



KubeCon



CloudNativeCon

Europe 2020

Virtual

Help! My Cluster Is On The Internet!

Container Security Fundamentals

Samuel Davidson

[linkedin.com/in/samuelbdavidson](https://www.linkedin.com/in/samuelbdavidson)

www.sambdavidson.com



About Me



Samuel Davidson

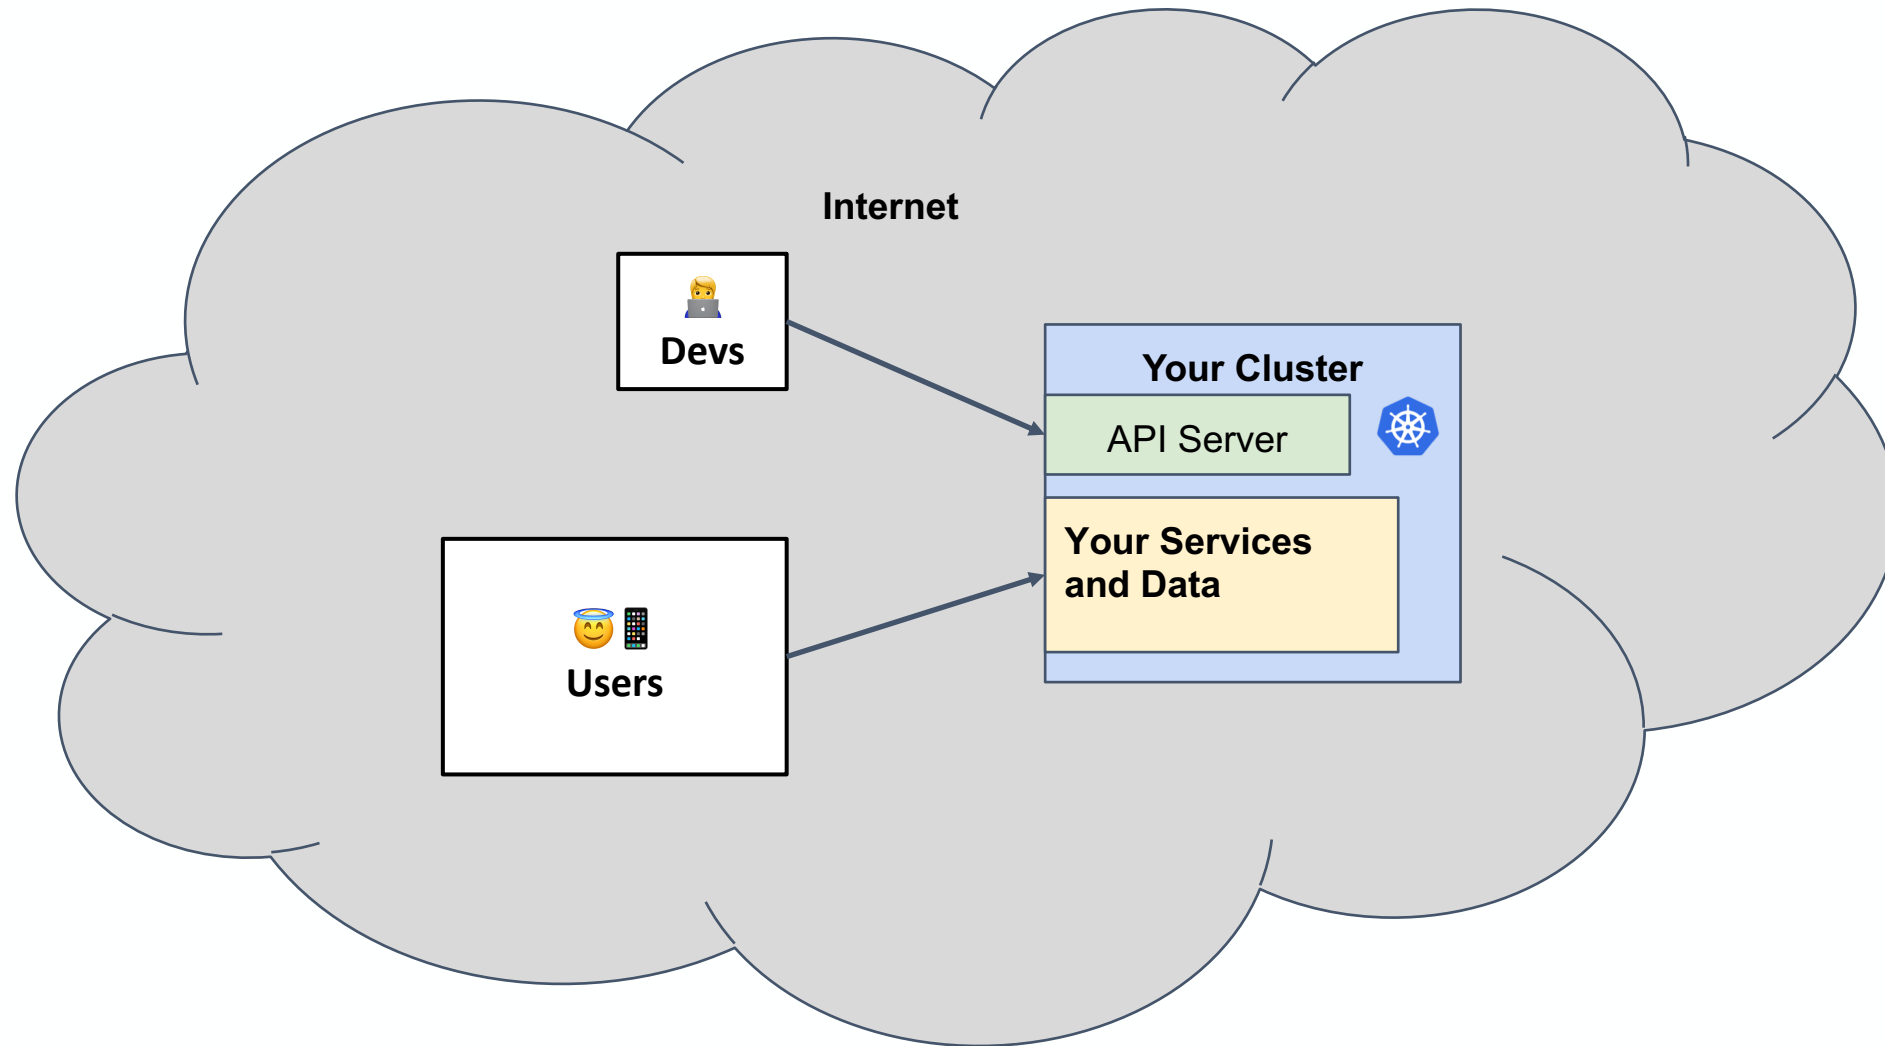
Google Kubernetes Engine (GKE) Security for 2.5 years.



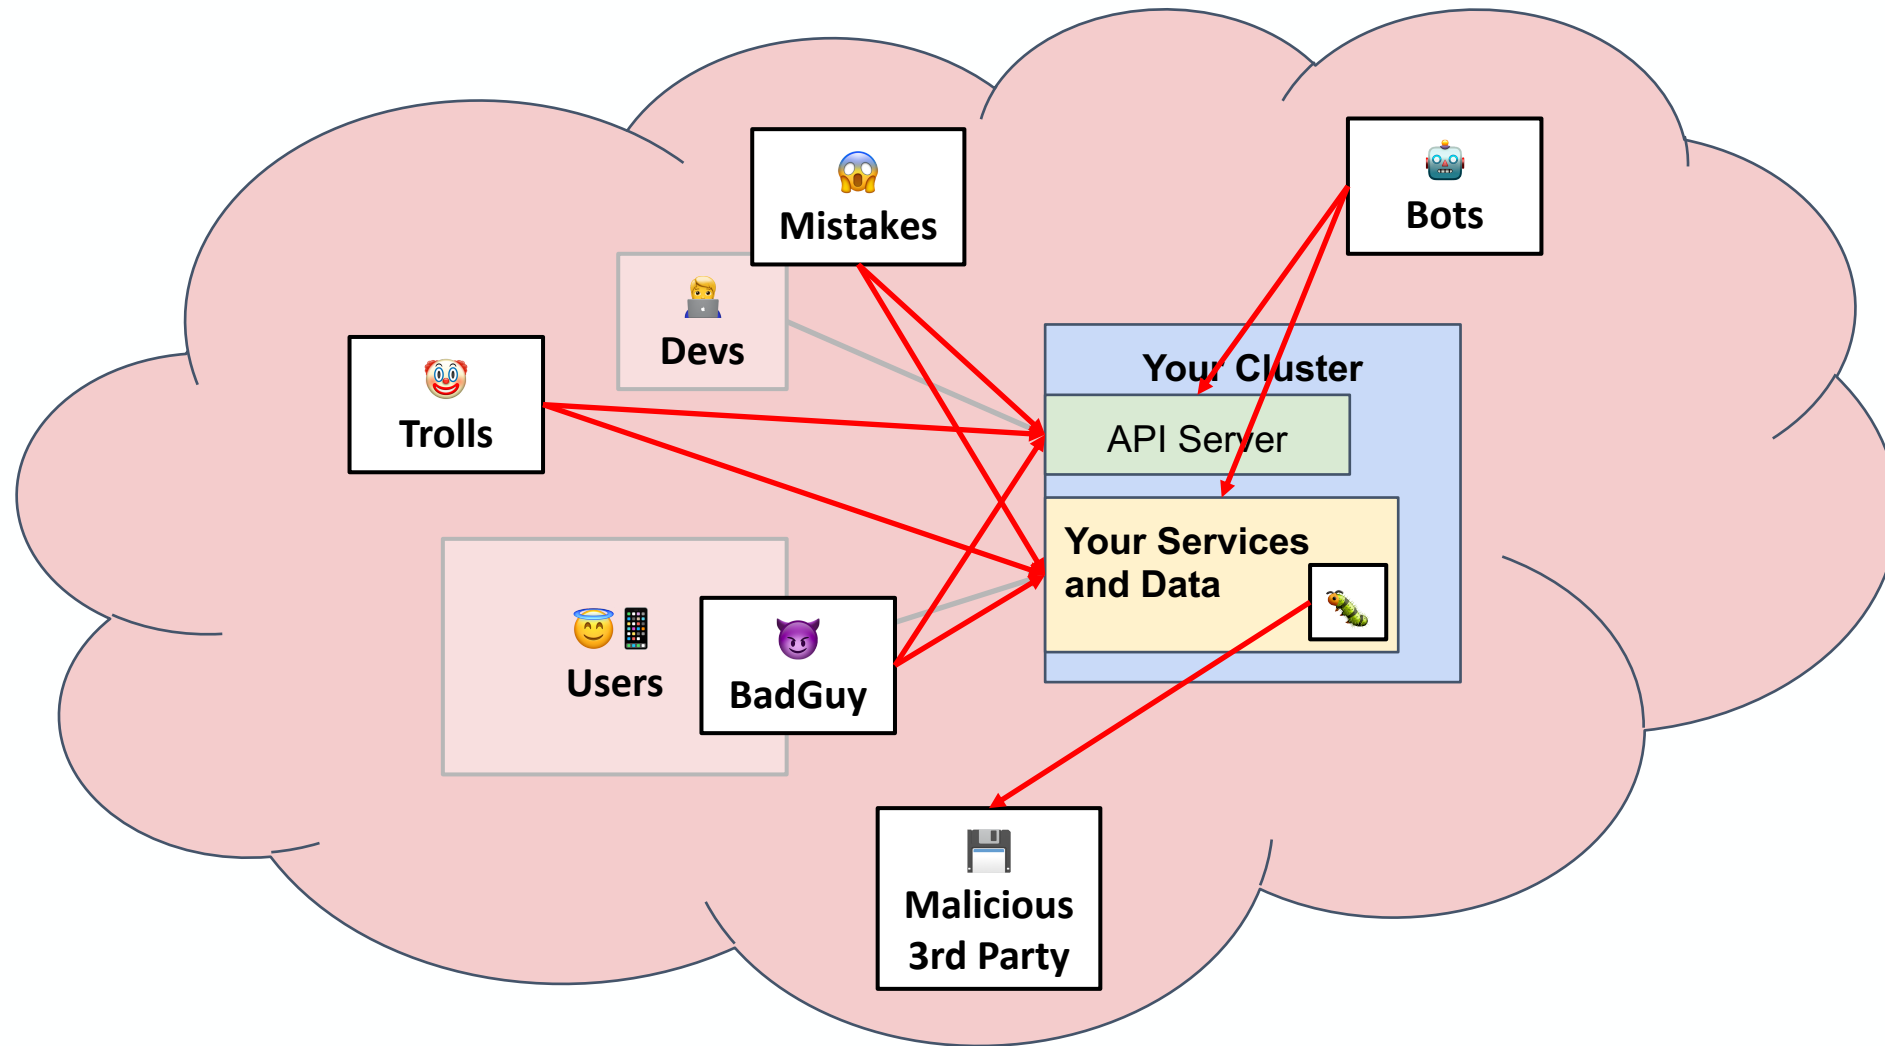
A lot of identity and authorization work for GKE.



Help! My Cluster Is On The Internet!



Help! My Cluster Is On The Internet!

