runtime-w --worker-threads 8
proc.name tokio-runtime-w
proc.pid 3622
proc.pname containerd-shim
user.loginuid -1
user.name root
T1083
container
filesystem
mitre_credential_access
mitre_discovery
process
2023/04/14 16:15:05:191
syscall
falco-n46xx
Warning
Read environment variable from /proc files
07:15:05.191205096: Warning Environment variables were retrieved from /proc files (user=root user_loginuid=-1 program=tokio-runtime-w command=tokio-runtime-w --worker-threads 8 pid=3622 file=/proc/32655/environ parent=containerd-shim gparent=systemd ggparent=<NA> gggparent=<NA> container_id=032bc2b72cbc image=732248494576.dkr.ecr.ap-northeast-2.amazonaws.com/aws-guardduty-agent) k8s.ns=amazon-guardduty k8s.pod=aws-guardduty-agent-tzqvk container=032bc2b72cbc
container.id 032bc2b72cbc
container.image.repository 732248494576.dkr.ecr.ap-northeast-2.amazonaws.com/aws-guardduty-agent
evt.time 1681456505191205000
fd.name /proc/32655/environ
k8s.ns.name amazon-guardduty
k8s.pod.name aws-guardduty-agent-tzqvk
proc.aname[2] systemd
proc.aname[3]
proc.aname[4]
proc.cmdline tokio-runtime-w --worker-threads 8
proc.name tokio-runtime-w
proc.pid 3622
proc.pname containerd-shim
user.loginuid -1
user.name root
T1083
container
filesystem
mitre_credential_access
mitre_discovery
process

