This year, it's about security

Brandon Baker & Maya Kaczorowski, Google Cloud
Dec 11 2018

Brandon Baker

Cloud Security Horizontal Lead,
Google Cloud

Maya Kaczorowski

Security PM, Google Cloud

@MayaKaczorowski

Google Cloud

What's happened this year

1. Kubernetes attacks in the wild
2. Developments in isolation
3. Software supply chain
4. Hardening and what's coming in 2019

Google Cloud

Kubernetes attacks in the wild

Threats seen in the wild

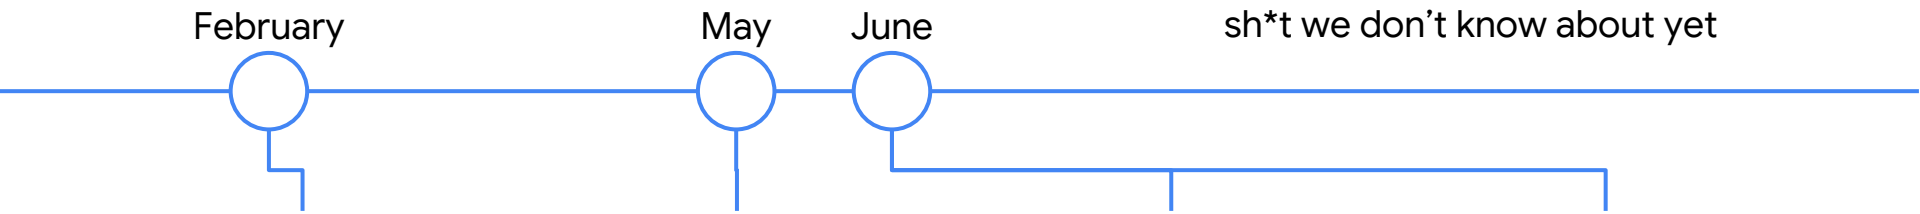

Tesla

Unsecured Kubernetes dashboard with cloud account credentials

Used to mine cryptocurrency

Shopify

Researcher could access and replay kubelet credentials

Not exploited

Weight Watchers

Unsecured Kubernetes dashboard with sensitive data, including credentials

Not exploited

Docker Hub

Public images with embedded cryptocurrency mining malware

Used to mine cryptocurrency

Google Cloud

Unsecured Kubernetes dashboard

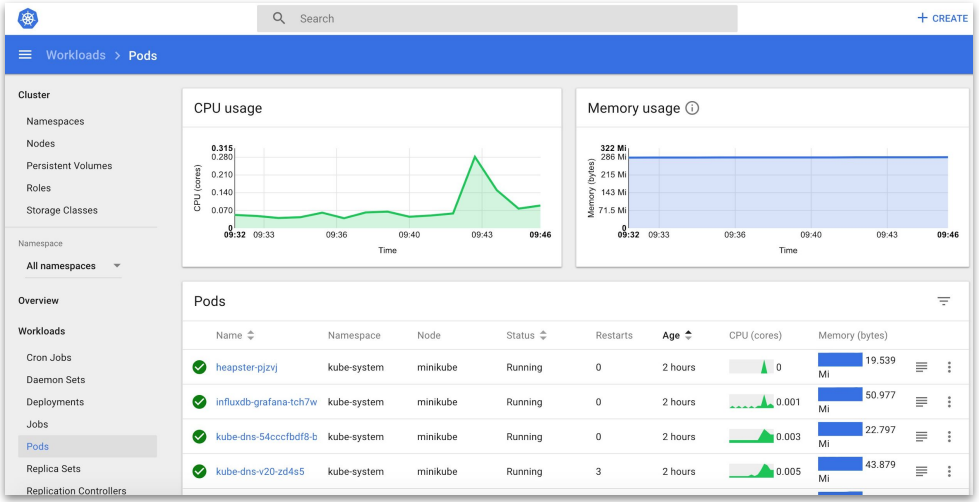

- Hackers accessed the Kubernetes console, which was **not password protected**
- Console contained **privileged cloud account credentials**
- Used credentials to access resources and **mine cryptocurrency**

Google Cloud

Shopify's cluster non-compromise

0xacb submitted a report to **Shopify**.

