

# A DevOps State of Mind: Continuous Security with DevSecOps + Containers

Chris Van Tuin  
Chief Technologist, NA West / Silicon Valley  
[cvantuin@redhat.com](mailto:cvantuin@redhat.com)



# SECURITY BREACH: BILLION DATA RECORDS

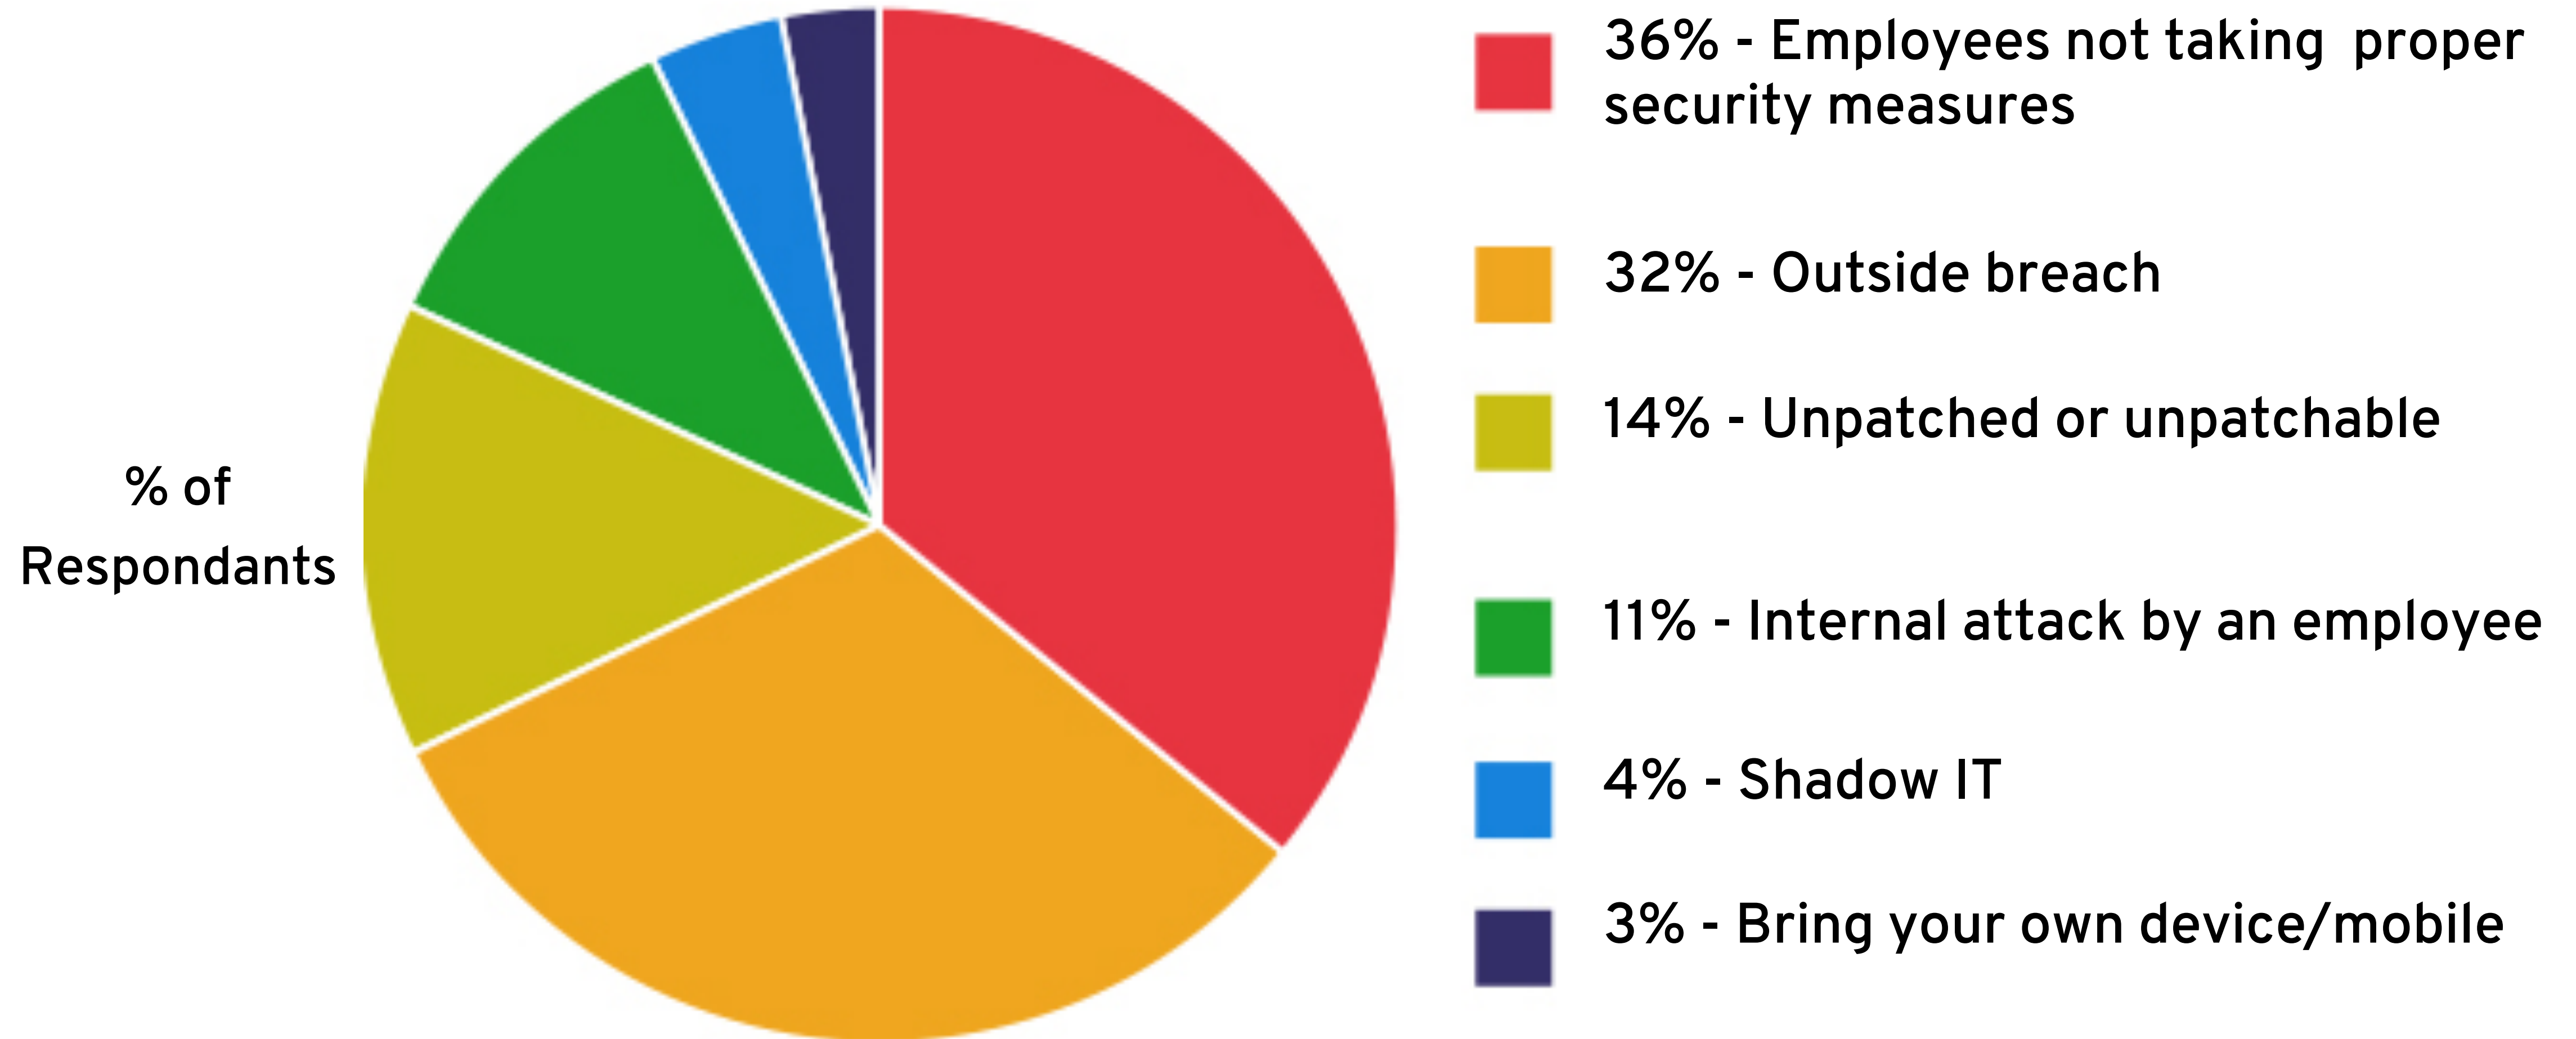


J.P.Morgan

Neiman Marcus



# WHAT IS THE GREATEST SECURITY RISK?



Source: Techvalidate/Red Hat

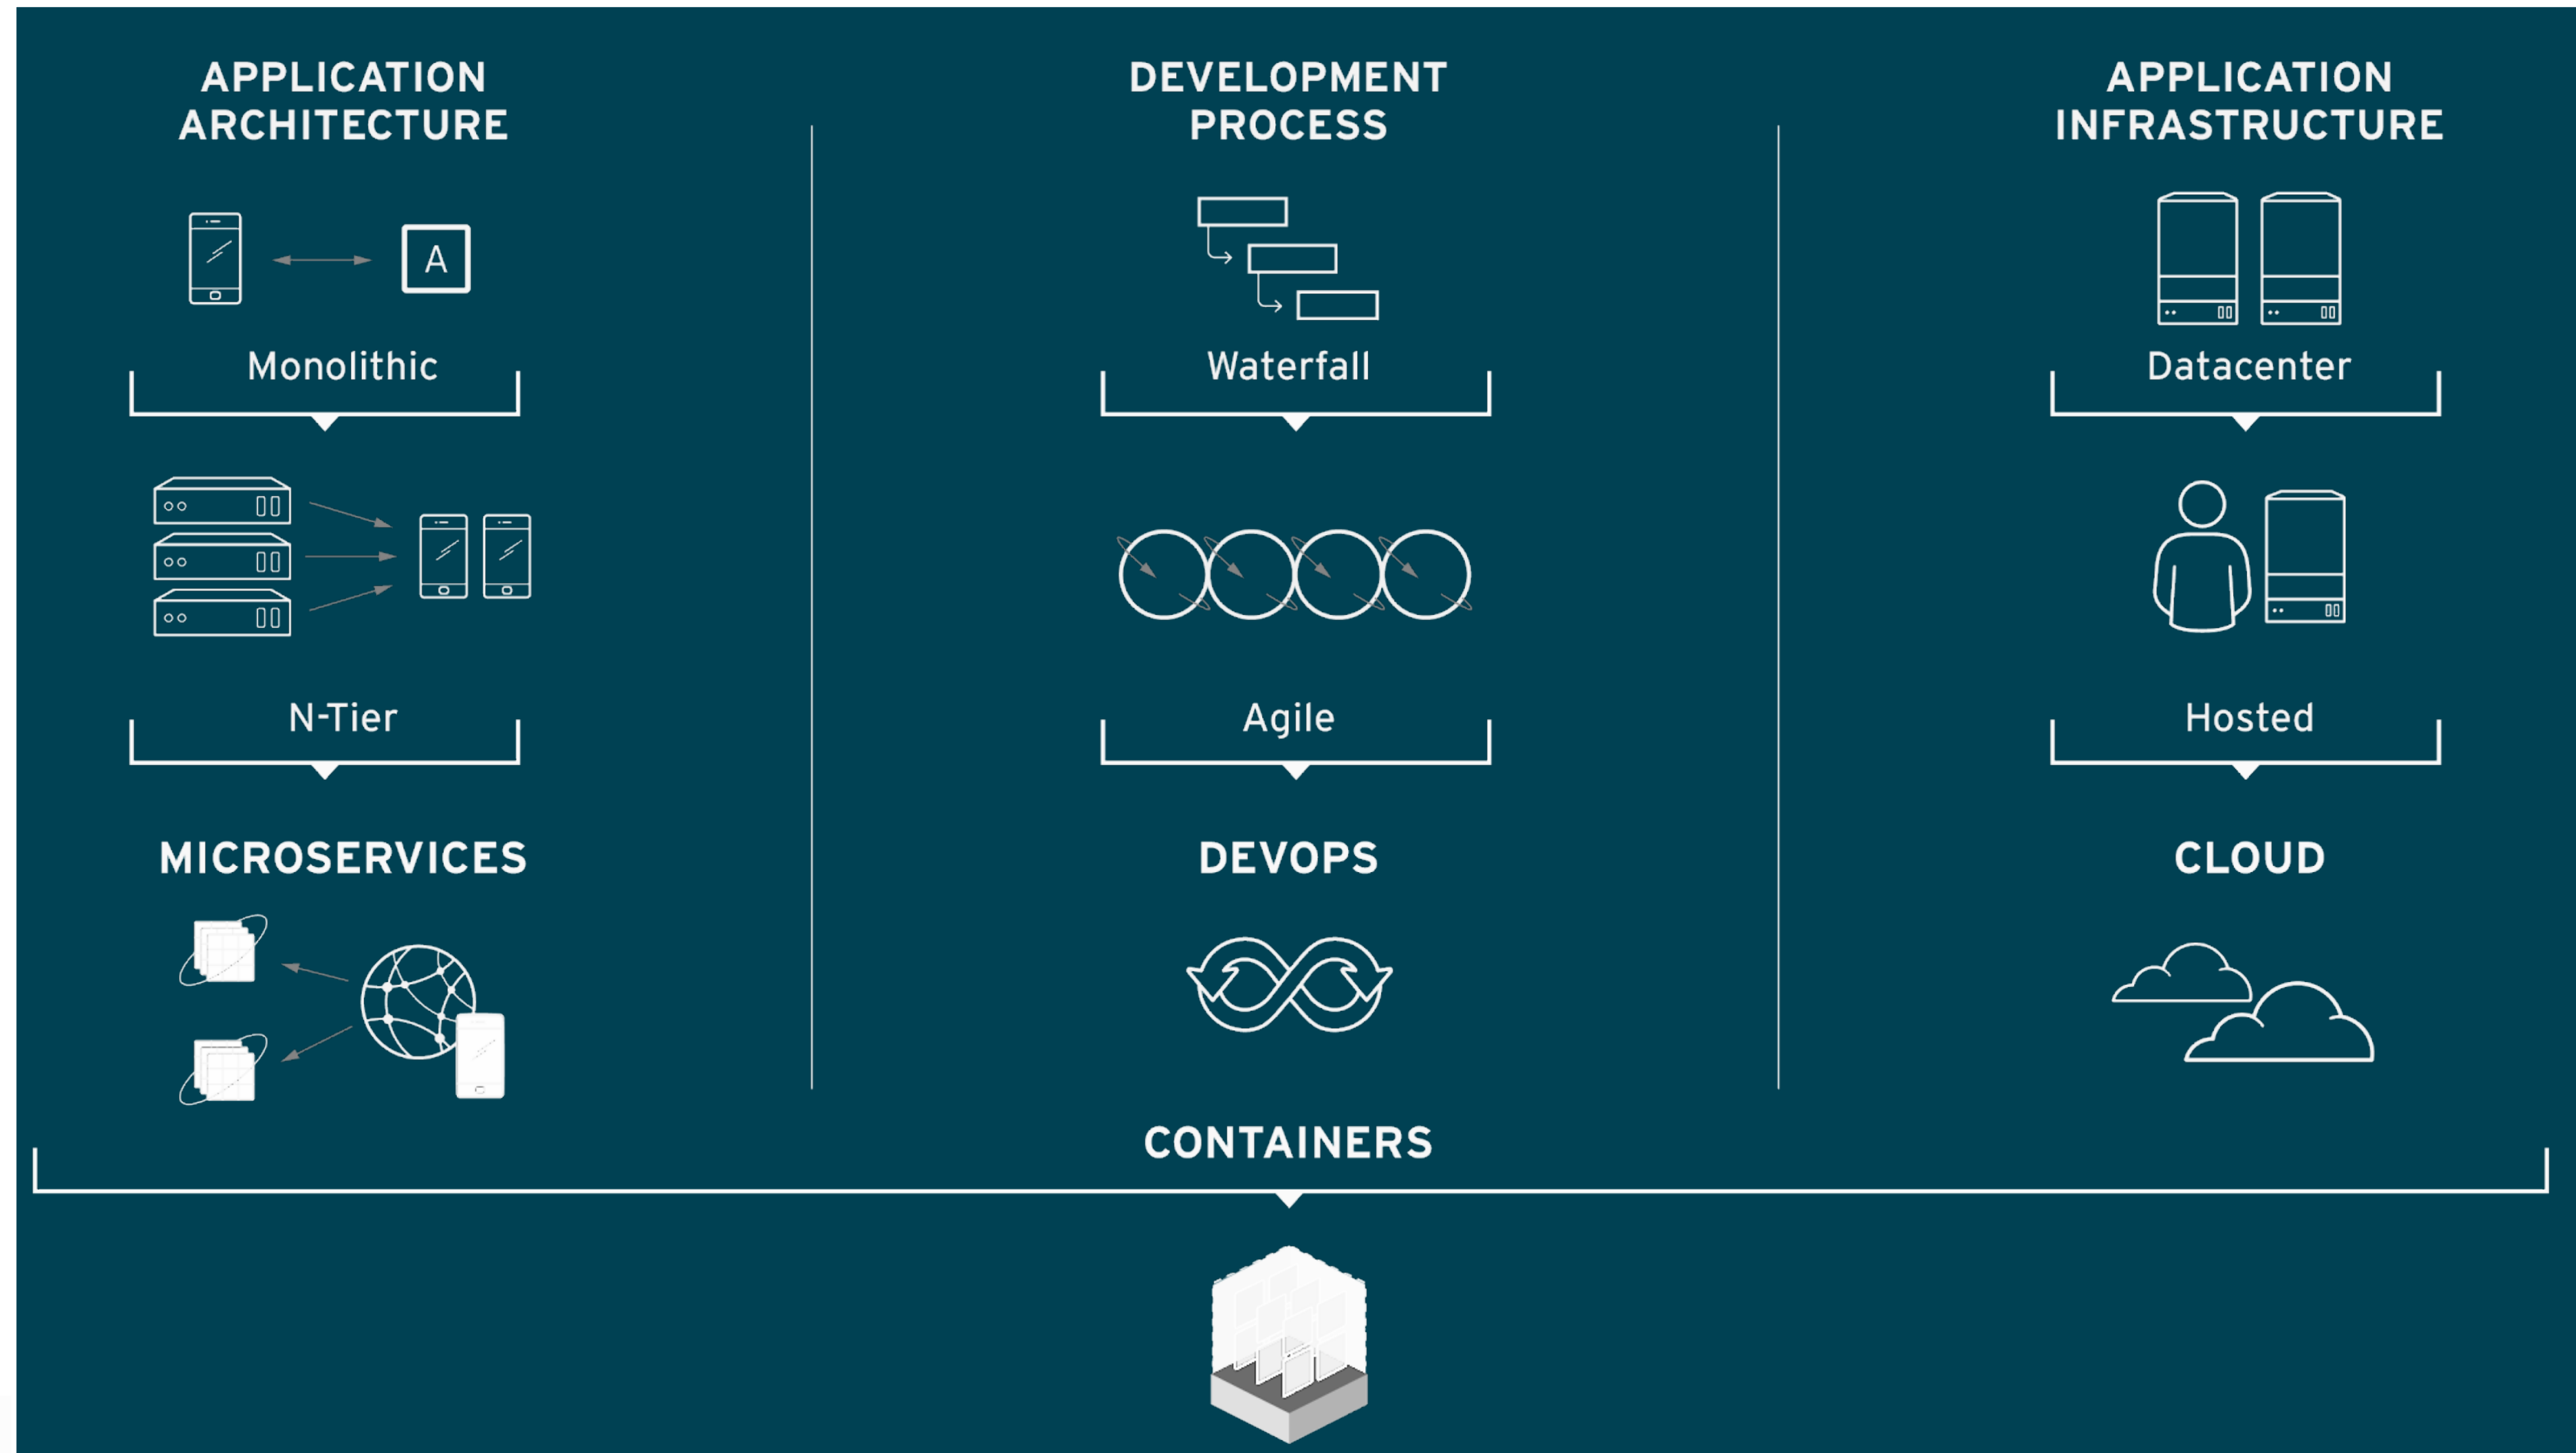
# “Only the paranoid survive”

- Andy Grove, 1996






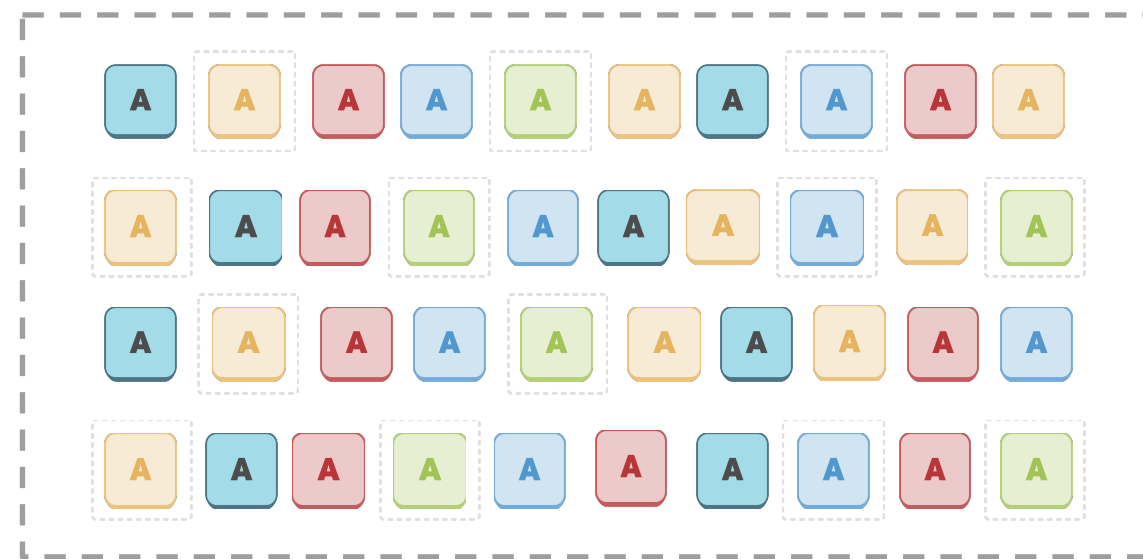
# SECURITY MUST EVOLVE & KEEP UP




# HYBRID CLOUD ENVIRONMENTS

ANY COMBINATION, WHETHER TRADITIONAL OR CONTAINERIZED


**LEGACY APPS**  
(1,000+)




**PRODUCTION**



**DEV/TEST**




BARE METAL



VIRTUAL




PRIVATE CLOUD



PUBLIC CLOUD

# DISTRIBUTED APPLICATIONS



**ON-PREMISE**



BARE METAL



VIRTUAL



PRIVATE CLOUD

**OFF-PREMISE**




PUBLIC CLOUD

# SECURITY IS AN AFTERTHOUGHT

“Patch?  
The servers are behind the firewall.”