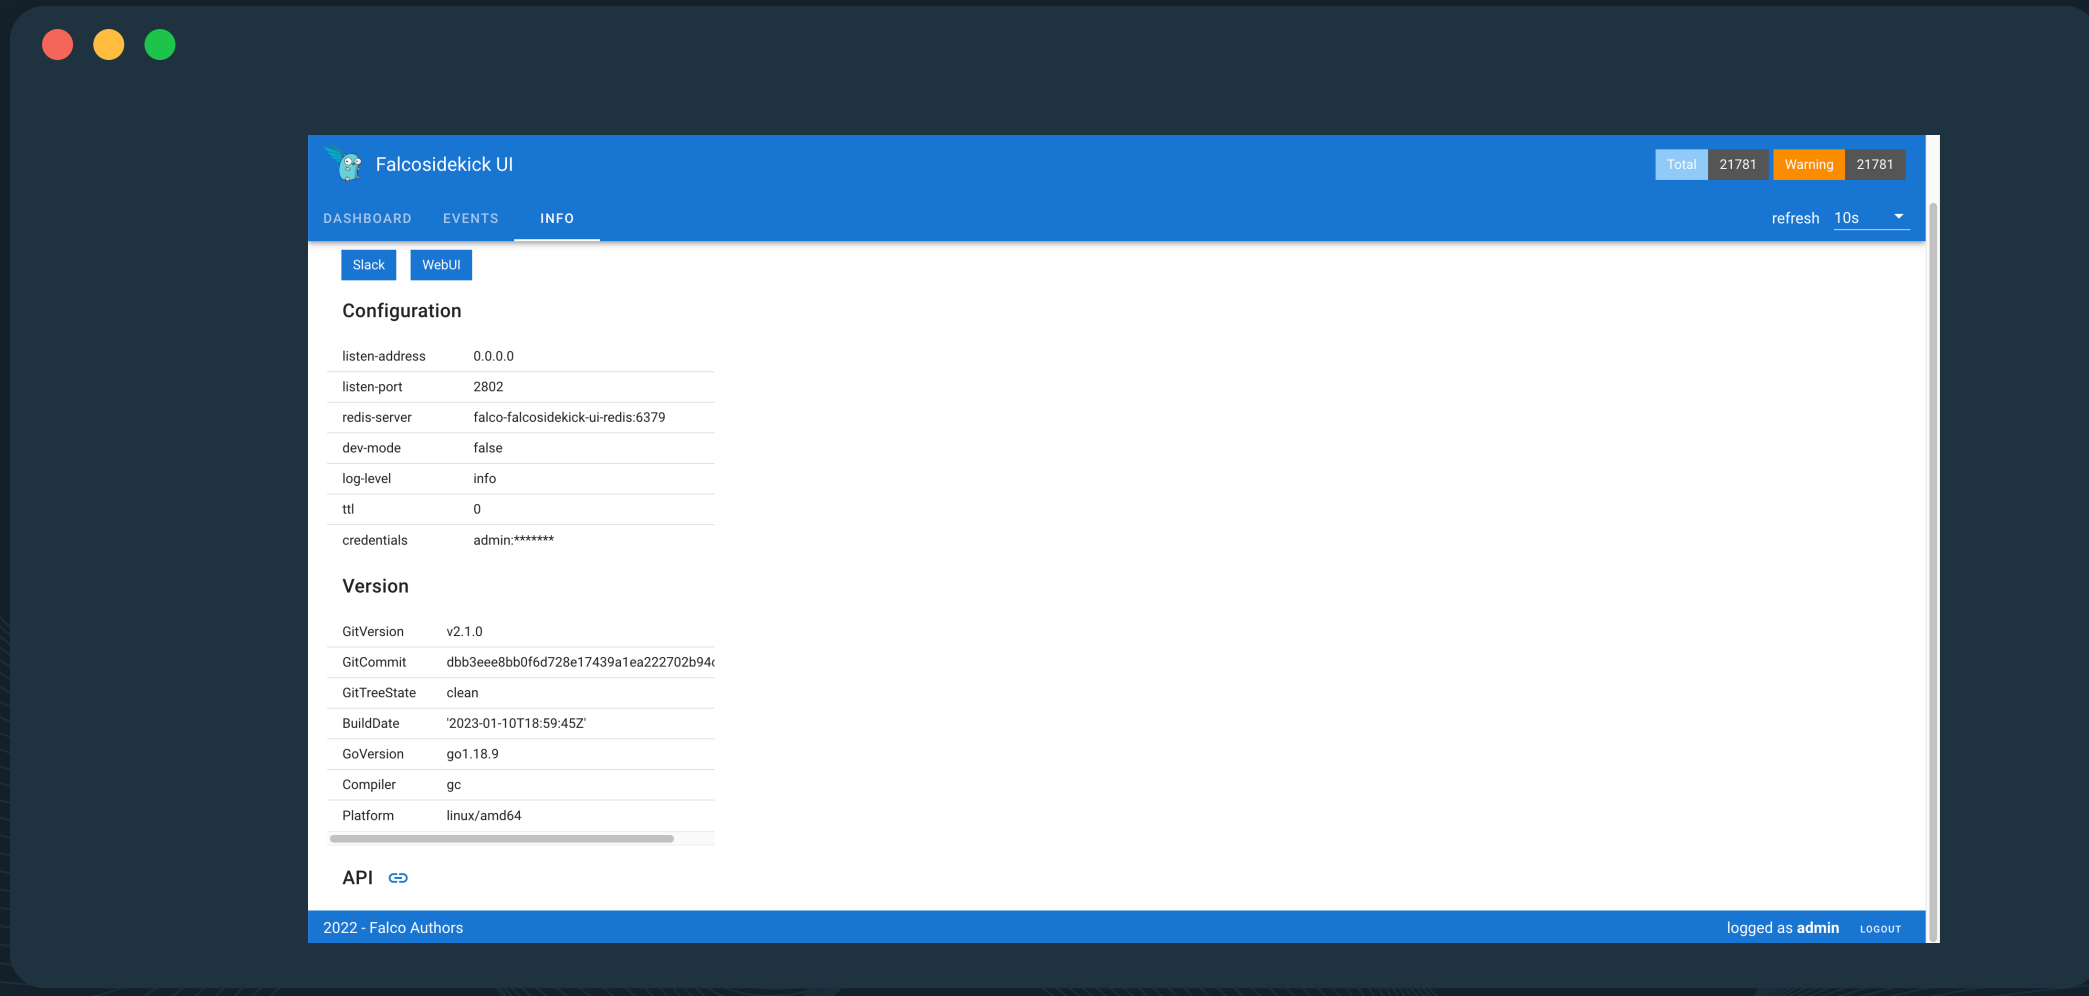
Falcosidekick UI
Total 21781
Warning 21781
DASHBOARD
EVENTS
INFO
refresh 10s
Slack
WebUI
Configuration
listen-address 0.0.0.0
listen-port 2802
redis-server falco-falcosidekick-ui-redis:6379
dev-mode false
log-level info
ttl 0
credentials admin:*******
Version
GitVersion v2.1.0
GitCommit dbb3eee8bb0f6d728e17439a1ea222702b94
GitTreeState clean
BuildDate '2023-01-10T18:59:45Z'
GoVersion go1.18.9
Compiler gc
Platform linux/amd64
API
2022 - Falco Authors
logged as admin
LOGOUT