The Exploit Chain - How to get root access on all Shopify instances

- Access Google Cloud metadata
- Dump kube-env
- Execute arbitrary commands using kubelet
- Profit

Learn more:
Thurs Dec 13th 4:30-5:05pm
https://sched.co/GrZf

Google Cloud

https://hackerone.com/reports/341876

Docker Hub cryptocurrency mining

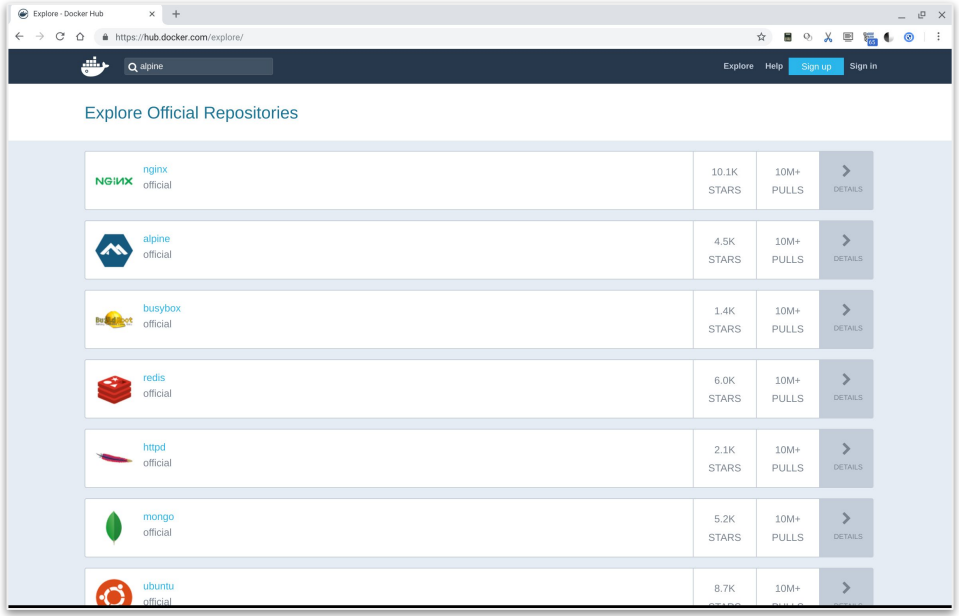

- Hackers made **17 malicious images** available on Docker Hub, like docker123321
- Malware included cryptomining software, netting **~$90k of Monero in ~1 year**

Google Cloud

What we are seeing: drive-by scanning

What we’re not seeing: container escape

Google Cloud

Developments in isolation

Threats can come from within

External to the container:

- DDoS, service disruption
- Data theft
- Cryptomining

What most people typically think about

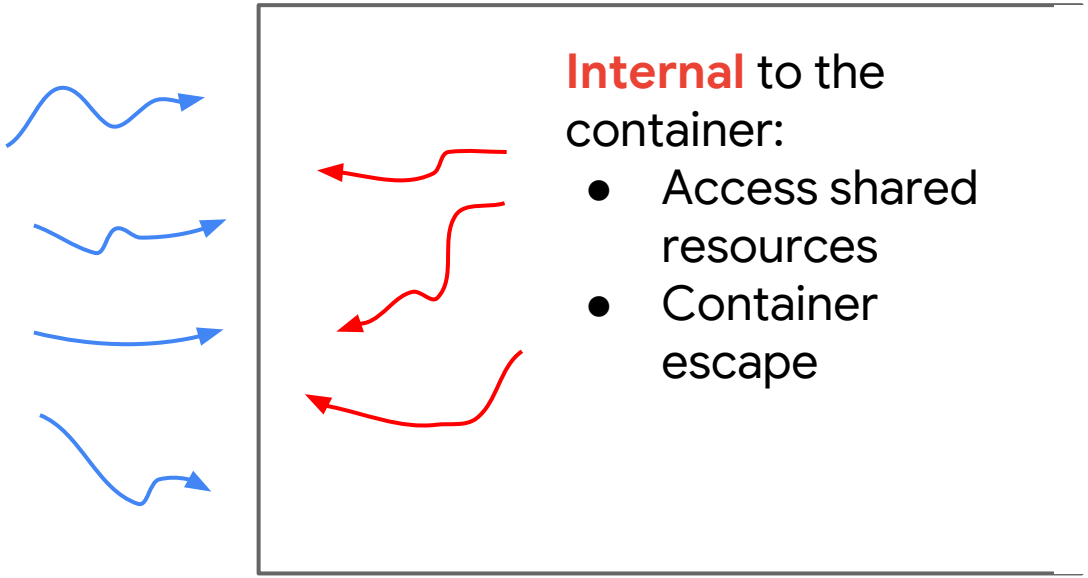

Internal to the container:

- Access shared resources
- Container escape

Layers of isolation in Kubernetes

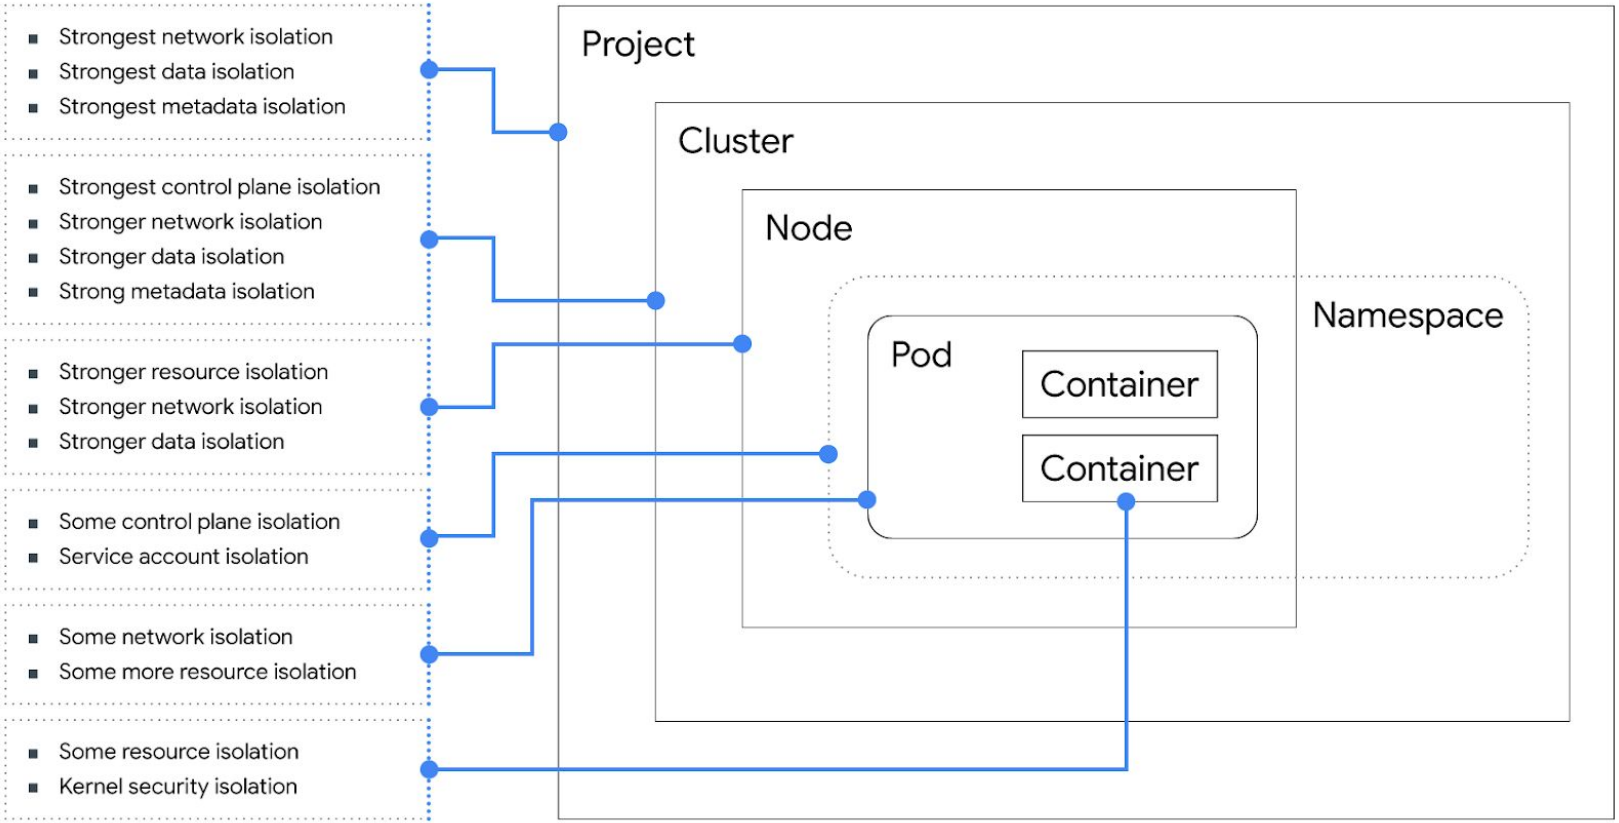

Developments in stronger isolation

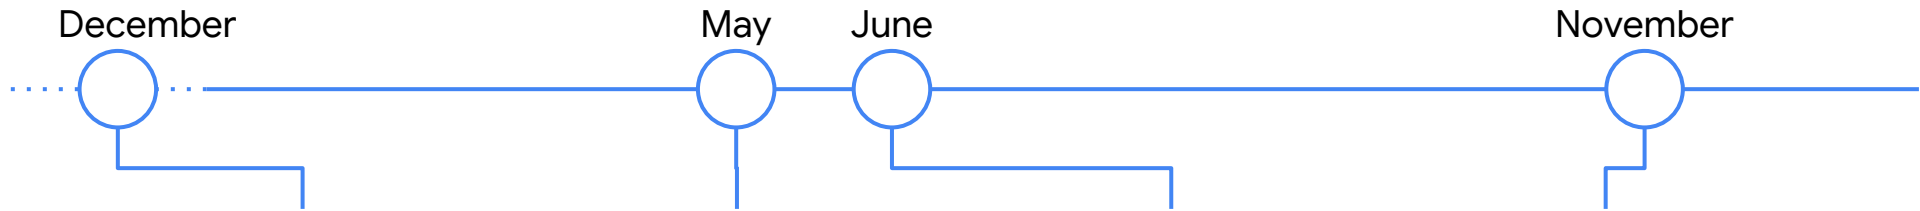

Kata Containers

Lightweight VM with stripped down guest kernel

Merger of Clear Containers + Hyper runV

gVisor

Intercepts system calls by acting as a guest kernel in user space

From Google

Nabla

Limits system calls using unikernel, and blocks the rest with seccomp

From IBM

Firecracker

Lightweight microVM meant for running in a non-virtualized environment

From AWS

Google Cloud

Sandboxing in Kubernetes

Kubernetes API to sandbox containers

Pod level isolation

Multiple containers in Pod

Two isolation boundaries

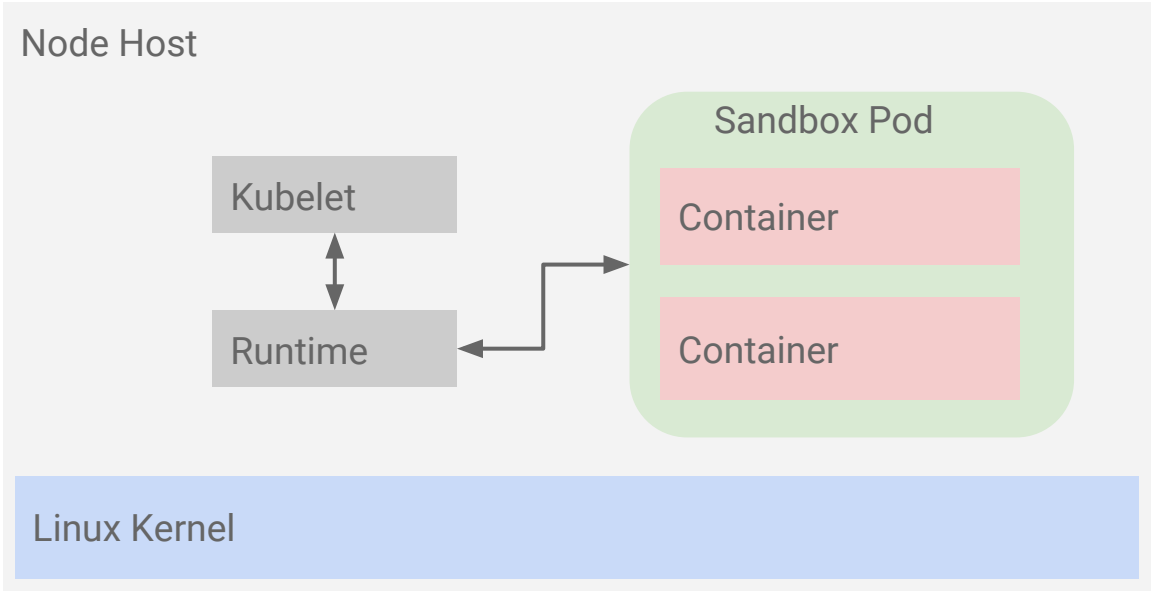

Google Cloud

RuntimeClass

RuntimeClass is a new API to specify runtimes

Specify the RuntimeClass in your Pod spec

```
apiVersion: v1
kind: RuntimeClass
metadata:
  name:gvisor
spec:
  parameters:
    io.containerd.runtime: gvisor
  support:
    linux:
      capabilities: [ '*' ]
      privileged: true
      namespaces:
        network: [ Pod ]
        PID: [ Pod, Container ]
        IPC: [ Pod ]
    ...
```

```
apiVersion: v1
kind: Pod
...
spec:
  ...
  runtimeClassName: gvisor
```

Google Cloud

Software supply chain

Ideal, security-hardened container supply chain

Base image	Code	Build	Application image	Deploy
Controlled base images	Static analysis	Hermetic	Vulnerability scanning	Admission control
Hash based addressing	Dependency analysis	Reproducible	Configuration scanning	Runtime configurations
		Rootless		

How could you have avoided the Docker Hub images with cryptocurrency mining?

Don’t pull images from public repositories directly

Scan your images for vulnerabilities, malware, and other security issues

Only deploy images you’ve scanned

Google Cloud

“Only deploy vulnerability-free images”