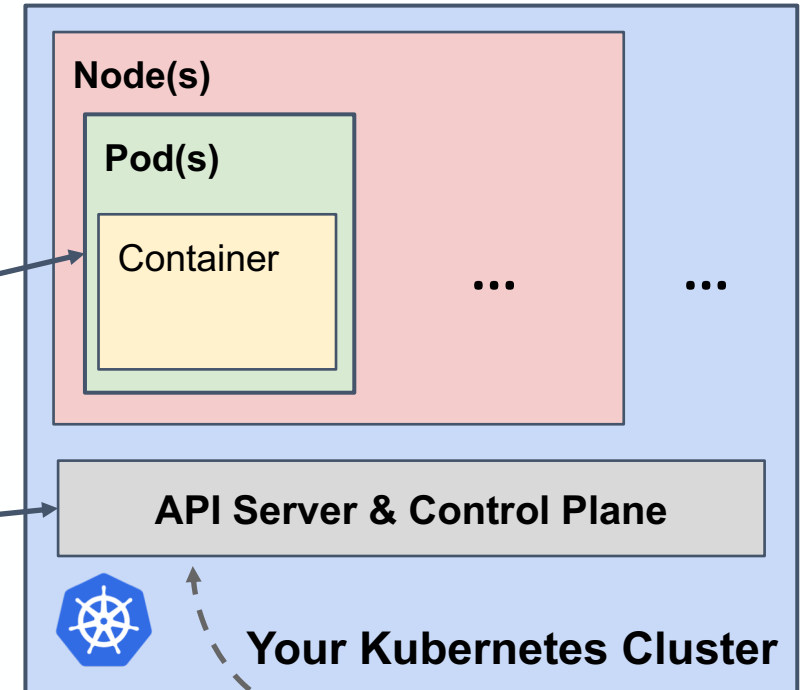


Structure Of This Talk


1. Workload Security

2. Cluster Security

3. User Security

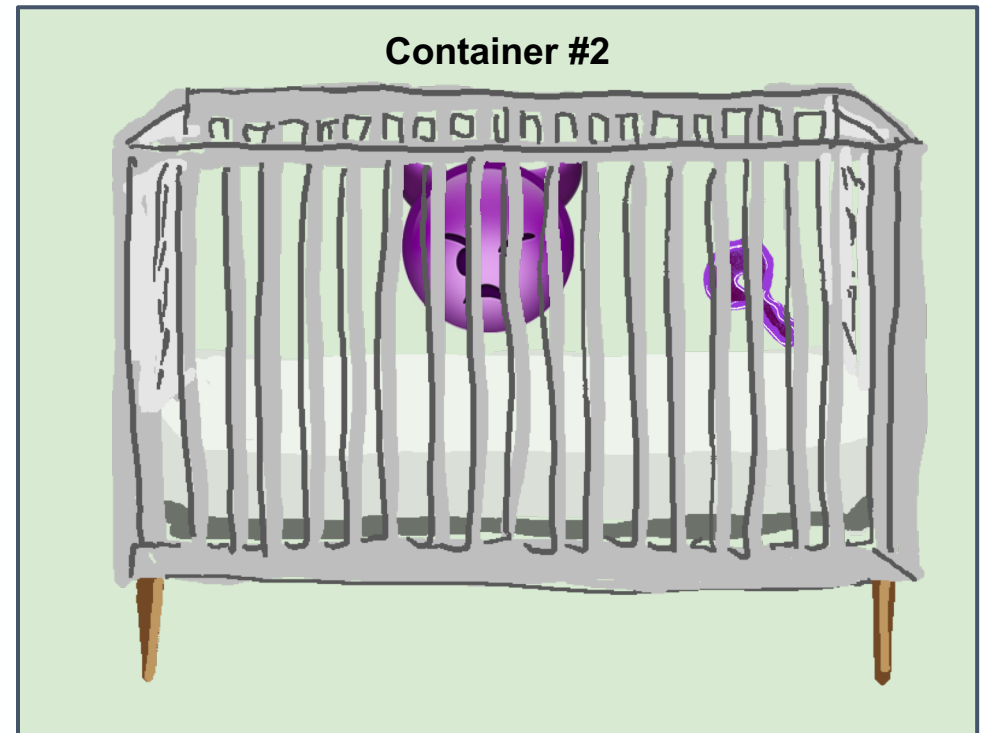



1. Workload Security



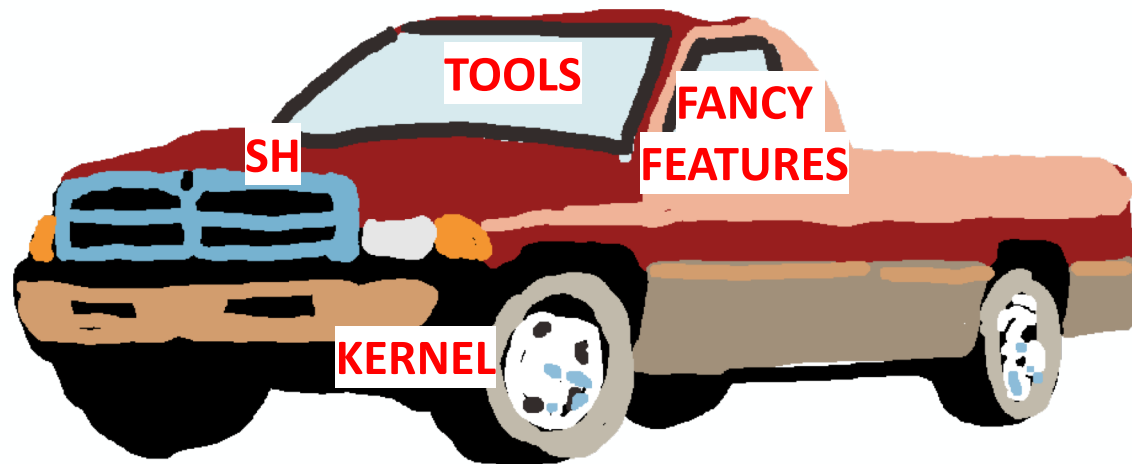
1. Workload Security - Containers

#1 - Assume that you will be OWNED.




1. Workload Security - Containers

#2 - Use a *distroless* base image.



Debian10



Distroless Debian10

1. Workload Security - Containers


#2 - Use a *distroless* base image.

```
FROM golang:1.13-buster as builder
WORKDIR /go/src/app
ADD . /go/src/app

RUN go get -d -v ./...

RUN go build -o /go/bin/app

# Now copy it into our base image.
FROM debian:10
COPY --from=builder /go/bin/app /
CMD ["/app"]
```



```
FROM golang:1.13-buster as builder
WORKDIR /go/src/app
ADD . /go/src/app

RUN go get -d -v ./...

RUN go build -o /go/bin/app

# Now copy it into our base image.
FROM gcr.io/distroless/base-debian10
COPY --from=builder /go/bin/app /
CMD ["/app"]
```

1. Workload Security - Containers



#2 - Use a *distroless* base image.

<https://github.com/GoogleContainerTools/distroless>

(bit.ly/39IU5i7)

or just search "distroless"

1. Workload Security - Containers



KubeCon



CloudNativeCon

Europe 2020



#3 - Containers are easy rebuild and deploy.

HELP! 🤖 I've got a critical vulnerability in my container!!

Just bump your dependencies and redeploy. 😊

 Agola Sorint.Lab ★ 605	 Argo Cloud Native Computing Foundation (CNCF) ★ 5,985	 BRIGADE Cloud Native Computing Foundation (CNCF) ★ 1,960	 Buildkite Buildkite ★ 489
 Concourse VMware ★ 5,067 MCap: \$60.1B	 flux Cloud Native Computing Foundation (CNCF) ★ 5,201	 GoCD Thoughtworks ★ 5,703 Funding: \$28M	 Jenkins Continuous Delivery Foundation (CDF) ★ 15,861
 JENKINSX Continuous Delivery Foundation (CDF) ★ 3,579	 keptn Cloud Native Computing Foundation (CNCF) ★ 442	 Razee IBM ★ 337 MCap: \$111.93B	 Spinnaker Continuous Delivery Foundation (CDF) ★ 7,239
 TEKTON Tekton Pipelines Continuous Delivery Foundation (CDF) ★ 5,241	 Travis CI Travis CI ★ 584	 weave flagger Weaveworks ★ 2,163 Funding: \$20M	 werf Flant ★ 1,781

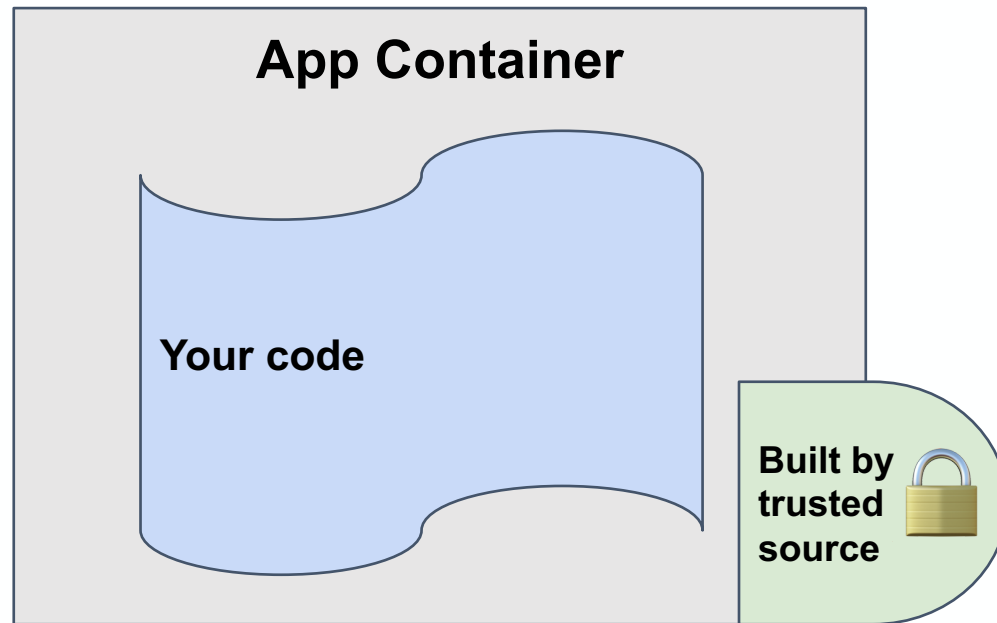
landscape.cncf.io

1. Workload Security - Containers

#4 - Trust your containers with **signatures!**

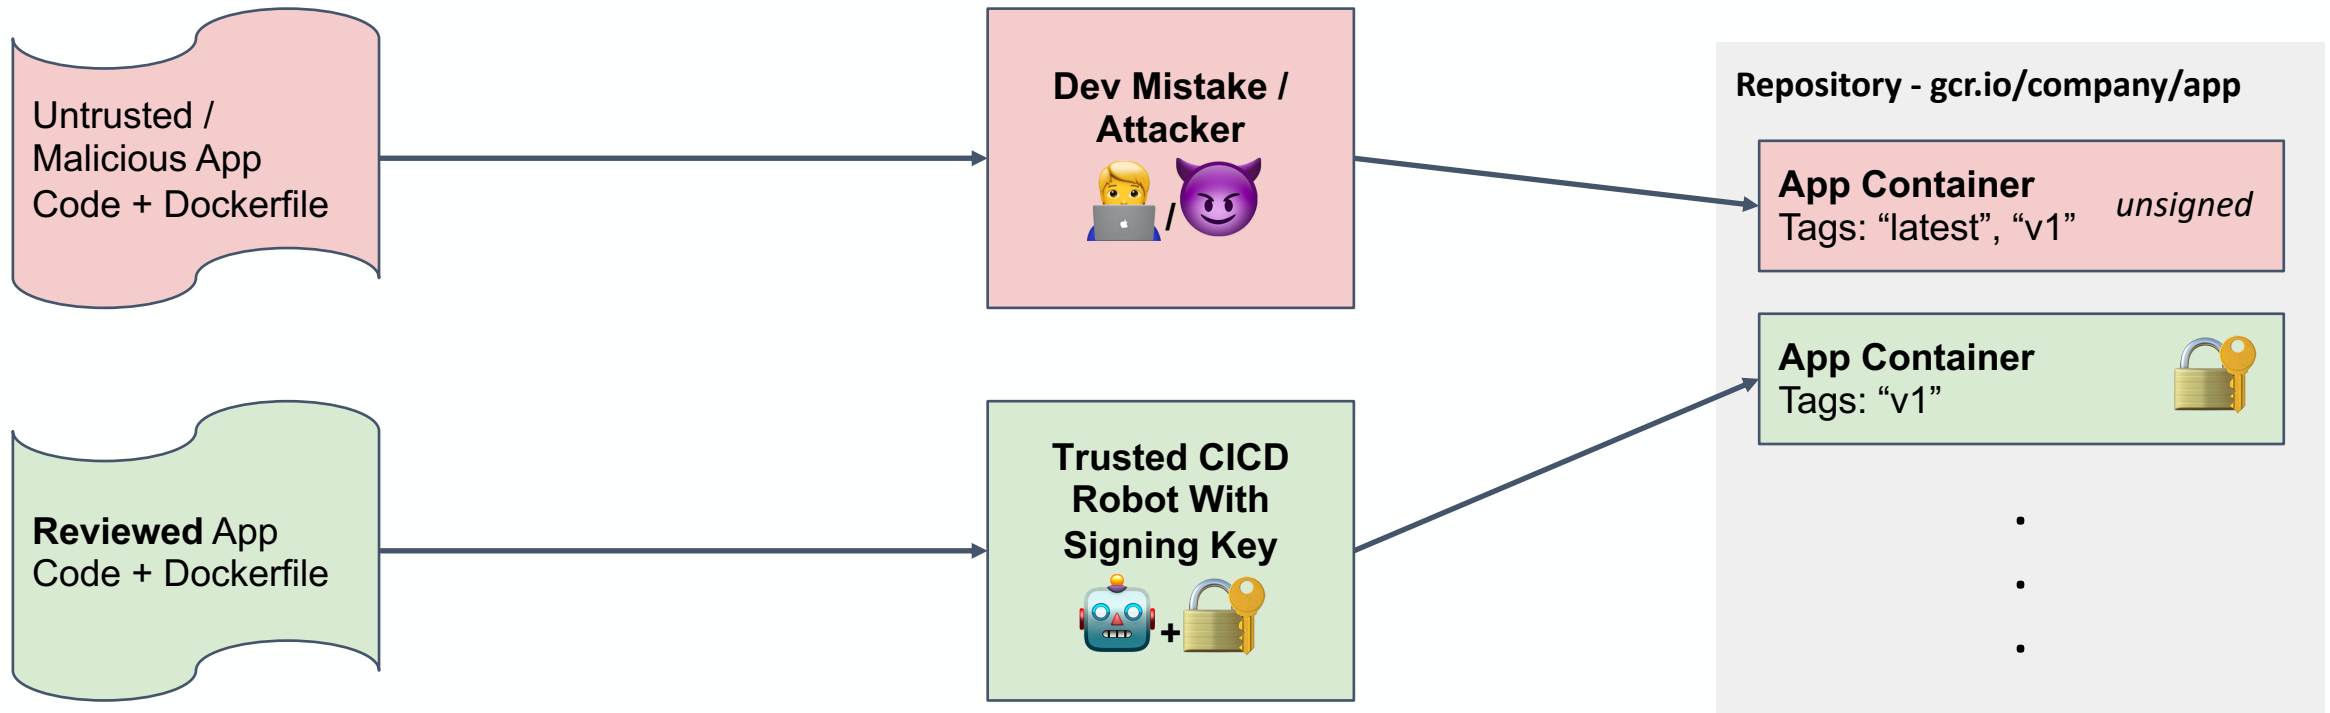
A.K.A:

Binary Authorization
Signed Containers
Image Signing
Binary Attestation
Content Trust



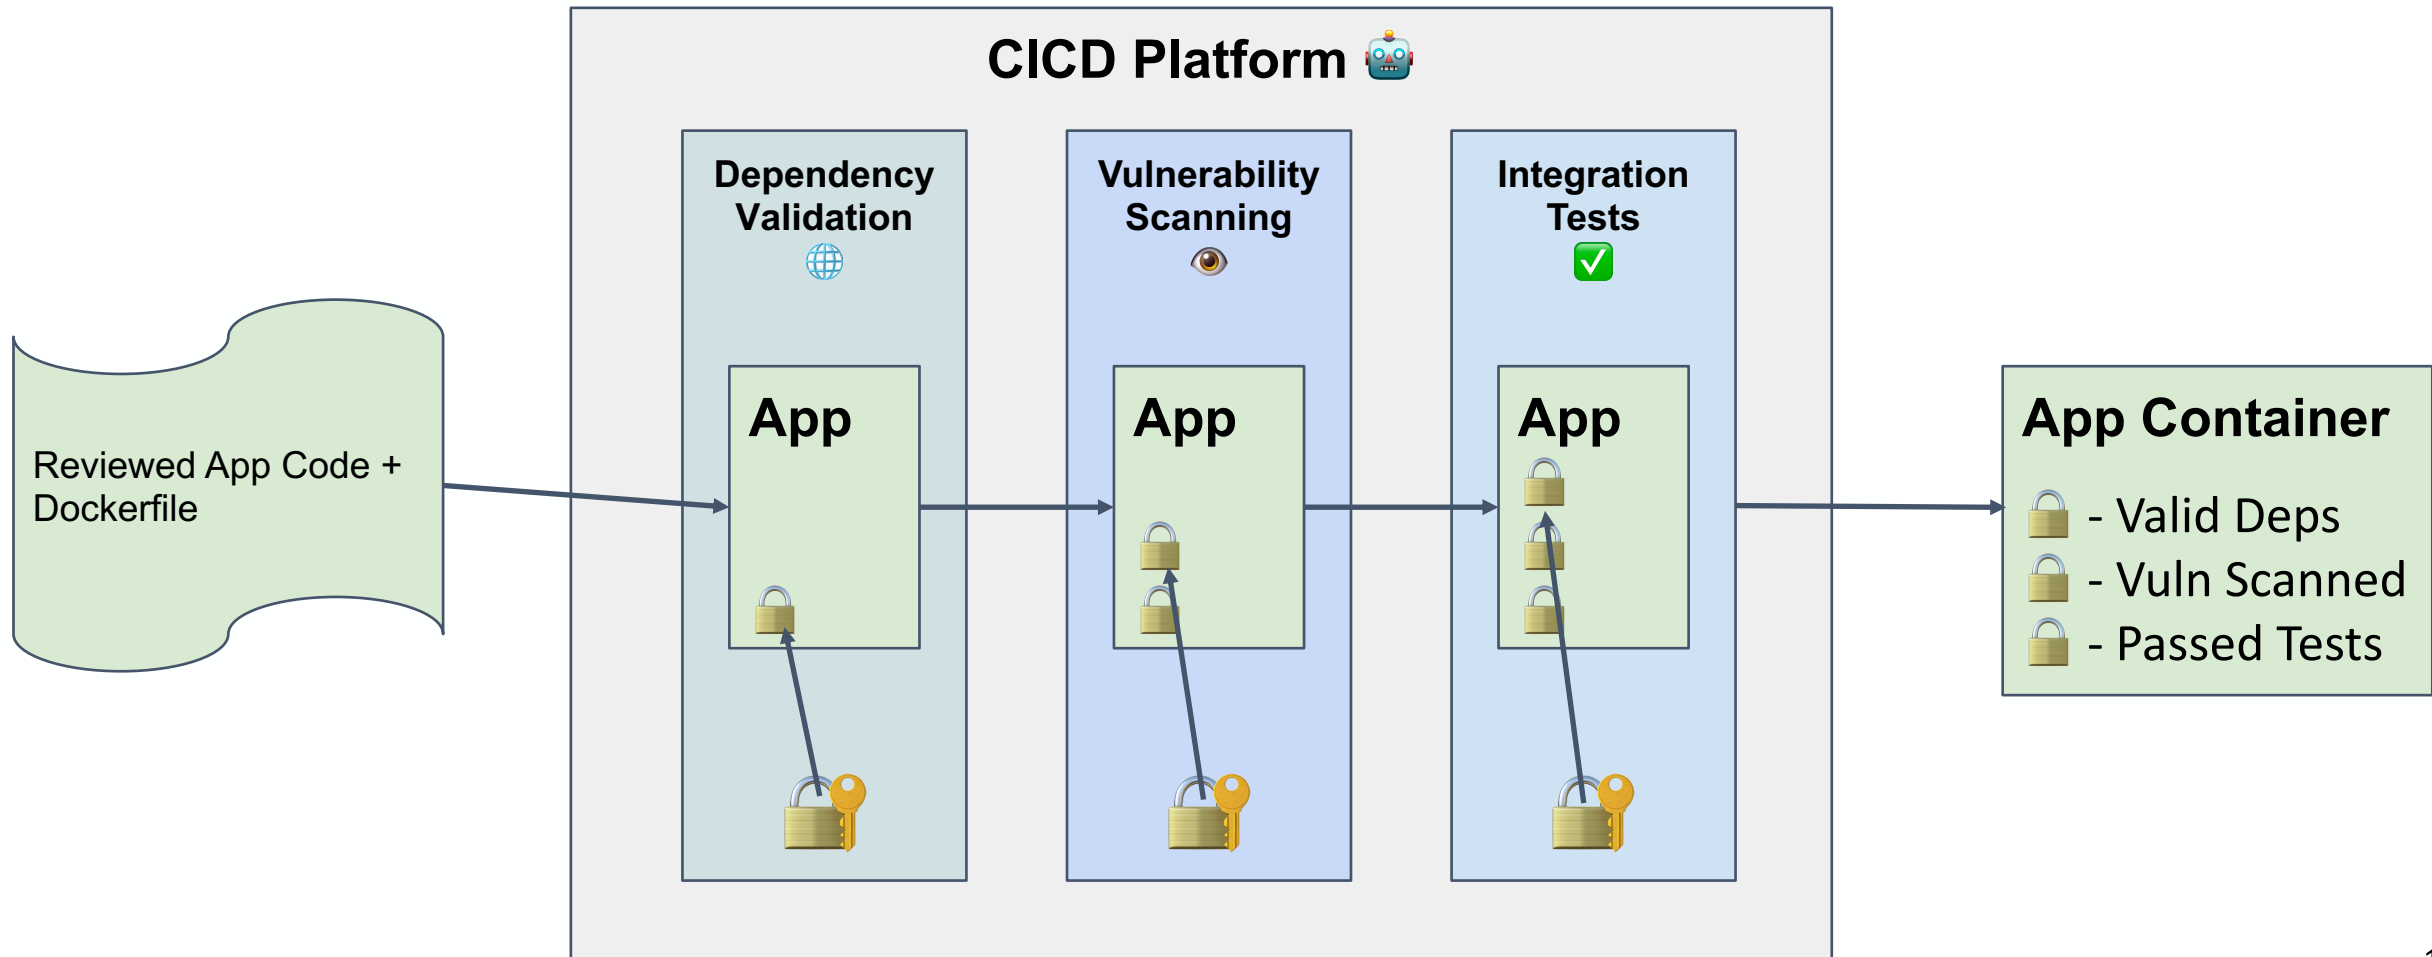
1. Workload Security - Containers

#4 - Trust your containers with **signatures!**



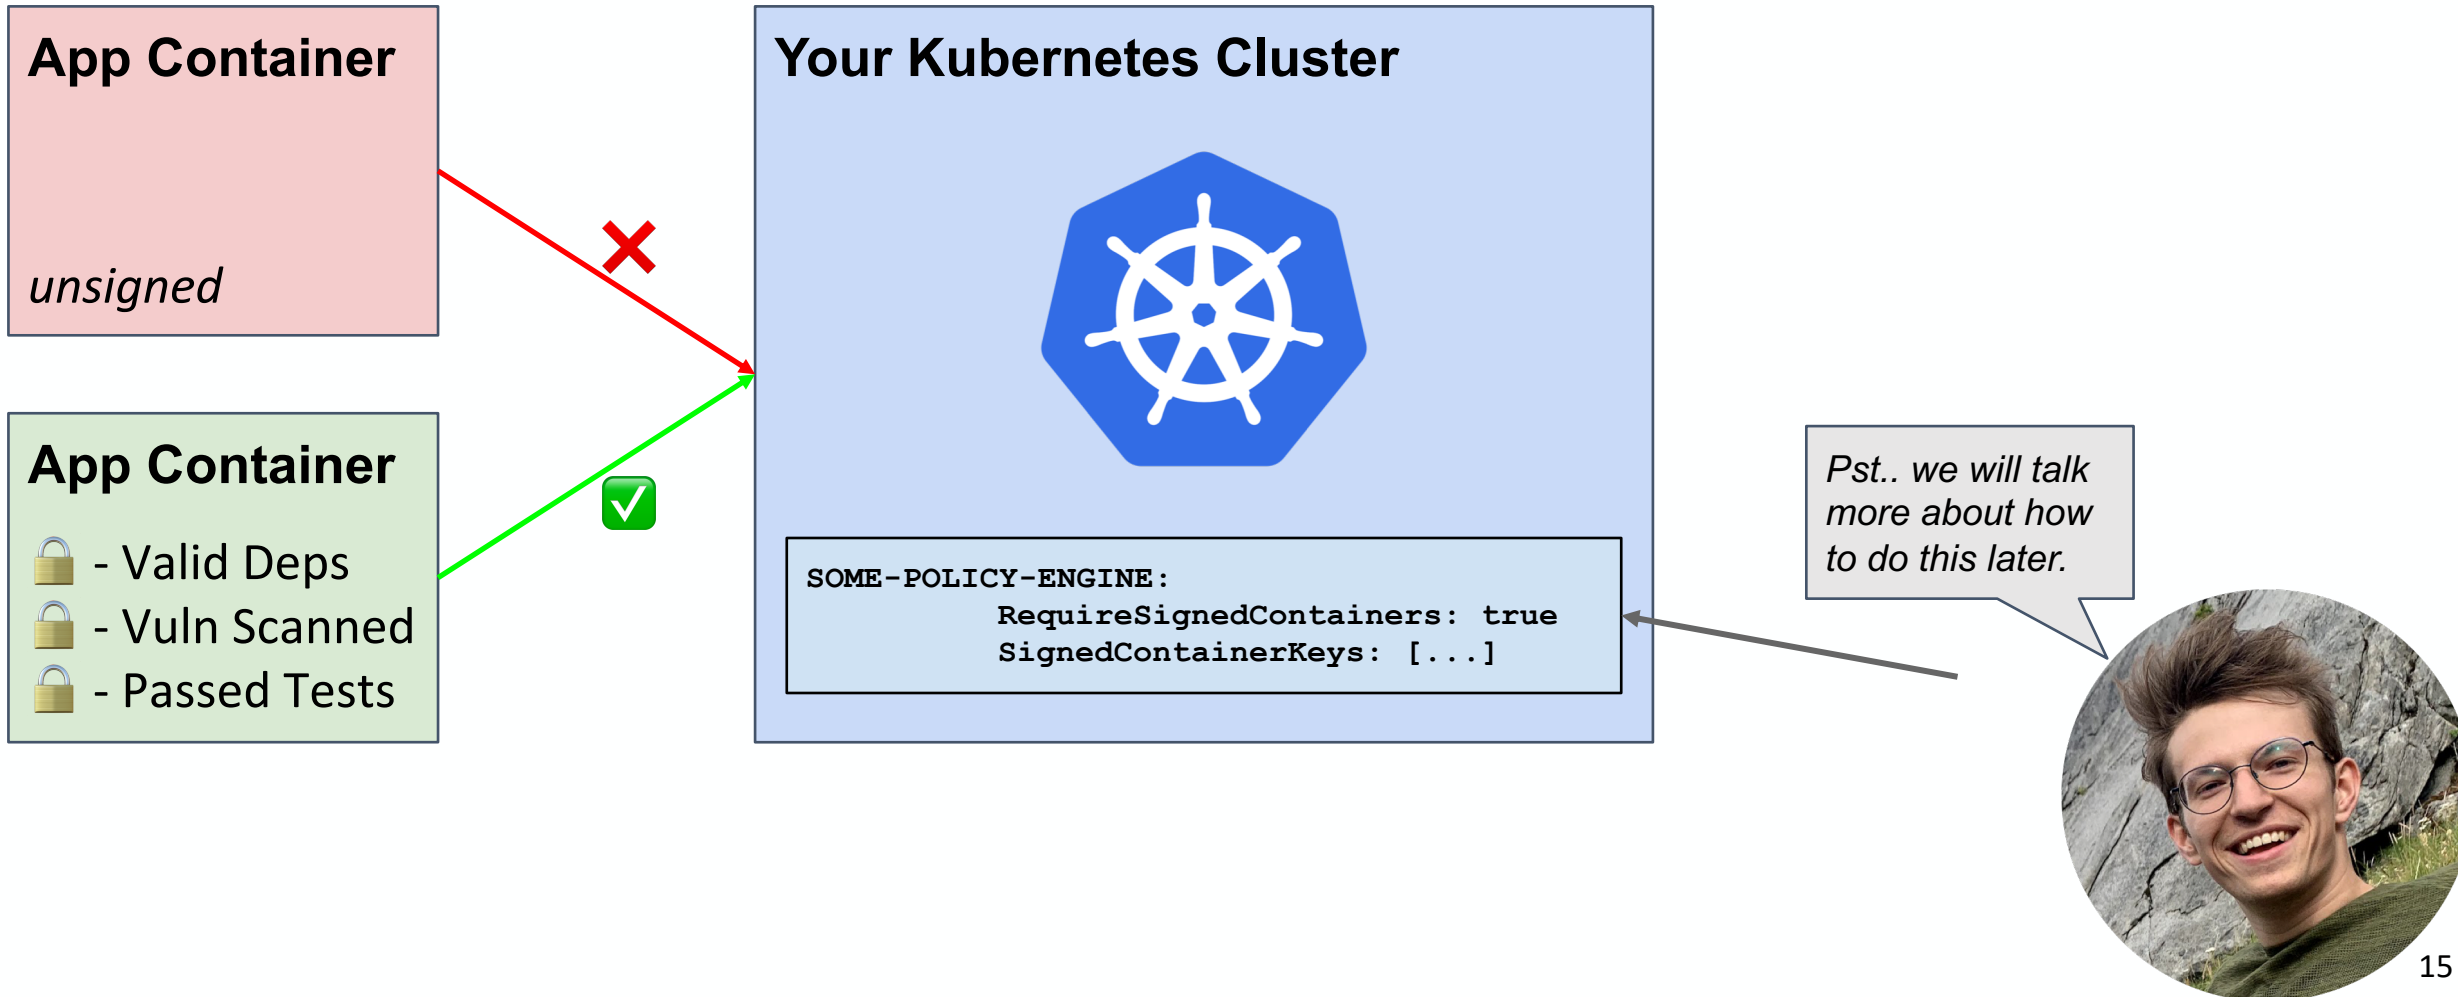
1. Workload Security - Containers

#4 - Trust your containers with **signatures!**



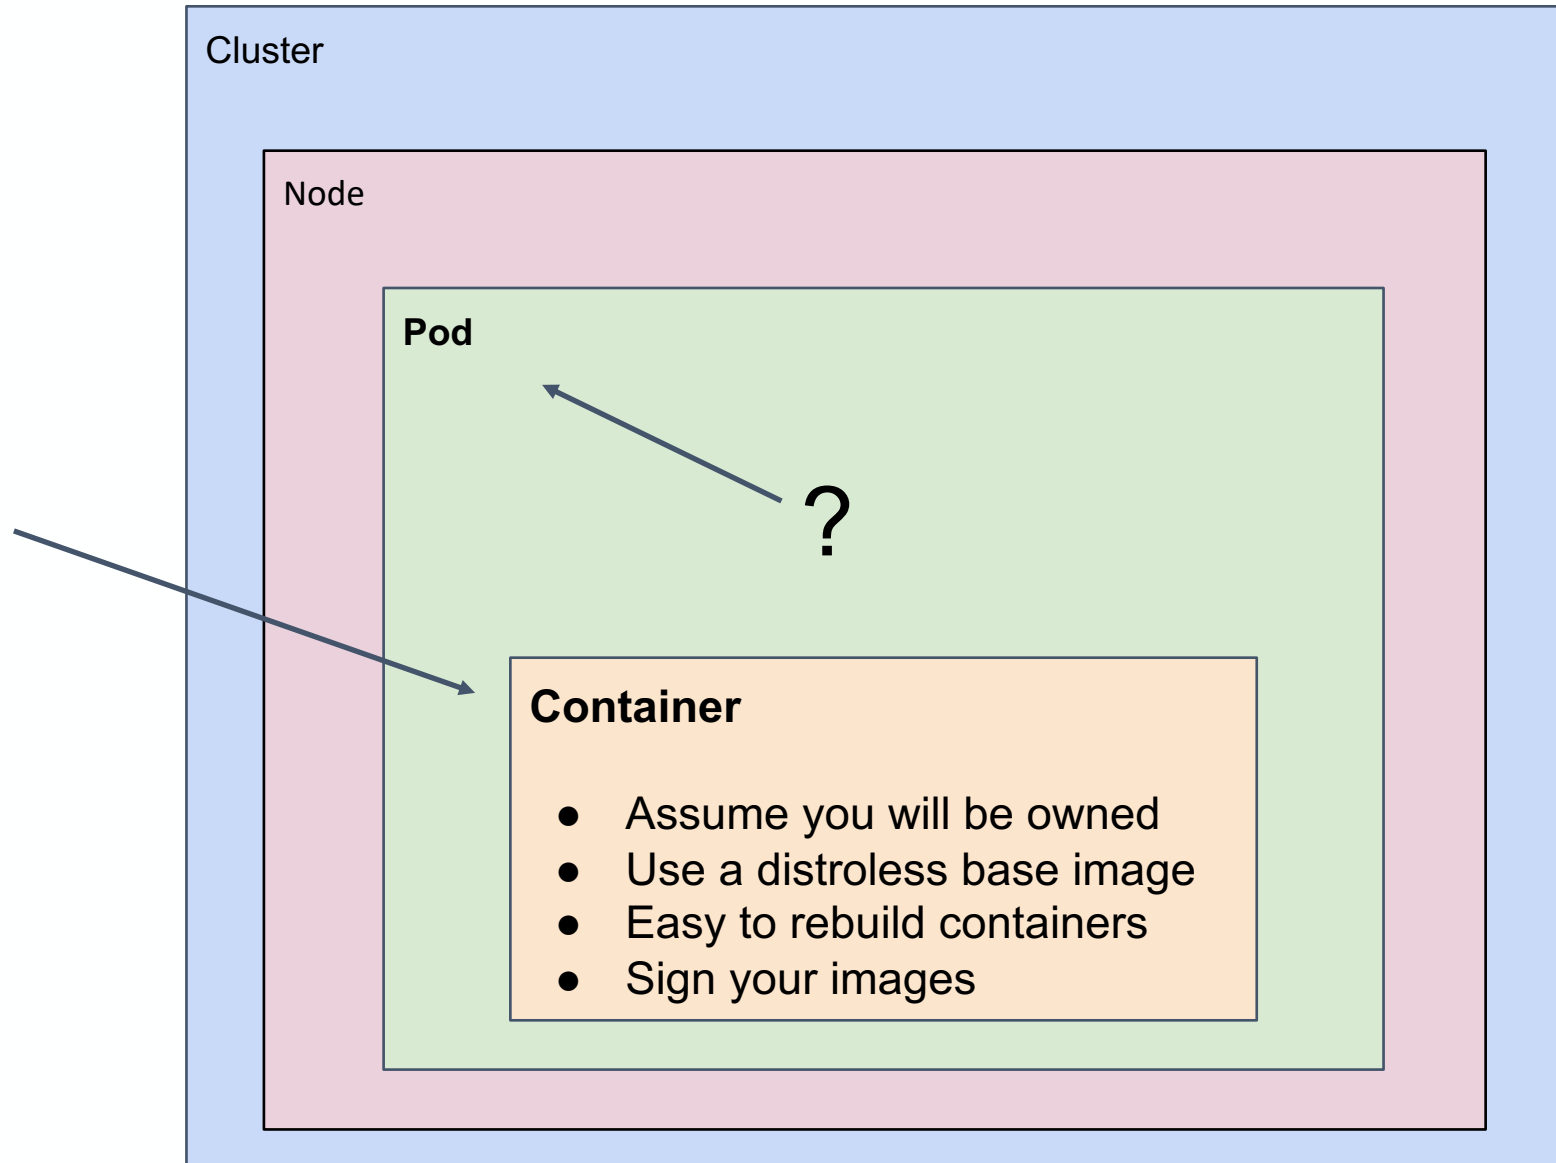
1. Workload Security - Containers

#4 - Trust your containers with **signatures!**



1. Workload Security

Recap:



1. Workload Security - Pods



KubeCon



CloudNativeCon

Europe 2020

Virtual

This concerns PodSpec configs *everywhere*.

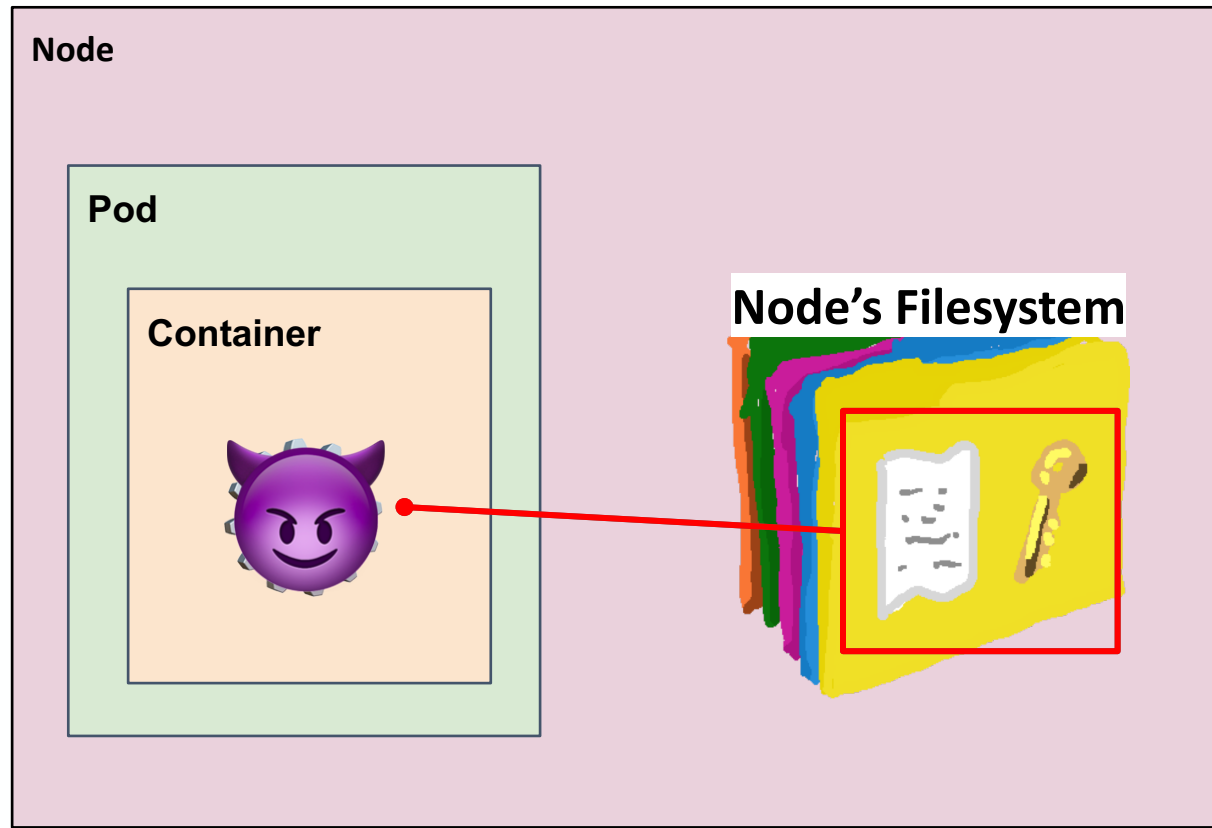
- **Pods**
- **DaemonSets**
- **Jobs**
- **CronJobs**
- **ReplicaSet**
- **StatefulSet**
- **etc...**

PodSpec API Reference

<https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.18/#podspec-v1-core>
(bit.ly/2CUKHM)

1. Workload Security - Pods

#1 - Don't use hostPath.



1. Workload Security - Pods

#1 - Don't use hostPath.