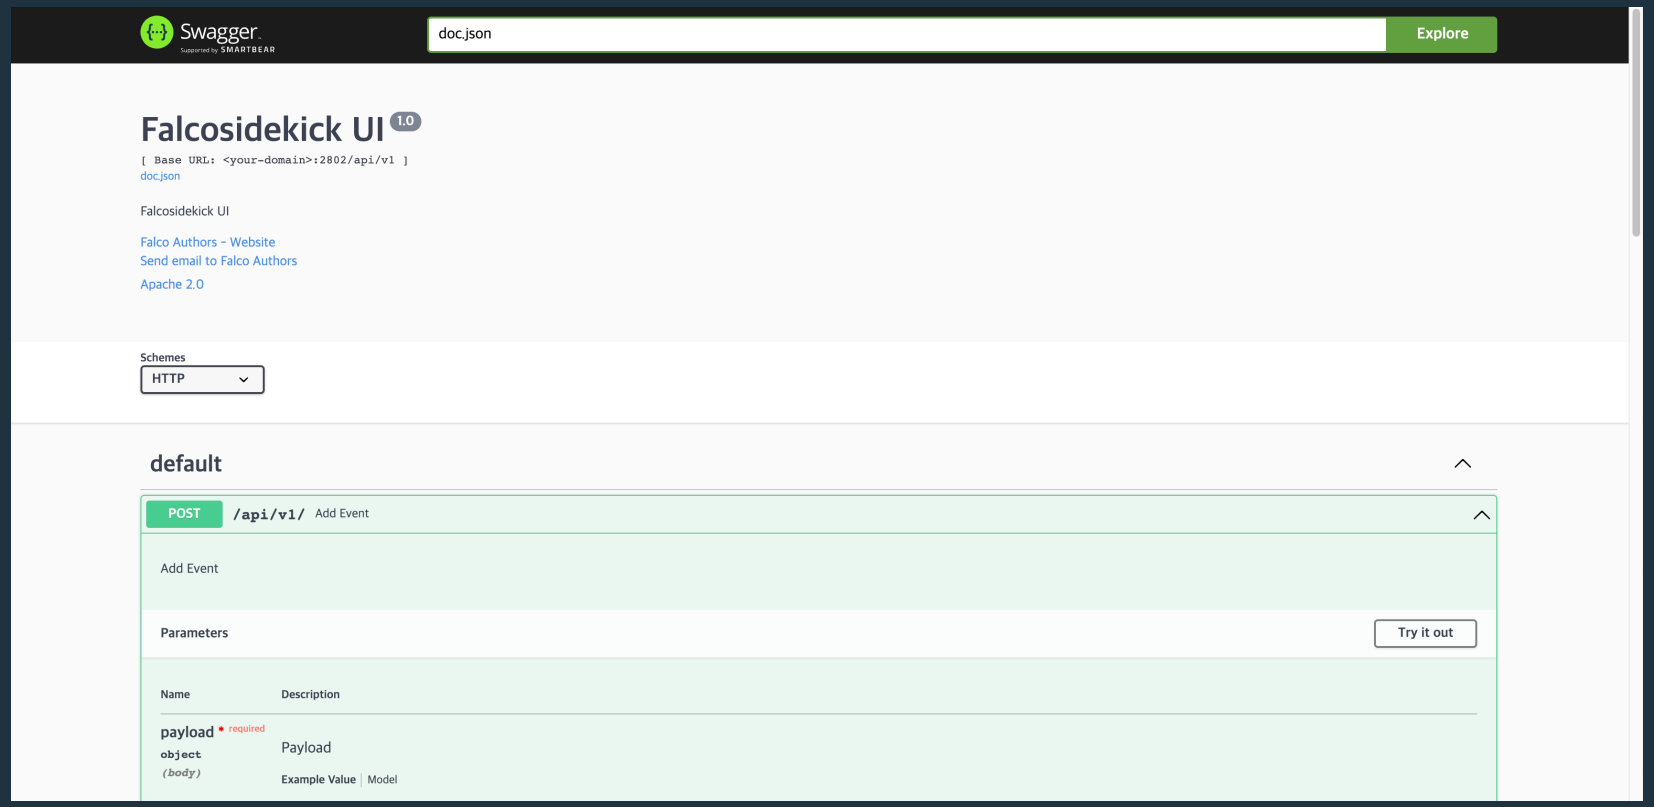
Swagger
Supported by SMARTBEAR
doc.json
Explore
Falcosidekick UI
1.0
[Base URL: <your-domain>:2802/api/v1]
doc.json
Falcosidekick UI
Falco Authors - Website
Send email to Falco Authors
Apache 2.0
Schemes
HTTP
default
POST
/api/v1/
Add Event
Add Event
Parameters
Try it out
Name
Description
payload
required
object
(body)
Payload
Example Value
Model

9:20

12:20:58.345731618: Informational Excessively capable container started (user=root user_loginuid=-1 command=sh -c node server.js pid=4705 k8s.ns=default k8s.pod=security-context-demo-4 container=82219f76c042 image=gcr.io/google-samples/node-hello:1.0 cap_permitted=CAP_CHOWN CAP_DAC_OVERRIDE CAP_FOWNER CAP_FSETID CAP_KILL CAP_SETGID CAP_SETUID CAP_SETPCAP CAP_NET_BIND_SERVICE CAP_NET_ADMIN CAP_NET_RAW CAP_SYS_CHROOT CAP_SYS_TIME CAP_MKNOD CAP_AUDIT_WRITE CAP_SETFCAP)

rule
Launch Excessively Capable Container

priority
Informational

source
syscall

hostname
falco-n46xx

tags
T1610, cis, container, mitre_lateral_movement, mitre_privilege_escalation

container.id
82219f76c042

container.image.repository
gcr.io/google-samples/node-hello

container.image.tag
1.0

k8s.ns.name
default

k8s.pod.name
security-context-demo-4

proc.cmdline
sh -c node server.js

thread.cap_permitted
CAP_CHOWN CAP_DAC_OVERRIDE CAP_FOWNER CAP_FSETID CAP_KILL CAP_SETGID CAP_SETUID CAP_SETPCAP CAP_NET_BIND_SERVICE CAP_NET_ADMIN CAP_NET_RAW CAP_SYS_CHROOT CAP_SYS_TIME CAP_MKNOD CAP_AUDIT_WRITE CAP_SETFCAP

