

SECURING CONTAINERS

What's so different?

Joel Lathrop

JOEL@DIDACTIC-SECURITY.COM

Triangle InfoSeCon | October 2018



DIDACTIC
SECURITY



Software development

Automation

DevOps

Consulting

Forensics

Pen testing

Malware reversing

Cryptography

“Containers don’t contain.”

– Dan Walsh


“... but they sure are handy!”

– your developers

GOALS


- ▶ What is a container?
- ▶ What are the security implications?
- ▶ How do I respond?

ANATOMY



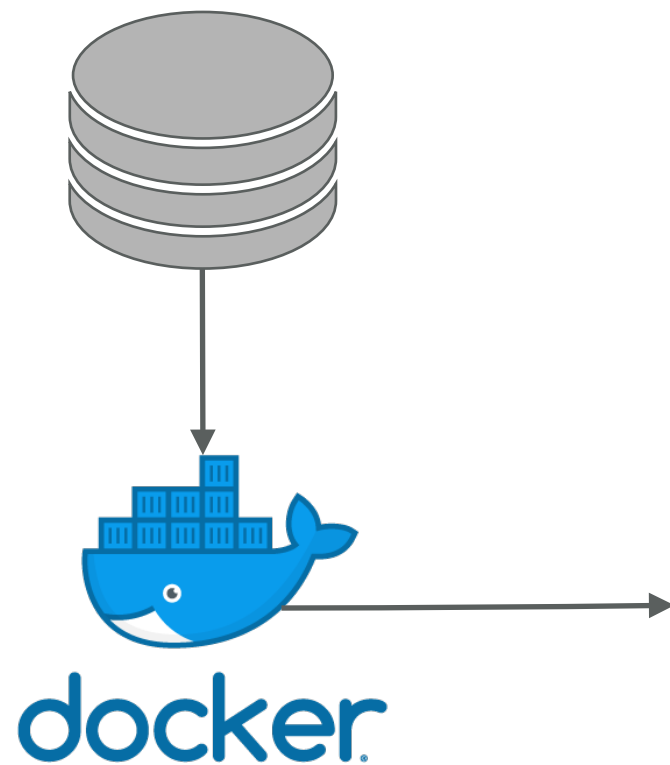
ANATOMY

Image Registry



ANATOMY

Image Registry



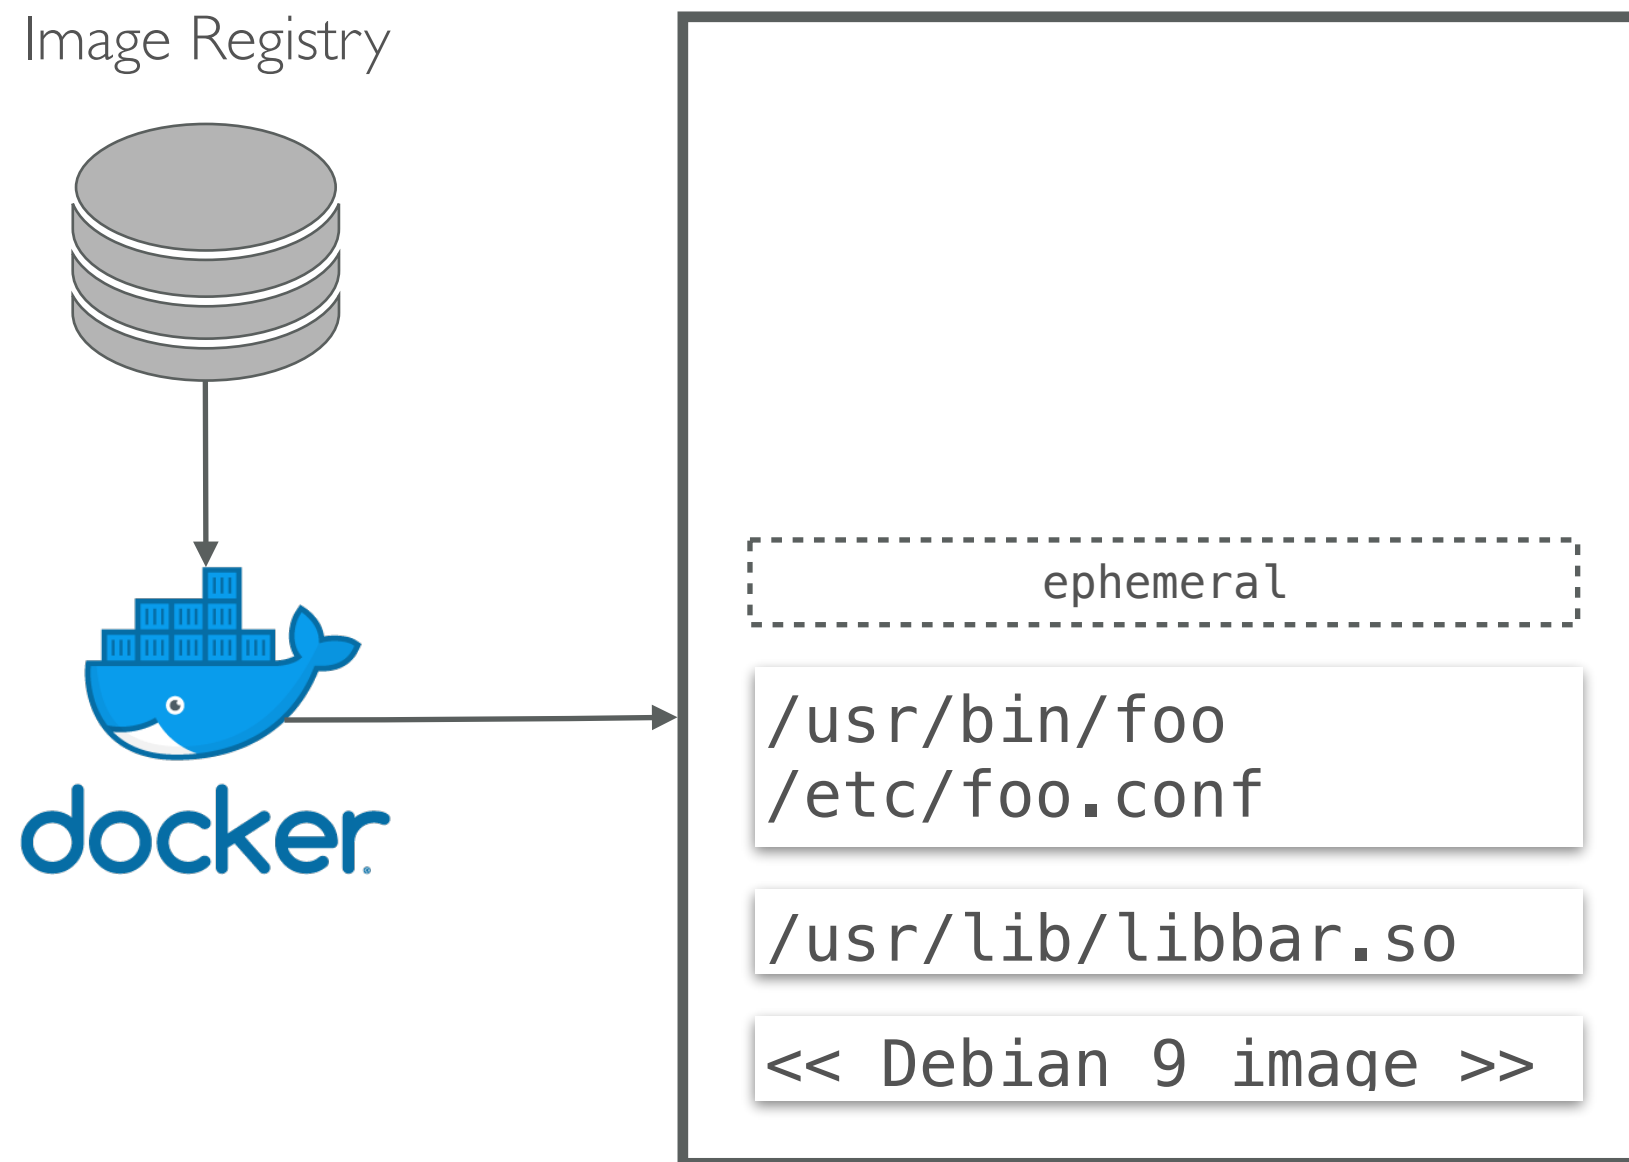
ephemeral

```
/usr/bin/foo  
/etc/foo.conf
```