- Anonymous (far too many to name), 2005 - ...




| SECURITY |

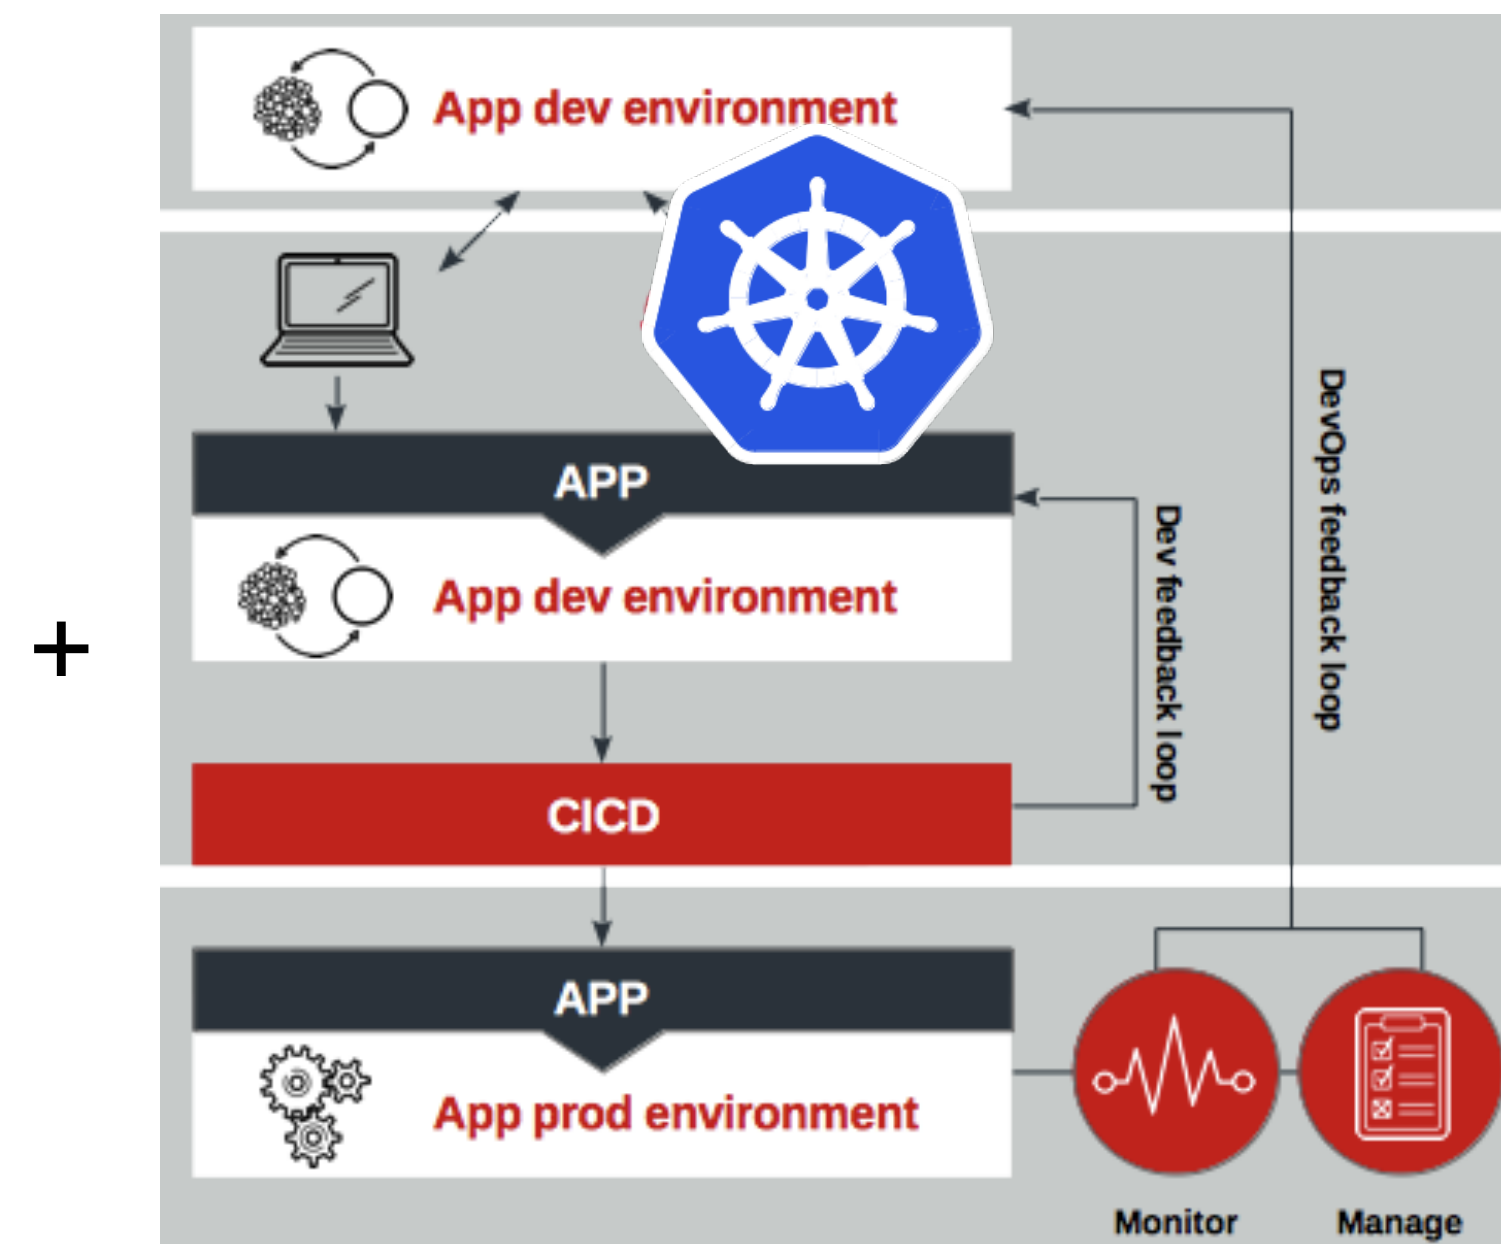


# DEVSECOPS

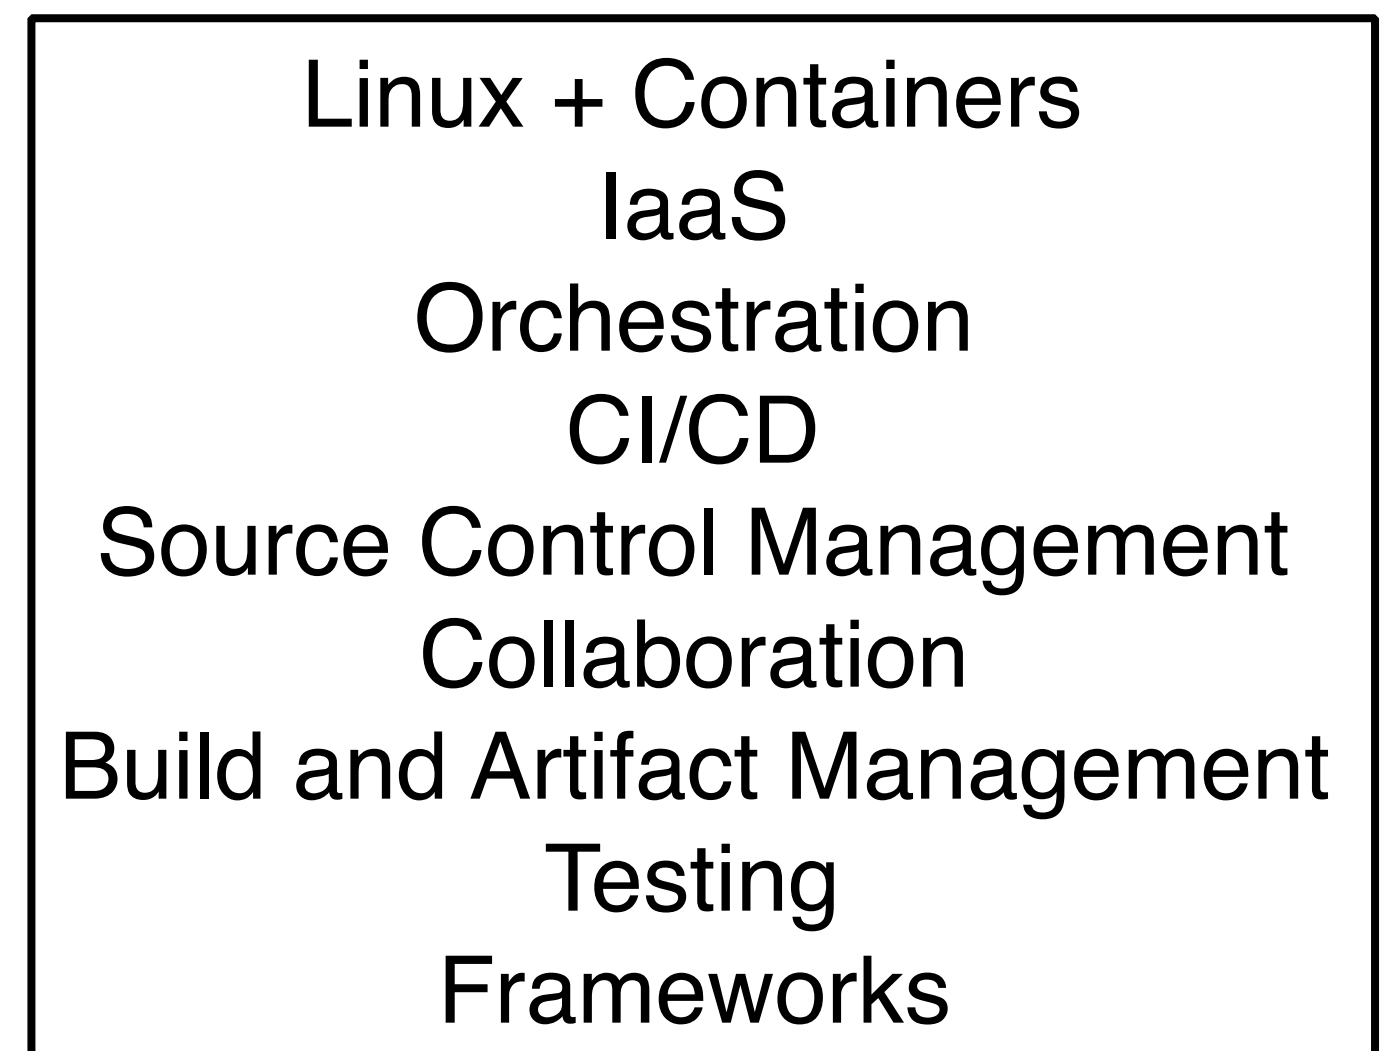
## Culture



## Process



## Technology



Open Source


← **End to End Security** →

# DEVSECOPS

**Reduce Risks, Lower Costs, Speed Delivery, Speed Reaction**



**Security  
Automation**



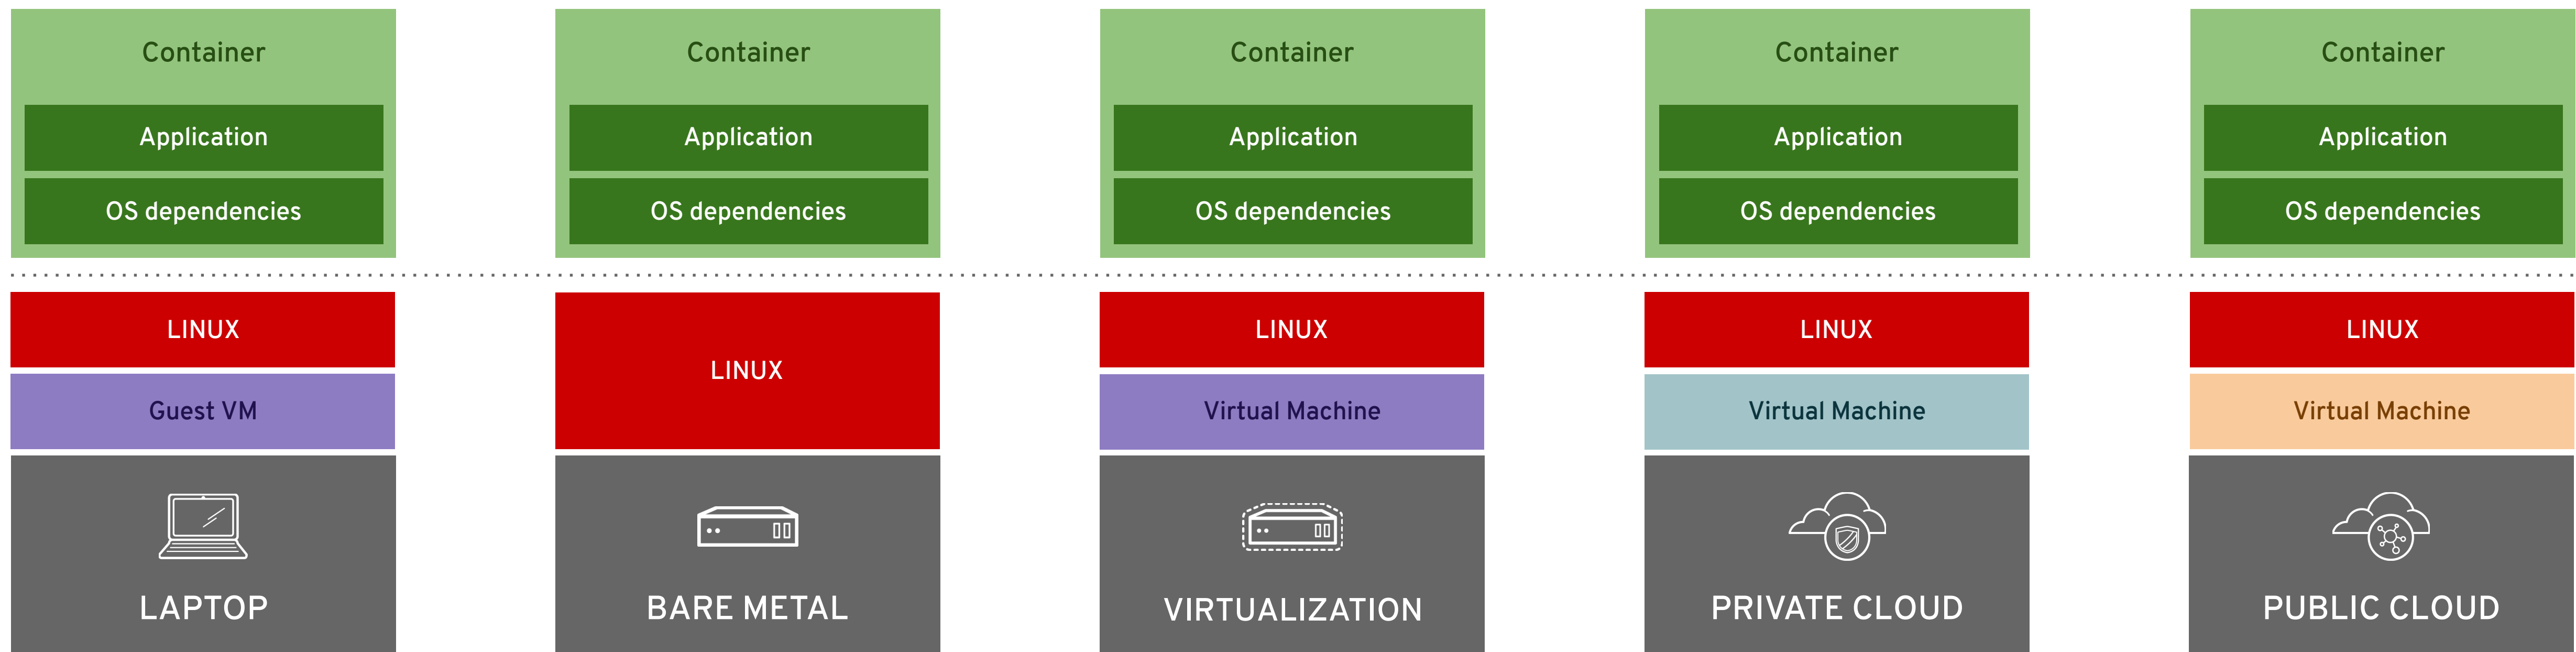
**Process  
Optimization**



**Continuous  
Security  
Improvement**

# CONTAINERS

# APPLICATION PORTABILITY WITH CONTAINERS




# CONTAINERS AT SCALE



# MORE THAN CONTAINERS...