user.name
root

time
2023-04-19 12:20:58.345731618 +0000 UTC

https://github.com/falcosecurity/falcosidekick

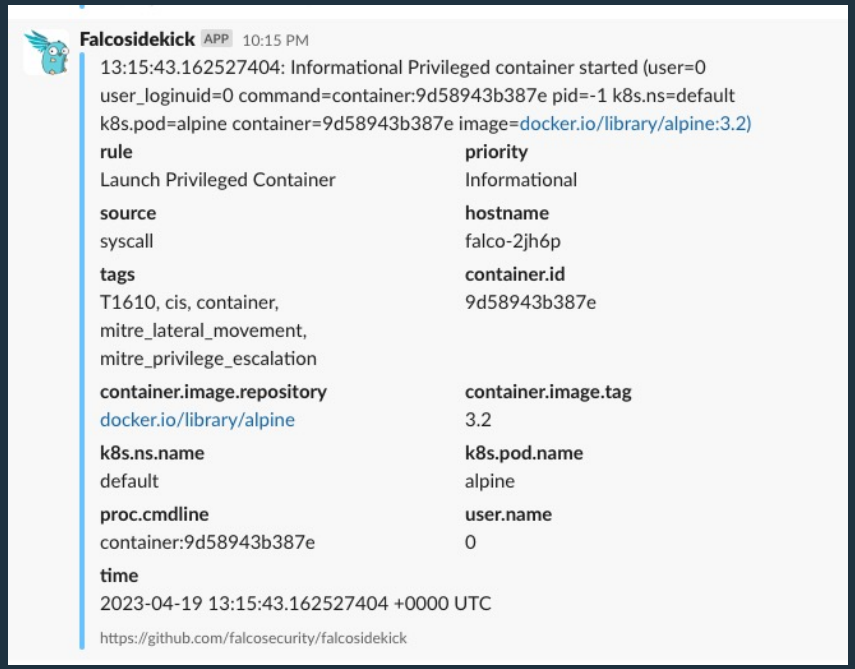
Falcosidekick APP 10:15 PM
13:15:43.162527404: Informational Privileged container started (user=0 user_loginuid=0 command=container:9d58943b387e pid=-1 k8s.ns=default k8s.pod=alpine container=9d58943b387e image=docker.io/library/alpine:3.2)
rule
Launch Privileged Container
priority
Informational
source
syscall
hostname
falco-2jh6p
tags
T1610, cis, container, mitre_lateral_movement, mitre_privilege_escalation
container.id
9d58943b387e
container.image.repository
docker.io/library/alpine
container.image.tag
3.2
k8s.ns.name
default
k8s.pod.name
alpine
proc.cmdline
container:9d58943b387e
user.name
0
time
2023-04-19 13:15:43.162527404 +0000 UTC
https://github.com/falcosecurity/falcosidekick

07 Falco Plugins

Audit from Event Sources

K8saudit
Cloudtrail
Docker
Okta
Github
Json

.
.
.

지원되는 플러그인이 계속 늘어나는 중

08 How to Manage Falco Rules

https://falco.org/docs/rules/

```
- rule: shell_in_container
  desc: notice shell activity within a container
  condition: >
    evt.type = execve and
    evt.dir = < and
    container.id != host and
    (proc.name = bash or
     proc.name = ksh)
  output: >
    shell in a container
    (user=%user.name container_id=%container.id container_name=%container.name
    shell=%proc.name parent=%proc.pname cmdline=%proc.cmdline)
  priority: WARNING
```

```
- list: shell_binaries
  items: [bash, csh, ksh, sh, tcsh, zsh, dash]
```

```
- macro: container
  condition: container.id != host
```

Rule, list, macro만 알면 끝 !

/etc/falco/falco_rules.yaml

```yaml
- list: my_programs
  items: [ls, cat, pwd]

- rule: my_programs_opened_file
  desc: track whenever a set of programs opens a file
  condition: proc.name in (my_programs) and (evt.type=open or evt.type=openat)
  output: a tracked program opened a file (user=%user.name command=%proc.cmdline file=%fd.name)
  priority: INFO
```

/etc/falco/falco_rules.local.yaml

```yaml
- list: my_programs
  append: true
  items: [cp]
```

기본 rule은 falco_rules.local.yaml에서 수정(기본 falco_rules.yaml은 수정X)

```
customRules:
  rules-traefik.yaml: |-
    - macro: traefik_consider_syscalls
      condition: (evt.num < 0)

    - macro: app_traefik
      condition: container and container.image startswith "traefik"

    # Restricting listening ports to selected set

    - list: traefik_allowed_inbound_ports_tcp
      items: [443, 80, 8080]

    - rule: Unexpected inbound tcp connection traefik
      desc: Detect inbound traffic to traefik using tcp on a port outside of expected set
      condition: inbound and evt.rawres >= 0 and not fd.sport in (traefik_allowed_inbound_ports_tcp) and app_traefik
      output: Inbound network connection to traefik on unexpected port (command=%proc.cmdline pid=%proc.pid connection=%f
      priority: NOTICE

    # Restricting spawned processes to selected set

    - list: traefik_allowed_processes
      items: ["traefik"]

    - rule: Unexpected spawned process traefik
      desc: Detect a process started in a traefik container outside of an expected set
      condition: spawned_process and not proc.name in (traefik_allowed_processes) and app_traefik
      output: Unexpected process spawned in traefik container (command=%proc.cmdline pid=%proc.pid user=%user.name %conta
      priority: NOTICE
```

helm install falco –f custom_rule.yaml

09 Guardduty EKS Runtime Monitoring

Falco Vs EKS Runtime Monitoring

[FALCO]