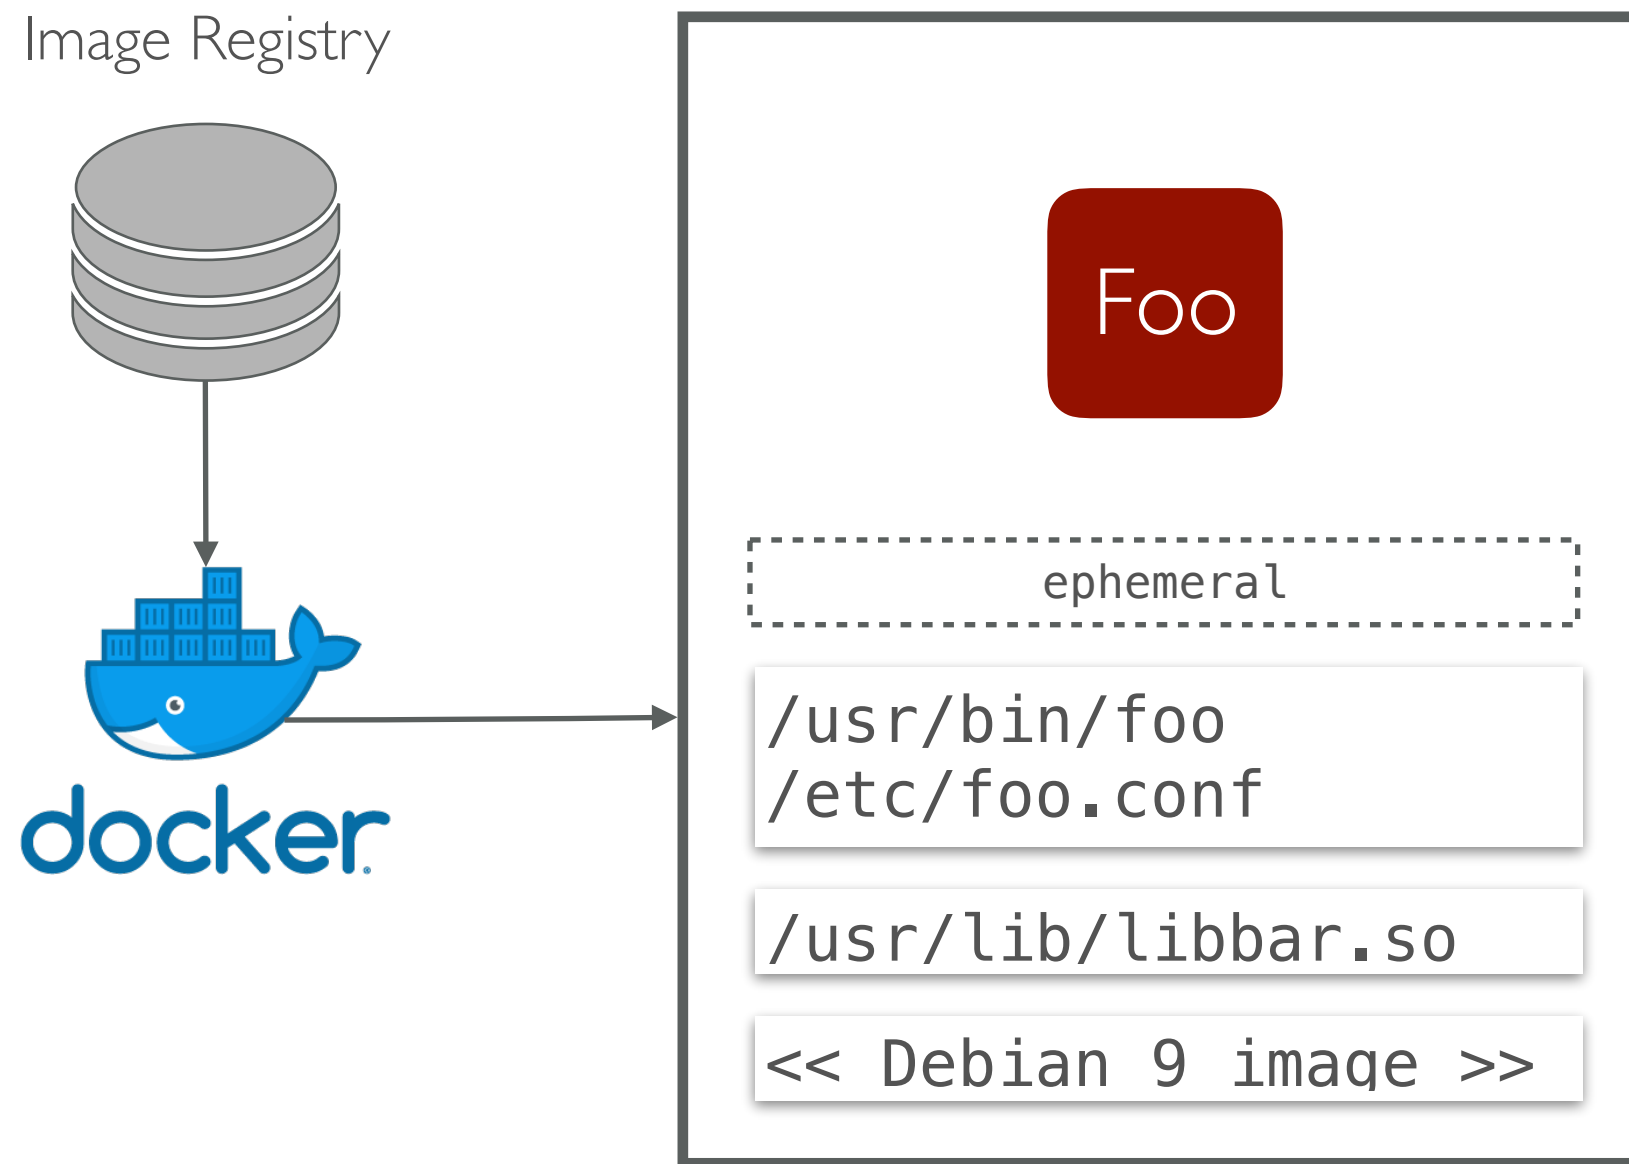
```
/usr/lib/libbar.so
```

```
<< Debian 9 image >>
```