**Scheduling**




**Lifecycle & health**



**Discovery**



**Monitoring**



**Scaling**



**Persistence**



**Aggregation**




**Security**



# kubernetes



BARE METAL



VIRTUAL



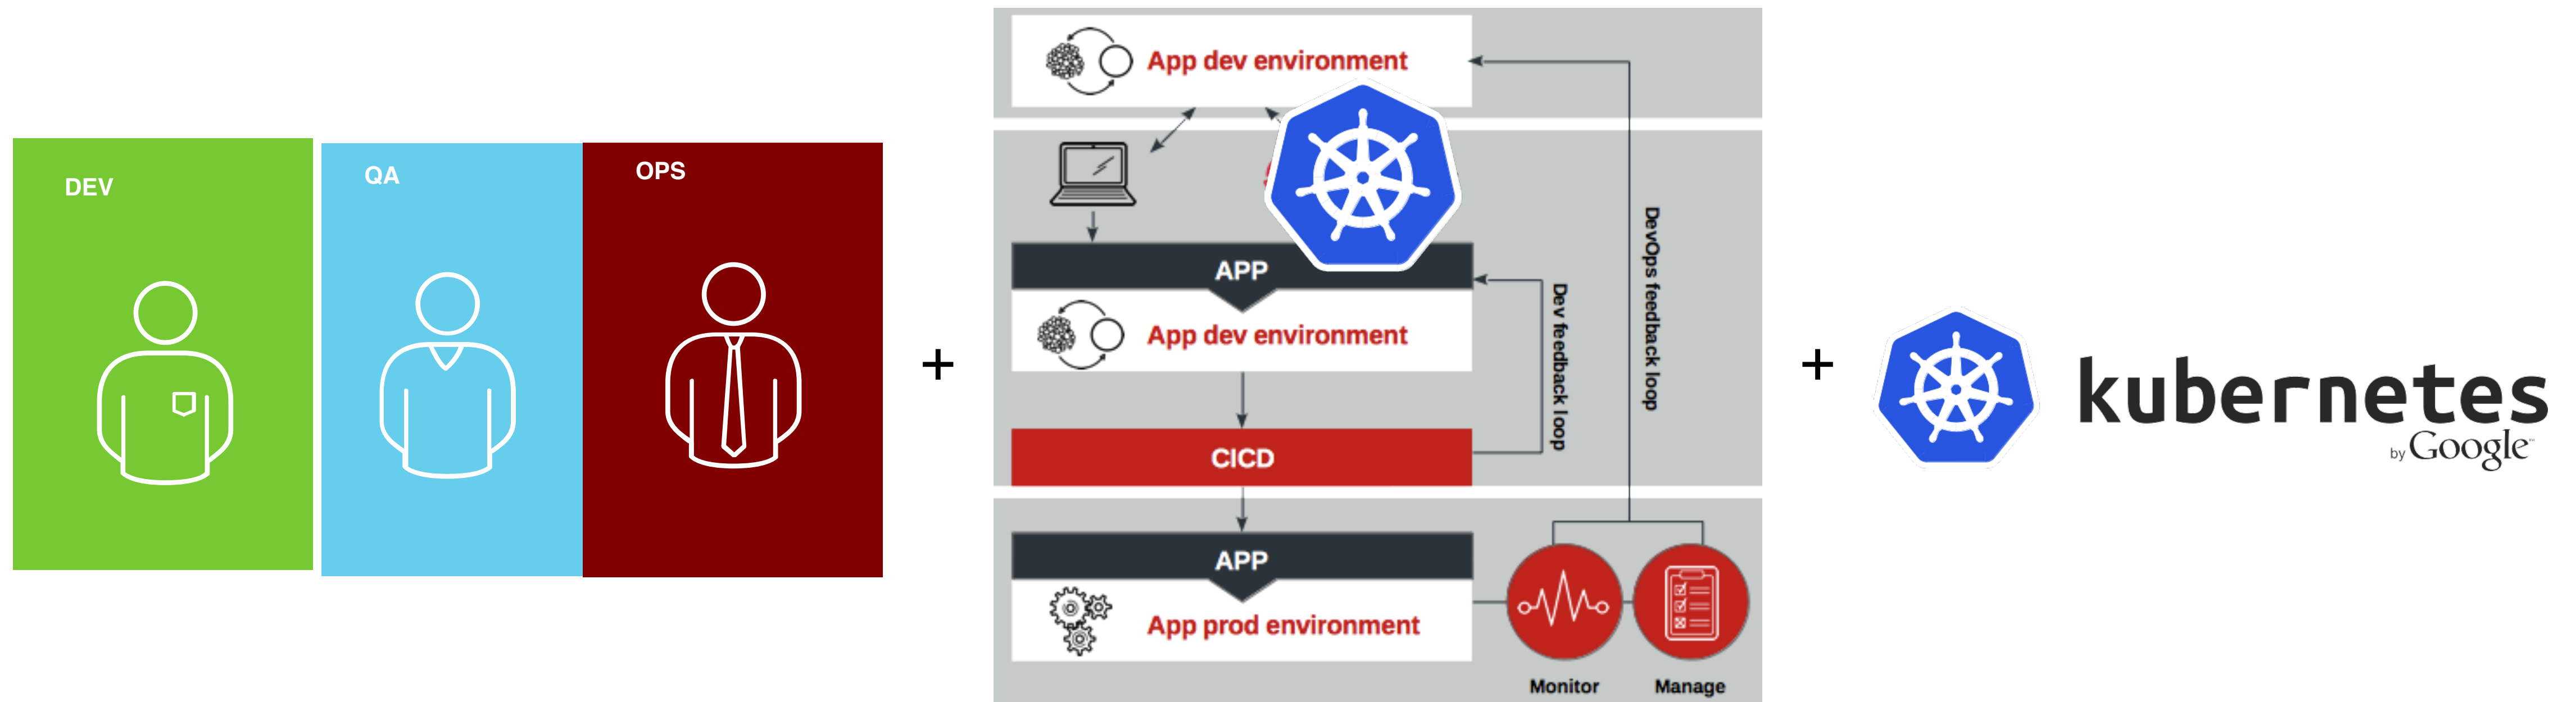
PRIVATE CLOUD



PUBLIC CLOUD

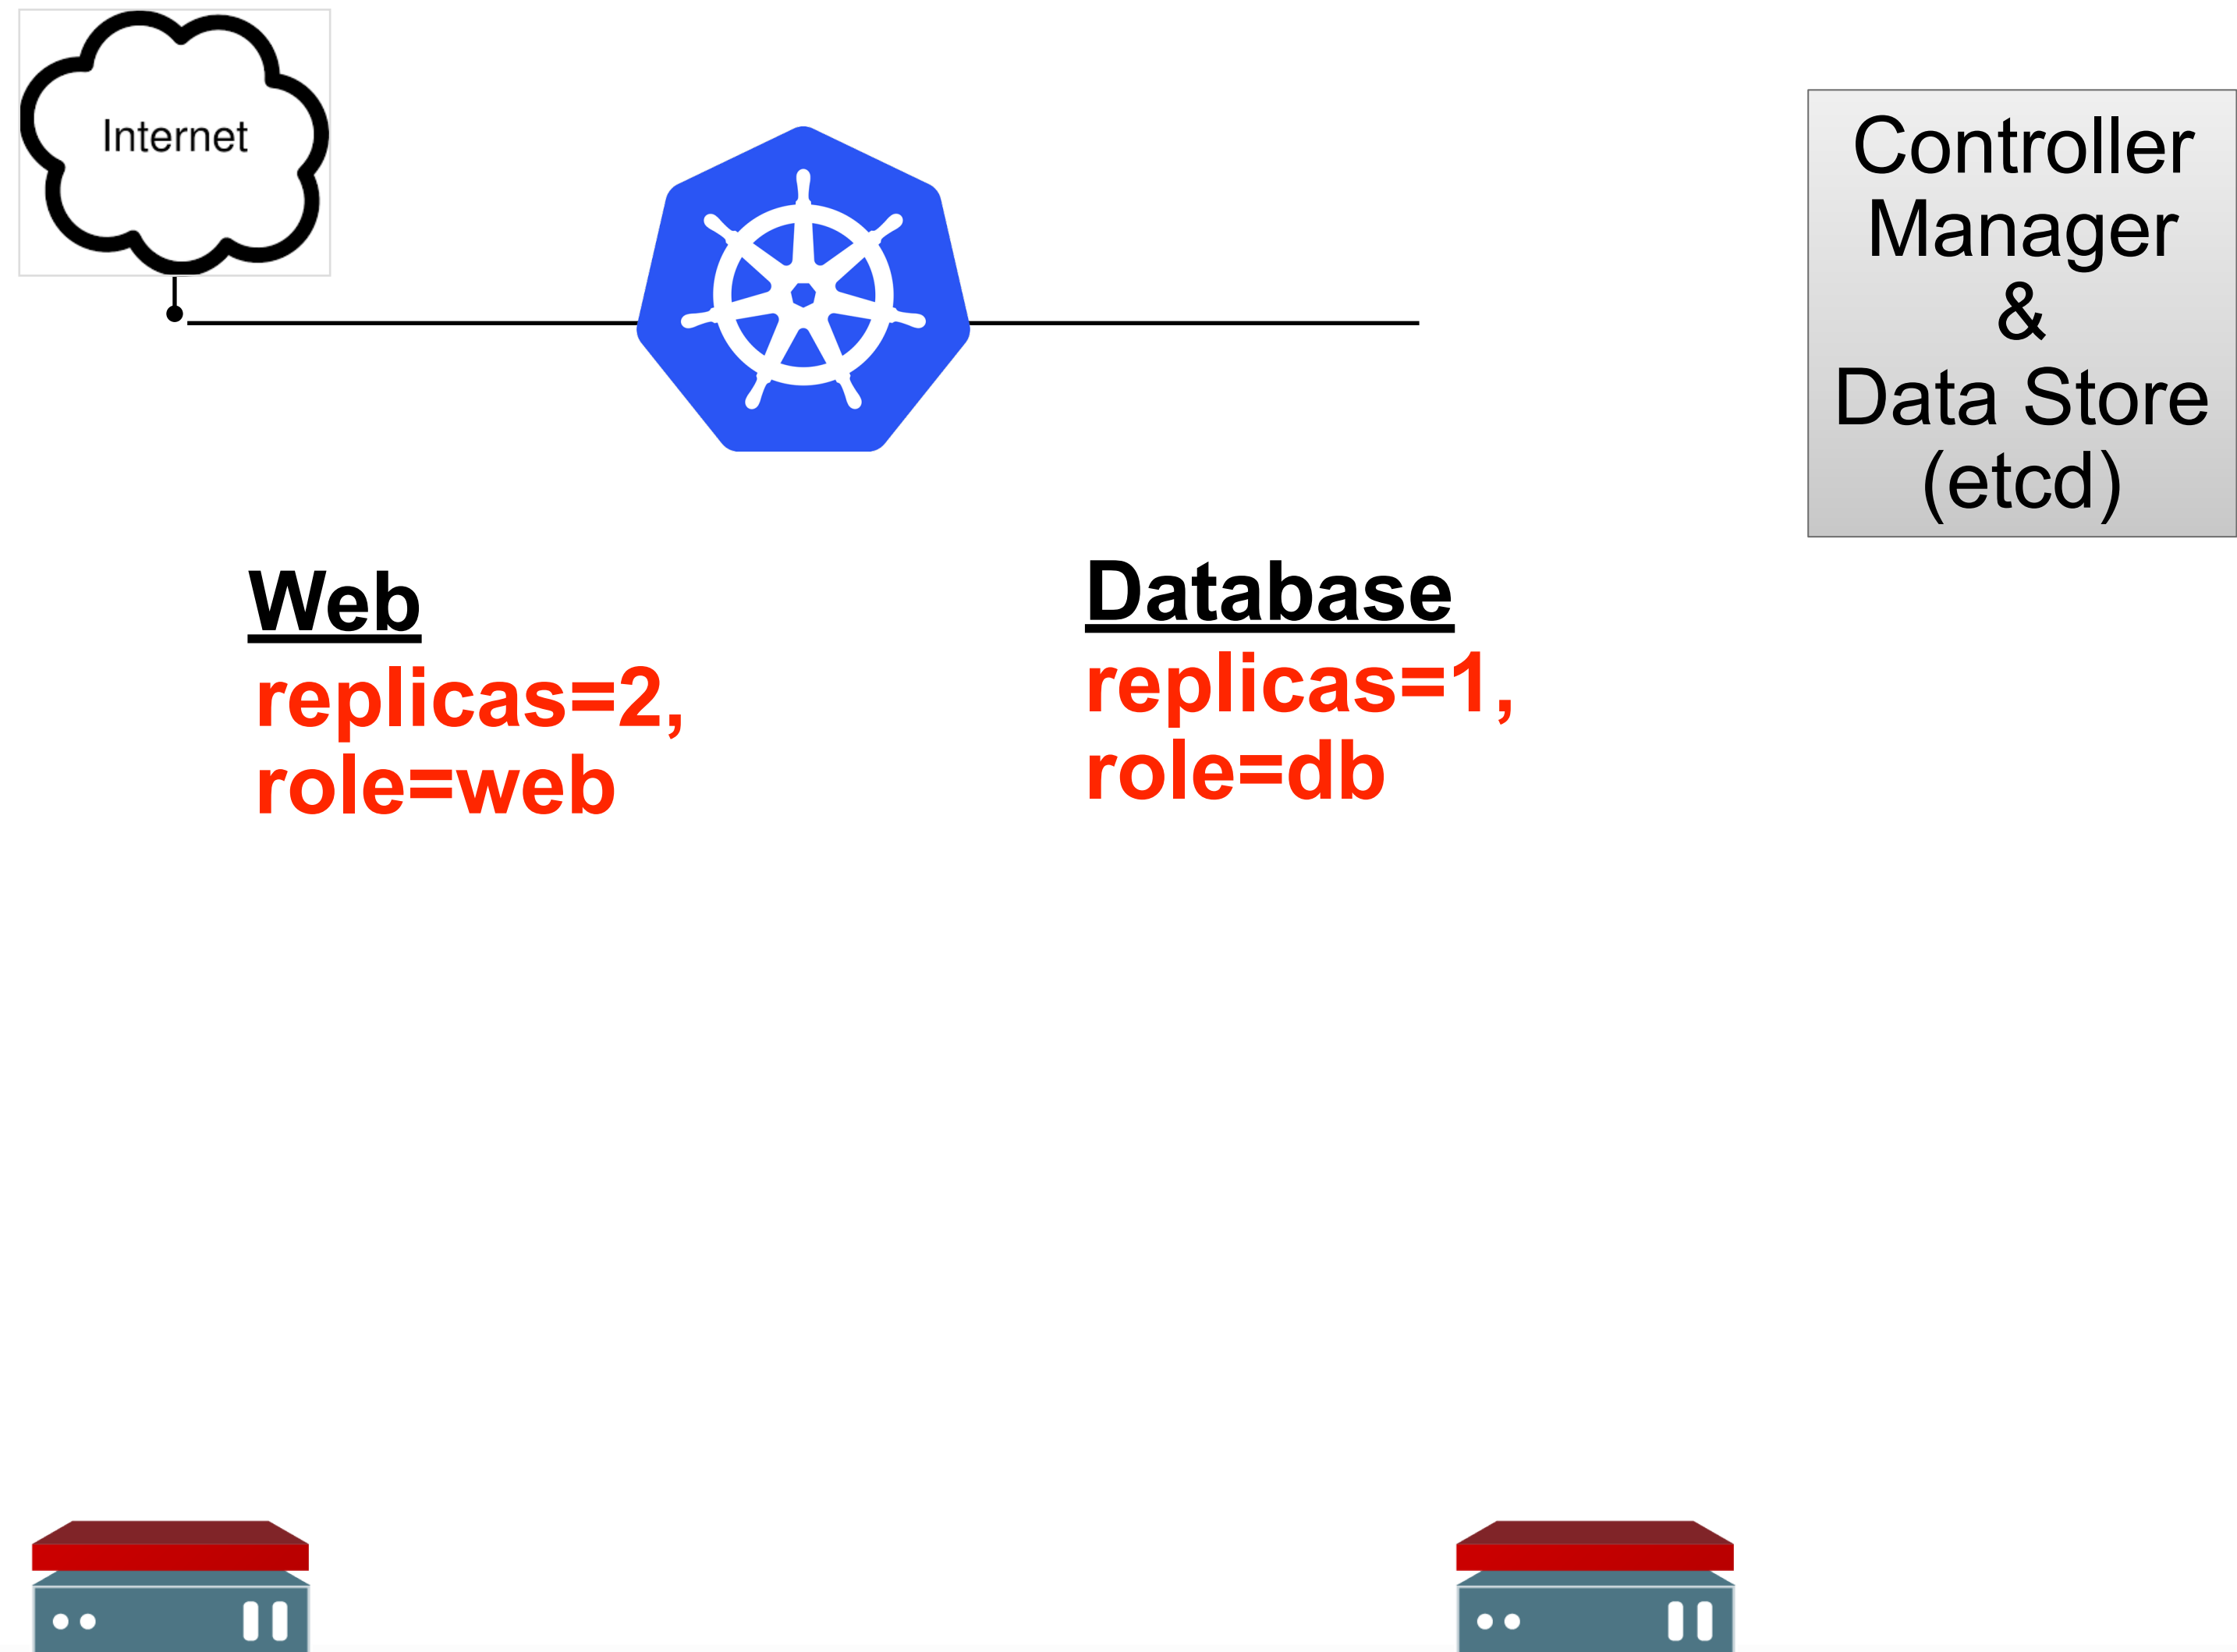
# DEVSECOPS

## End to End Security




# ORCHESTRATION

Deployment, Declarative



# ORCHESTRATION

Schedule + Provision Pods (Compute/Storage/Network)



Web  
replicas=2,  
role=web  
ReplicaSet

**Pods**



role=web



role=web

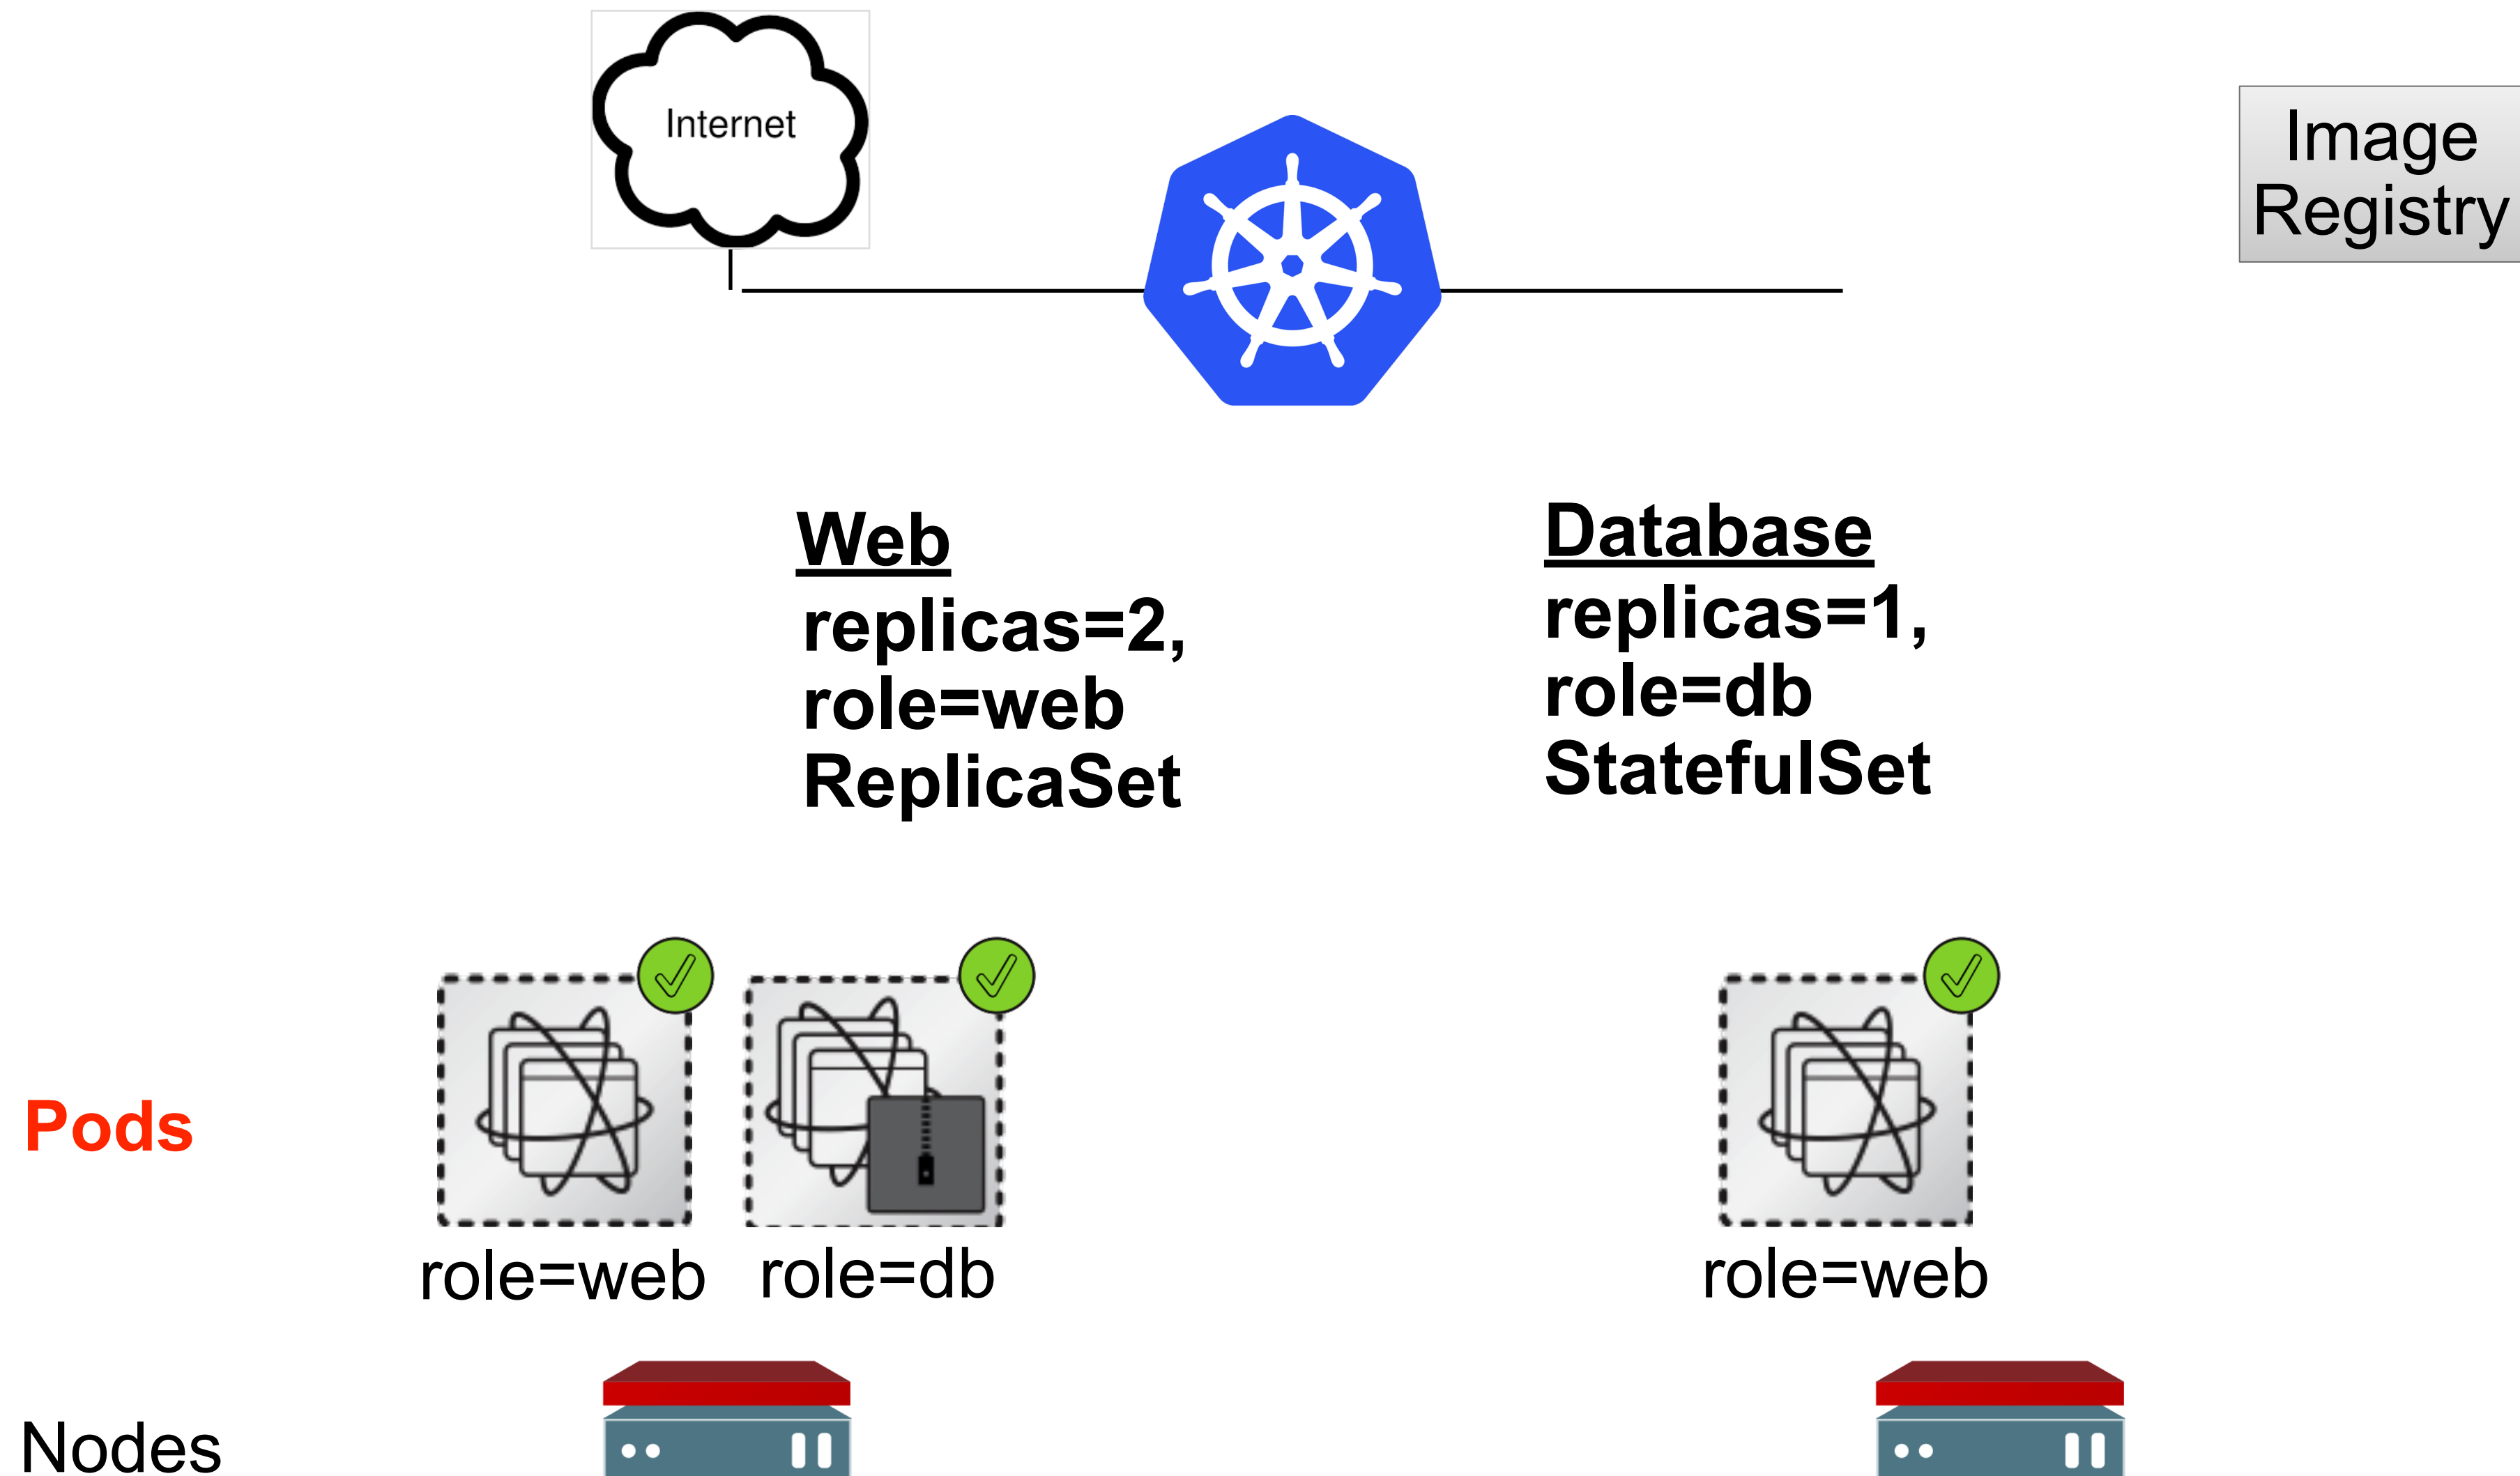
**Nodes**





# ORCHESTRATION

Schedule + Provision Pods (Compute/Storage/Network)