`.volumes.hostPath`

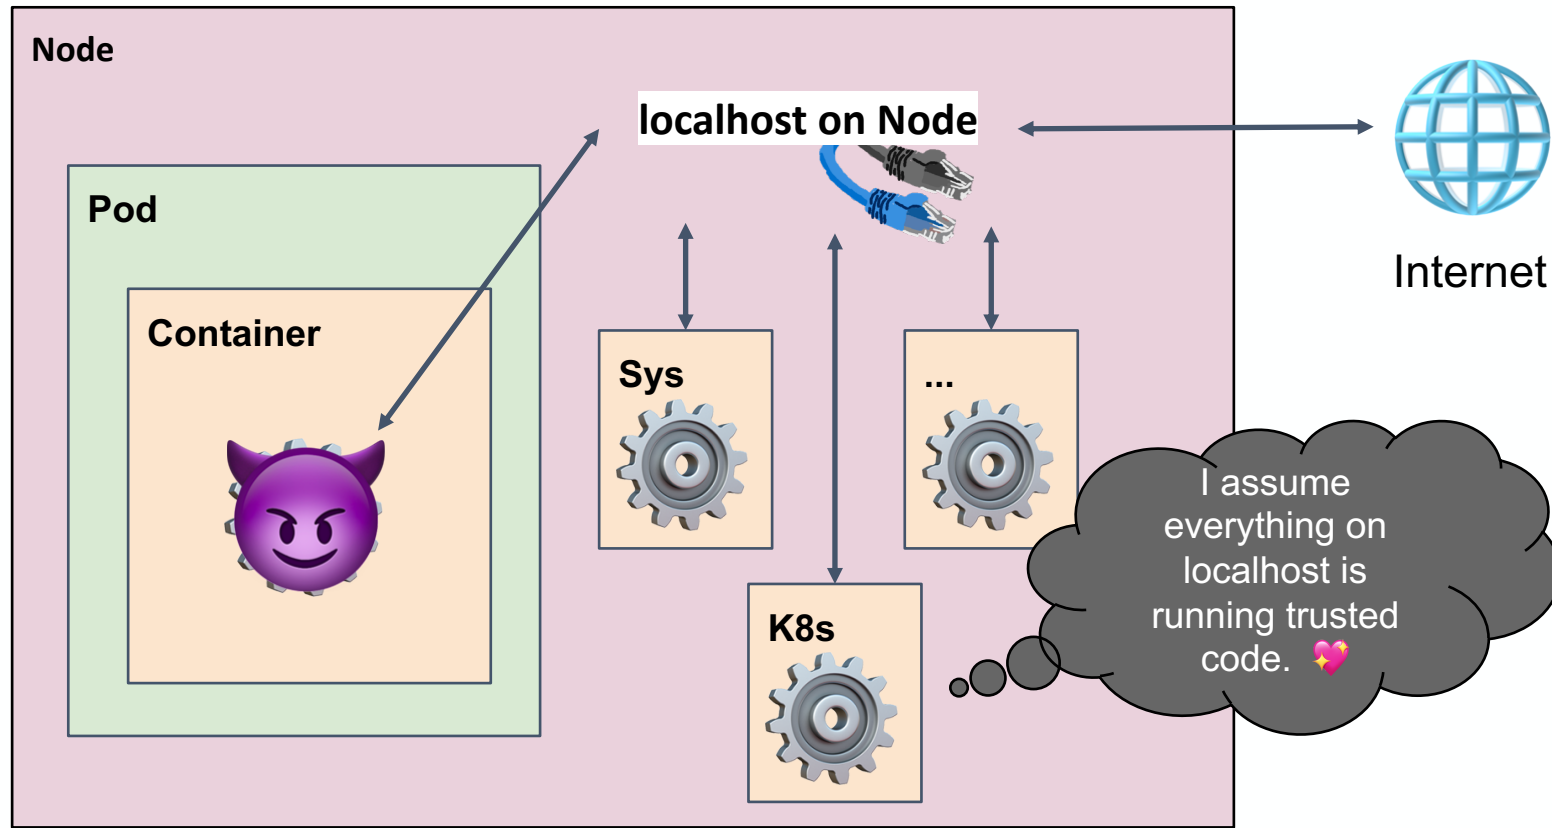
example:

```
apiVersion: v1
kind: Pod
metadata:
  name: example
spec:
  containers:
  - image: gcr.io/org/app
    name: example
    volumeMounts:
    - mountPath: /keys
      name: keys-volume
  volumes:
  - name: keys-volume
    hostPath: /etc/pod_data/
```

*Sure, maybe there is nothing scary in /etc/pod_data right now...
But what about a year from now?
Do all devs know the danger of this folder?*

1. Workload Security - Pods

#2 - Don't use hostNetwork.



1. Workload Security - Pods

#2 - Don't use hostNetwork.

.hostNetwork

example:

```
apiVersion: v1
kind: Pod
metadata:
  name: example
spec:
  containers:
  - image: gcr.io/org/app
    name: example
    ports:
    - 4444
  hostNetwork: true
```

just don't include this line

1. Workload Security - Pods

#3 - Be conscious of your pod's Service Account.

simple pod spec:

```
apiVersion: v1
kind: Pod
metadata:
  name: simple
spec:
  containers:
  - image: gcr.io/org/app
    name: simple
```

Are there any SA credentials bound to this pod?

Actually yes! Every pod is bound to a SA.

If no SA is specified the SA named "default" is used.

Pod is in "default" namespace so the SA loaded is:
`/api/v1/namespaces/default/serviceaccounts/default`

Mounted at:

`/var/run/secrets/kubernetes.io/serviceaccount/`



1. Workload Security - Pods

#3 - Be conscious of your pod's Service Account.

Some easy recommendations:

- Bind a different SA, that is unique to its use-case.
- Put the pod in a different namespace.
- **Set `automountServiceAccountToken` to false.**

If your workload doesn't need Kubernetes API server access, just do this!

```
apiVersion: v1
kind: Pod
metadata:
  name: simple
spec:
  automountServiceAccountToken: false
  containers:
  - image: gcr.io/org/app
    name: simple
```

1. Workload Security - Pods



#3 - Be conscious of your pod's Service Account.

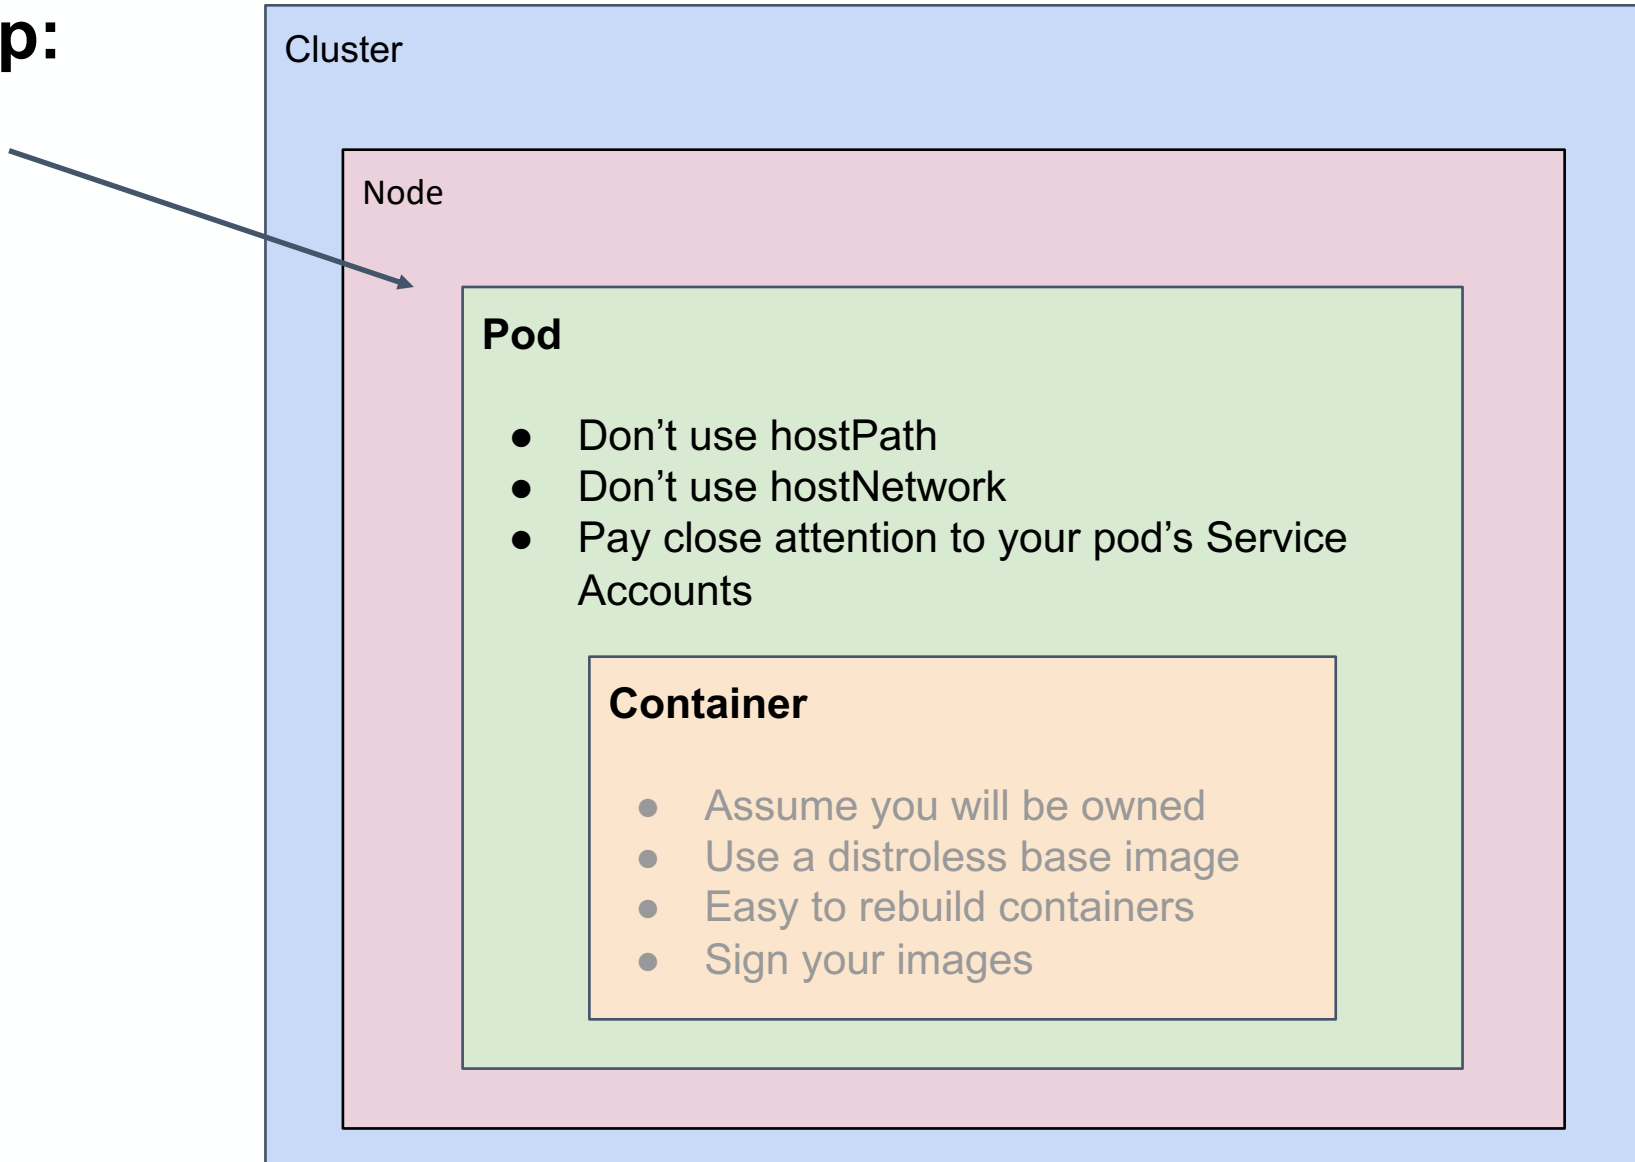
Learn More!