- Free
- Customize rules
- Plugin을 통한 다양한 소스에서 위협 탐지

[EKS Runtime Monitoring]

- AWS Native Service
- 해당 기능만 단독으로 사용 불가
- 가드듀티 기본 기능을 반드시 사용해야 하므로 비용 부담

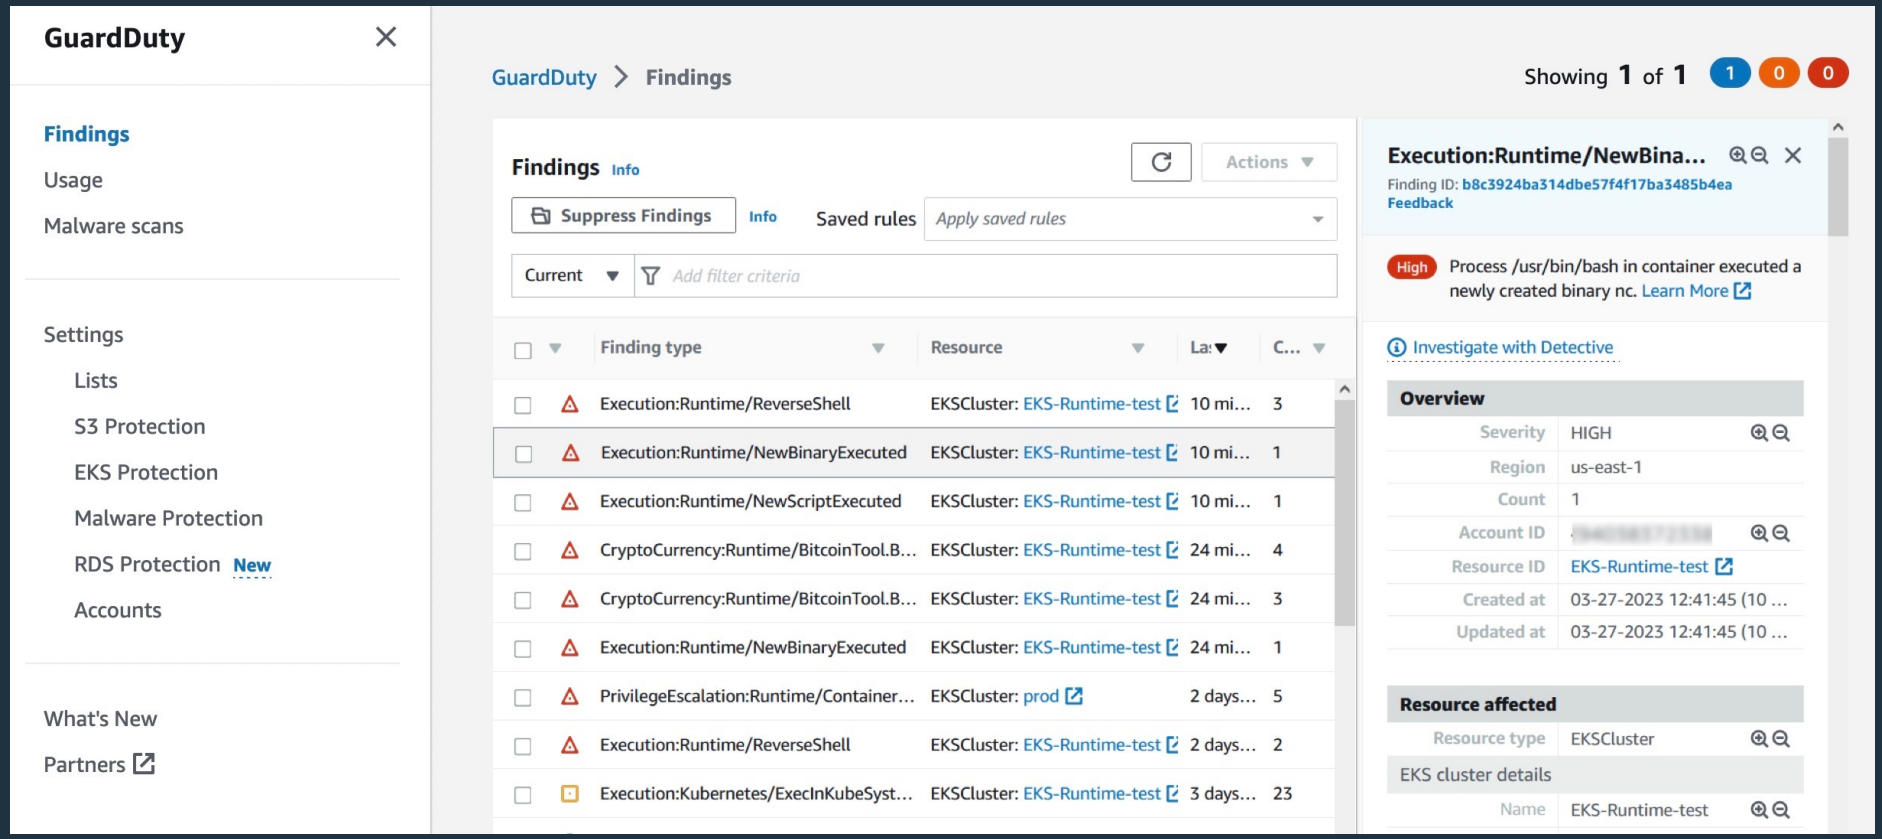
GuardDuty
Findings
Usage
Malware scans
Settings
Lists
S3 Protection
EKS Protection
Malware Protection
RDS Protection New
Accounts
What's New
Partners
GuardDuty > Findings
Showing 1 of 1
1 0 0
Findings Info
Actions
Suppress Findings Info
Saved rules
Apply saved rules
Current
Add filter criteria
Finding type
Resource
La:
C...
Execution:Runtime/ReverseShell EKSCluster: EKS-Runtime-test 10 mi... 3
Execution:Runtime/NewBinaryExecuted EKSCluster: EKS-Runtime-test 10 mi... 1
Execution:Runtime/NewScriptExecuted EKSCluster: EKS-Runtime-test 10 mi... 1
CryptoCurrency:Runtime/BitcoinTool.B... EKSCluster: EKS-Runtime-test 24 mi... 4
CryptoCurrency:Runtime/BitcoinTool.B... EKSCluster: EKS-Runtime-test 24 mi... 3
Execution:Runtime/NewBinaryExecuted EKSCluster: EKS-Runtime-test 24 mi... 1
PrivilegeEscalation:Runtime/Container... EKSCluster: prod 2 days... 5
Execution:Runtime/ReverseShell EKSCluster: EKS-Runtime-test 2 days... 2
Execution:Kubernetes/ExecInKubeSyst... EKSCluster: EKS-Runtime-test 3 days... 23
Execution:Runtime/NewBina...
Finding ID: b8c3924ba314dbe57f4f17ba3485b4ea
Feedback
High Process /usr/bin/bash in container executed a newly created binary nc. Learn More
Investigate with Detective
Overview
Severity HIGH
Region us-east-1
Count 1
Account ID
Resource ID EKS-Runtime-test
Created at 03-27-2023 12:41:45 (10 ...
Updated at 03-27-2023 12:41:45 (10 ...
Resource affected
Resource type EKSCluster
EKS cluster details
Name EKS-Runtime-test