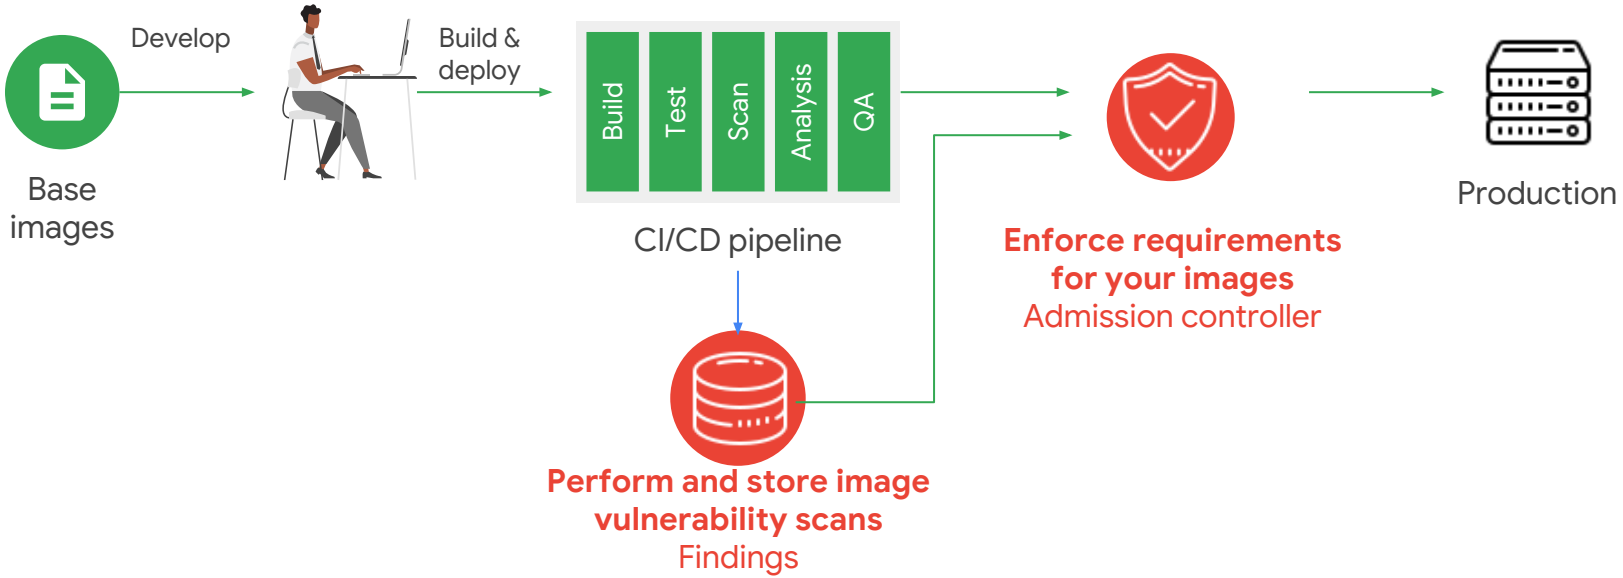

Google Cloud

Grafeas

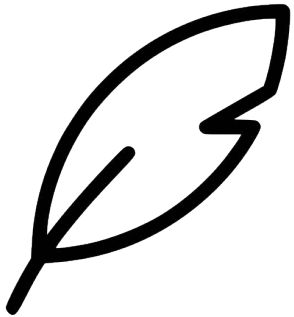

- Structured artifact metadata repository
  - Meant to be used as part of a container registry
- Spec includes multiple kinds of metadata
  - Package, Vulnerabilities, Discovery, Builds, Image basis, Deployment history, Attestation
- Can use multiple metadata providers
  - Providers include other scanning companies, e.g., JFrog, Red Hat, IBM, Black Duck, Twistlock, and Aqua

Kritis

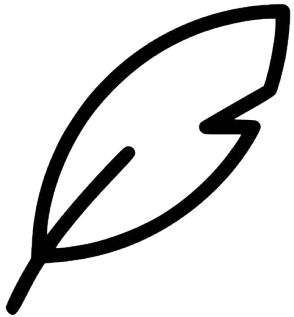

- Signing and deploy enforcement tool for Kubernetes
  - Implemented as a Kubernetes admission controller
  - Integrates with Grafeas attestation metadata APIs
- Generate attestations based on your requirements
  - Build provenance
  - Vulnerability findings