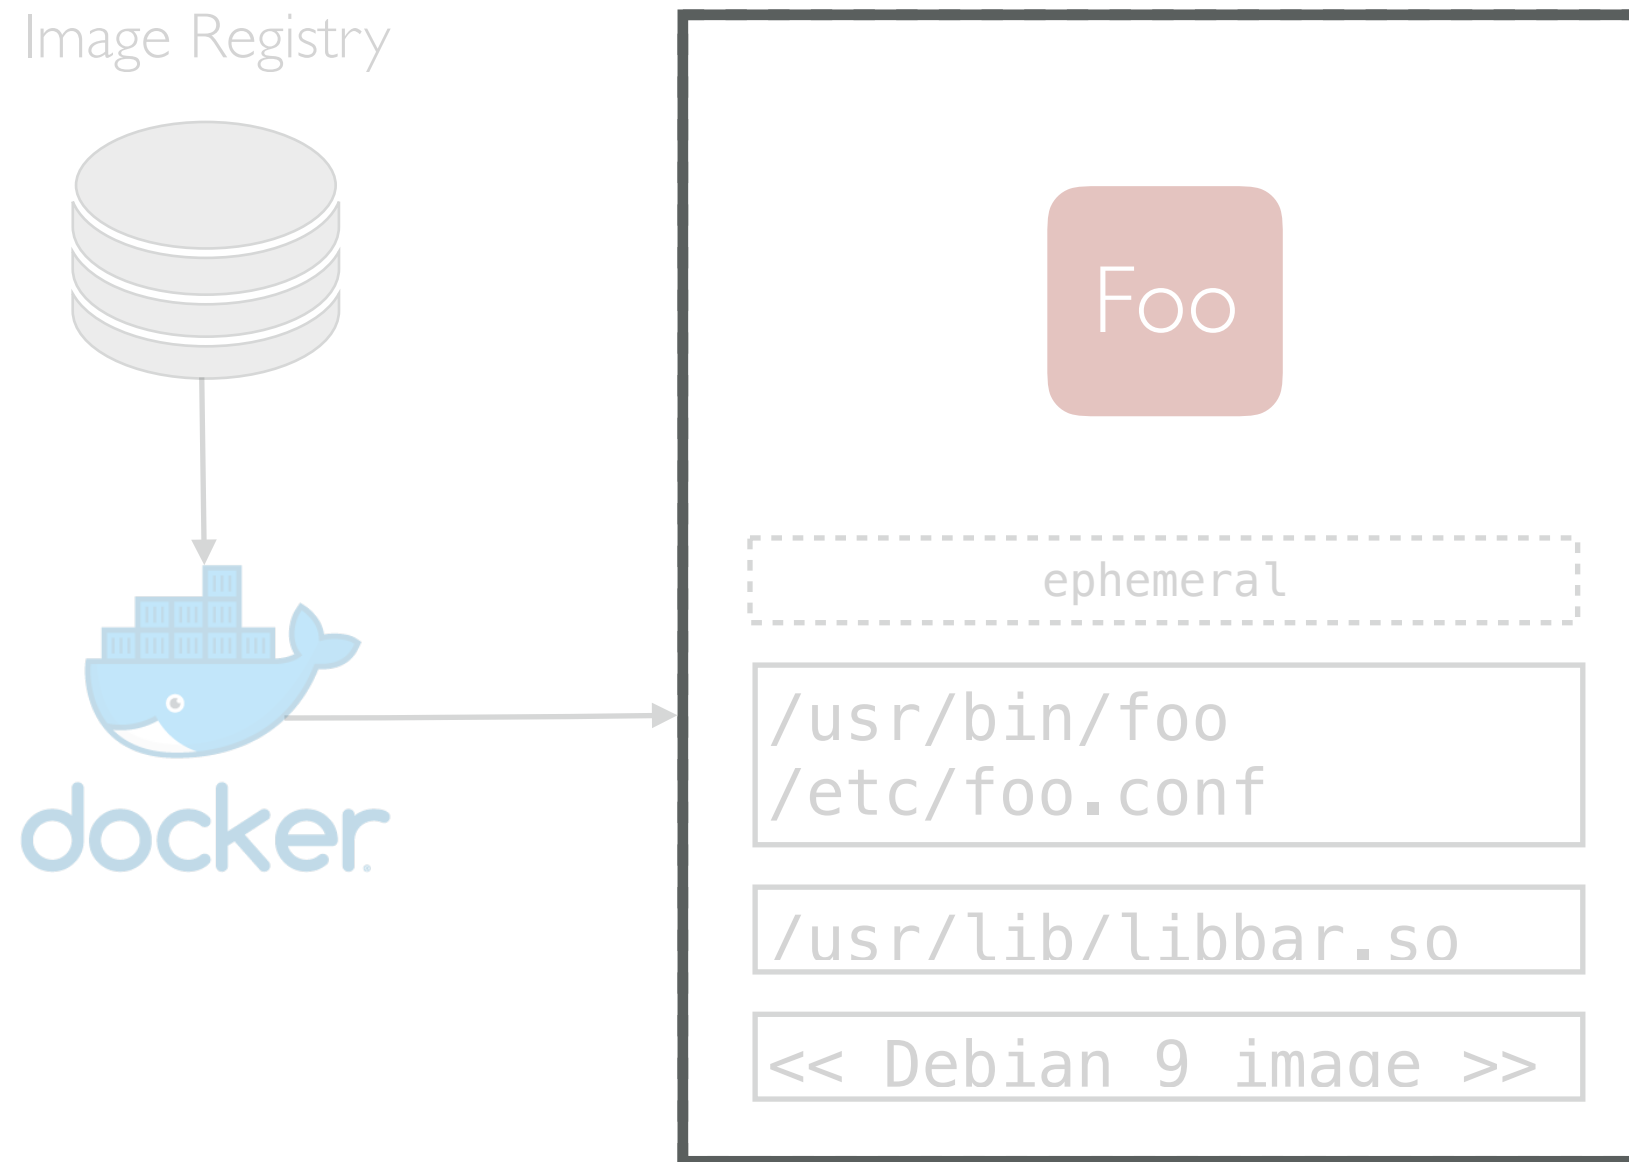
ANATOMY



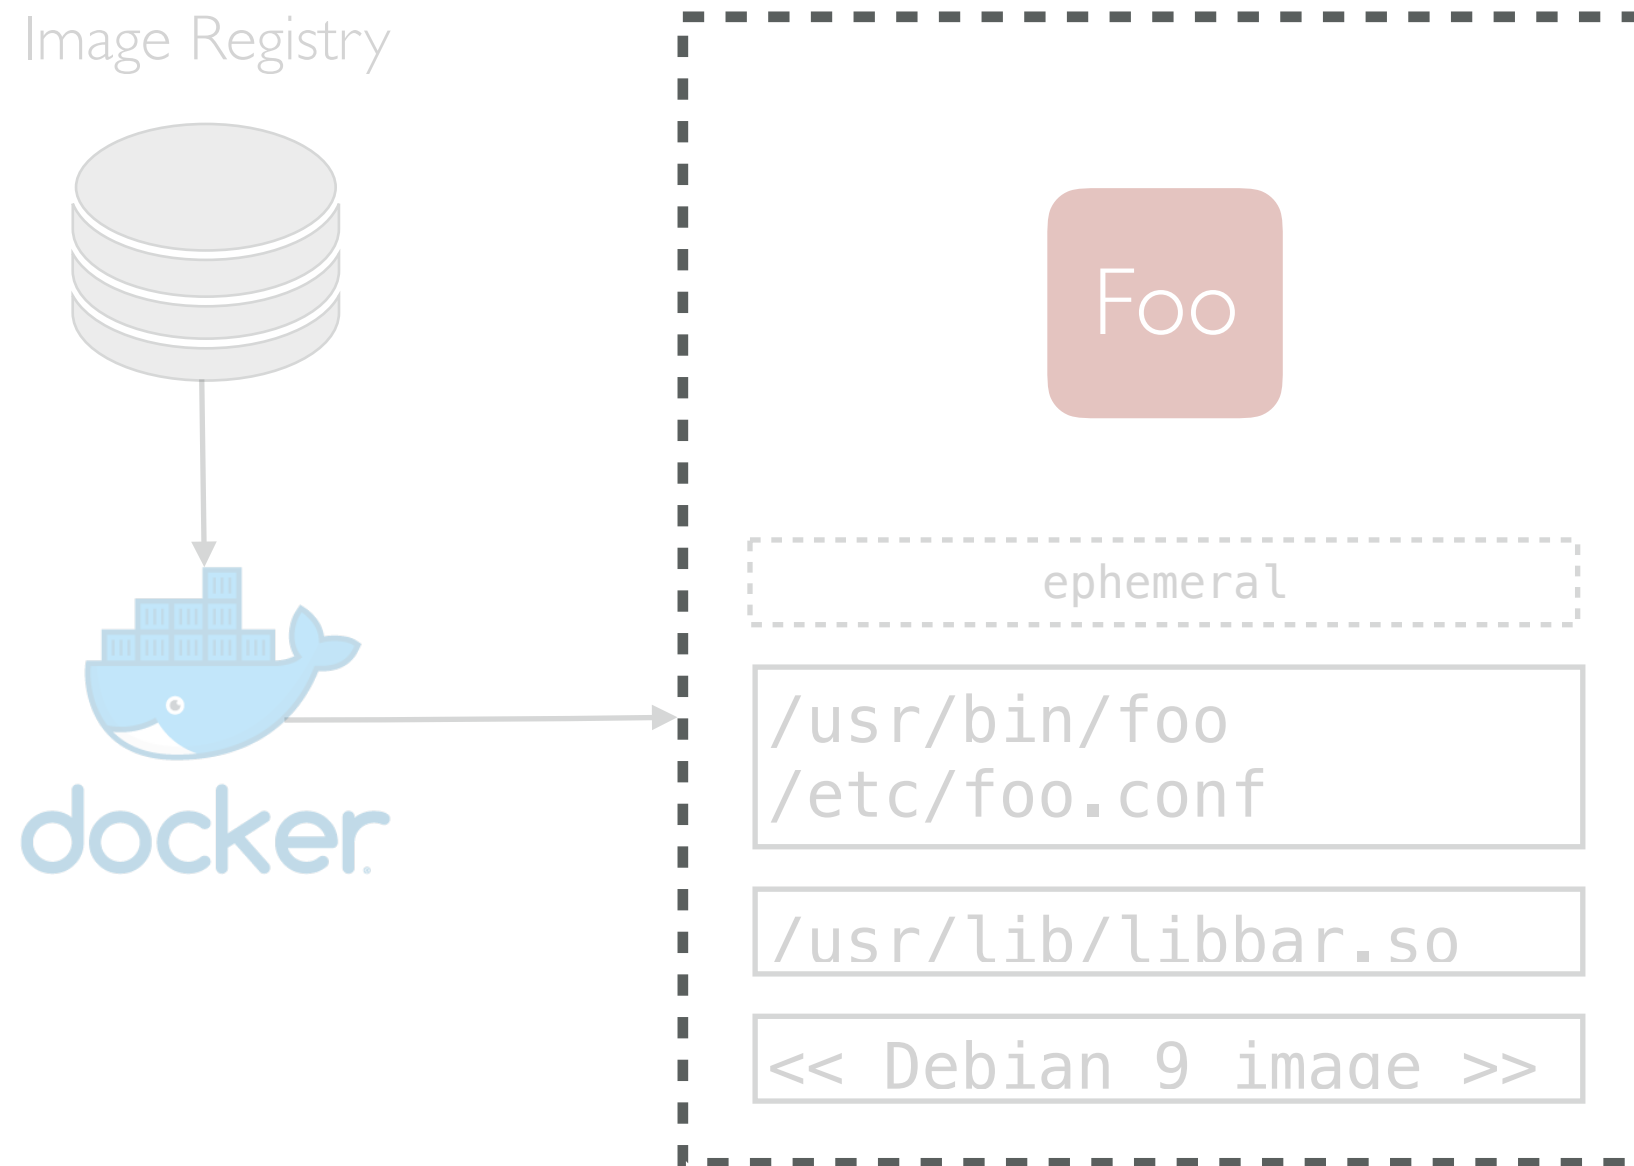
ANATOMY




KERNEL ESCAPES



KERNEL ESCAPES



KERNEL ESCAPES



- ▶ Kernel design is open by default
 - Therefore, trying to close holes
 - Not everything is namespaced
- ▶ Kernel bugs → Container escapes


MINDSET: CONTAINER ISOLATION

A container is **not** like a VM.


A container is like an **application package** running in a *slightly leaky* sandbox.

POLICY: OPERATING SYSTEMS

Use a **container-focused** host operating system.