Kubernetes Docs on Service Accounts

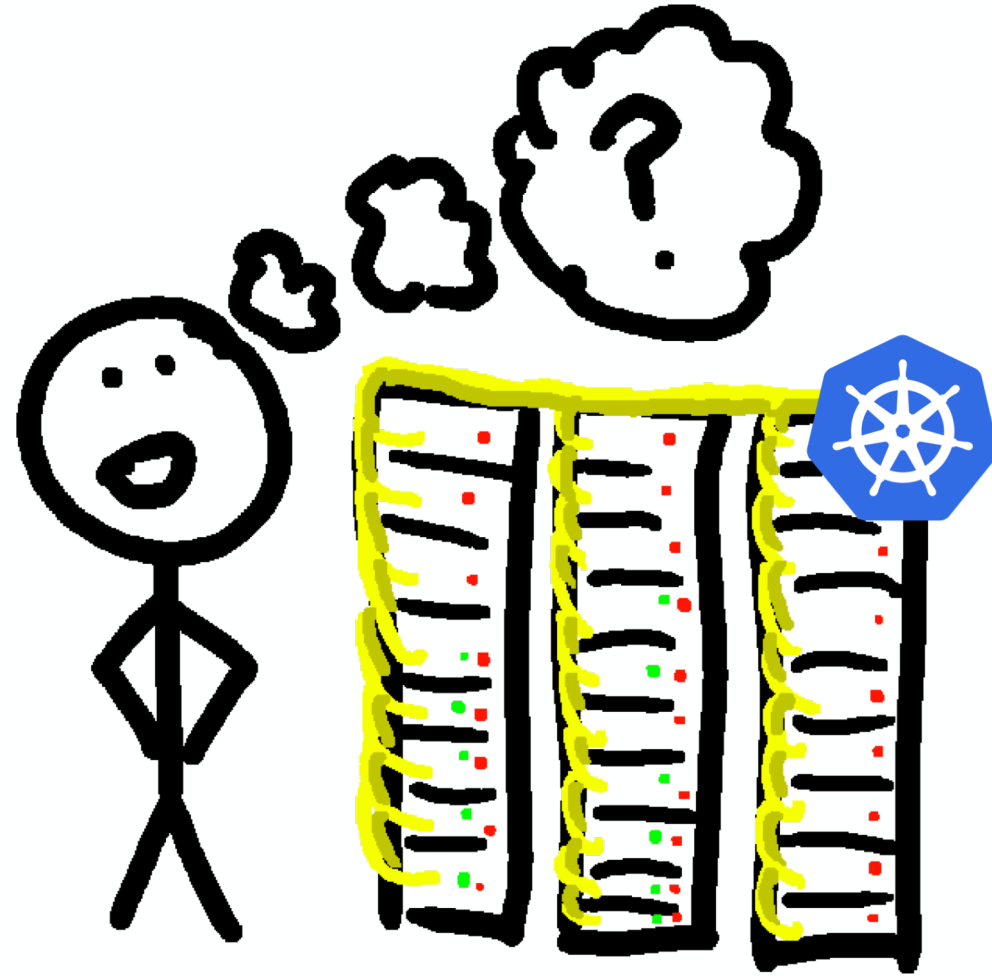
<https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/>
(bit.ly/30C3rIP)

1. Workload Security

Recap:



2. Cluster Security



2. Cluster Security

#1 - Keep your cluster up to date.

Bugs and vulnerabilities are fixed all the time!



v1.16.0 has been working just fine for us, I don't want to rock the boat.

github 1.16 post-release bugfix PRs
<https://bit.ly/2OPsoVA>

Consider this! Since 1.16.0 there have been **174 bugfix PRs** into the release branch. The latest patch version is **1.16.14**.*

* as of 7/29/2020

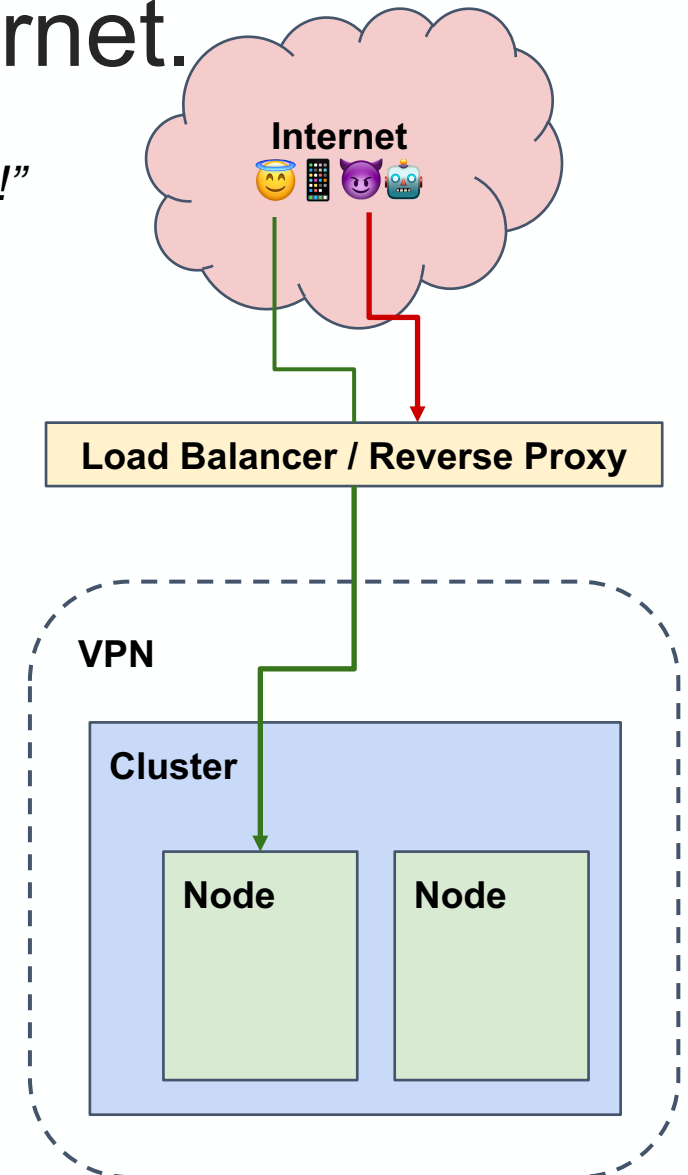


2. Cluster Security

#2 - Isolate your cluster from the internet.

“Help! My Cluster Is On The Internet!”

- Ideally **the entire cluster** is in a private network (VPN, auth-proxy, etc).
 - No public IPs for any cluster VMs.
- Solutions to common needs:
 - Devs/bots need API access?
 - Log them into the network
 - Users on internet need access to services/pods?
 - External load balancer that can forward traffic to nodes
 - Cluster needs internet access?
 - Egress only internet access from private network.



#3 - For your secrets use Secrets.

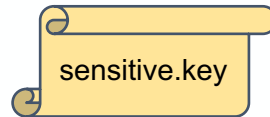
- Great for Access Keys, Passwords, Tokens, etc.
- Stored in memory, never saved to a node.
- Only loaded as-needed by pods.
- Easy authorization policy with RBAC.
- Not great for non-sensitive or lengthy configs, documents, large files.
 - Use ConfigMaps or other storage.

Kubernetes Docs on Secrets

<https://kubernetes.io/docs/concepts/configuration/secret/>
(bit.ly/3064n2E)

2. Cluster Security

#3 - For your secrets use Secrets.



```
$ kubectl create secret generic sensitive-key --from-file=./sensitive.key --namespace=app-sensitive
```

Some namespace for the workload.

```
apiVersion: v1  
kind: Secret  
metadata:  
  name: sensitive-key  
  namespace: app-sensitive  
type: Opaque  
data:  
  sensitive.key: ...BASE64...
```

2. Cluster Security

#3 - For your secrets use Secrets.

```
apiVersion: v1
kind: Secret
metadata:
  name: sensitive-key
  namespace: app-sensitive
type: Opaque
data:
  sensitive.key: ...BASE64...
```

```
apiVersion: v1
kind: Pod
metadata:
  name: pod-with-secret
  namespace: app-sensitive
spec:
  containers:
  - image: gcr.io/org/app
    name: app-with-secret
  volumeMounts:
  - name: keys
    mountPath: "/etc/key"
    readOnly: true
  volumes:
  - name: keys
    secret:
      secretName: sensitive-key
```

2. Cluster Security

Recap:

Cluster

- Keep your cluster up to date
- Isolate your cluster from the internet
- Use Secrets

Node

- *Isolate your nodes from the internet*


Pod

- Don't use hostPath
- Don't use hostNetwork
- Pay close attention to your pod's Service Accounts

Container

- Assume you will be owned
- Use a distroless base image
- Easy to rebuild containers
- Sign your images

3. User Security

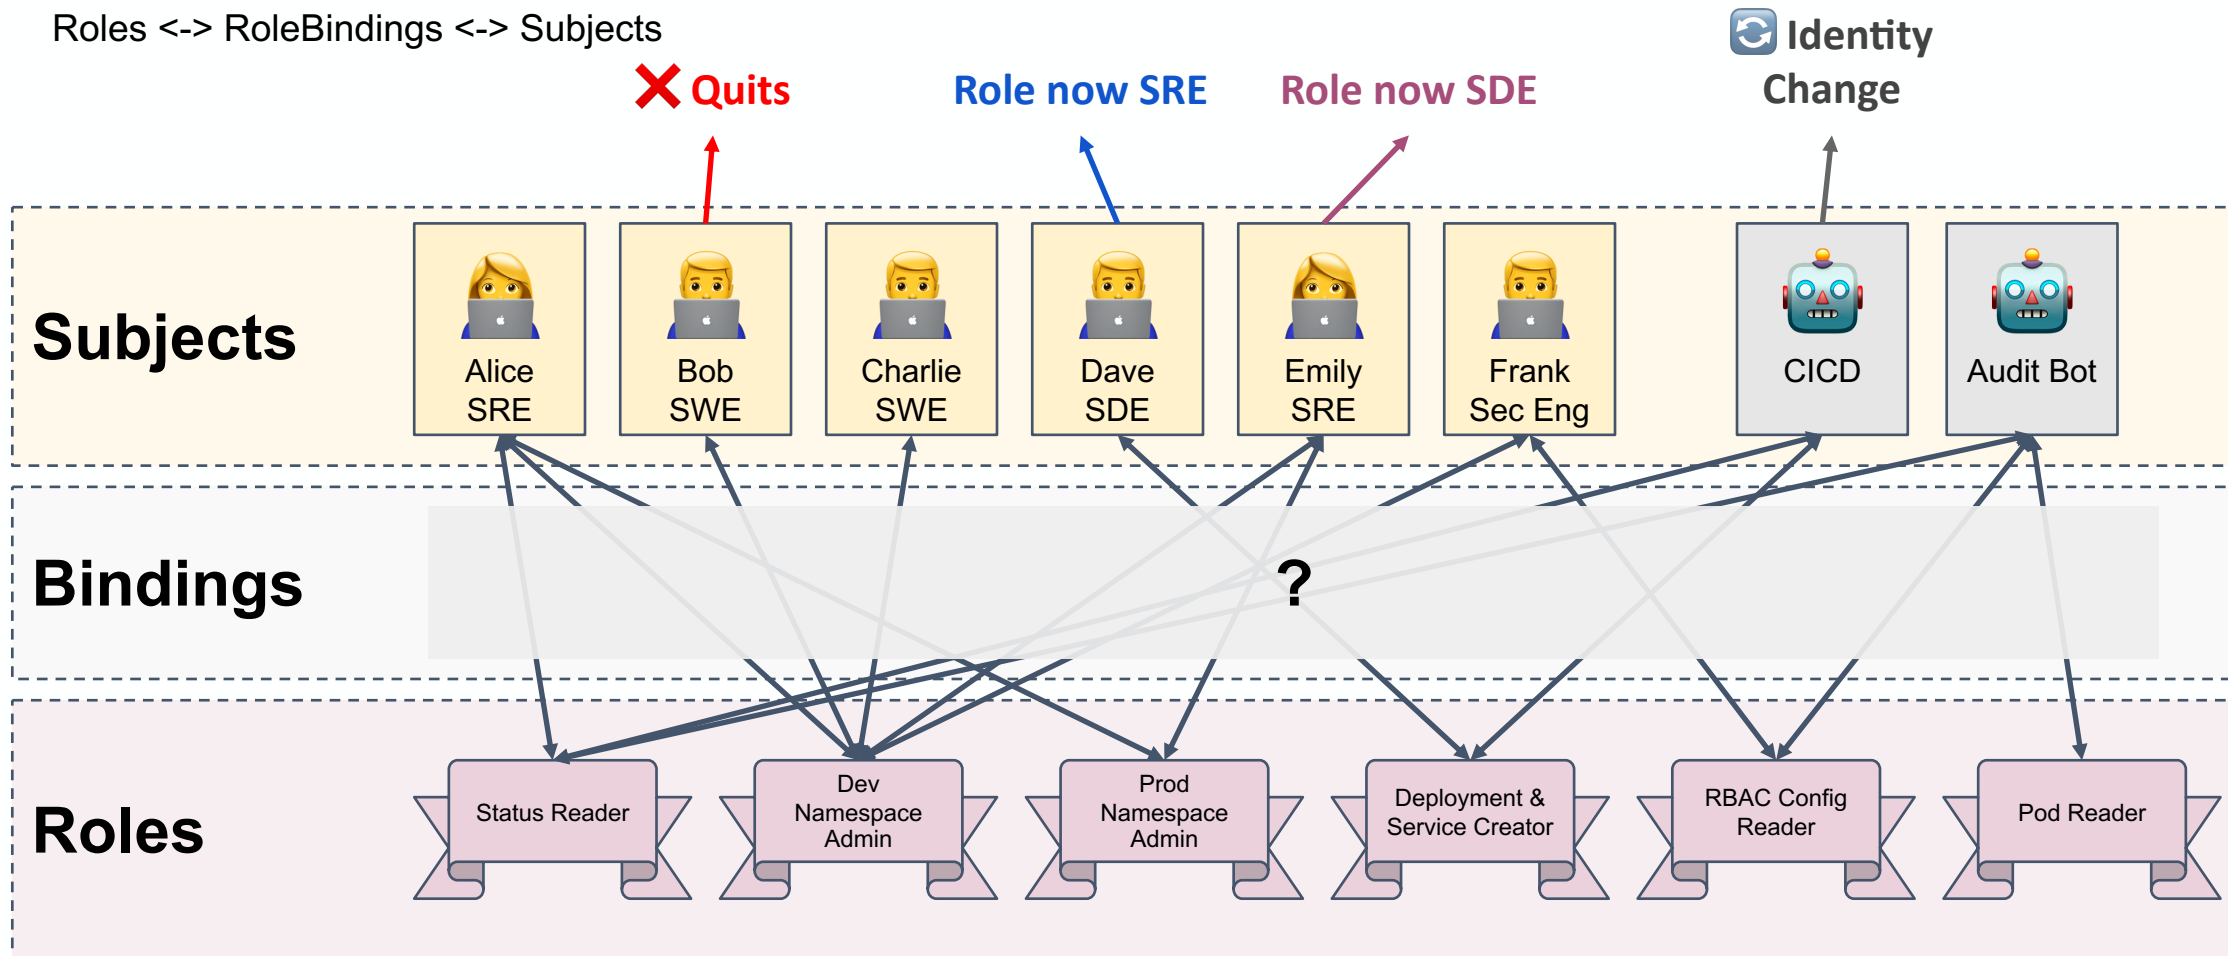


3. User Security

#1 - Use RBAC and groups.