Google Cloud

Open source: Grafeas & Kritis

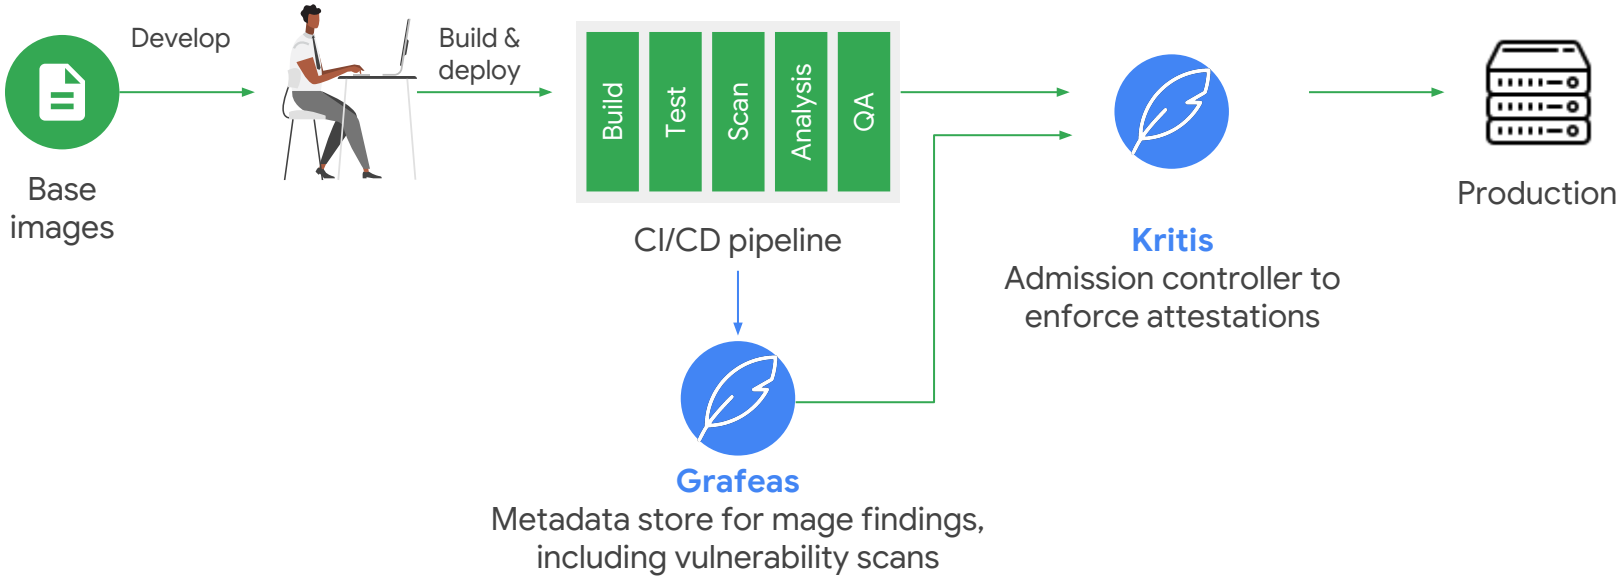

Google Cloud

https://github.com/grafeas/kritis/blob/master/docs/tutorial.md

Google Cloud: GCR Vulnerability Scanning and Binary Authorization