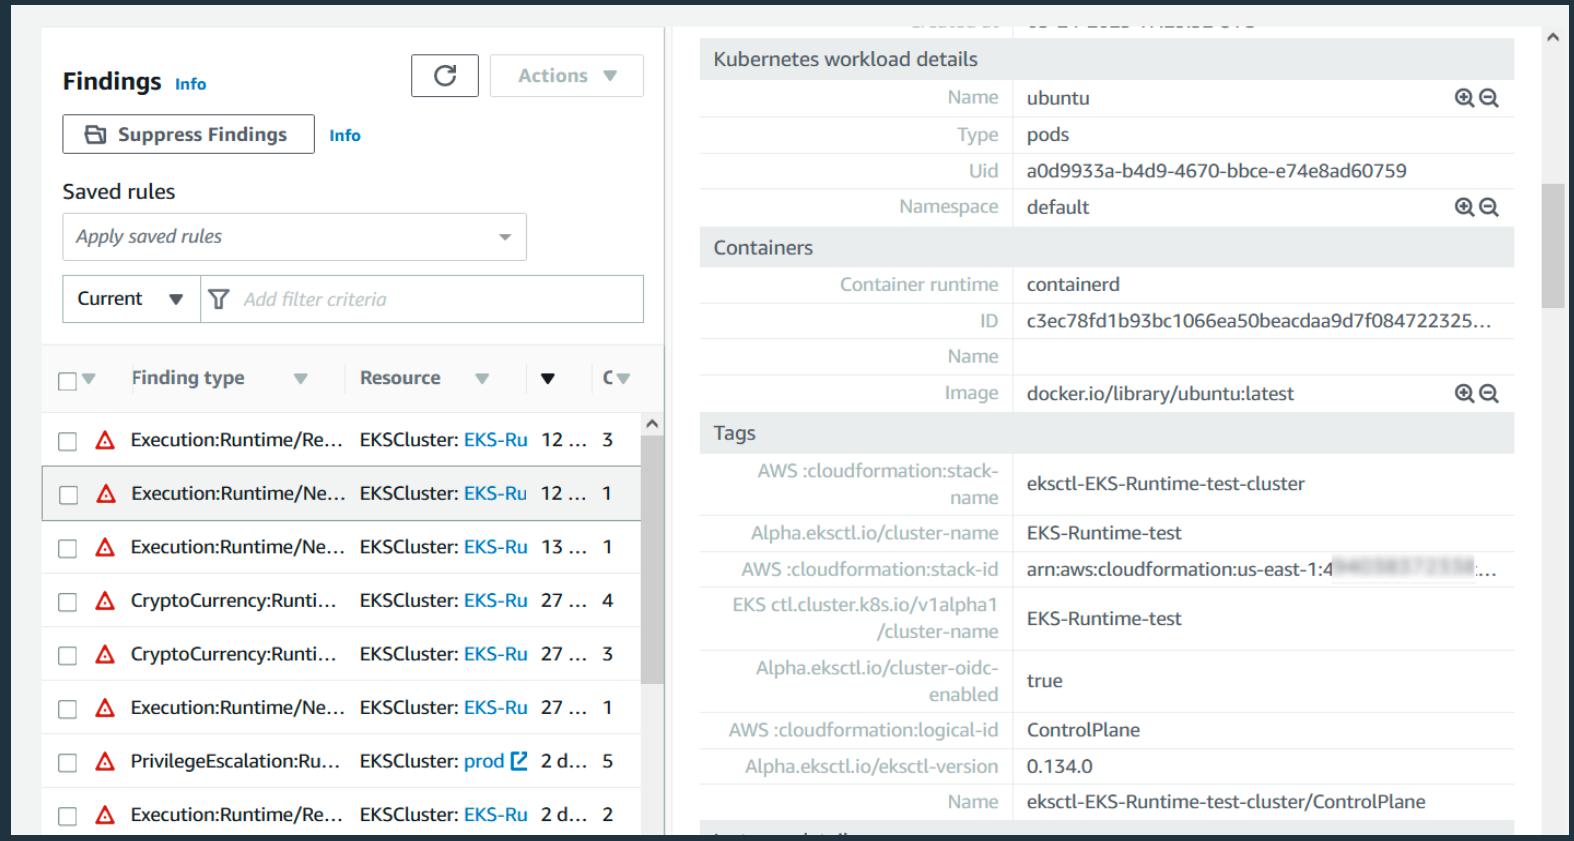
Findings Info
Actions
Suppress Findings
Info
Saved rules
Apply saved rules
Current
Add filter criteria
Finding type
Resource
Execution:Runtime/Re... EKSCluster: EKS-Ru 12 ... 3
Execution:Runtime/Ne... EKSCluster: EKS-Ru 12 ... 1
Execution:Runtime/Ne... EKSCluster: EKS-Ru 13 ... 1
CryptoCurrency:Runti... EKSCluster: EKS-Ru 27 ... 4
CryptoCurrency:Runti... EKSCluster: EKS-Ru 27 ... 3
Execution:Runtime/Ne... EKSCluster: EKS-Ru 27 ... 1
PrivilegeEscalation:Ru... EKSCluster: prod 2 d... 5
Execution:Runtime/Re... EKSCluster: EKS-Ru 2 d... 2
Kubernetes workload details
Name ubuntu
Type pods
Uid a0d9933a-b4d9-4670-bbce-e74e8ad60759
Namespace default
Containers
Container runtime containerd
ID c3ec78fd1b93bc1066ea50beacdaa9d7f084722325...
Name
Image docker.io/library/ubuntu:latest
Tags
AWS :cloudformation:stack-name eksctl-EKS-Runtime-test-cluster
Alpha.eksctl.io/cluster-name EKS-Runtime-test
AWS :cloudformation:stack-id arn:aws:cloudformation:us-east-1:4...
EKS ctl.cluster.k8s.io/v1alpha1/cluster-name EKS-Runtime-test
Alpha.eksctl.io/cluster-oidc-enabled true
AWS :cloudformation:logical-id ControlPlane
Alpha.eksctl.io/eksctl-version 0.134.0
Name eksctl-EKS-Runtime-test-cluster/ControlPlane