CoreOS



Project Atomic



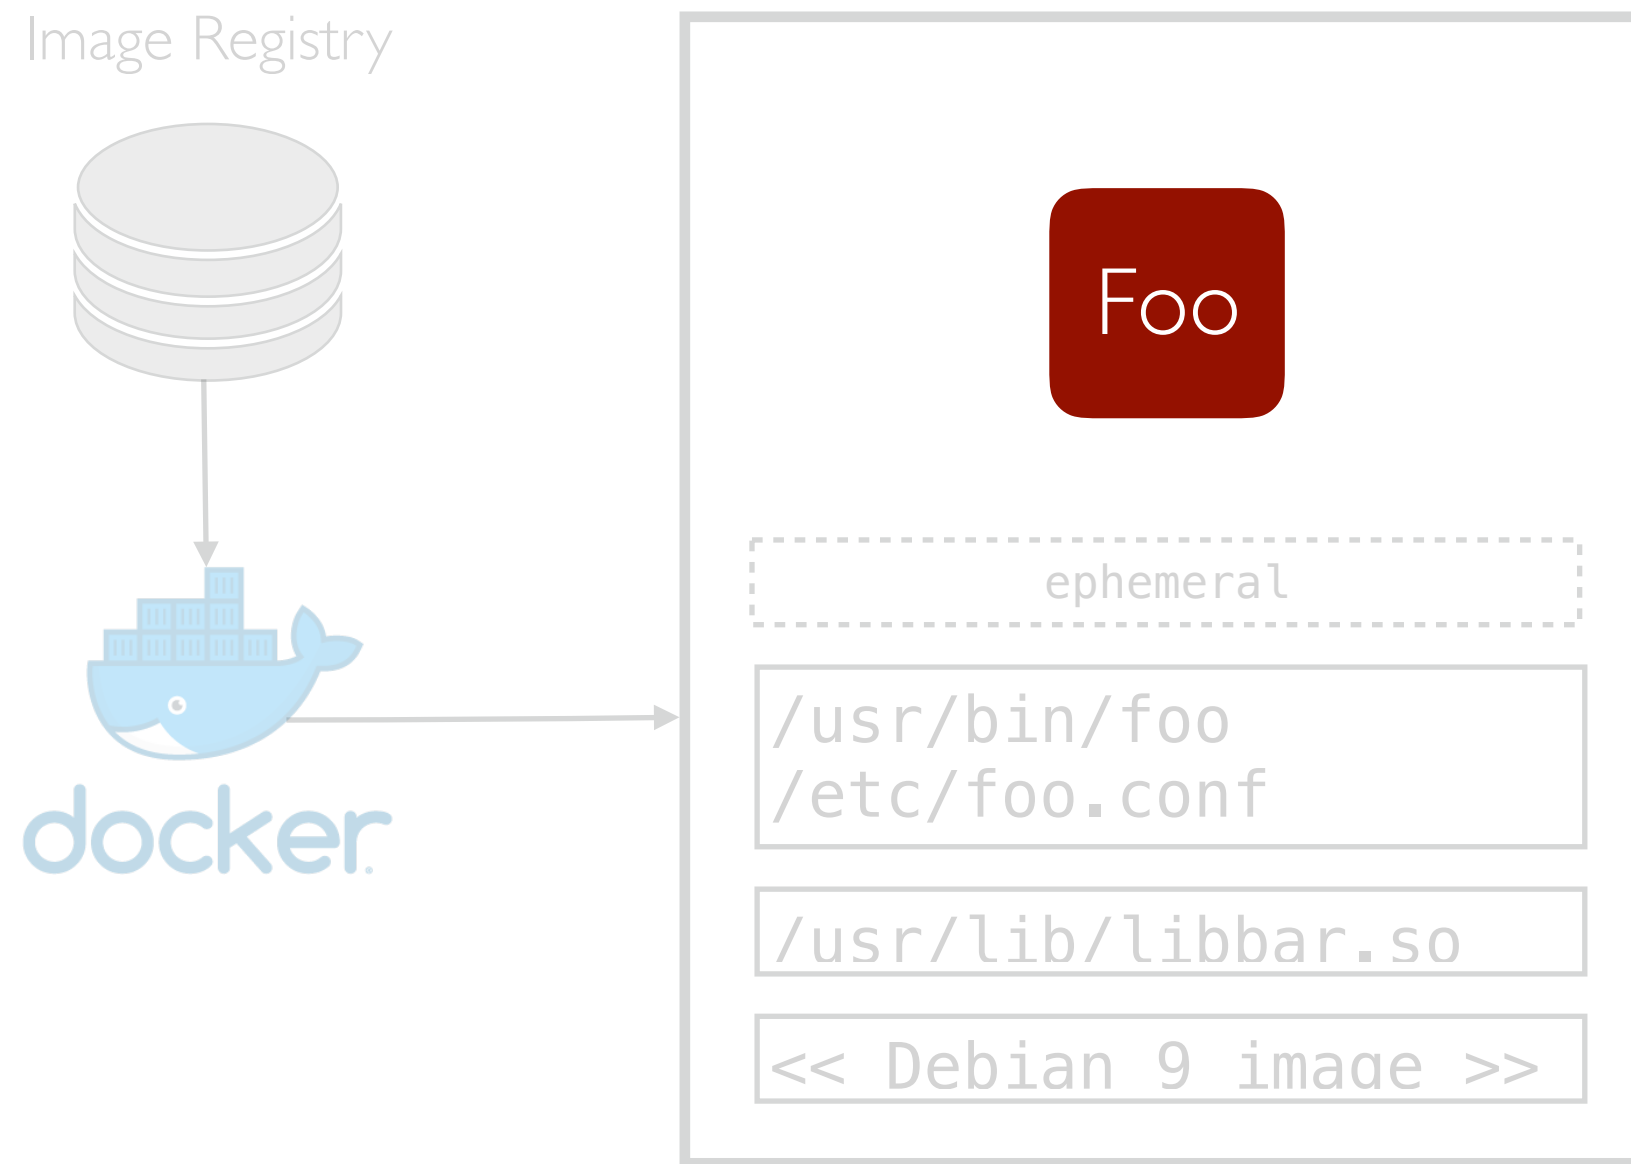
Clear Linux

etc. etc.

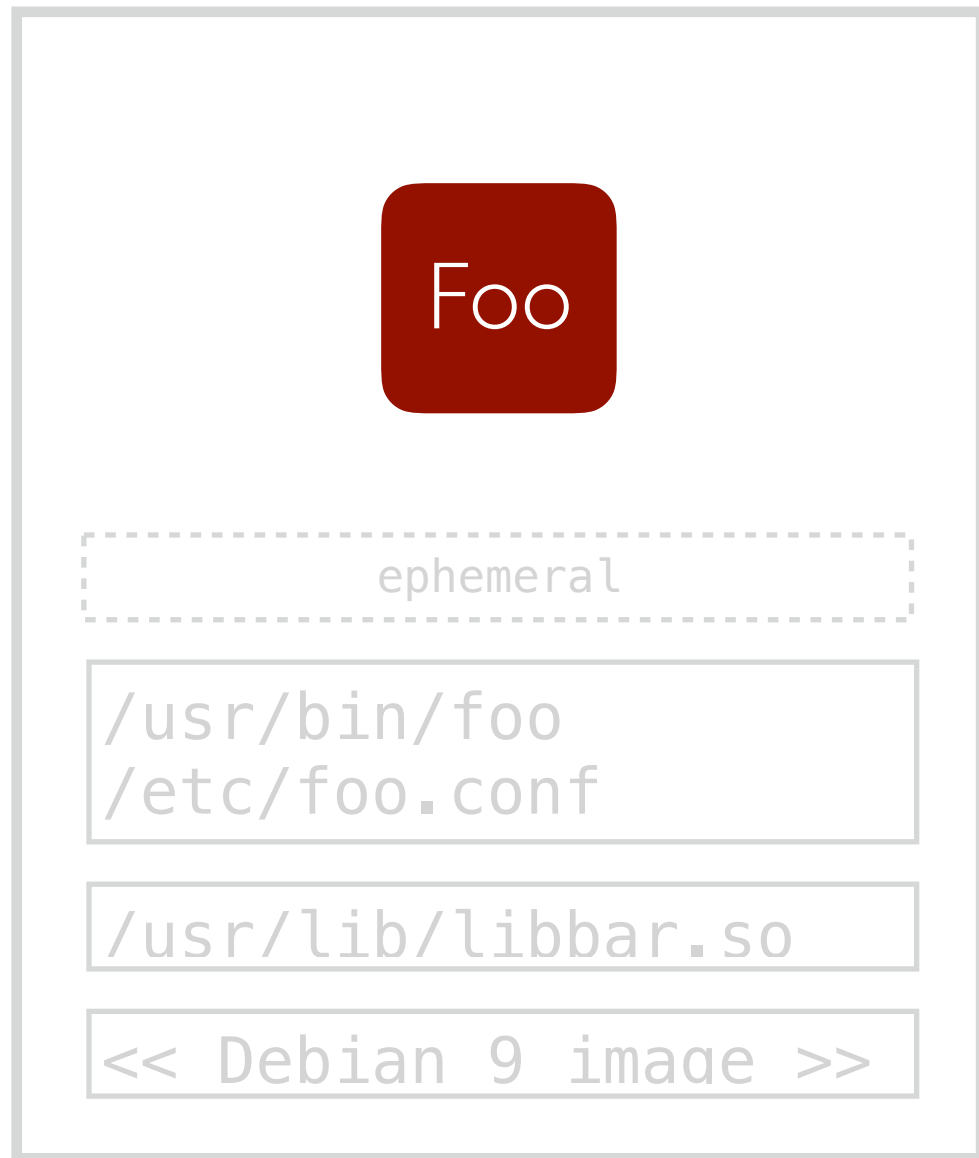
(Google for more)

* All brands are the property of their respective owners.