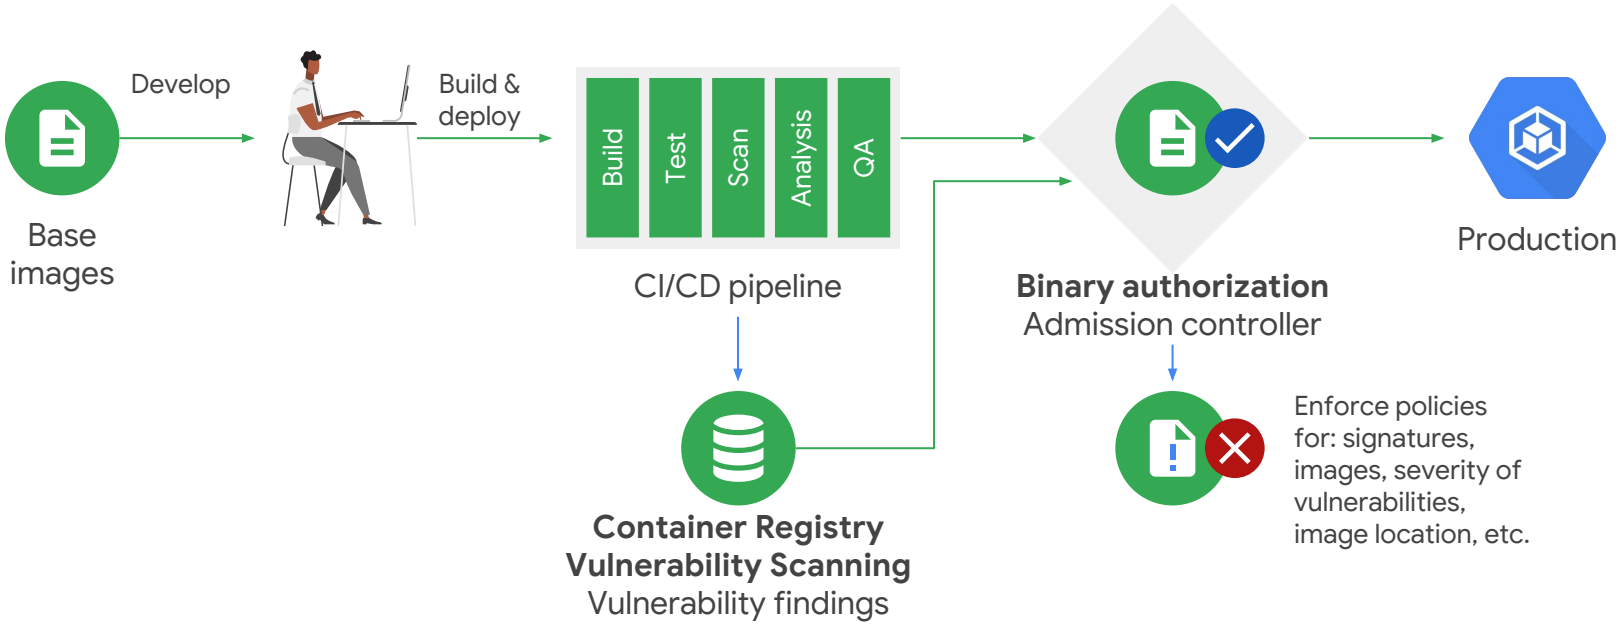

Google Cloud

Enforced governance

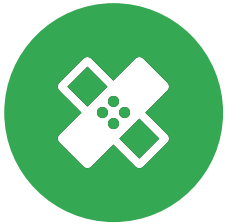

Containers are short lived and frequently re-deployed, **you can constantly be patching**.