# ORCHESTRATION

Service (Load Balancer)

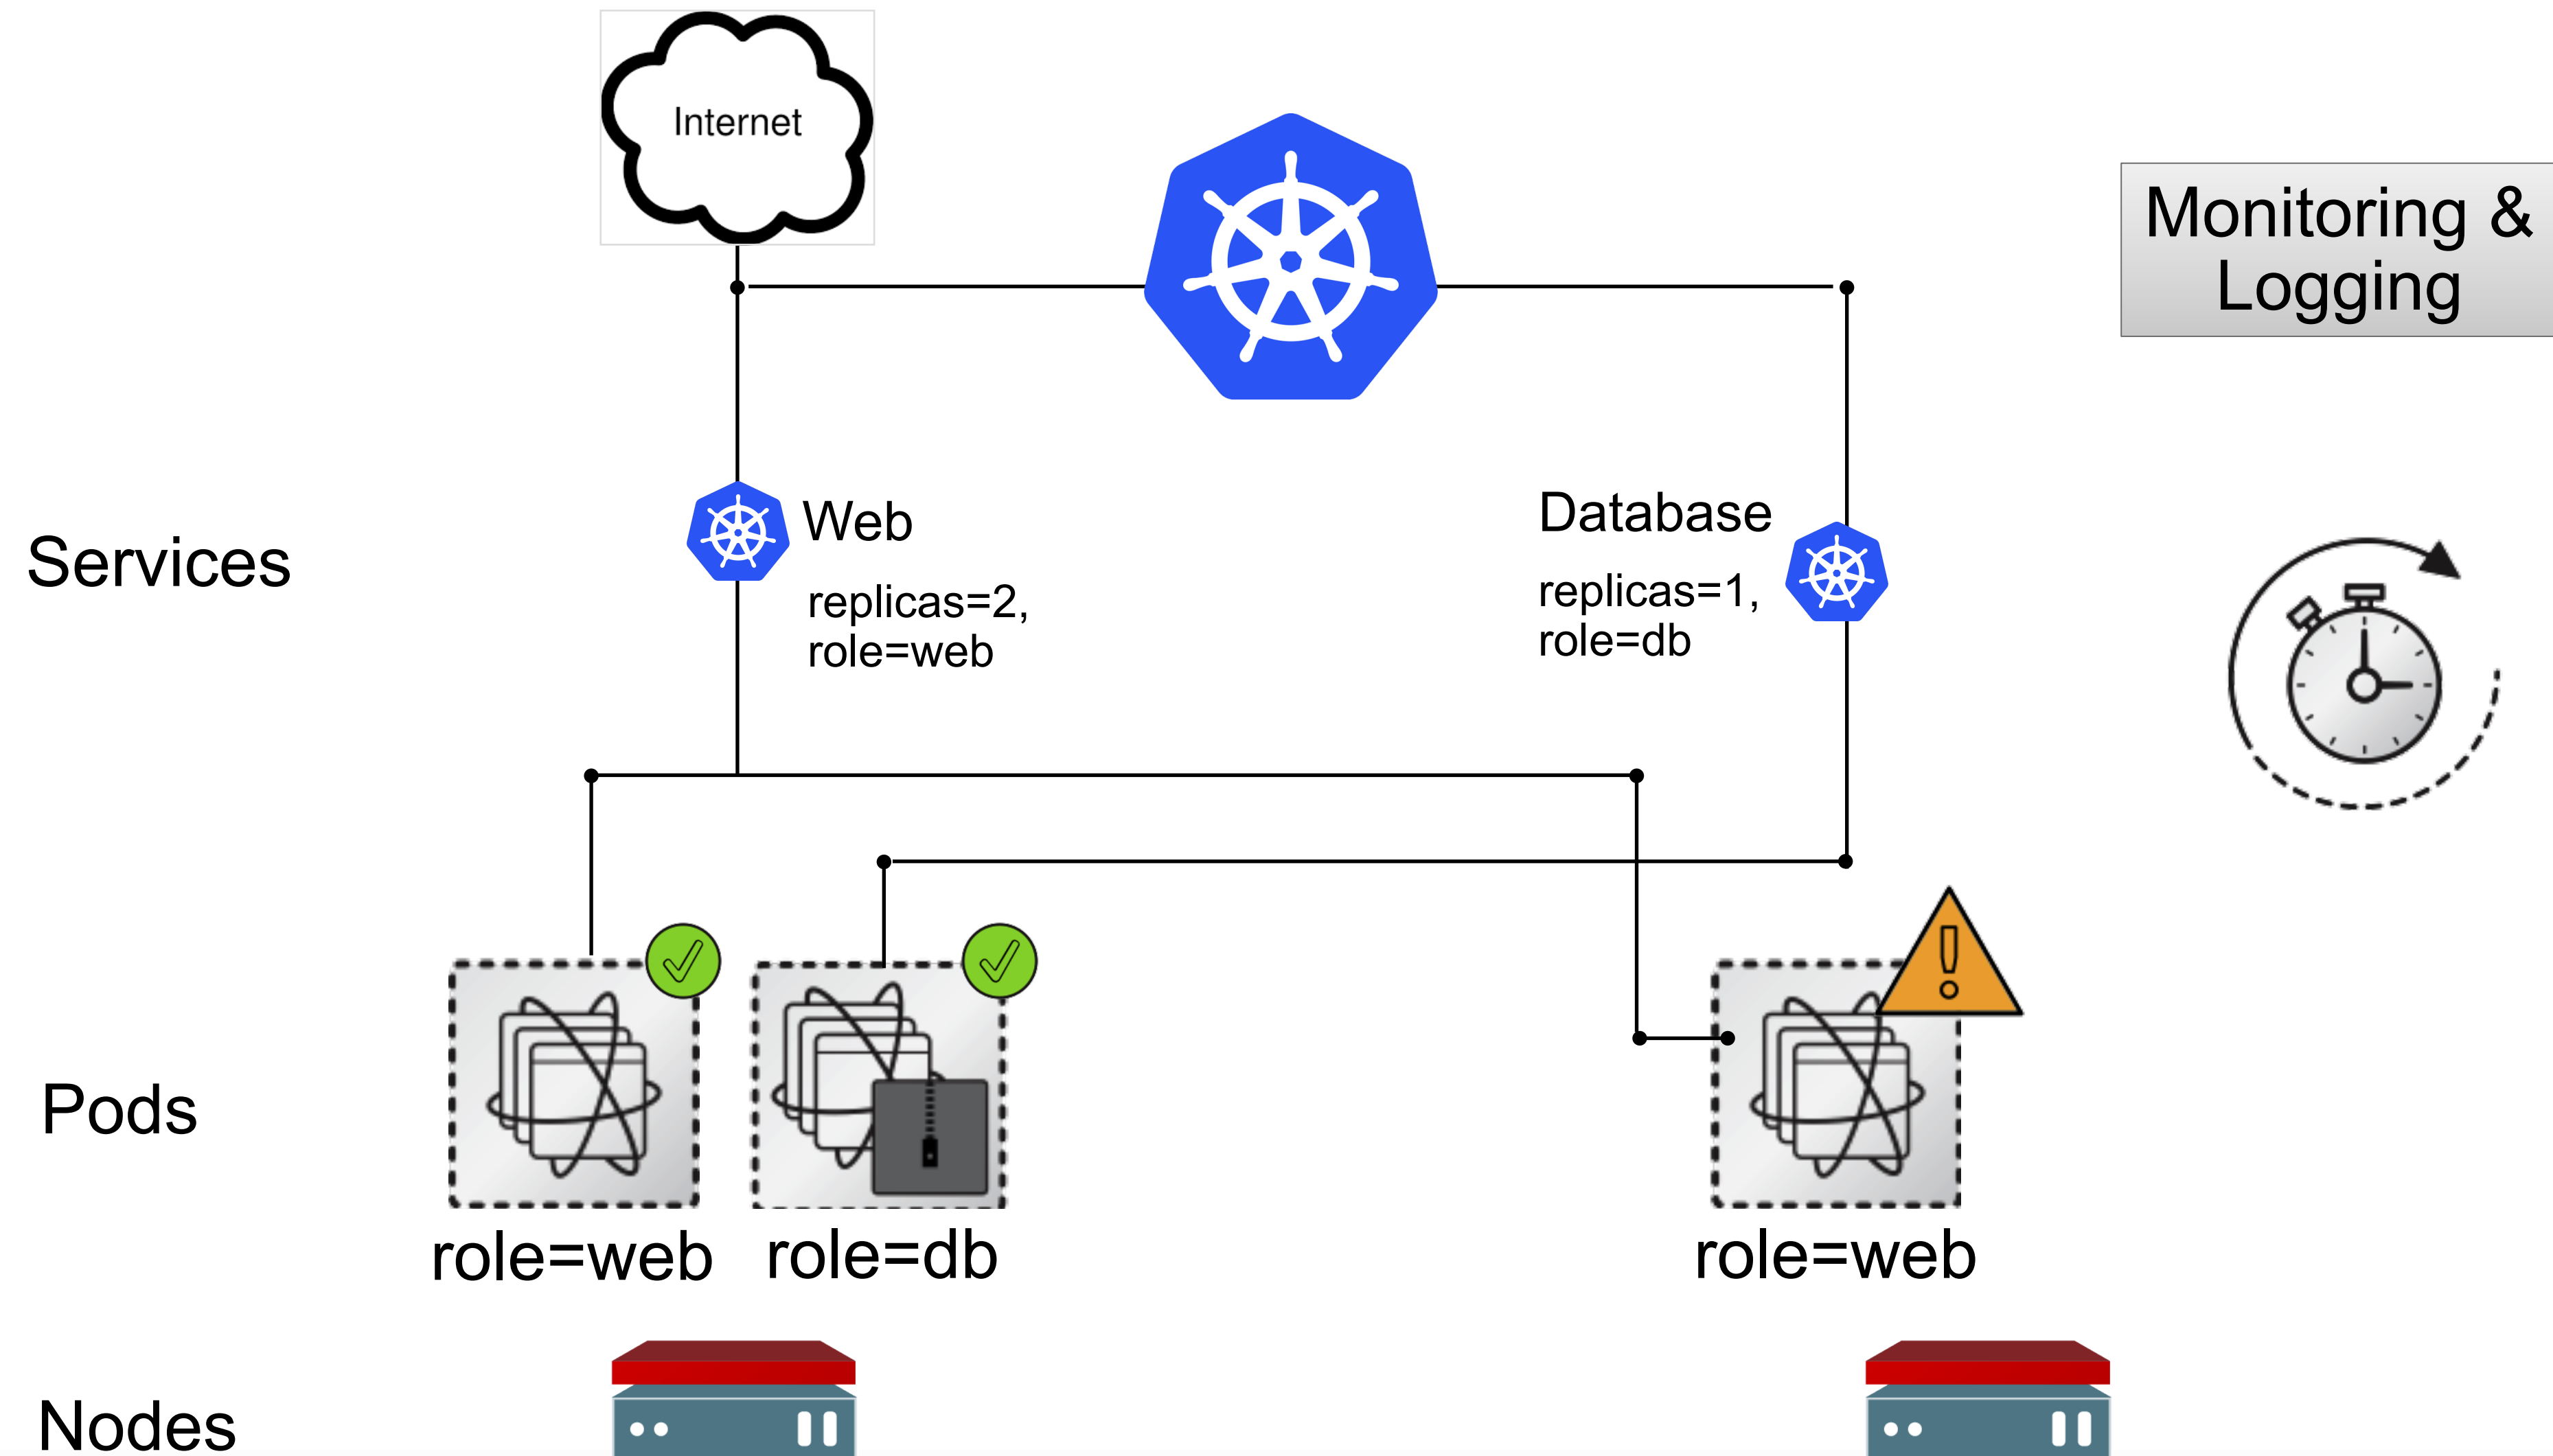
Services

Pods

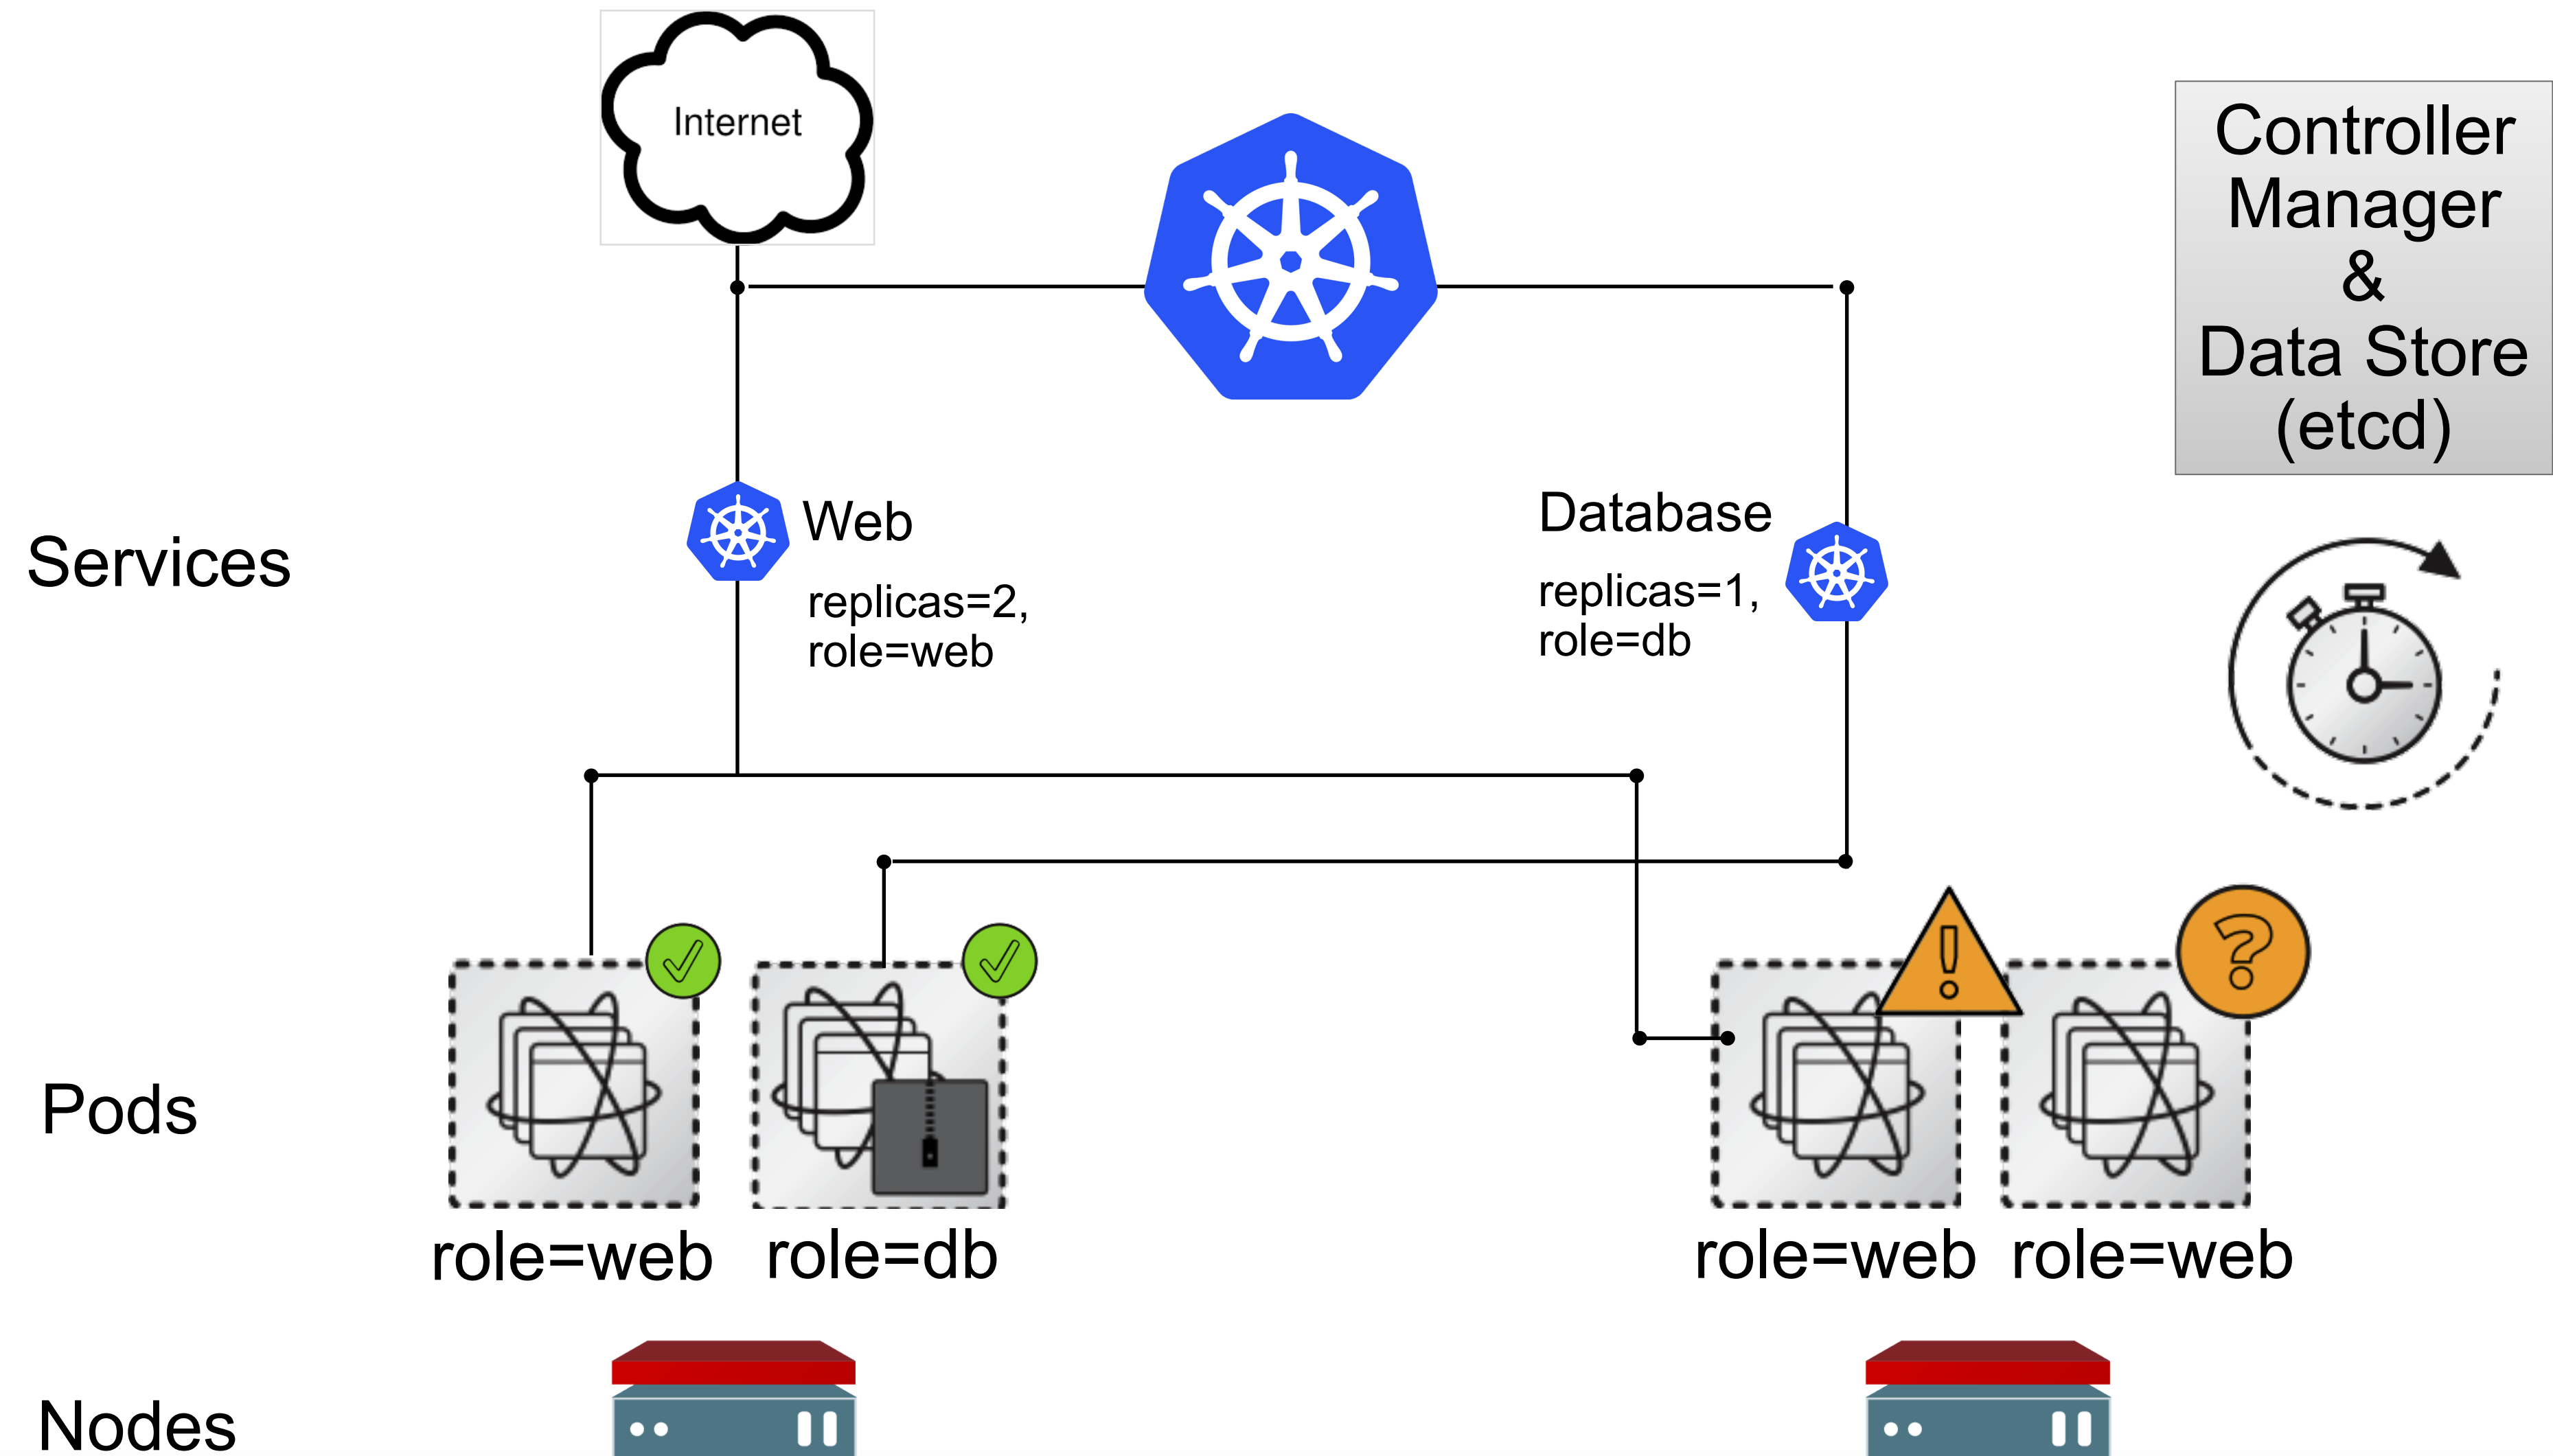
Nodes



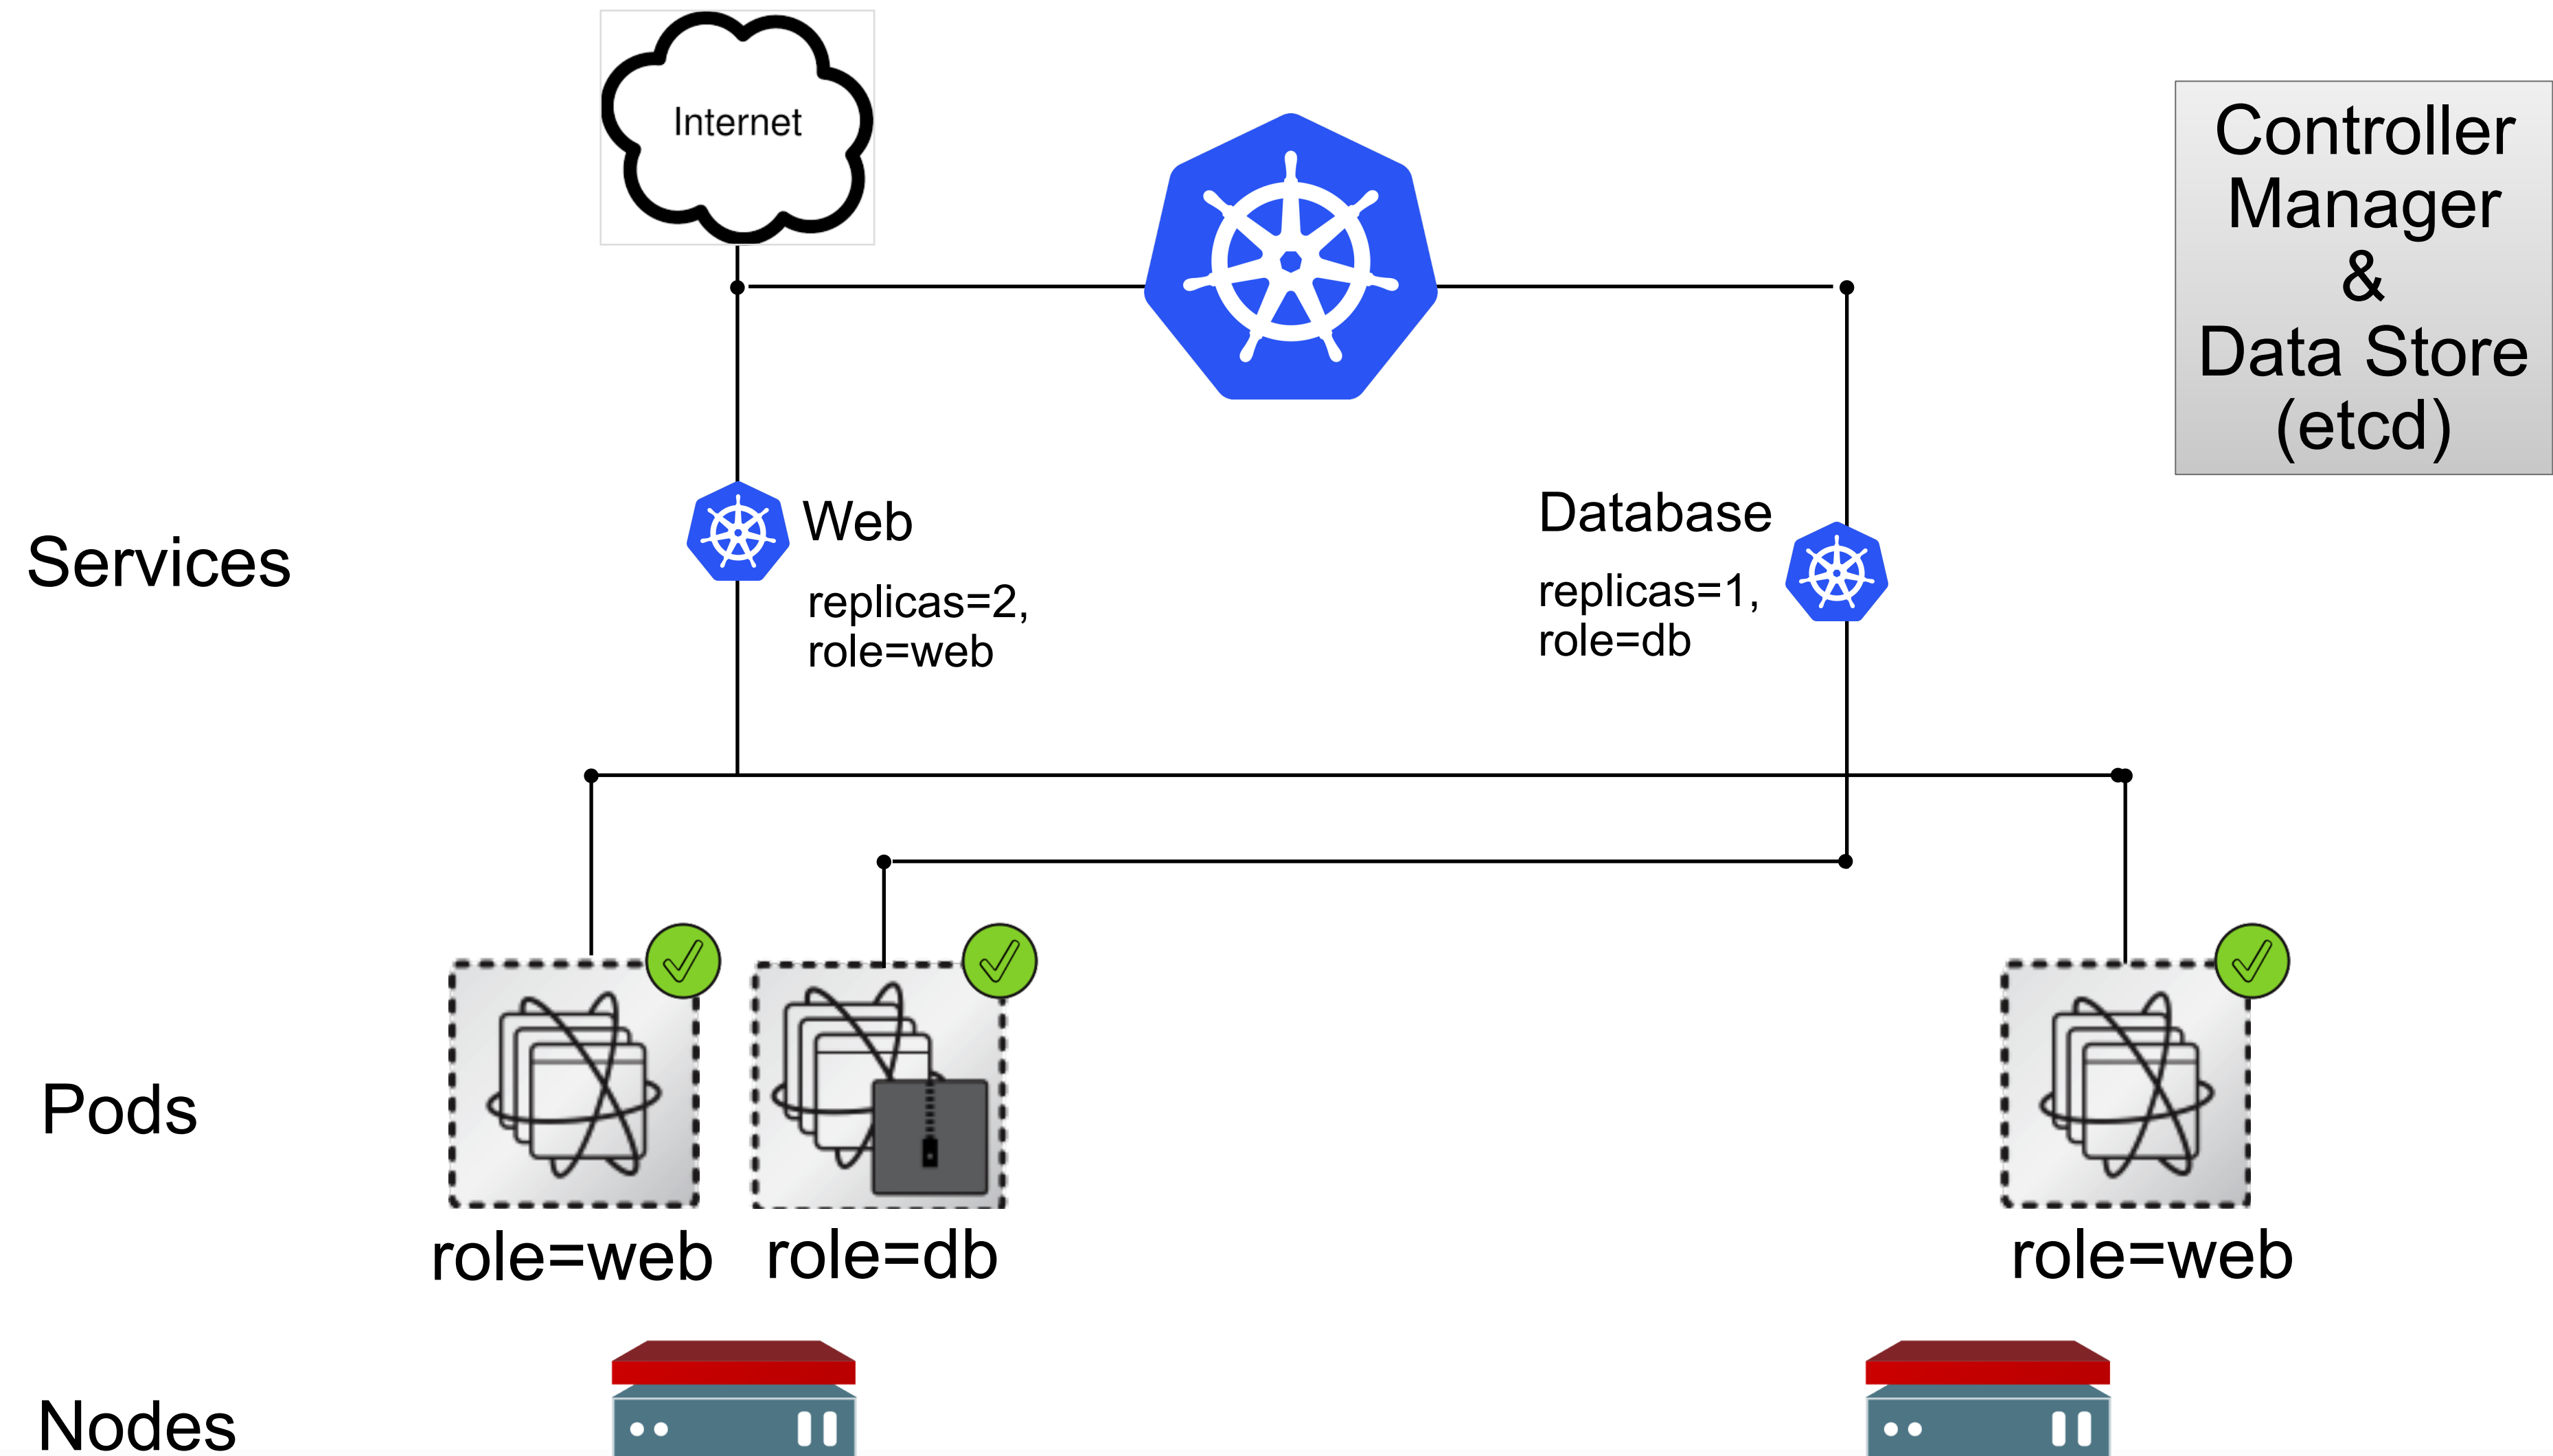
# HEALTH CHECK



# HEALTH CHECK

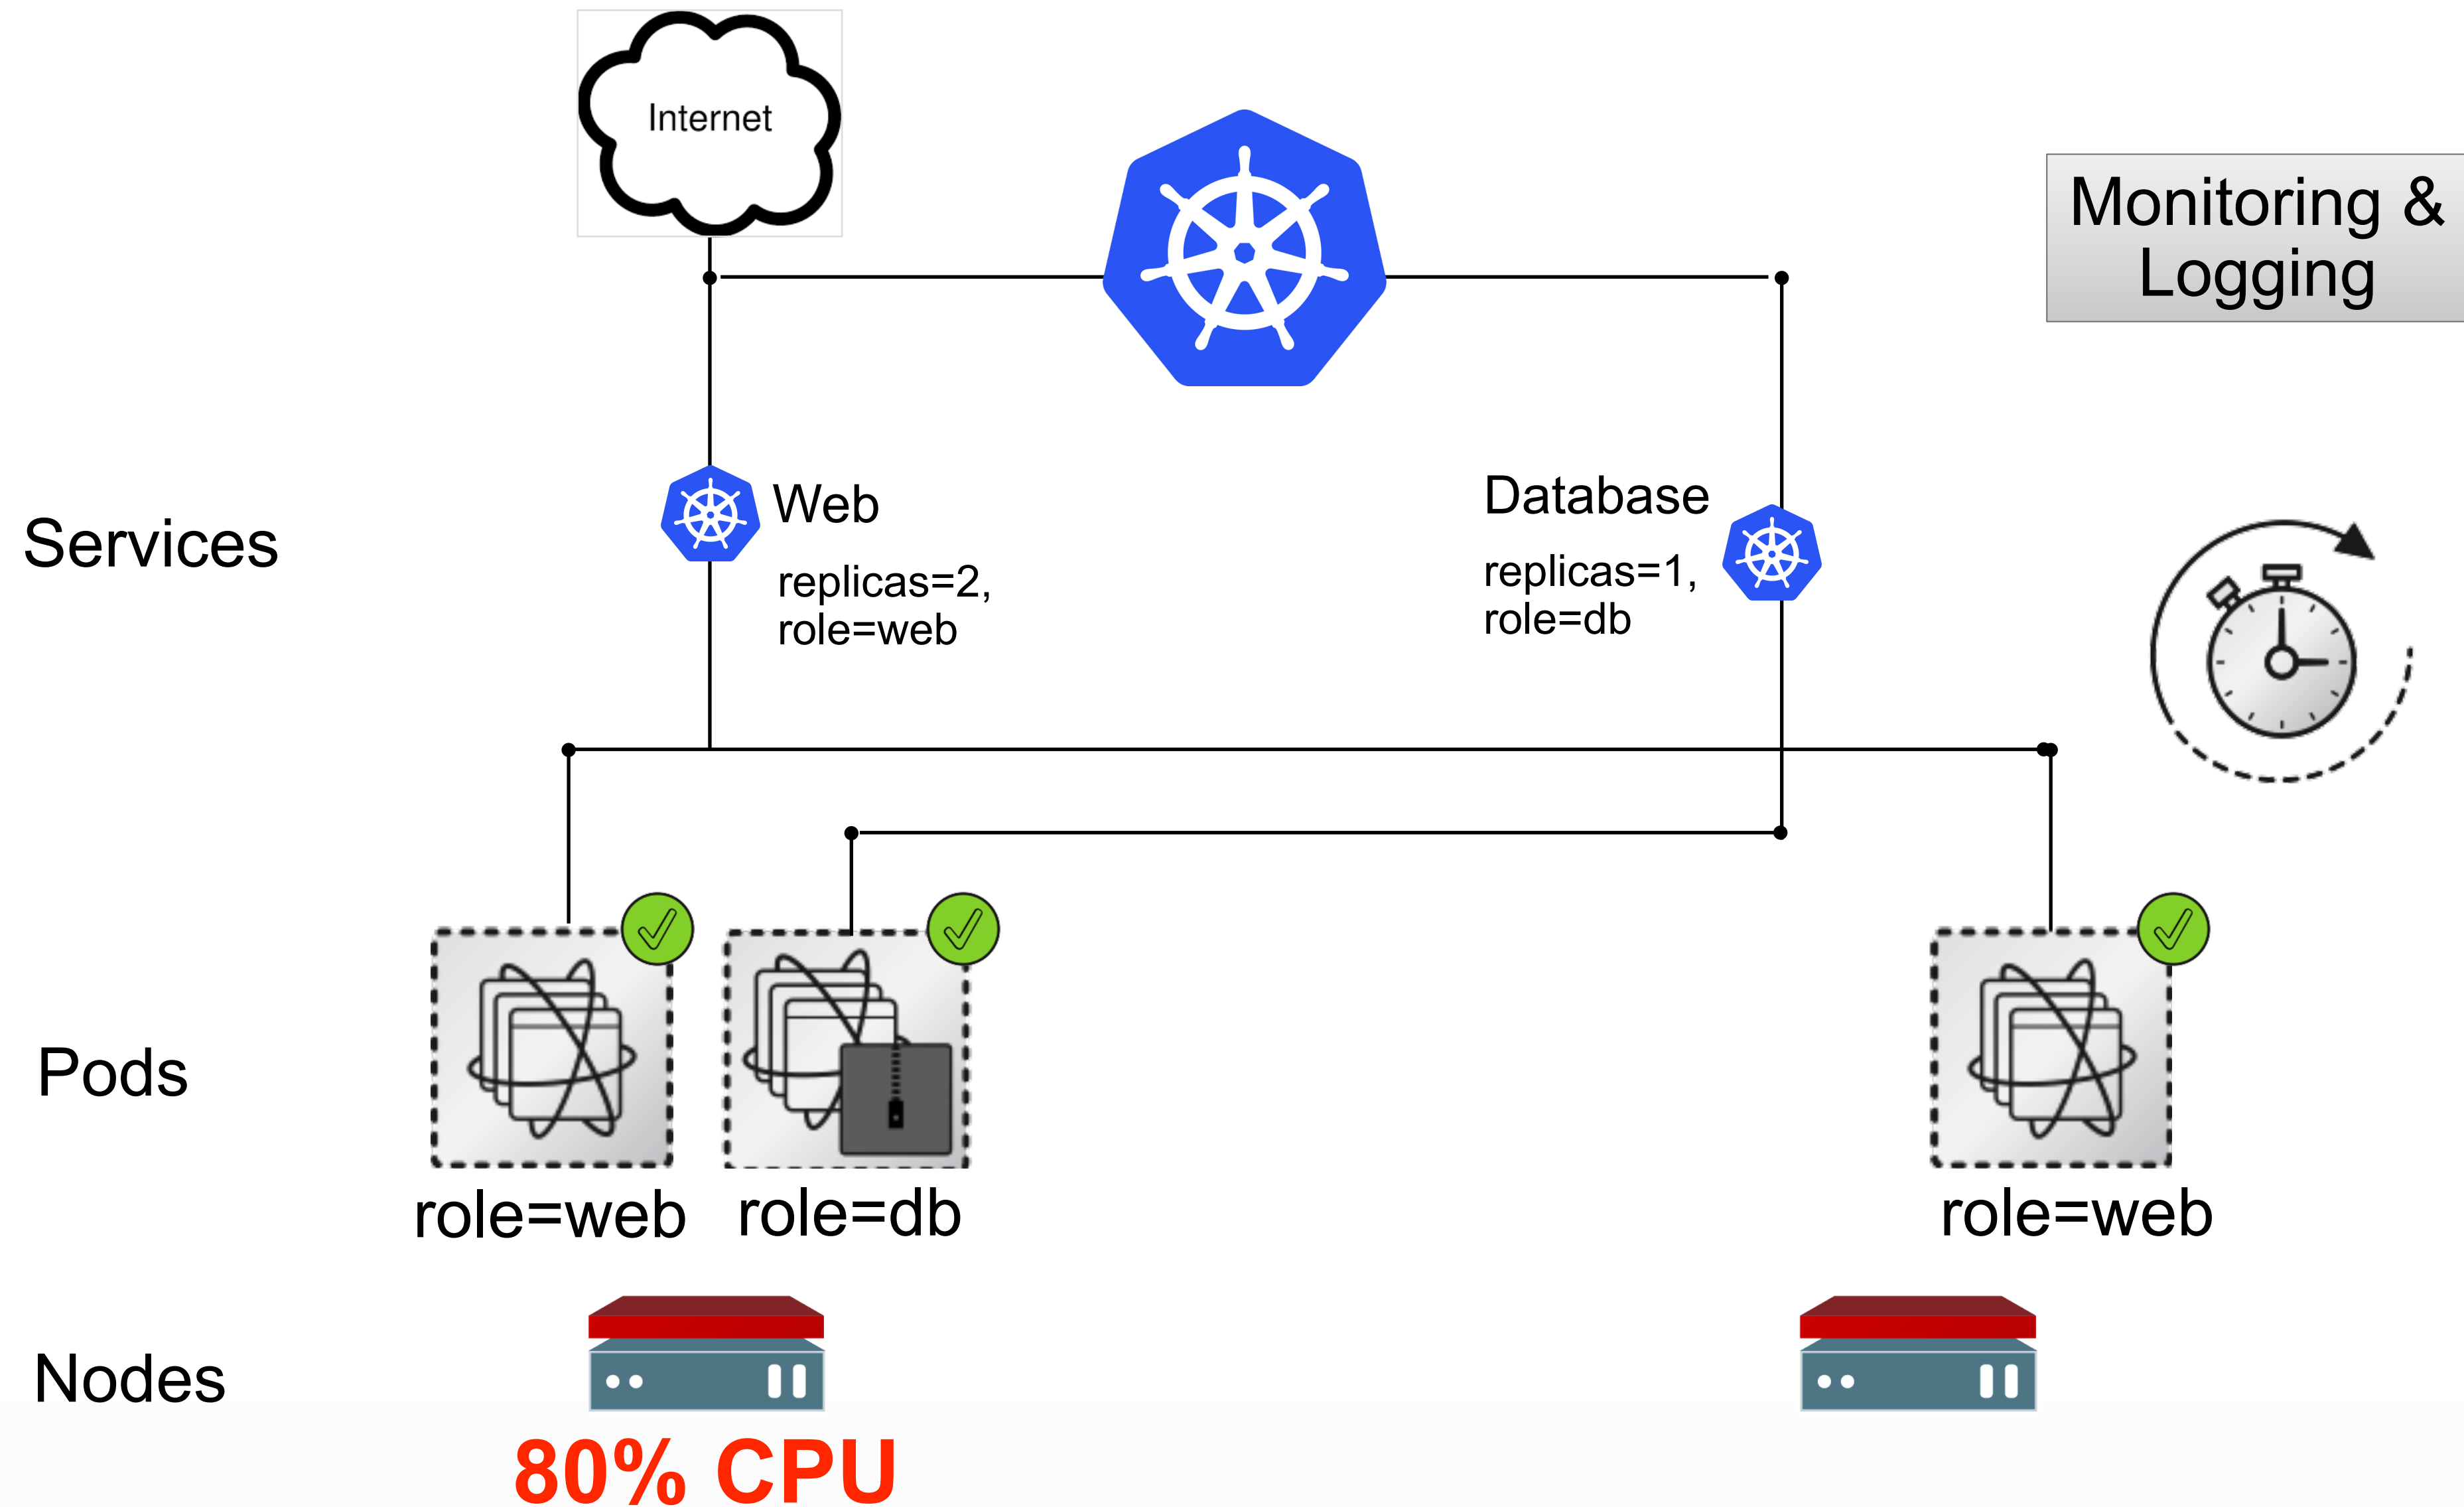


# HEALTH CHECK

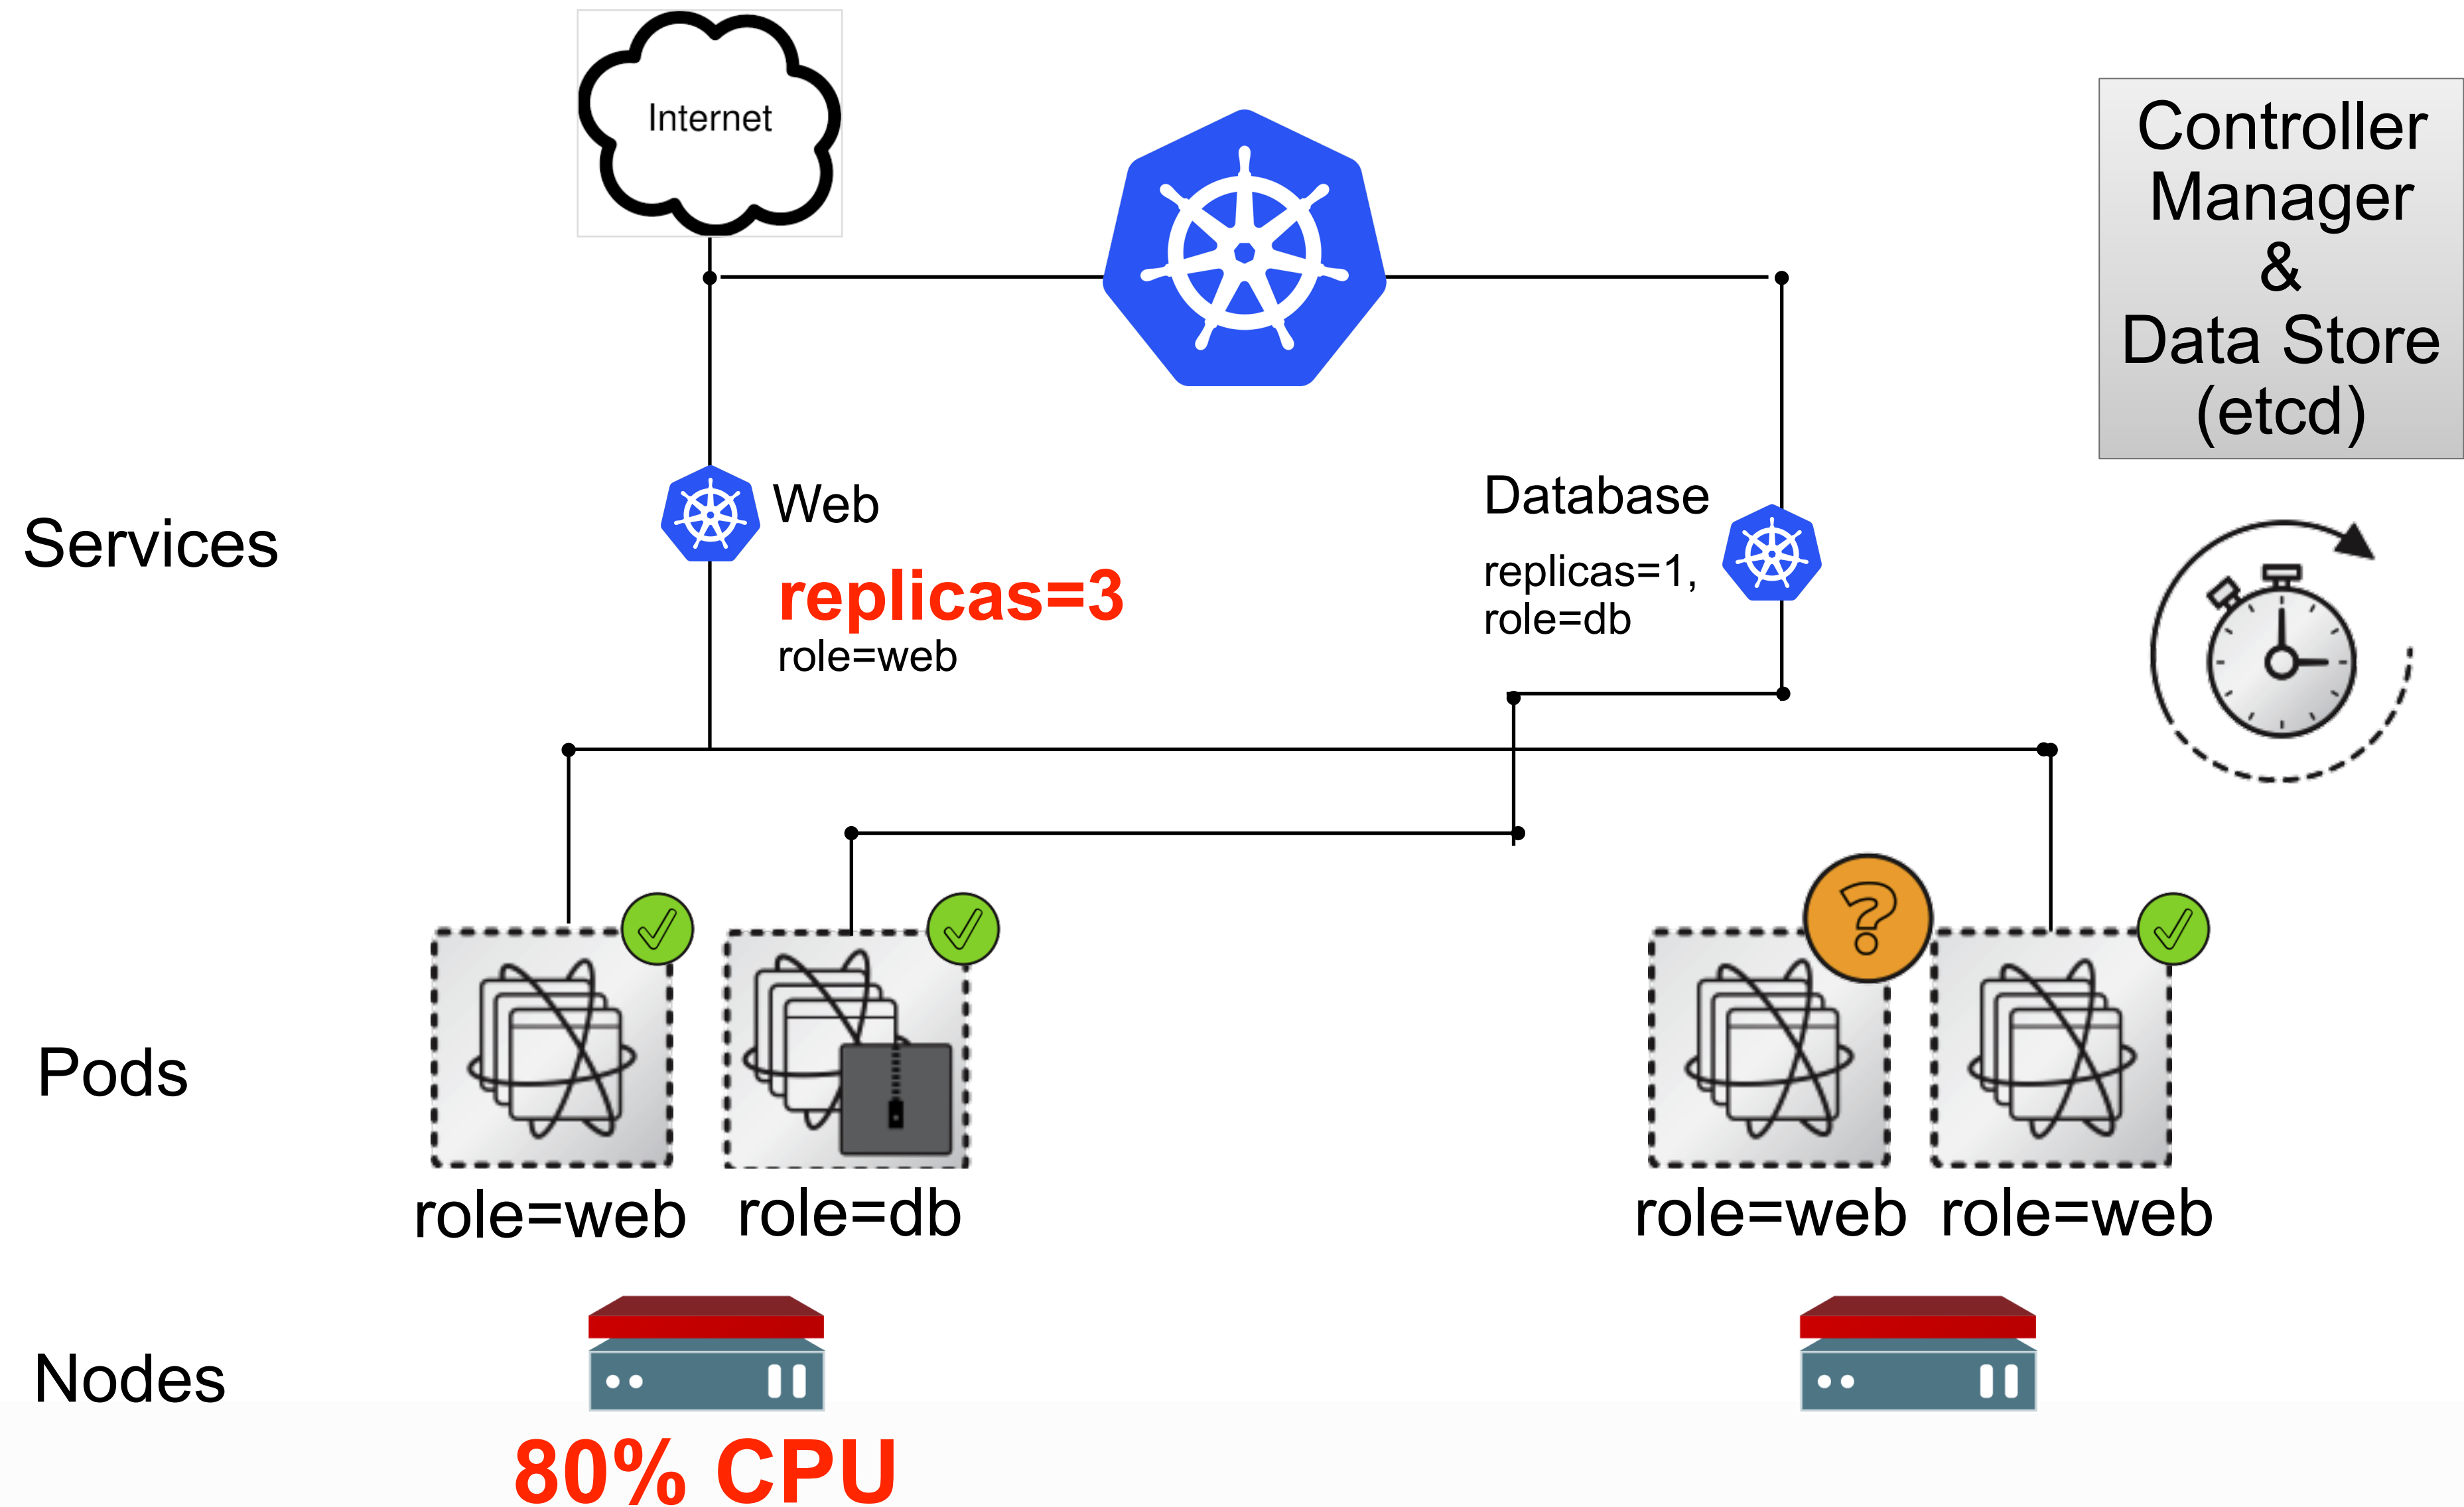