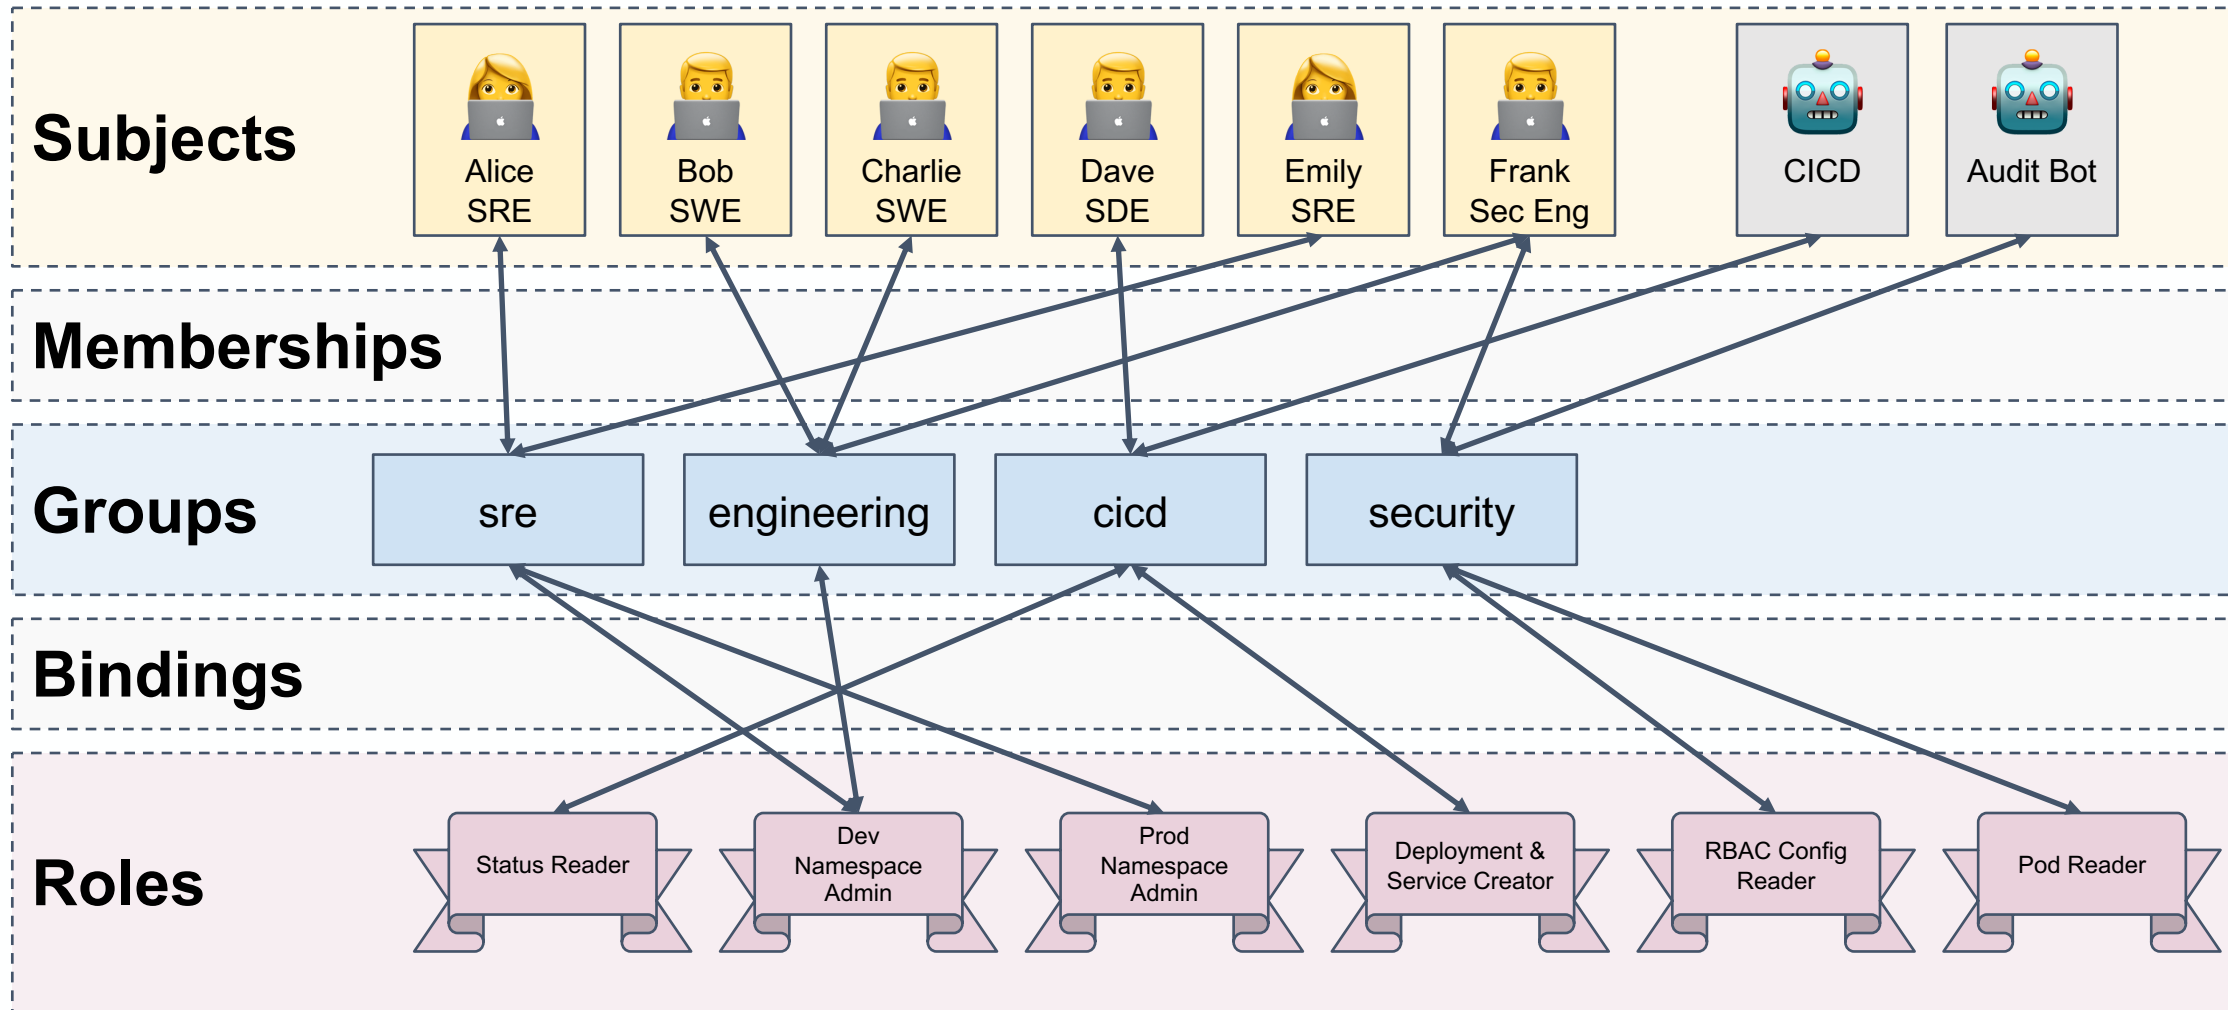
“Role Based Access Control”:

Roles <-> RoleBindings <-> Subjects



3. User Security

#1 - Use RBAC and groups.



3. User Security



#1 - Use RBAC and groups.

RBAC Docs

<https://kubernetes.io/docs/reference/access-authn-authz/rbac/>
(bit.ly/30GcGRR)

RBAC API



<https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.18/#role-v1-rbac-authorization-k8s-io>
(bit.ly/3hrhAJj)

3. User Security

#2 - Use a policy agent to protect you cluster.

- Typically a Kubernetes *AdmissionController* which selectively allows/denies Kubernetes resource requests based on rules (or policies).

RBAC



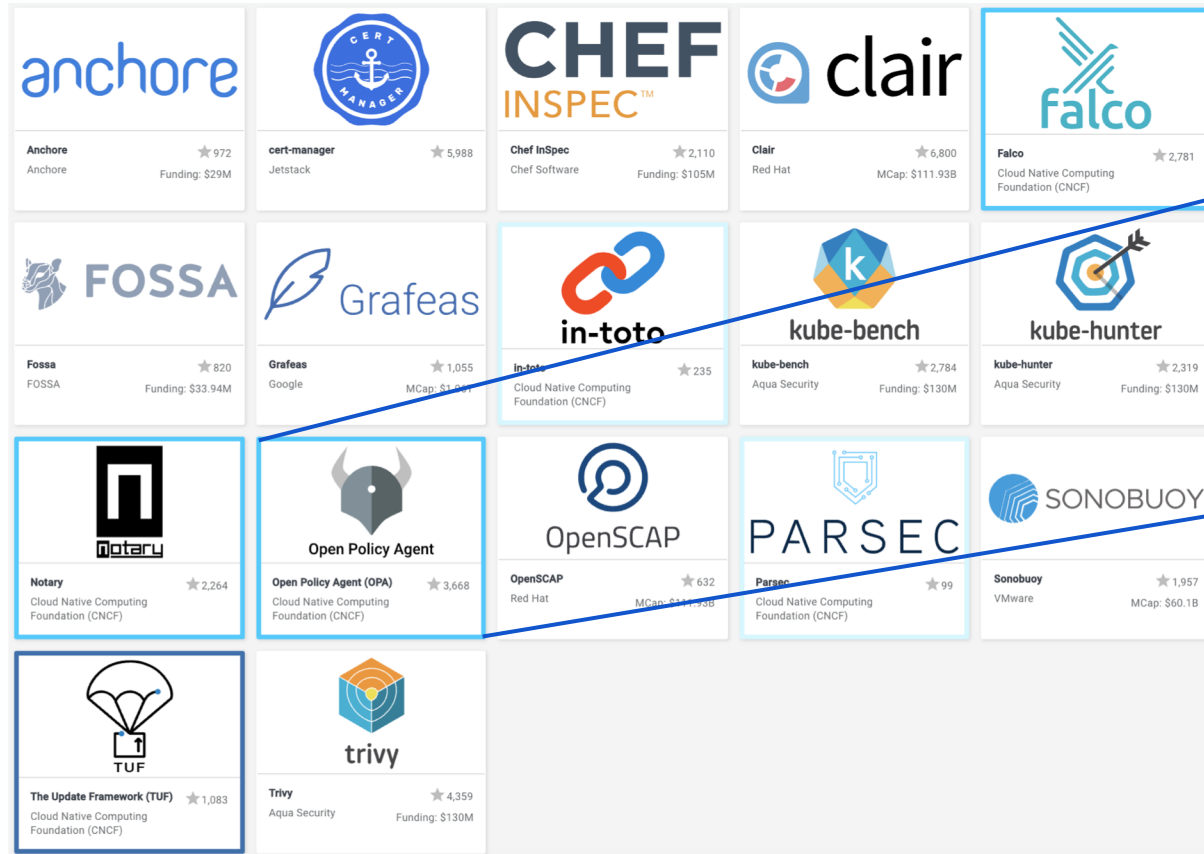
Admission Controllers

#2 - Use a policy agent to protect you cluster.