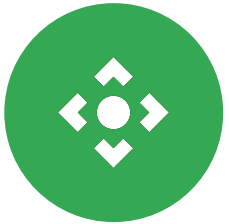

Containers are immutable, **you can control what is deployed in your environment**.

Google Cloud

Hardening and what's coming in 2019

Vulnerabilities in many layers affect your Kubernetes distribution

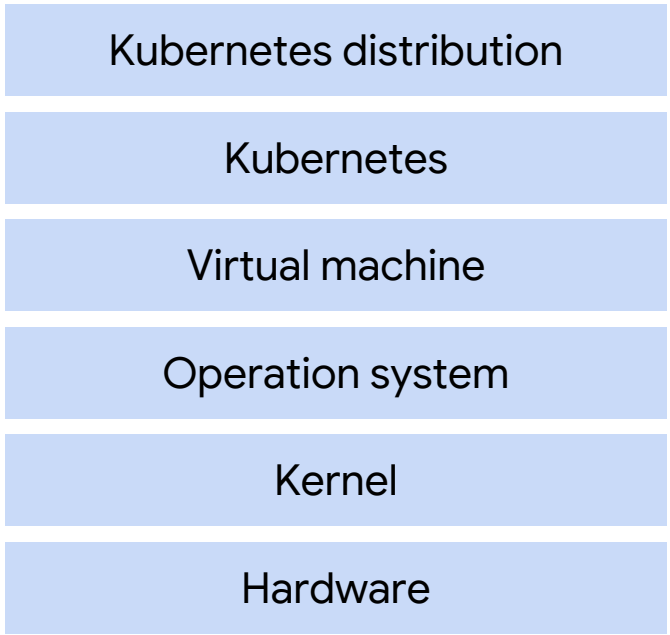

Google Cloud


The simplest thing you can do to improve your security is...
Keep your Kubernetes version up to date!

Best practices to harden your clusters

Set up a cluster

- Restrict access to kubectl
- Use RBAC
- Use a Network Policy
- Use namespaces
- Bootstrap TLS

Prevent known attacks

- Disable dashboard
- Disable default service account token
- Protect node metadata
- Scan images for known vulnerabilities

Follow security hygiene

- Keep Kubernetes updated
- Use a minimal OS
- Use minimal IAM roles
- Use private IPs on your nodes
- Monitor access with audit logging
- Verify binaries that are deployed

Prevent/limit impact of microservice compromise

- Set a Pod Security Policy
- Protect secrets
- Consider sandboxing
- Limit the identity used by pods
- Use a service mesh for authentication & encryption

Google Cloud

2019 prediction:

More hardening

Another 2019 prediction: more attacks

- Container-specific attacks
  - Container escape in the wild?
  - Continued supply chain attacks
  - Better detection: IDS/IPS-like solutions for containers

Learn more

cloud.google.com/containers/security
g.co/gke/security
g.co/gke/hardening

Q&A

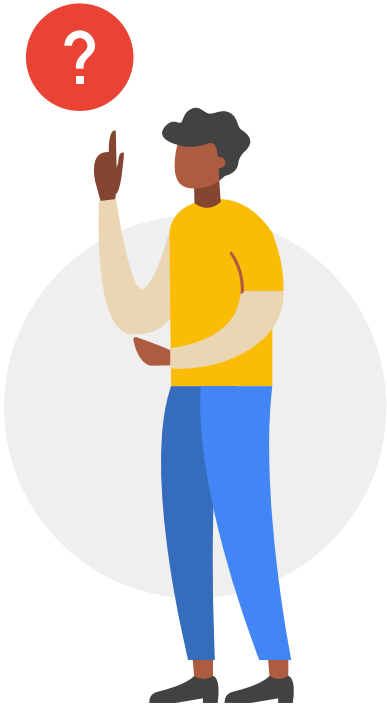

Google Cloud