# AUTO-SCALE



# AUTO-SCALE




# AUTO-SCALE




# CONTAINER SECURITY


# SECURING CONTAINERS




**Builds**




**Images**




**Registry**




**CI/CD**




**Container  
host**



**Network  
isolation**




**Storage**



**API & Platform  
access**



**Monitoring &  
Logging**




**Federated  
clusters**



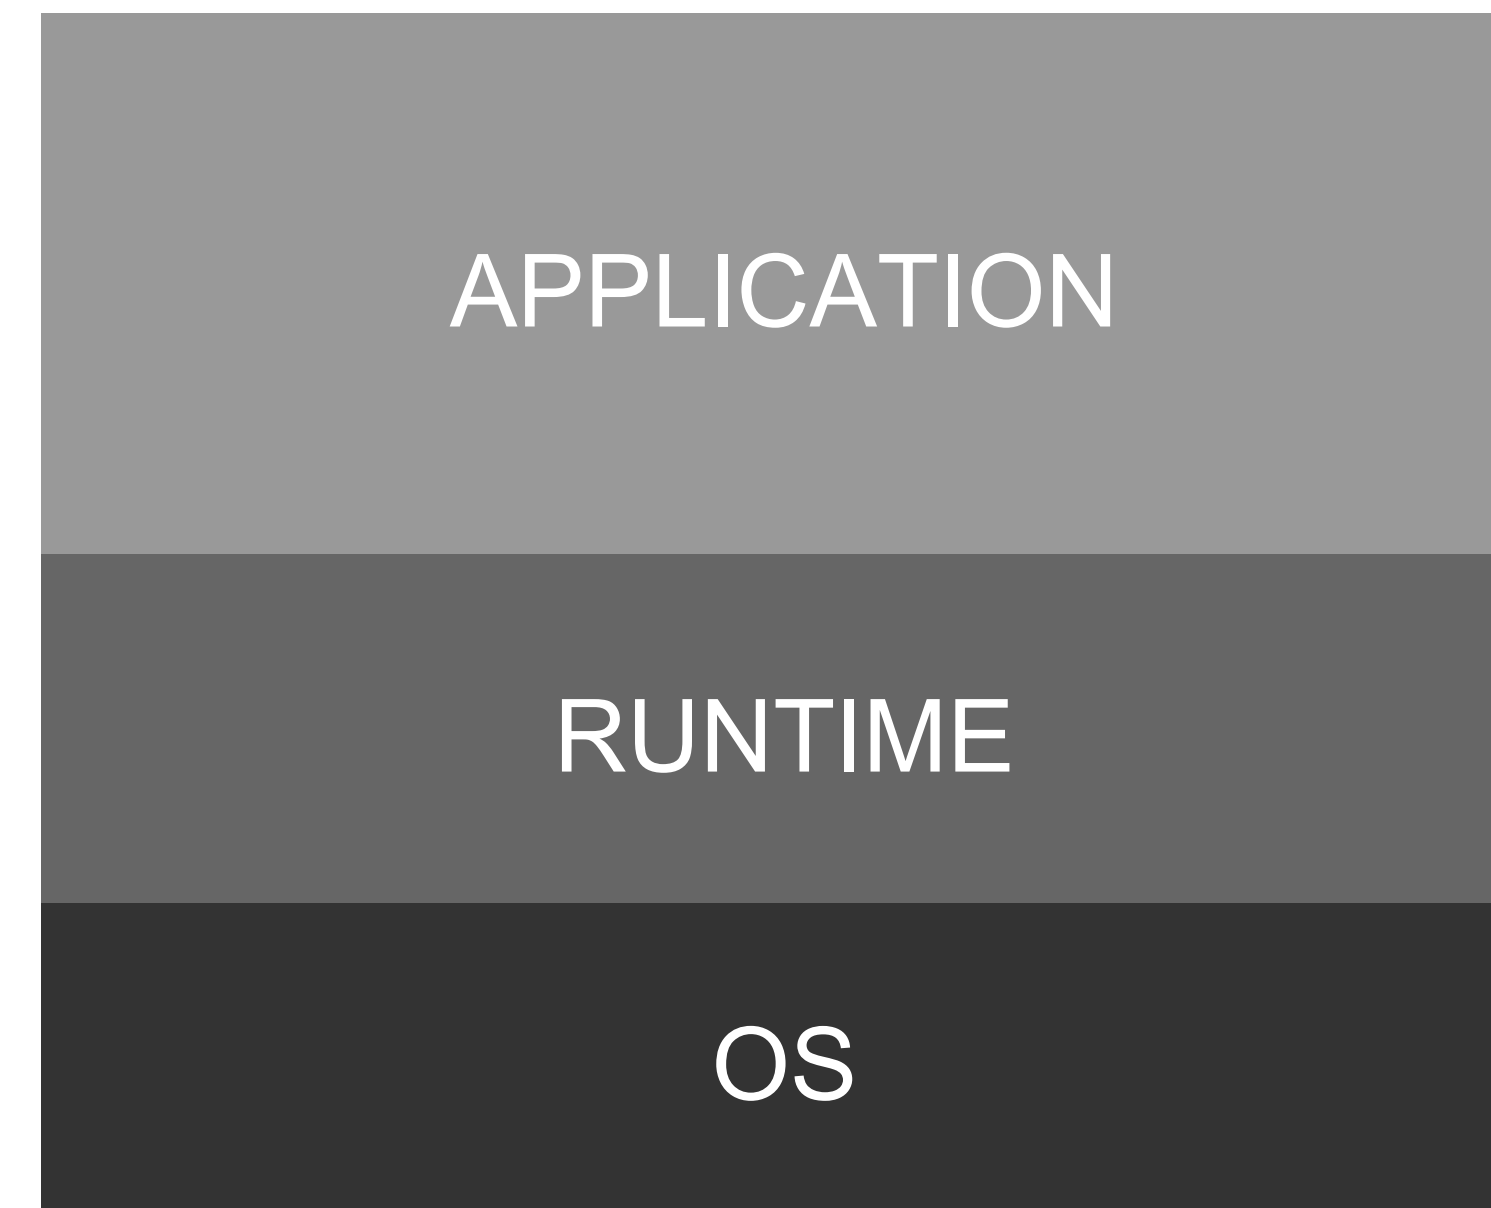
# CONTAINER BUILDS

# CONTENT: EACH LAYER MATTERS

JAR