OVER-PRIVILEGED CONTAINERS



OVER-PRIVILEGED CONTAINERS



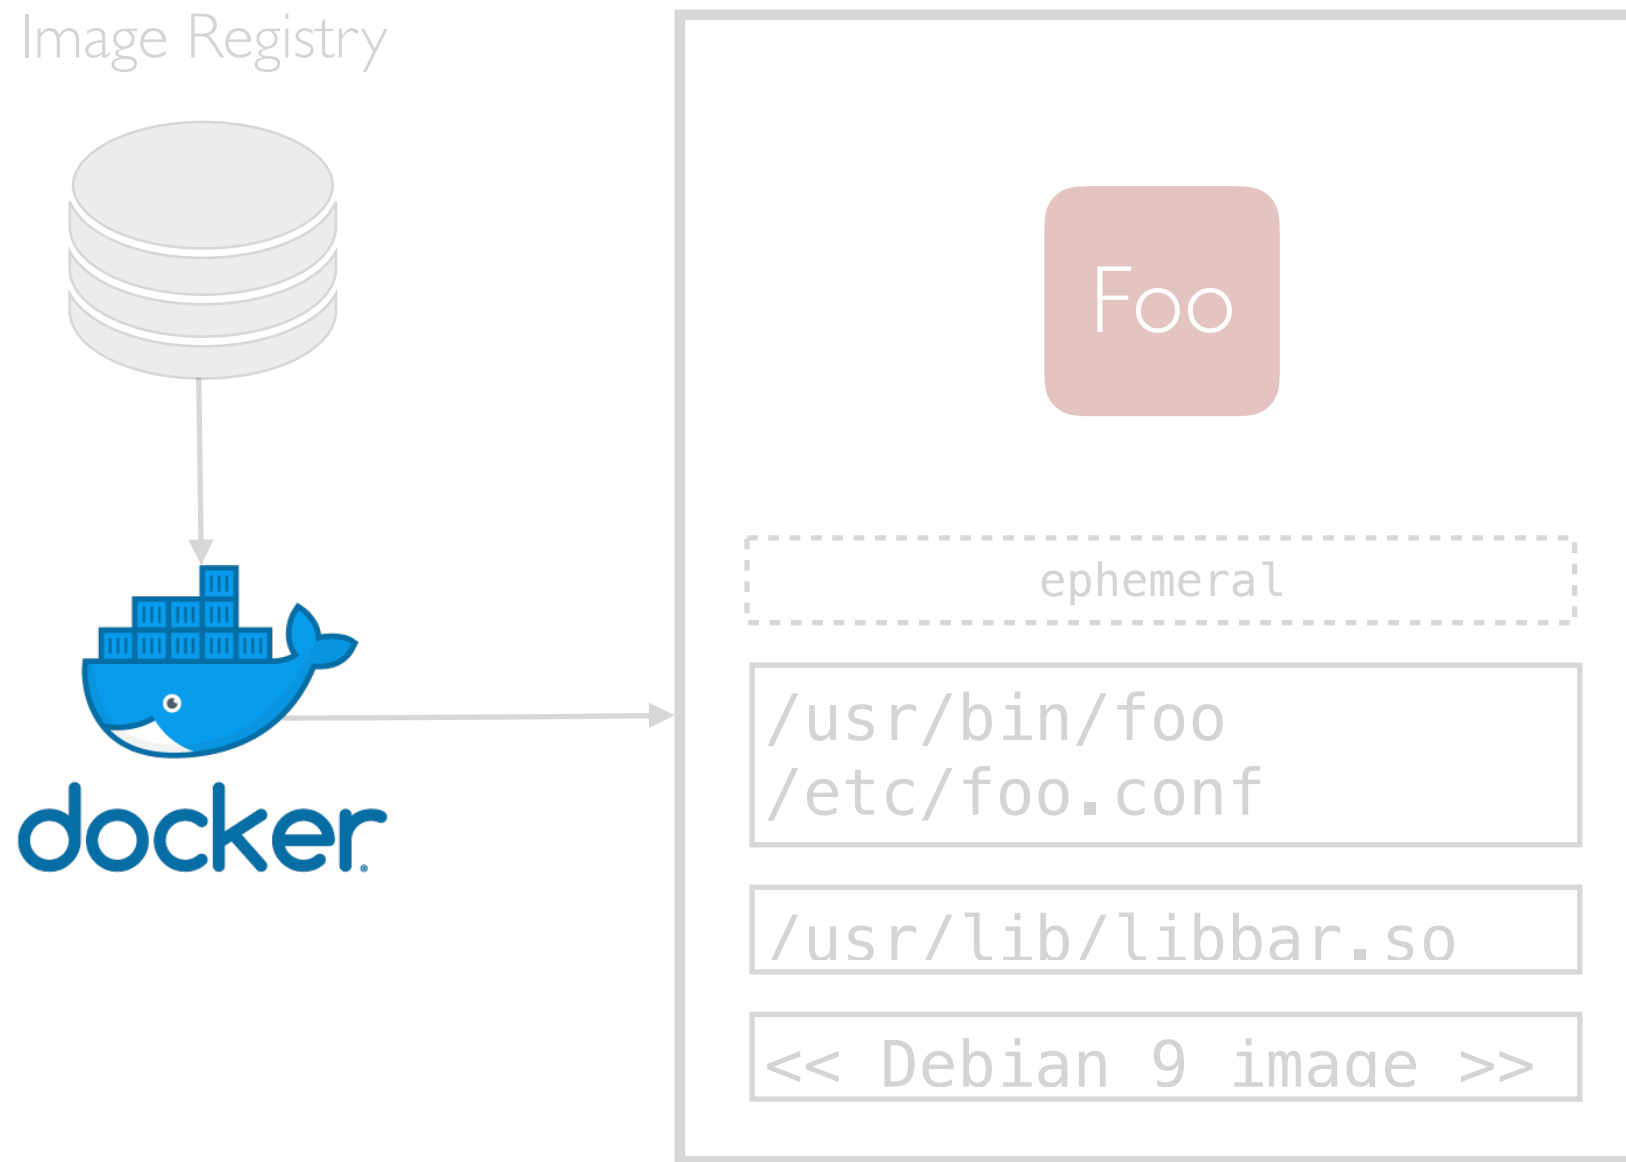
- ▶ The `--privileged` option
- ▶ Exposing the Docker socket
- ▶ Over-reaching volume mounts
- ▶ Unlimited resource allocation

MINDSET: LEAST PRIVILEGE


Don't give a container what you wouldn't give a regular application.

- ▶ Don't run as `root` inside the container
- ▶ Drop as many privileges as you can
- ▶ Minimize what volume mounts expose
- ▶ Set resource constraints

DAEMON DANGERS



DAEMON DANGERS



- ▶ Underlying socket API
 - No authorization required!
 - Can be exposed to the network
- ▶ Run a container \Rightarrow root on the host

MINDSET: LEAST PRIVILEGE

Minimize access to the Docker daemon.

- ▶ Avoid exposing to the network
- ▶ Think of `docker run` like `sudo`

IMAGE DISTRIBUTION

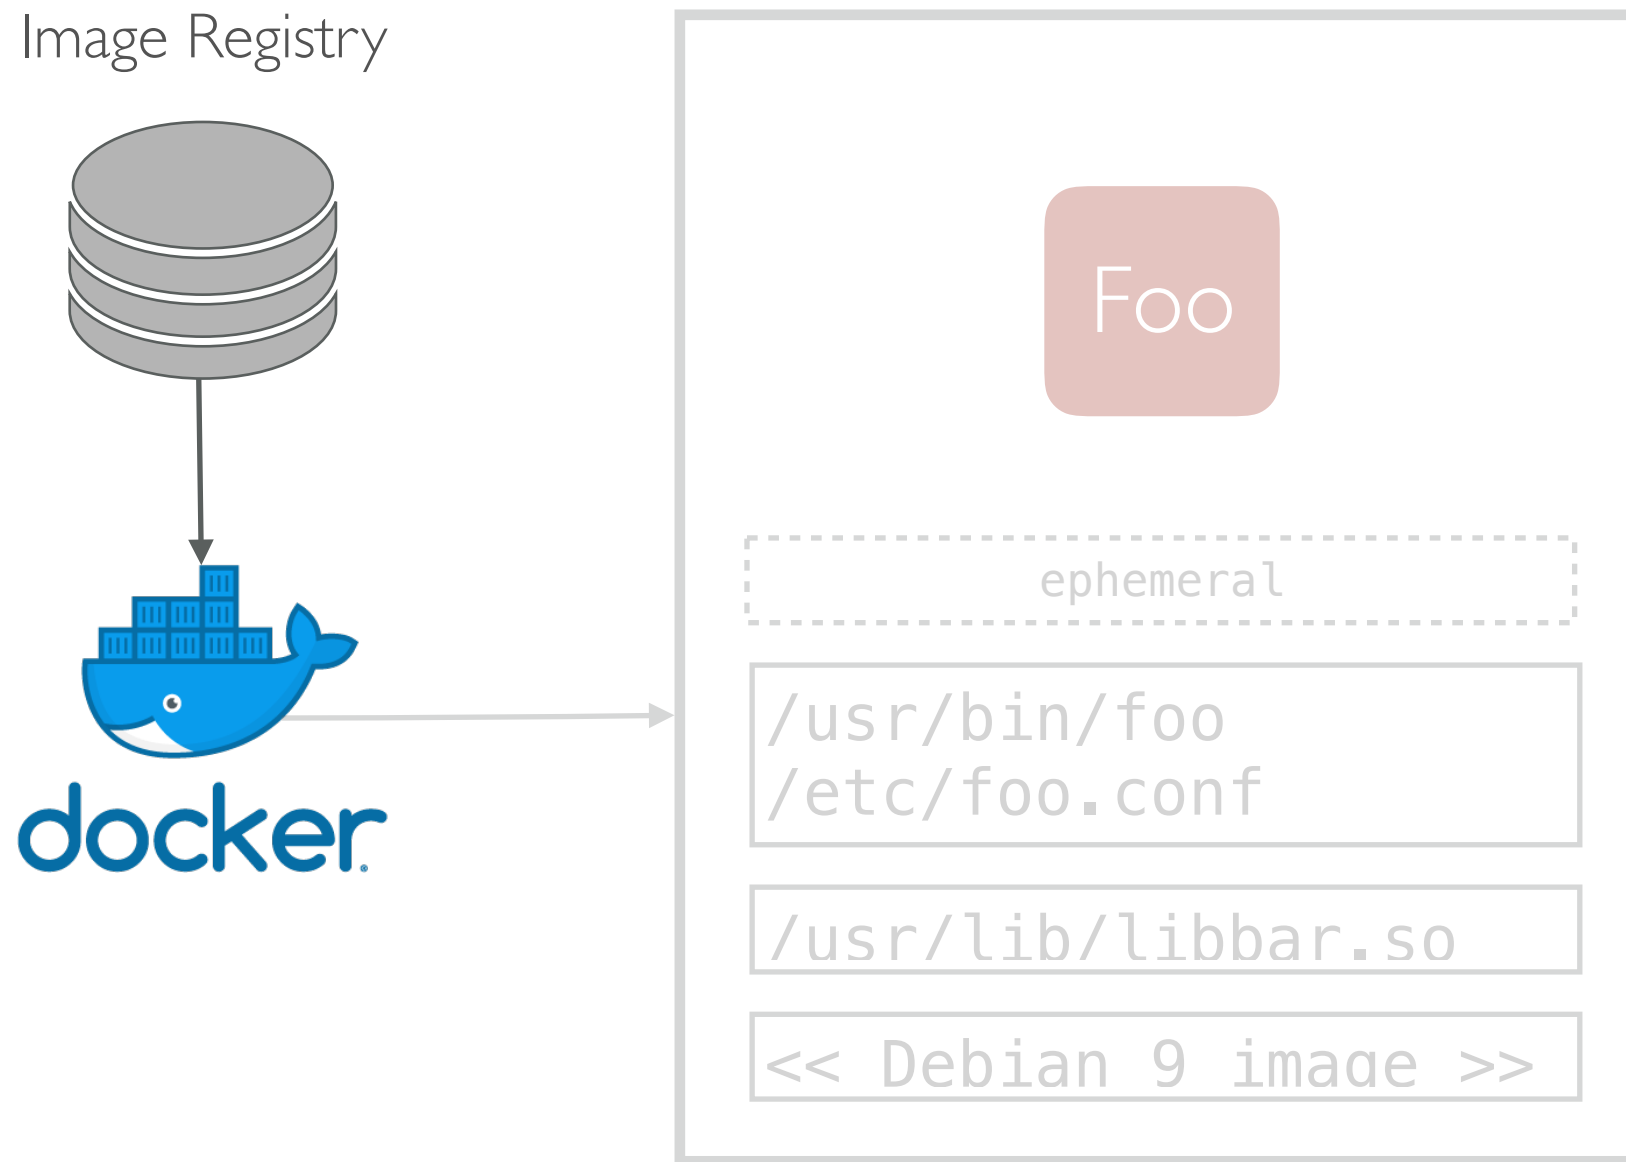
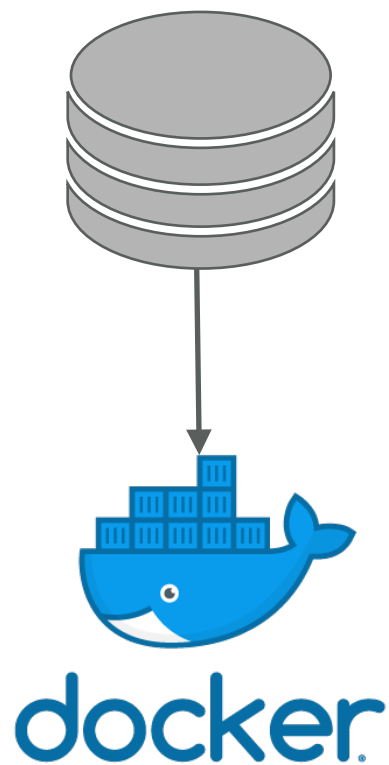