- Can enforce all kinds of best practices at runtime
 - No **hostPath**, **hostNetwork**
 - Default SAs
 - Allow/block images
 - Signatures
 - No Keys
 - RBAC
 - Labels (owner)
 - ... plus much, MUCH more
- Audits of existing resources

3. User Security

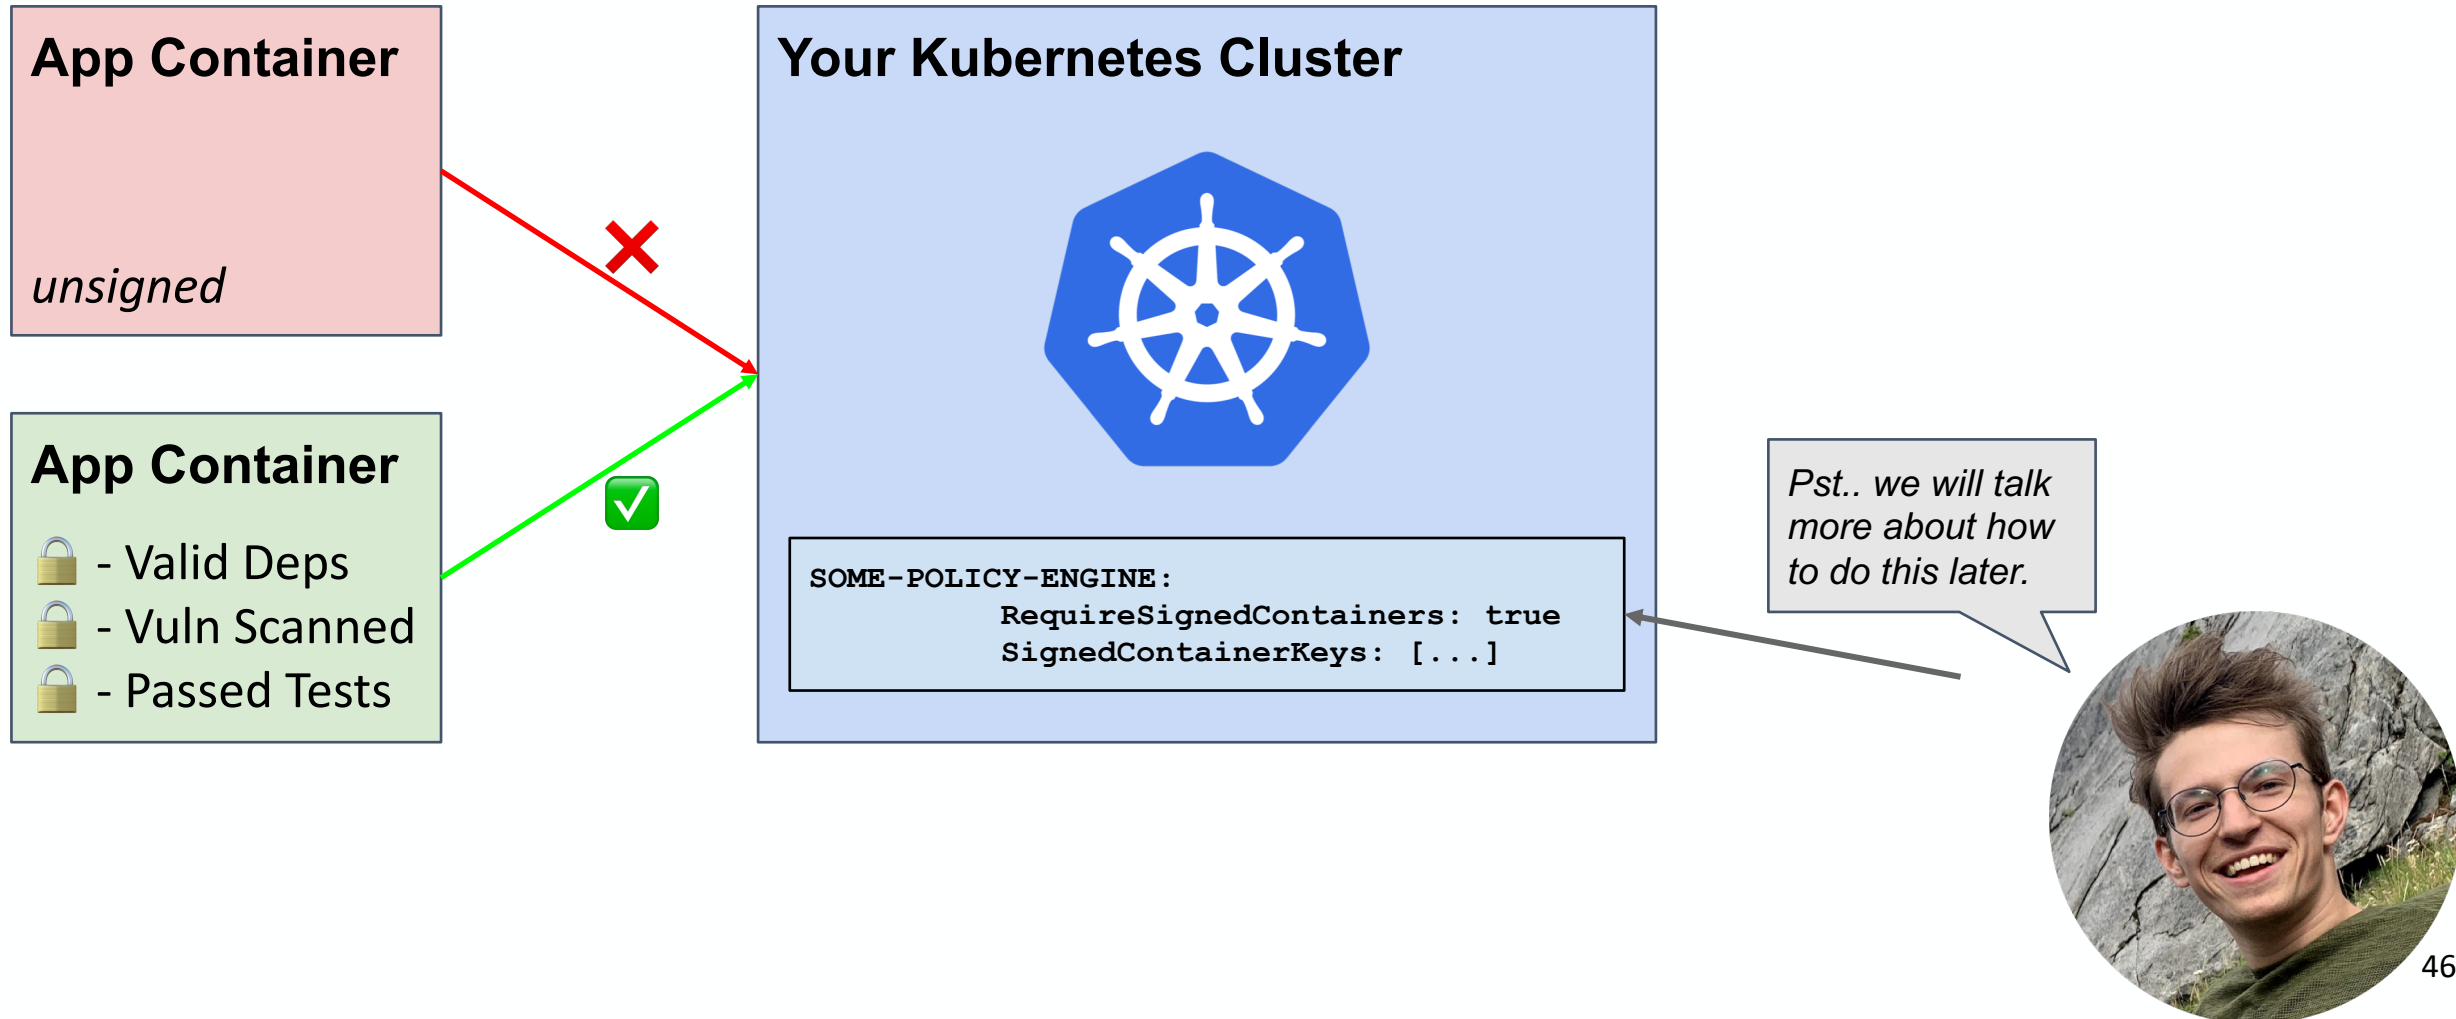
#2 - Use a policy agent to protect you cluster.



<https://github.com/open-policy-agent/gatekeeper>
(bit.ly/2WNIEM3)

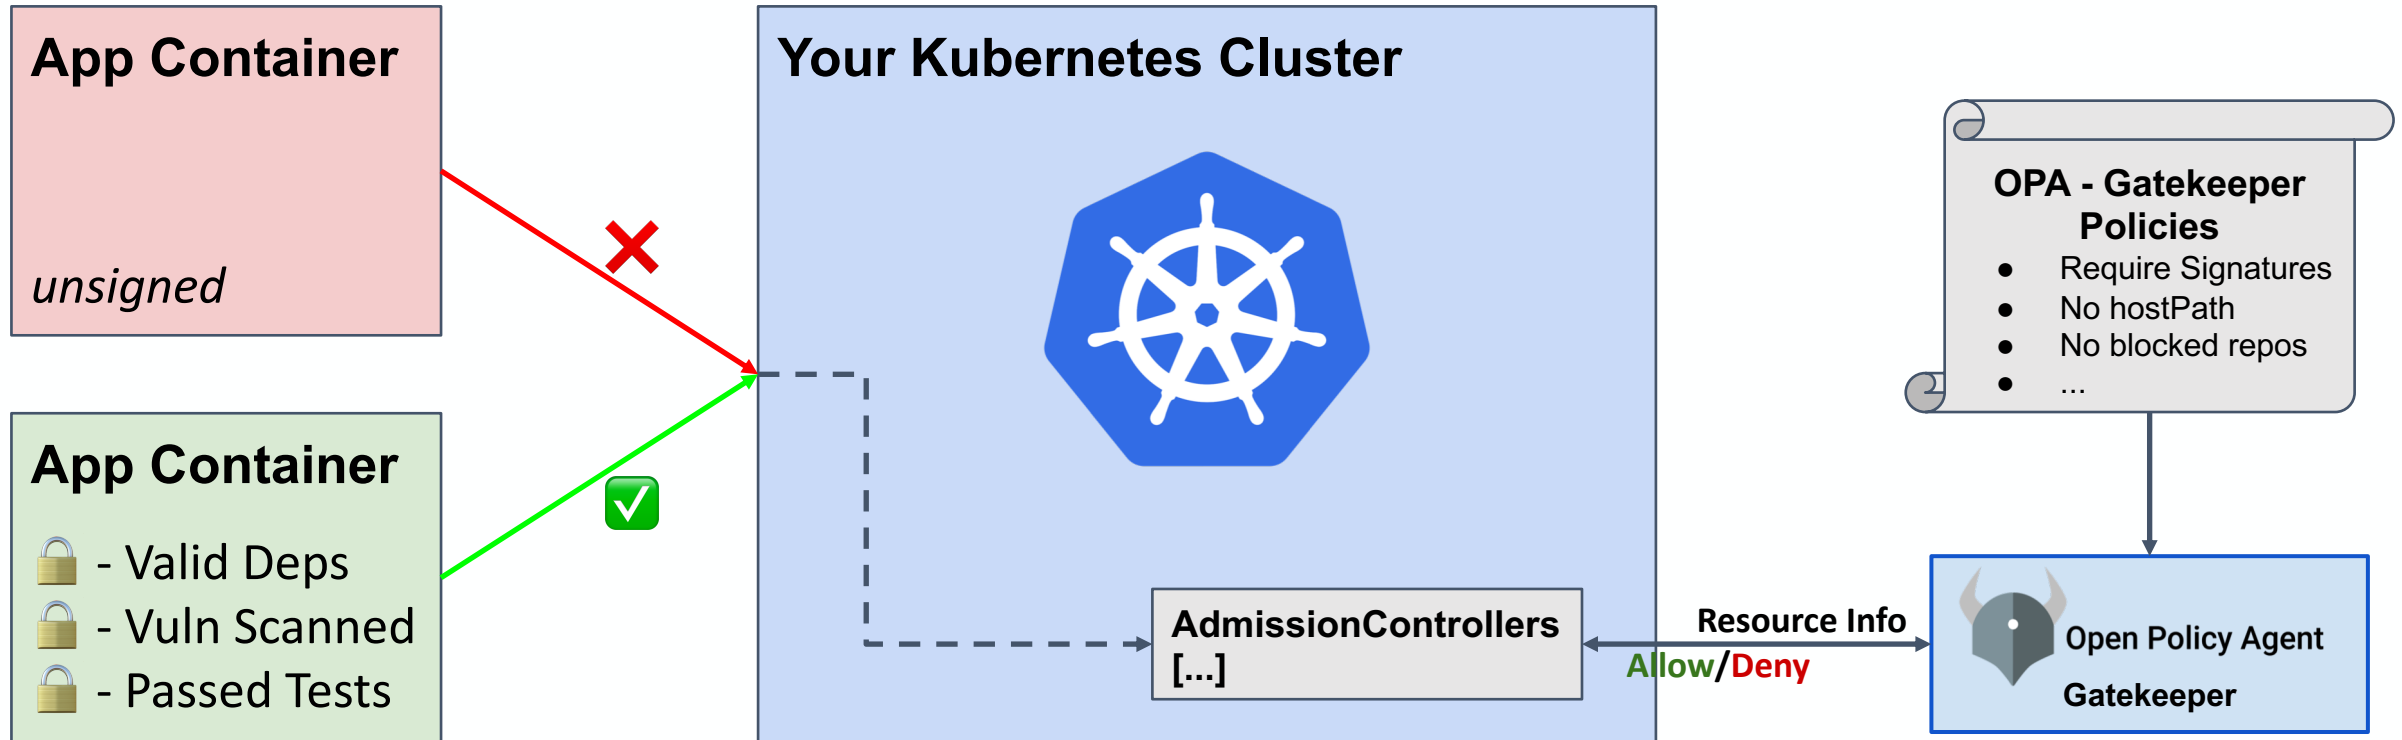
3. User Security

#2 - Use a policy agent to protect you cluster.



3. User Security

#2 - Use a policy agent to protect you cluster.



3. User Security

Recap:

Developers & Automation

- Use RBAC and groups
- Use a policy agent

Cluster

- Keep your cluster up to date
- Isolate your cluster from the internet
- Use Secrets
- Don't use Basic Auth

Node

- *Isolate your nodes from the internet*

Pod

- Don't use hostPath
- Don't use hostNetwork
- Pay close attention to your pod's Service Accounts


Container

- Assume you will be owned
- Use a distroless base image
- Easy to rebuild containers
- Sign your images

3. Cluster Security

Your workloads, cluster, and developers enjoy a much more secure Kubernetes experience. You are recognized for your efforts and compensated handsomely. You are filled with a sense of satisfaction.

:)





I can't remember all that!!

I made a doc!

- All the tips and tricks
- Lots of links and reading
- Plus other stuff not covered
 - *Don't run pods as root*
 - *Trusting your nodes*
 - *Disable basic auth*
 - *Namespace isolation*
 - *Identity developers and robots*
 - *Other authorization engines*
 - *Pod security policy (PSP)*



bit.ly/SamK8sSec



KubeCon



CloudNativeCon

Europe 2020



Virtual



KEEP CLOUD NATIVE

CONNECTED