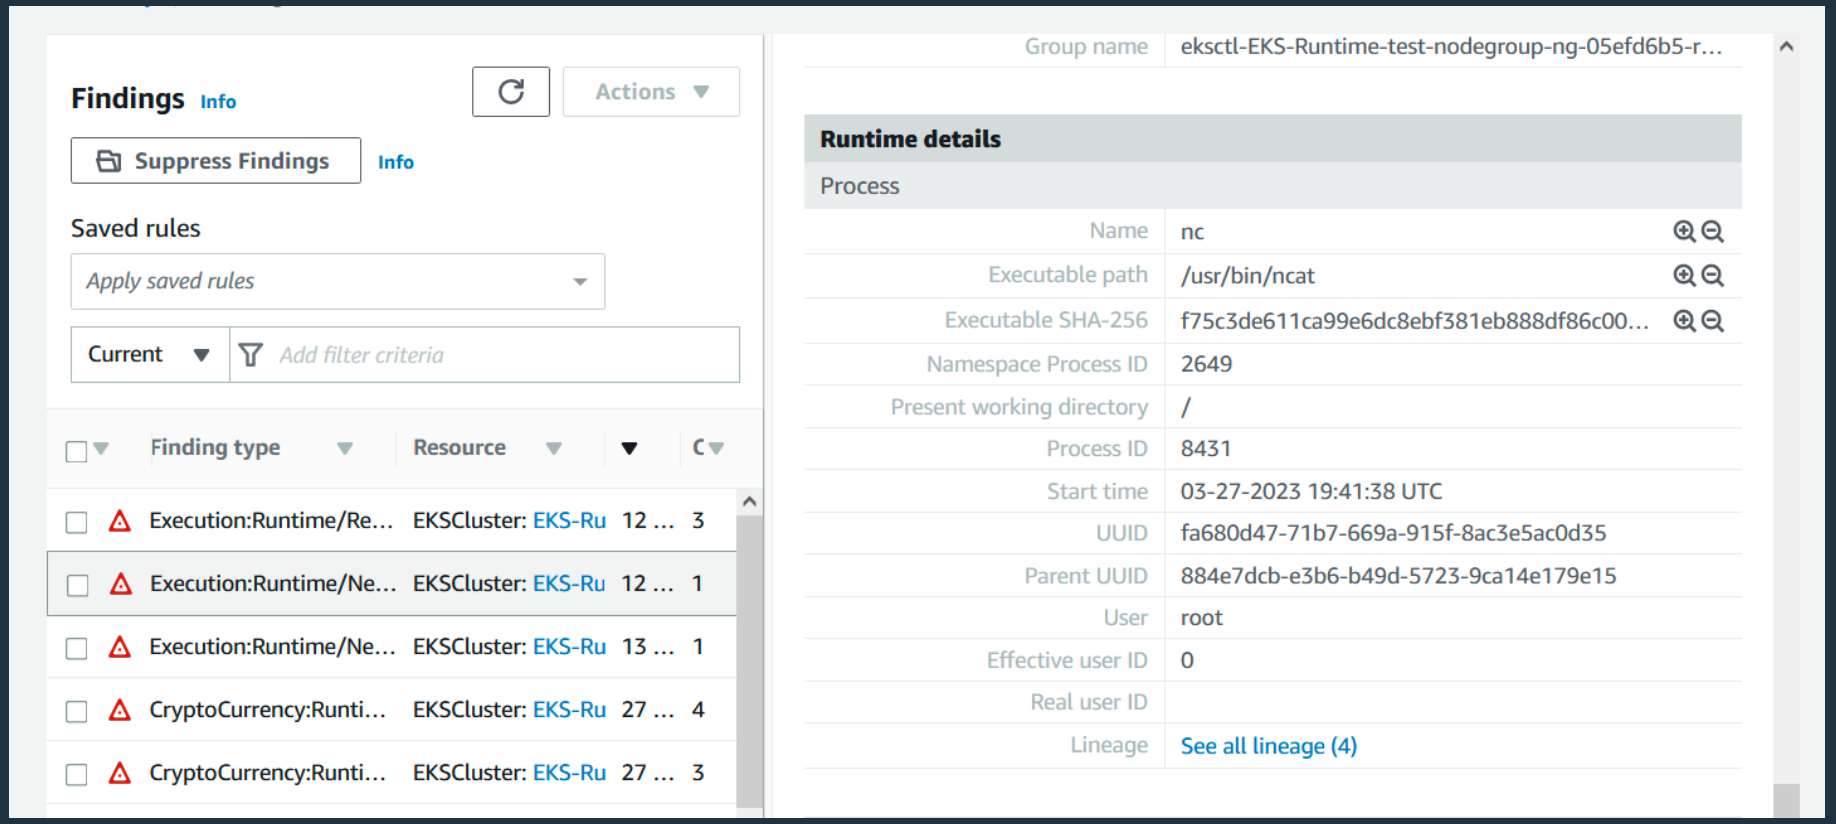
Findings Info
Actions
Suppress Findings
Info
Saved rules
Apply saved rules
Current
Add filter criteria
Finding type
Resource
Execution:Runtime/Re…
EKSCluster: EKS-Ru
Execution:Runtime/Ne…
EKSCluster: EKS-Ru
Execution:Runtime/Ne…
EKSCluster: EKS-Ru
CryptoCurrency:Runti…
EKSCluster: EKS-Ru
CryptoCurrency:Runti…
EKSCluster: EKS-Ru
Group name
eksctl-EKS-Runtime-test-nodegroup-ng-05efd6b5-r…
Runtime details
Process
Name
nc
Executable path
/usr/bin/ncat
Executable SHA-256
f75c3de611ca99e6dc8ebf381eb888df86c00…
Namespace Process ID
2649
Present working directory
/
Process ID
8431
Start time
03-27-2023 19:41:38 UTC
UUID
fa680d47-71b7-669a-915f-8ac3e5ac0d35
Parent UUID
884e7dcb-e3b6-b49d-5723-9ca14e179e15
User
root
Effective user ID
0
Real user ID
Lineage
See all lineage (4)

THANK YOU!

Do you have any questions?

질문은 언제든

AWS KRUG Slack 채널
#Security