IMAGE DISTRIBUTION

Image Registry



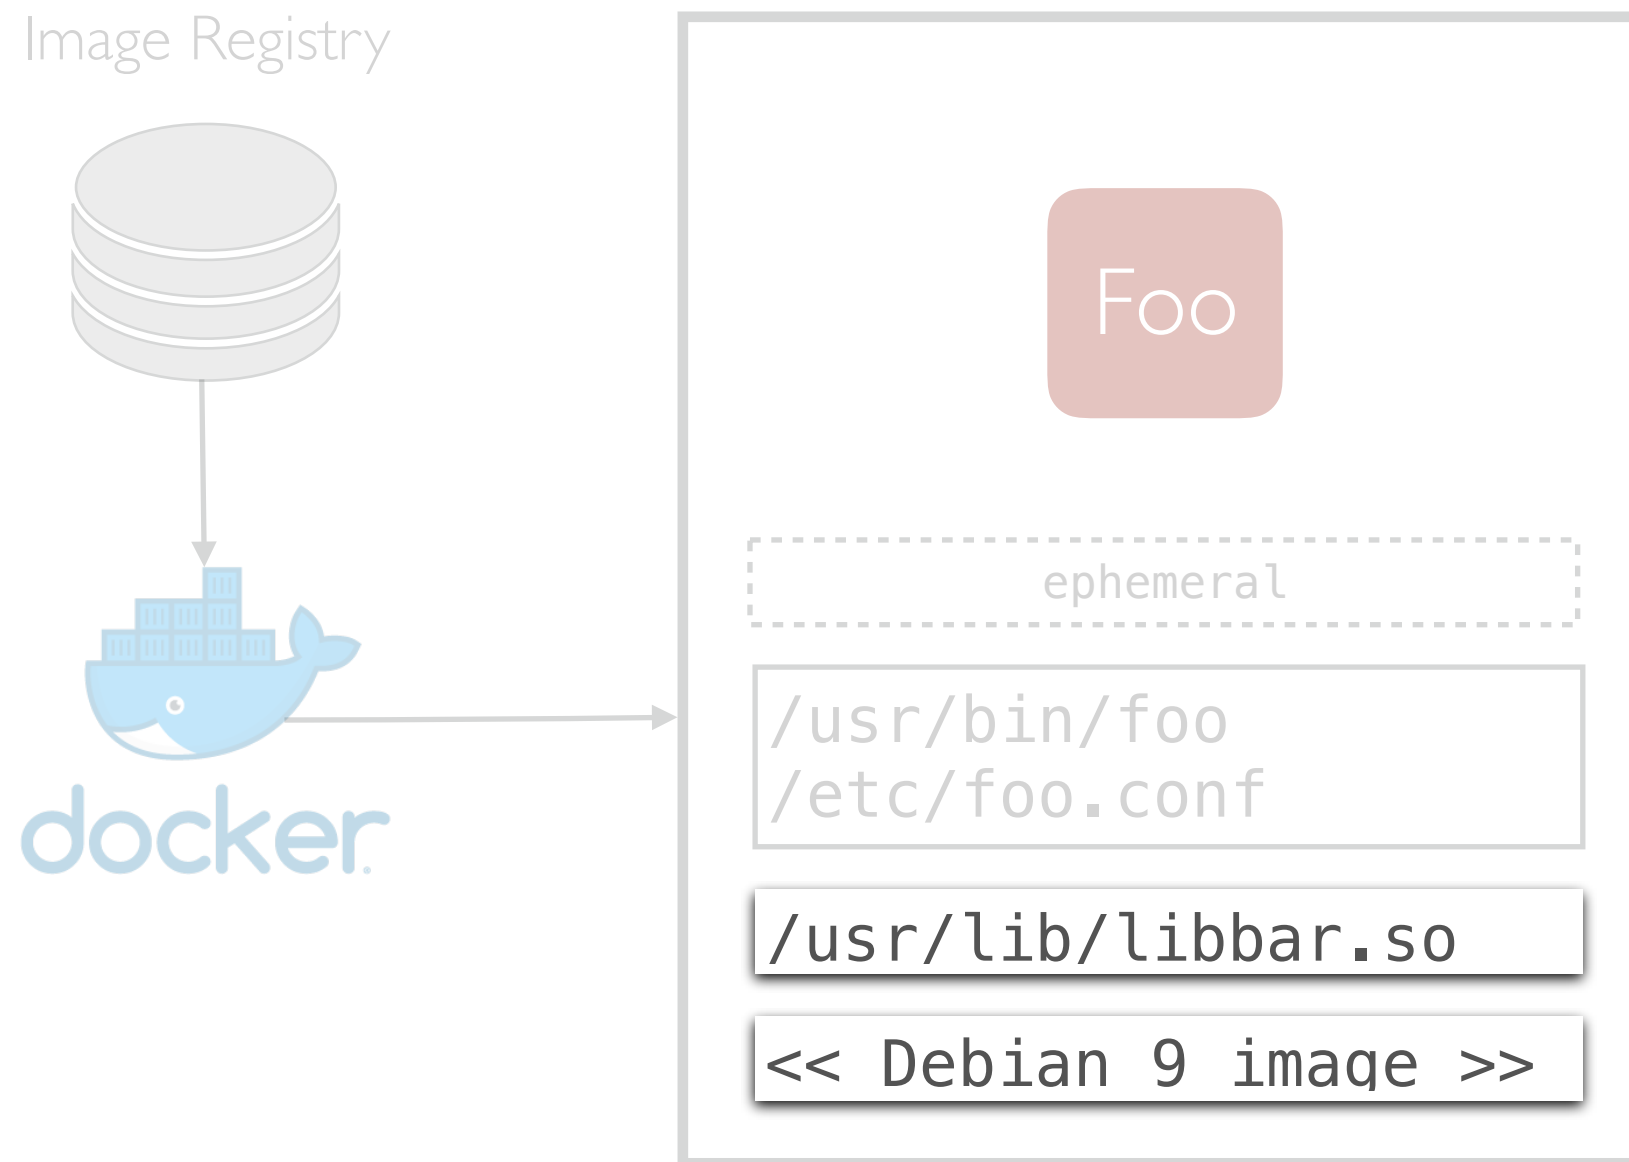
- ▶ Malicious source?
- ▶ Tampering in the registry?
- ▶ Tampering in transit?
- ▶ Tampering on the host system?

POLICY: SIGNED ARTIFACTS

We already do this for regular software updates.

- ▶ Unfortunately, competing approaches
- ▶ Some implementations raise issues
 - Docker Content Trust uses ToFU
- ▶ Consider running an in-house registry

VULNERABLE DEPENDENCIES



VULNERABLE DEPENDENCIES



ephemeral

```
/usr/bin/foo  
/etc/foo.conf
```

```
/usr/lib/libbar.so
```

```
<< Debian 9 image >>
```

- ▶ In direct dependencies
- ▶ In the base image
- ▶ Cannot idiomatically patch at run-time

MINDSET: RE-BUILD IMAGES

The old way was...

- ▶ Install a patch
- ▶ Restart the application

The new way is...

- ▶ Rebuild the image
- ▶ Re-deploy the container

MINDSET: INVENTORY & SCANS

Simpler to scan **images**, not hosts.

(Run a private image registry.)

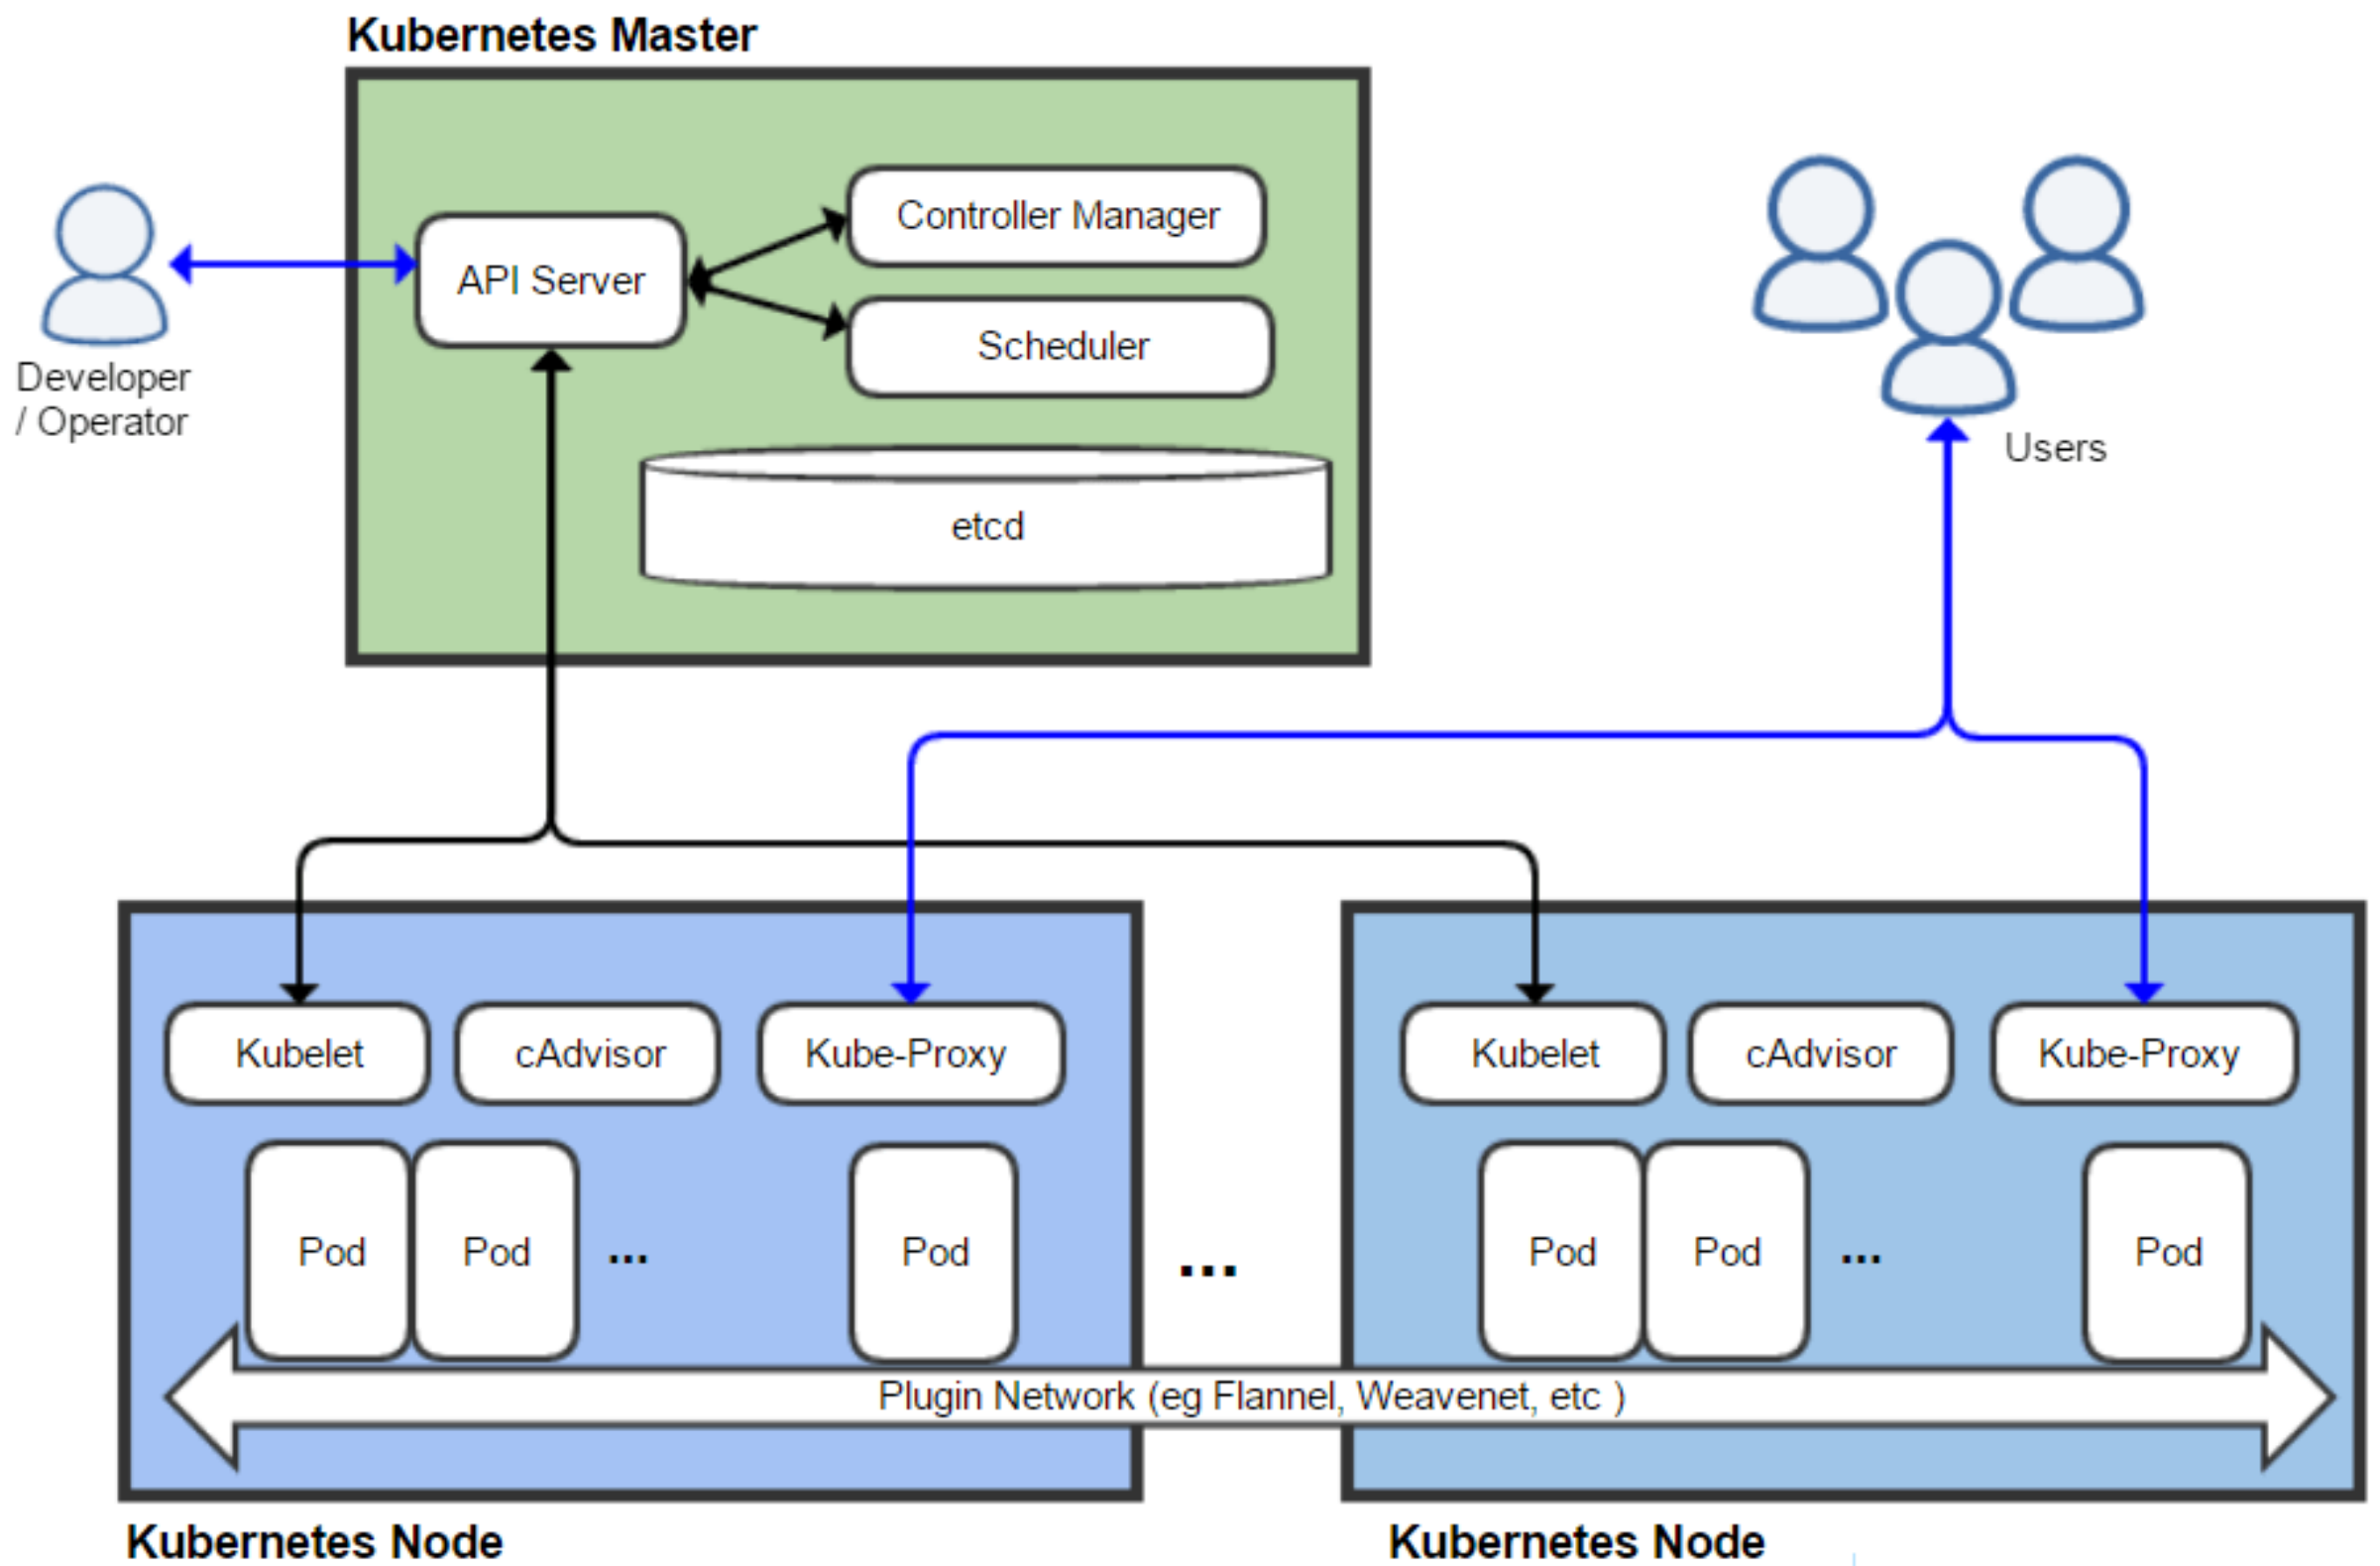
Asset tracking of where images are running.

SUMMARY

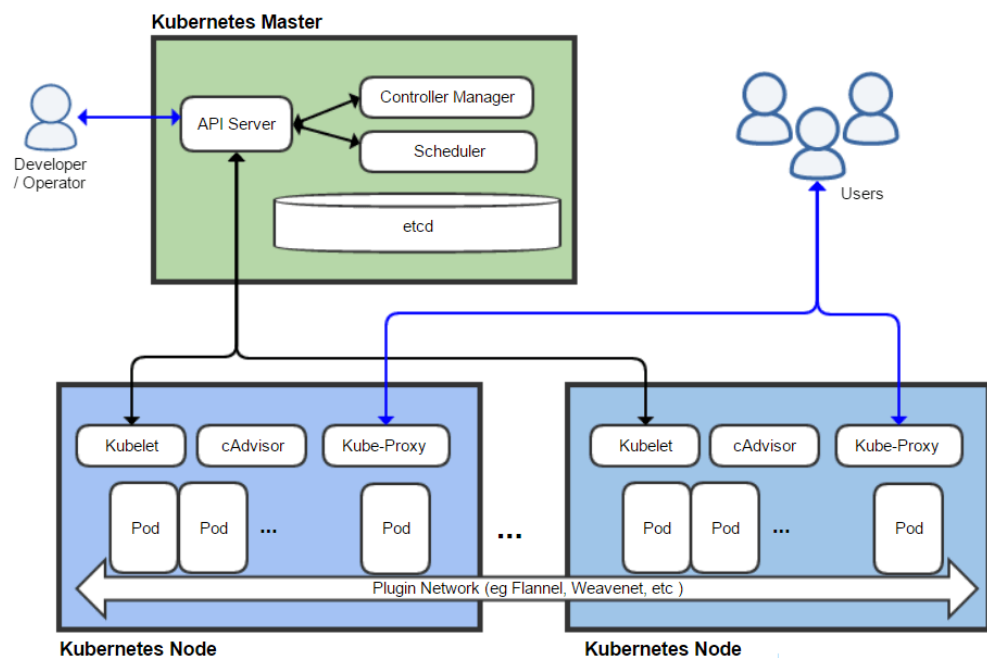
- ▶ *“Container’s don’t contain.”*
 - ▶ They’re not VM’s
 - ▶ It helps if you use a container-focused host OS
- ▶ Don’t give a container what you wouldn’t give a regular host application
- ▶ Restrict access to the Docker daemon
- ▶ Verify image signatures / Have a plan for distribution
- ▶ Scan images for vulns & re-deploy from fresh images
 - ▶ Use your own private image registry



kubernetes



HARDENING BASICS



CC BY-SA: Wikipedia user Khtan66

- ▶ API Server: TLS + Auth + RBAC
- ▶ etcd: Mutual TLS auth
- ▶ kubelet: Disable anonymous auth

Read the docs: *“Securing a Cluster”*

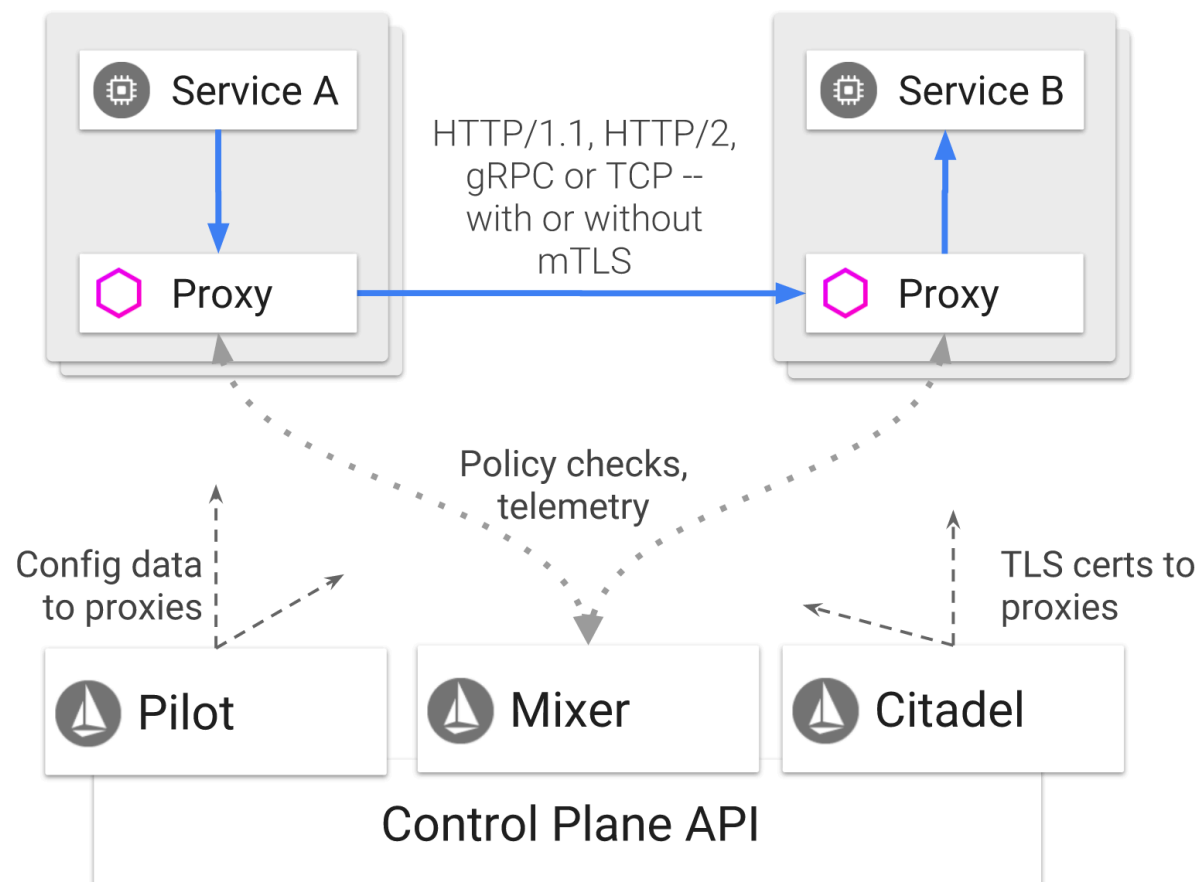
Test your config:





Istio

IMPLICATIONS



Source: <https://istio.io/docs/concepts/what-is-istio/>

- ▶ Separates **data plane** from **control plane**
- ▶ Pro: Mutual TLS everywhere
- ▶ Con: Single point of failure

See the interactive landscape at l.cncf.io

Greyed logos are not open source

App Definition and Development

Database

Streaming & Messaging

Source Code Management

Application Definition & Image Build

Continuous Integration & Delivery

Platform

Observability and Analysis

Orchestration & Management

Scheduling & Orchestration

Coordination & Service Discovery

Remote Procedure Call

Service Proxy

API Gateway

Service Mesh

Cloud-Native Storage

Container Runtime

Cloud-Native Network

Runtime

Automation & Configuration

Container Registries

Security & Compliance

Key Management

Provisioning

Public

Cloud

Kubernetes Certified Service Provider

Kubernetes Training Partner

Special

Monitoring

Logging

Tracing

Chaos Engineering

Serverless

Certified Kubernetes - Distribution

Certified Kubernetes - Hosted

Certified Kubernetes - Installer

Non-Certified Kubernetes

PaaS/Container Service

Cloud Native Landscape
This landscape is intended as a map through the previously uncharted terrain of cloud native technologies. There are many routes to deploying a cloud native application, with CNCF Projects representing a particularly well-traveled path.

l.cncf.io

CLOUD NATIVE Landscape
CLOUD NATIVE COMPUTING FOUNDATION
Partners: Redpoint, Amplify



Questions?