CONTAINER




# CONTAINER BUILDS

## Build

```
FROM fedora:1.0  
CMD echo "Hello"
```

## Build file

## Ship



## Image

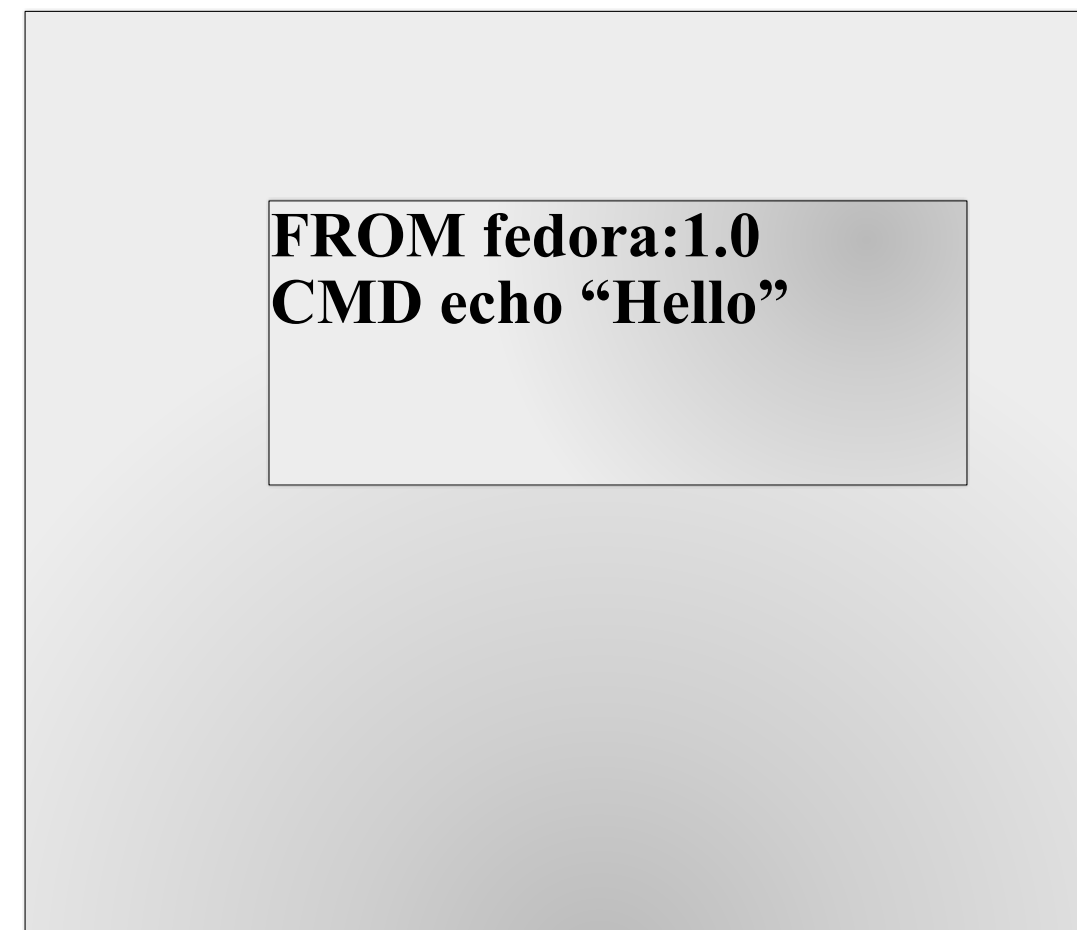
## Run



## Container

# CONTAINER BUILDS

## Build

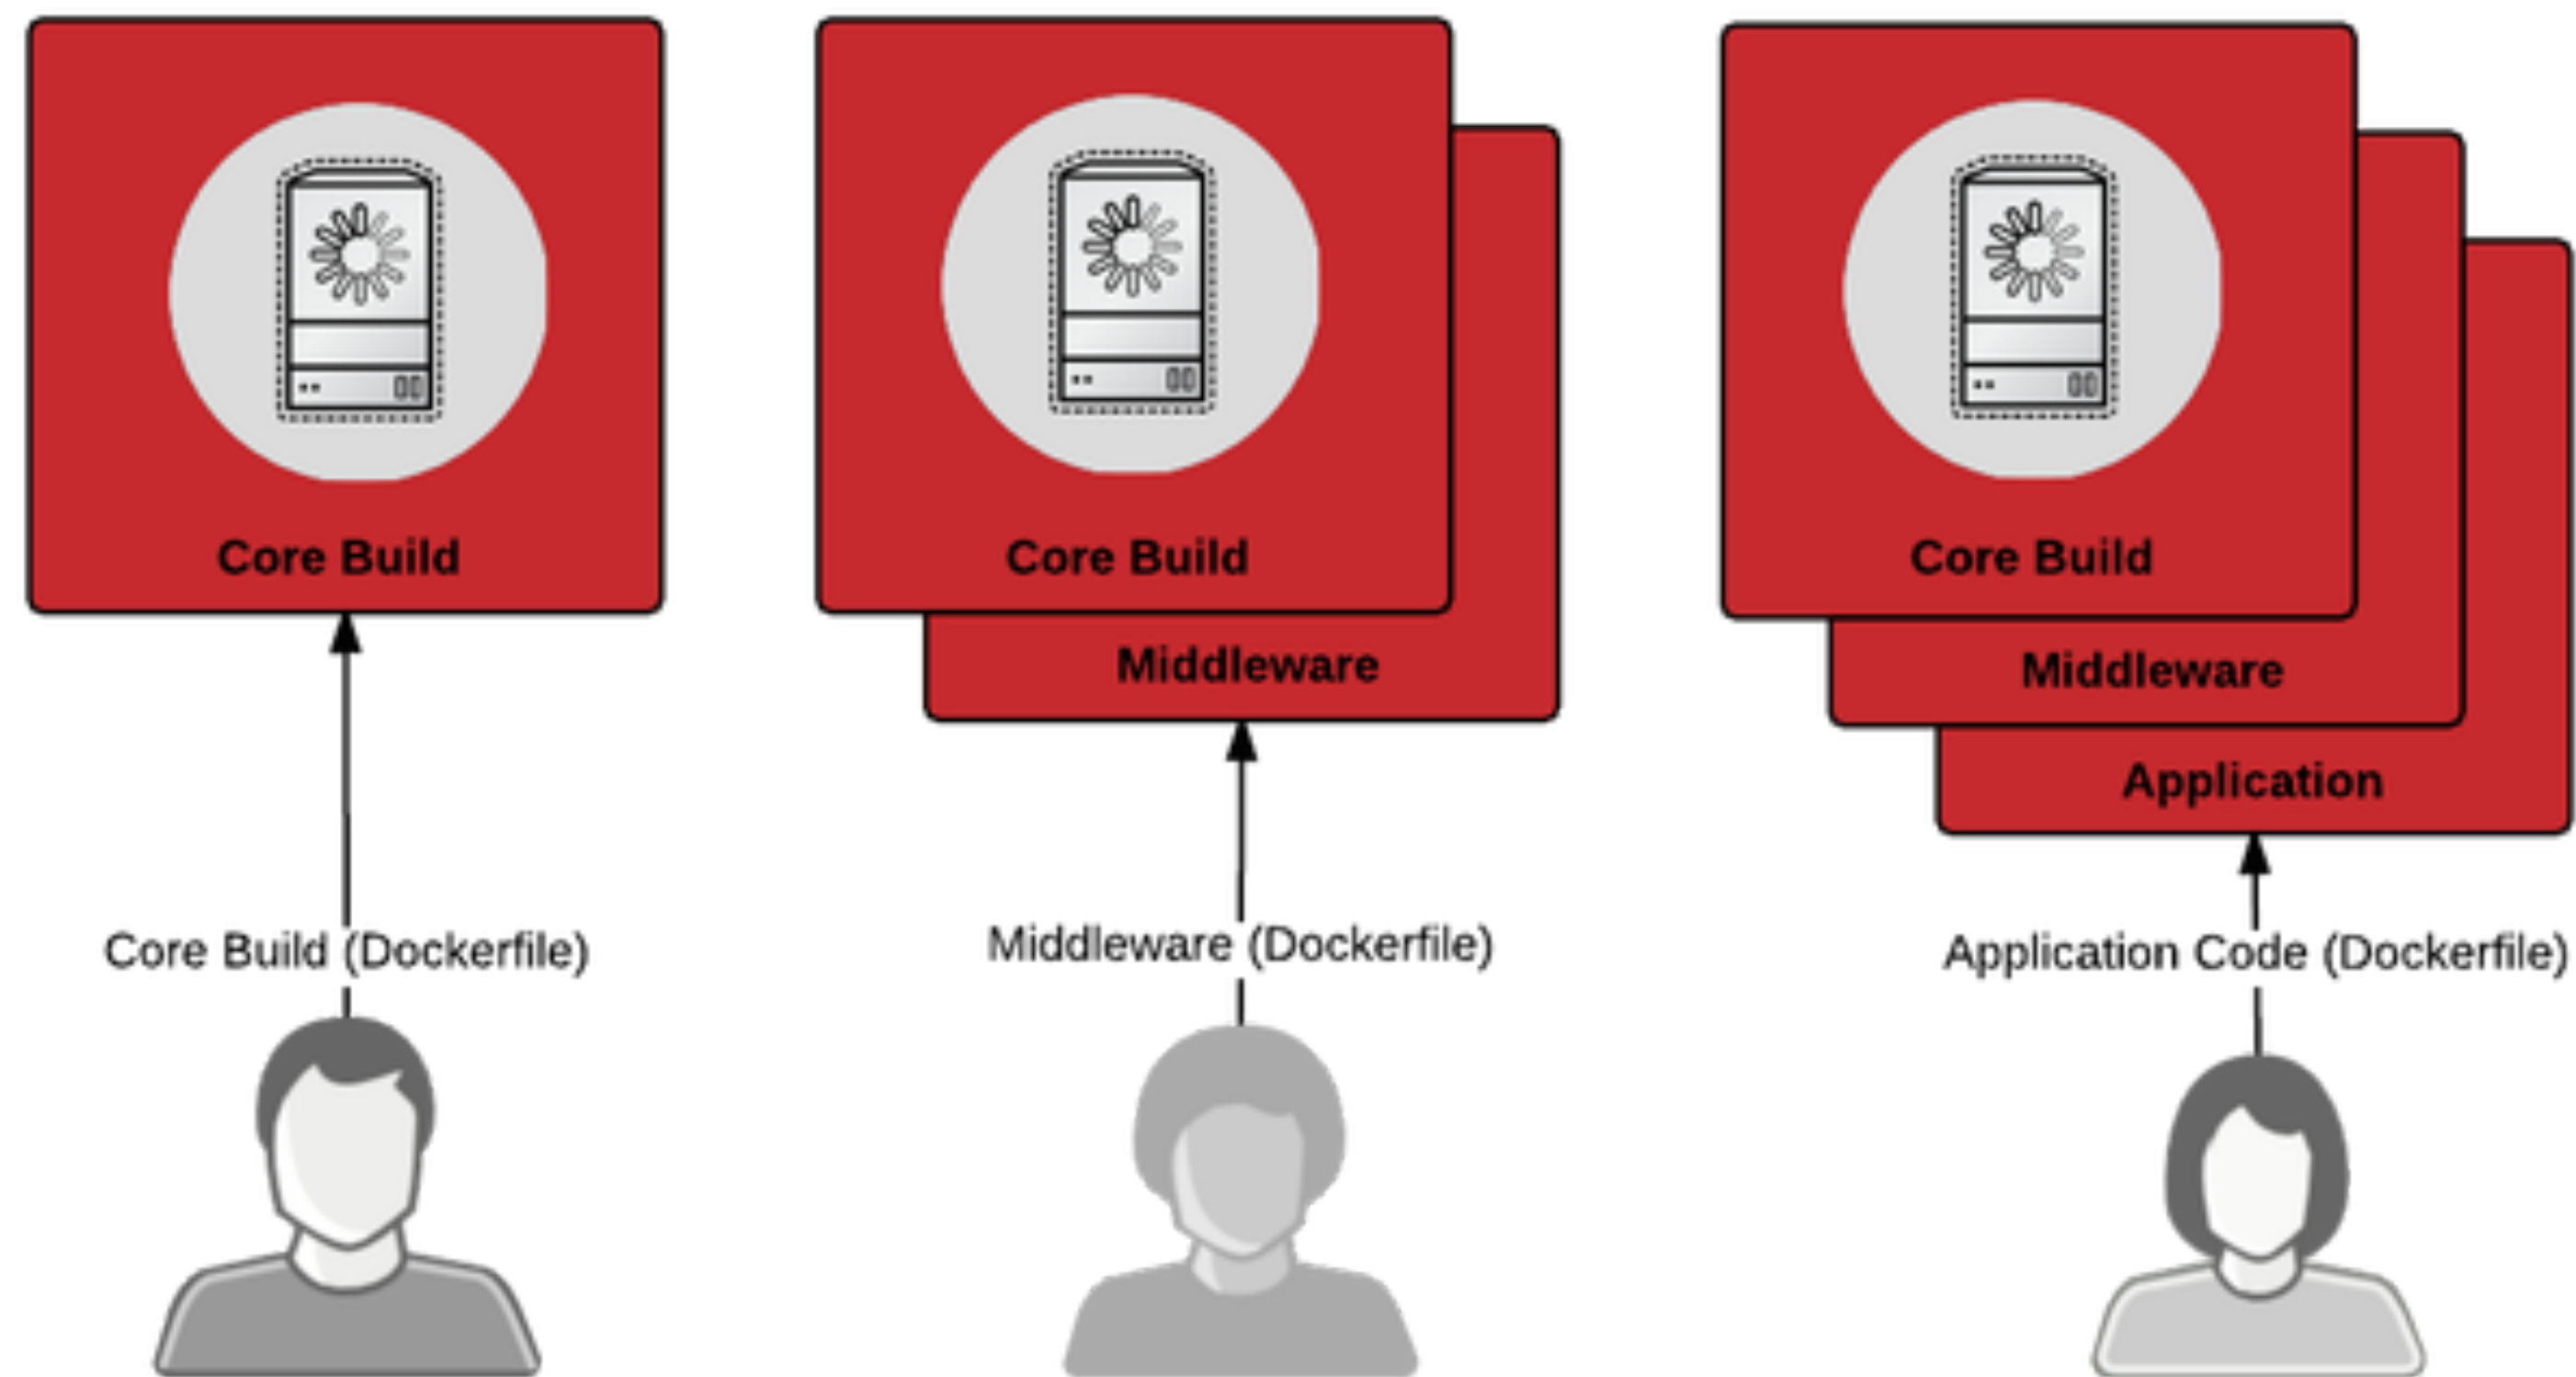


## Build file

## Best Practices

- Treat as a Blueprint
- Specify a user, defaults to root
- Don't login to build/configure
- Version control build file
- Be explicit with versions, not latest
- Each Run creates a new layer

# A CONVERGED SOFTWARE SUPPLY CHAIN




# CONTAINER IMAGE SECURITY



# WHAT'S INSIDE THE CONTAINER MATTERS

64% of official images in Docker Hub  
contain **high** priority security vulnerabilities



examples:

**ShellShock (bash)**

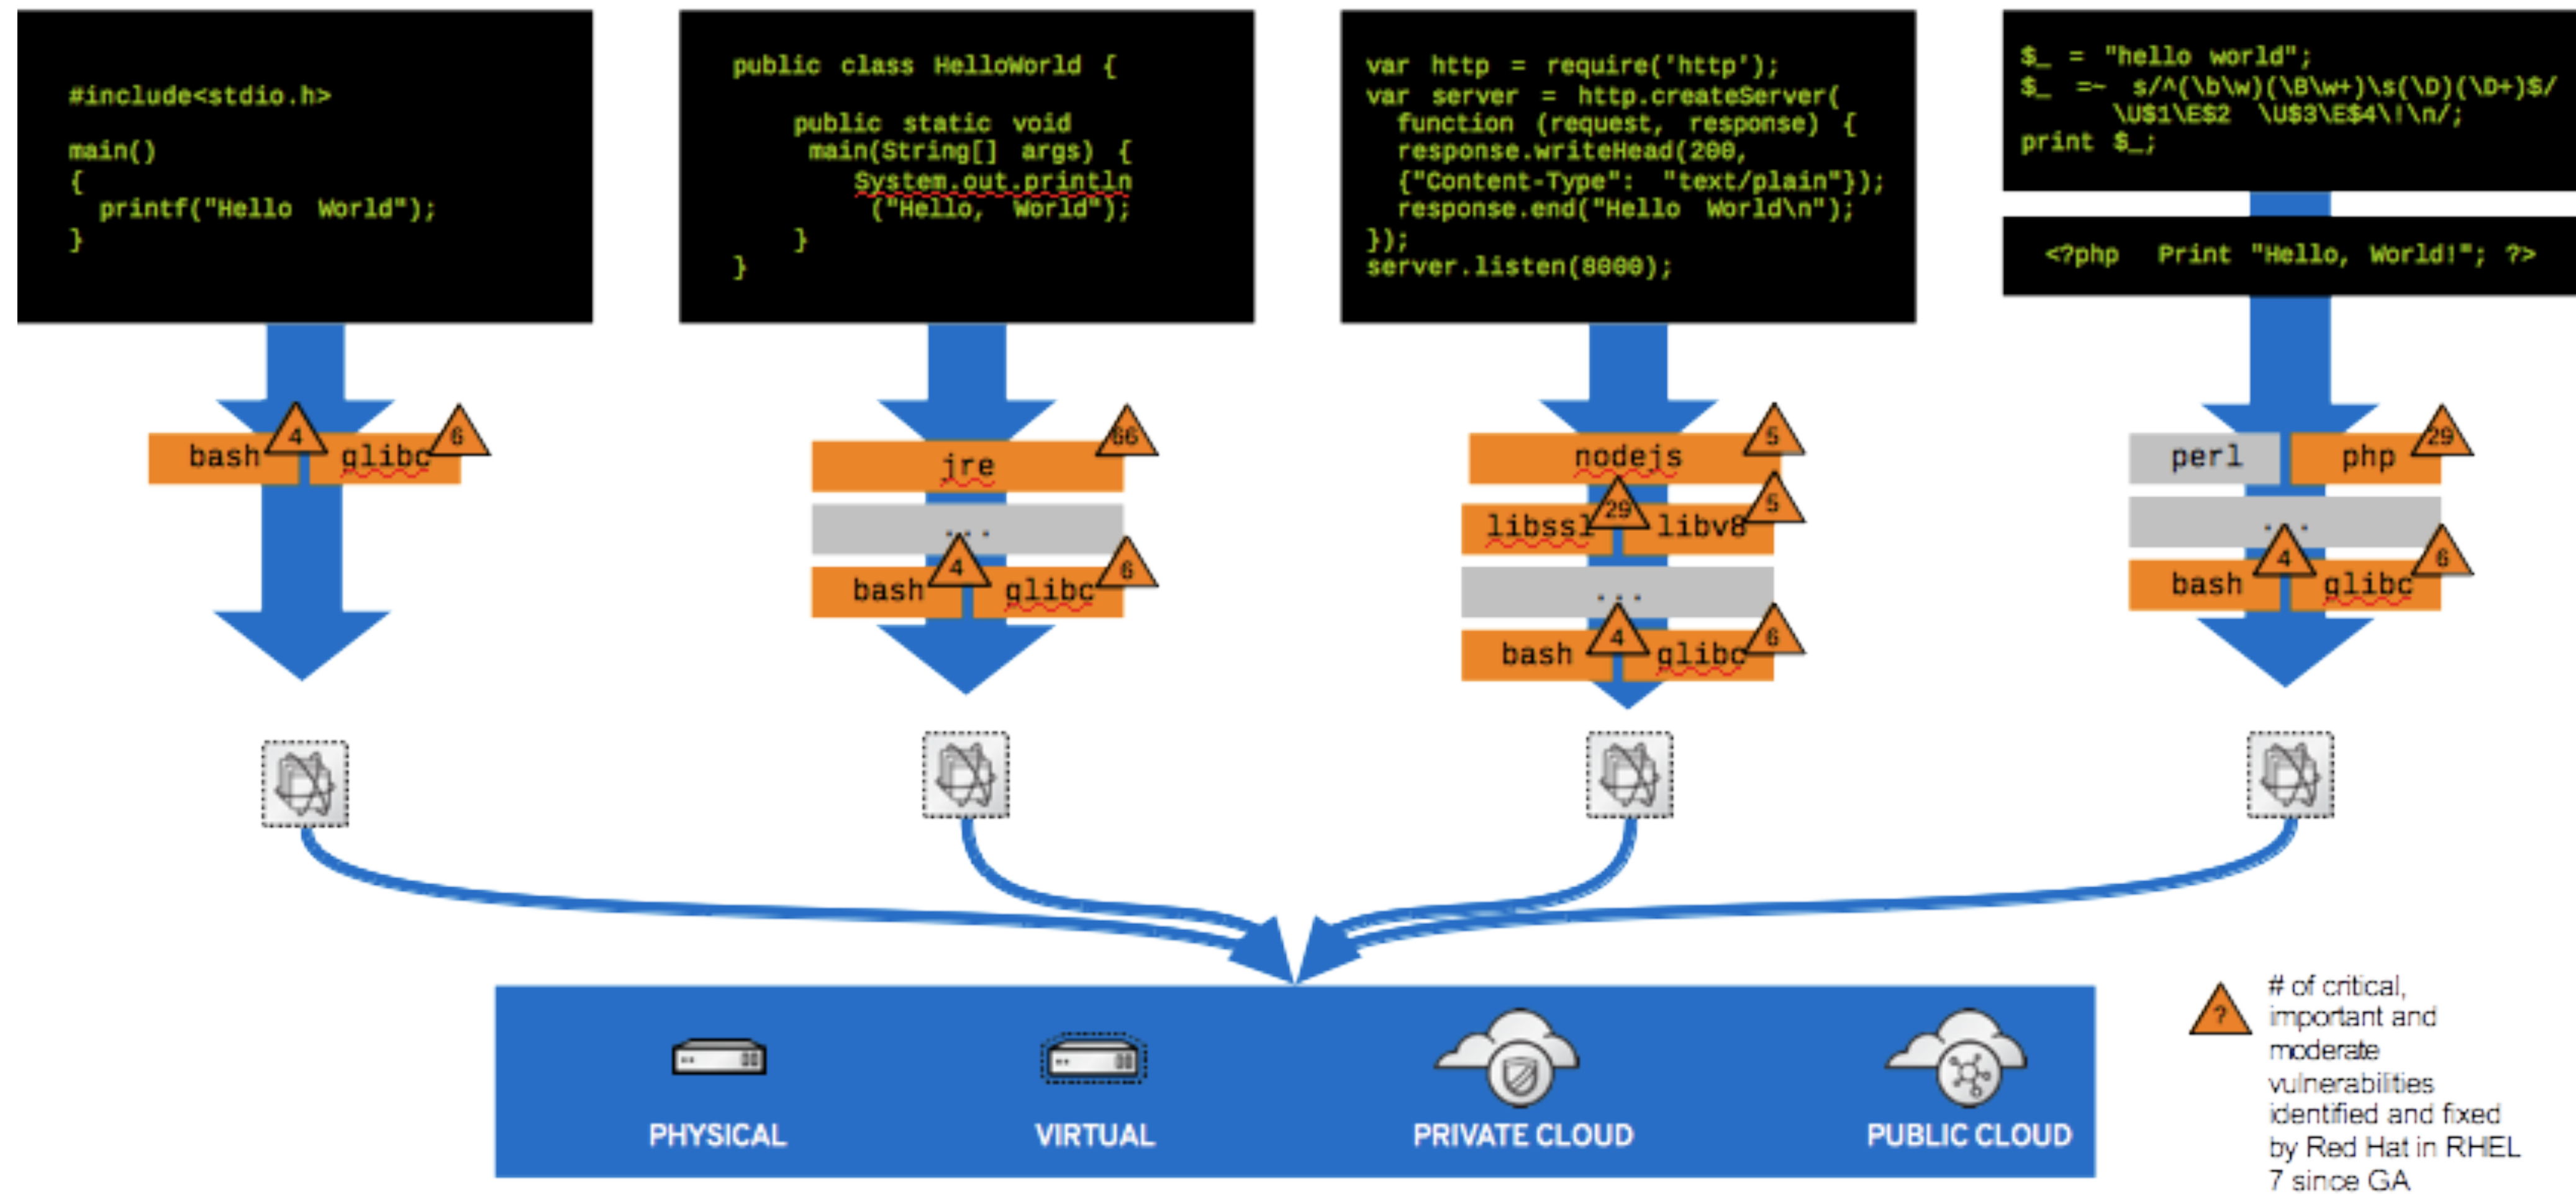
**Heartbleed (OpenSSL)**

**Poodle (OpenSSL)**

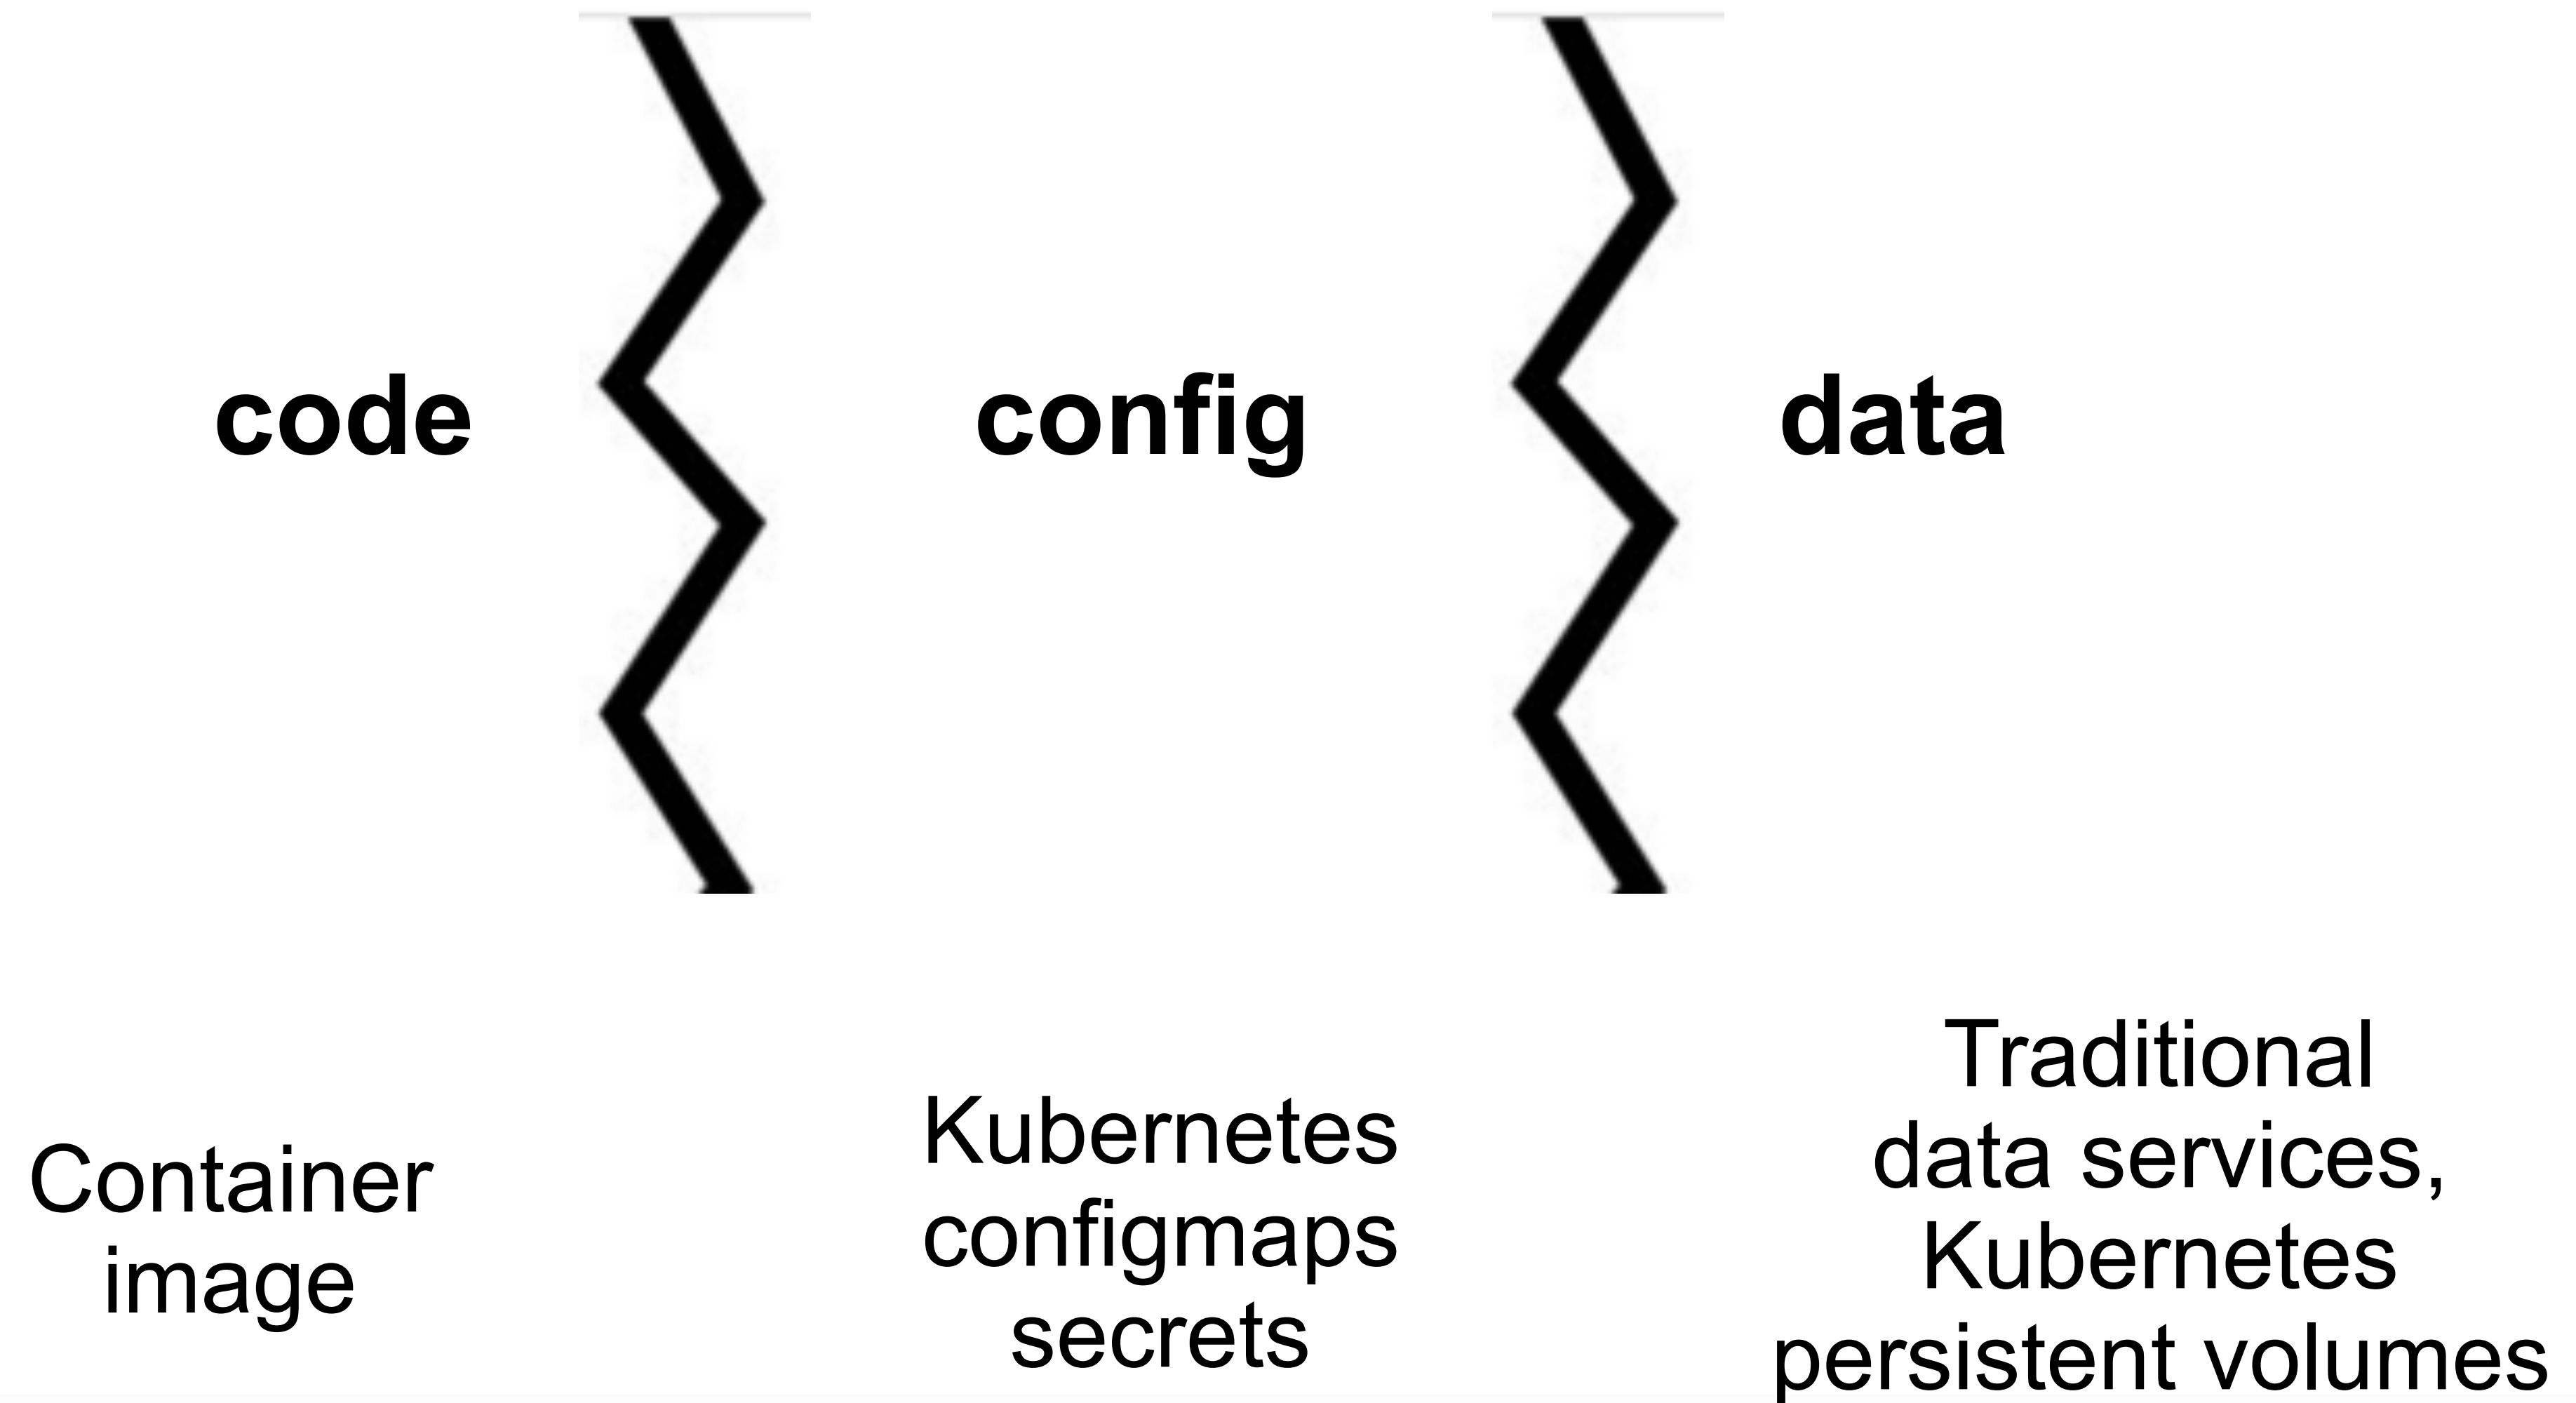
Source: *Over 30% of Official Images in Docker Hub Contain High Priority Security Vulnerabilities*, Jayanth Gummaraju, Tarun Desikan, and Yoshio Turner, BanyanOps, May 2015 (<http://www.banyanops.com/pdf/BanyanOps-AnalyzingDockerHub-WhitePaper.pdf>)

# SECURITY IMPLICATIONS

## What's inside matters...



# TREAT CONTAINERS AS IMMUTABLE




# CONTAINER REGISTRY SECURITY



# PRIVATE REGISTRY

REGISTRY

 project/test:latest [Show all images](#)

### Image

Description These are images in the test project, more description data goes here

Source <http://project.example.com/test>

Author Stef Walter <stefw@redhat.com>

Built 3 days ago

Digest sha256:91e54dfb11794fad694460162bf0cb0a4fa710cfa3f60979c177

Tags project/test:latest project/test:0.5

```
$ sudo docker pull 172.30.0.0/project/test:latest
```

### Configuration

```
# /bin/registry "/etc/docker/registry/config.yml"
```

Run as	root	Environment	PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin
Working Directory	/		
Stop with	SIGTERM		
Architecture	amd64		

---

Ports	5000 (TCP)	Volumes	/var/lib/registry
-------	------------	---------	-------------------

### Metadata




Image Layers

RUN yum -y update && yum install -y httpd

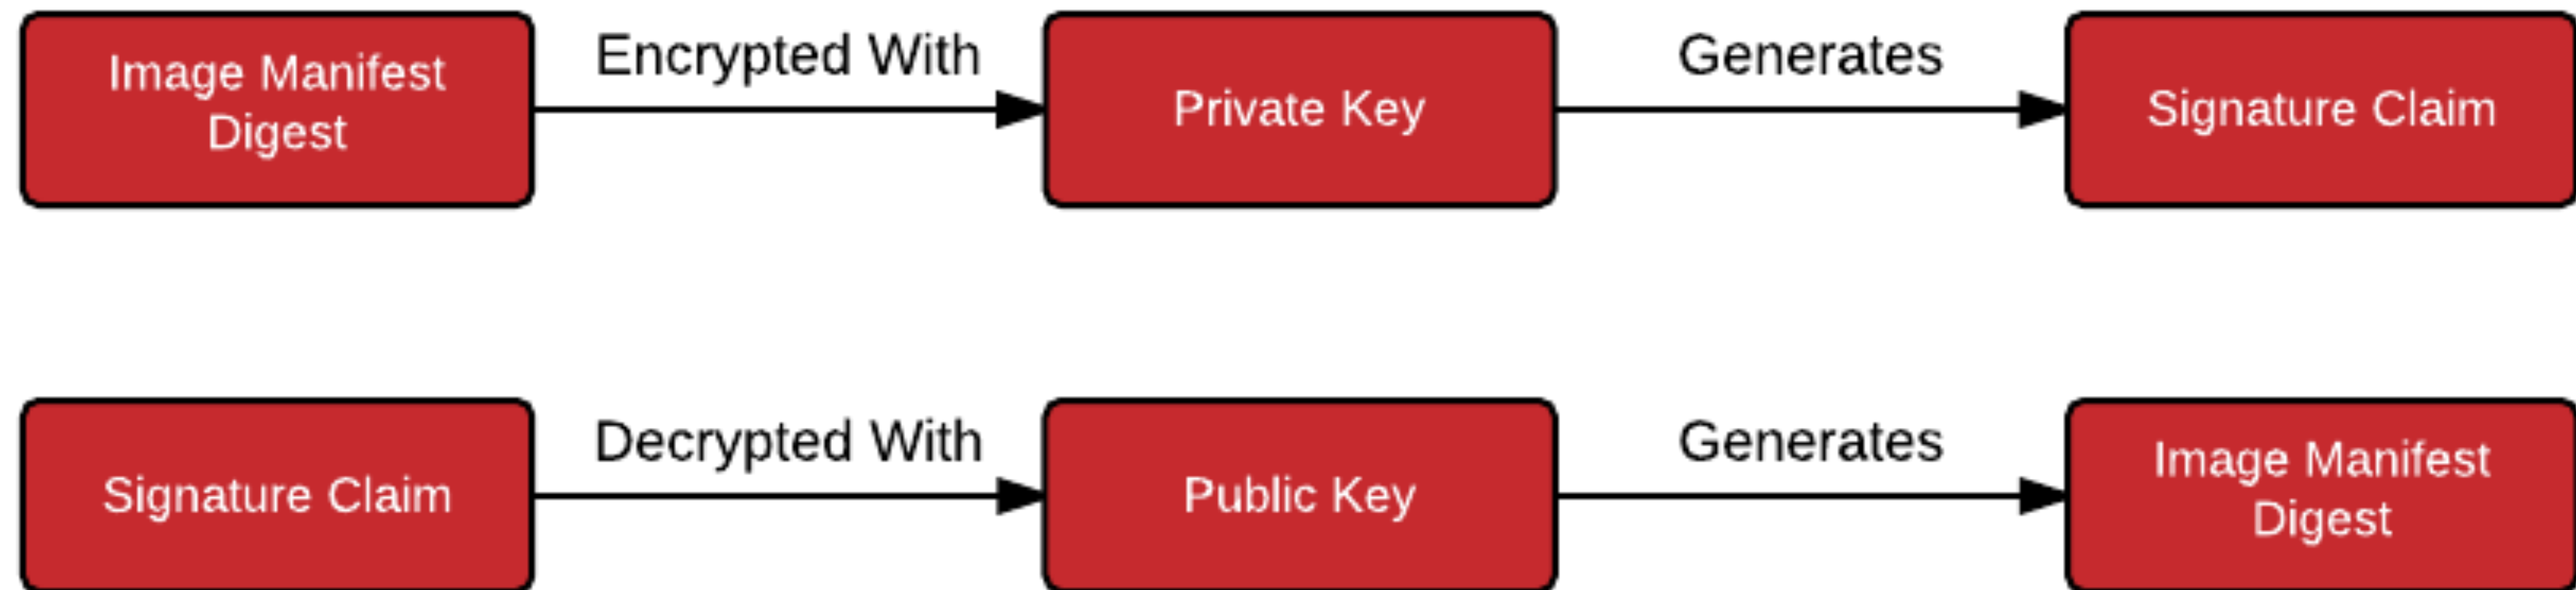
Labels INSTALL docker run -v /srv/registry:/var/lib/registry --privileged IMAGE /install

Annotations openshift.blah: Blah blah blah

Docker Version 1.7.4

# IMAGE SIGNING

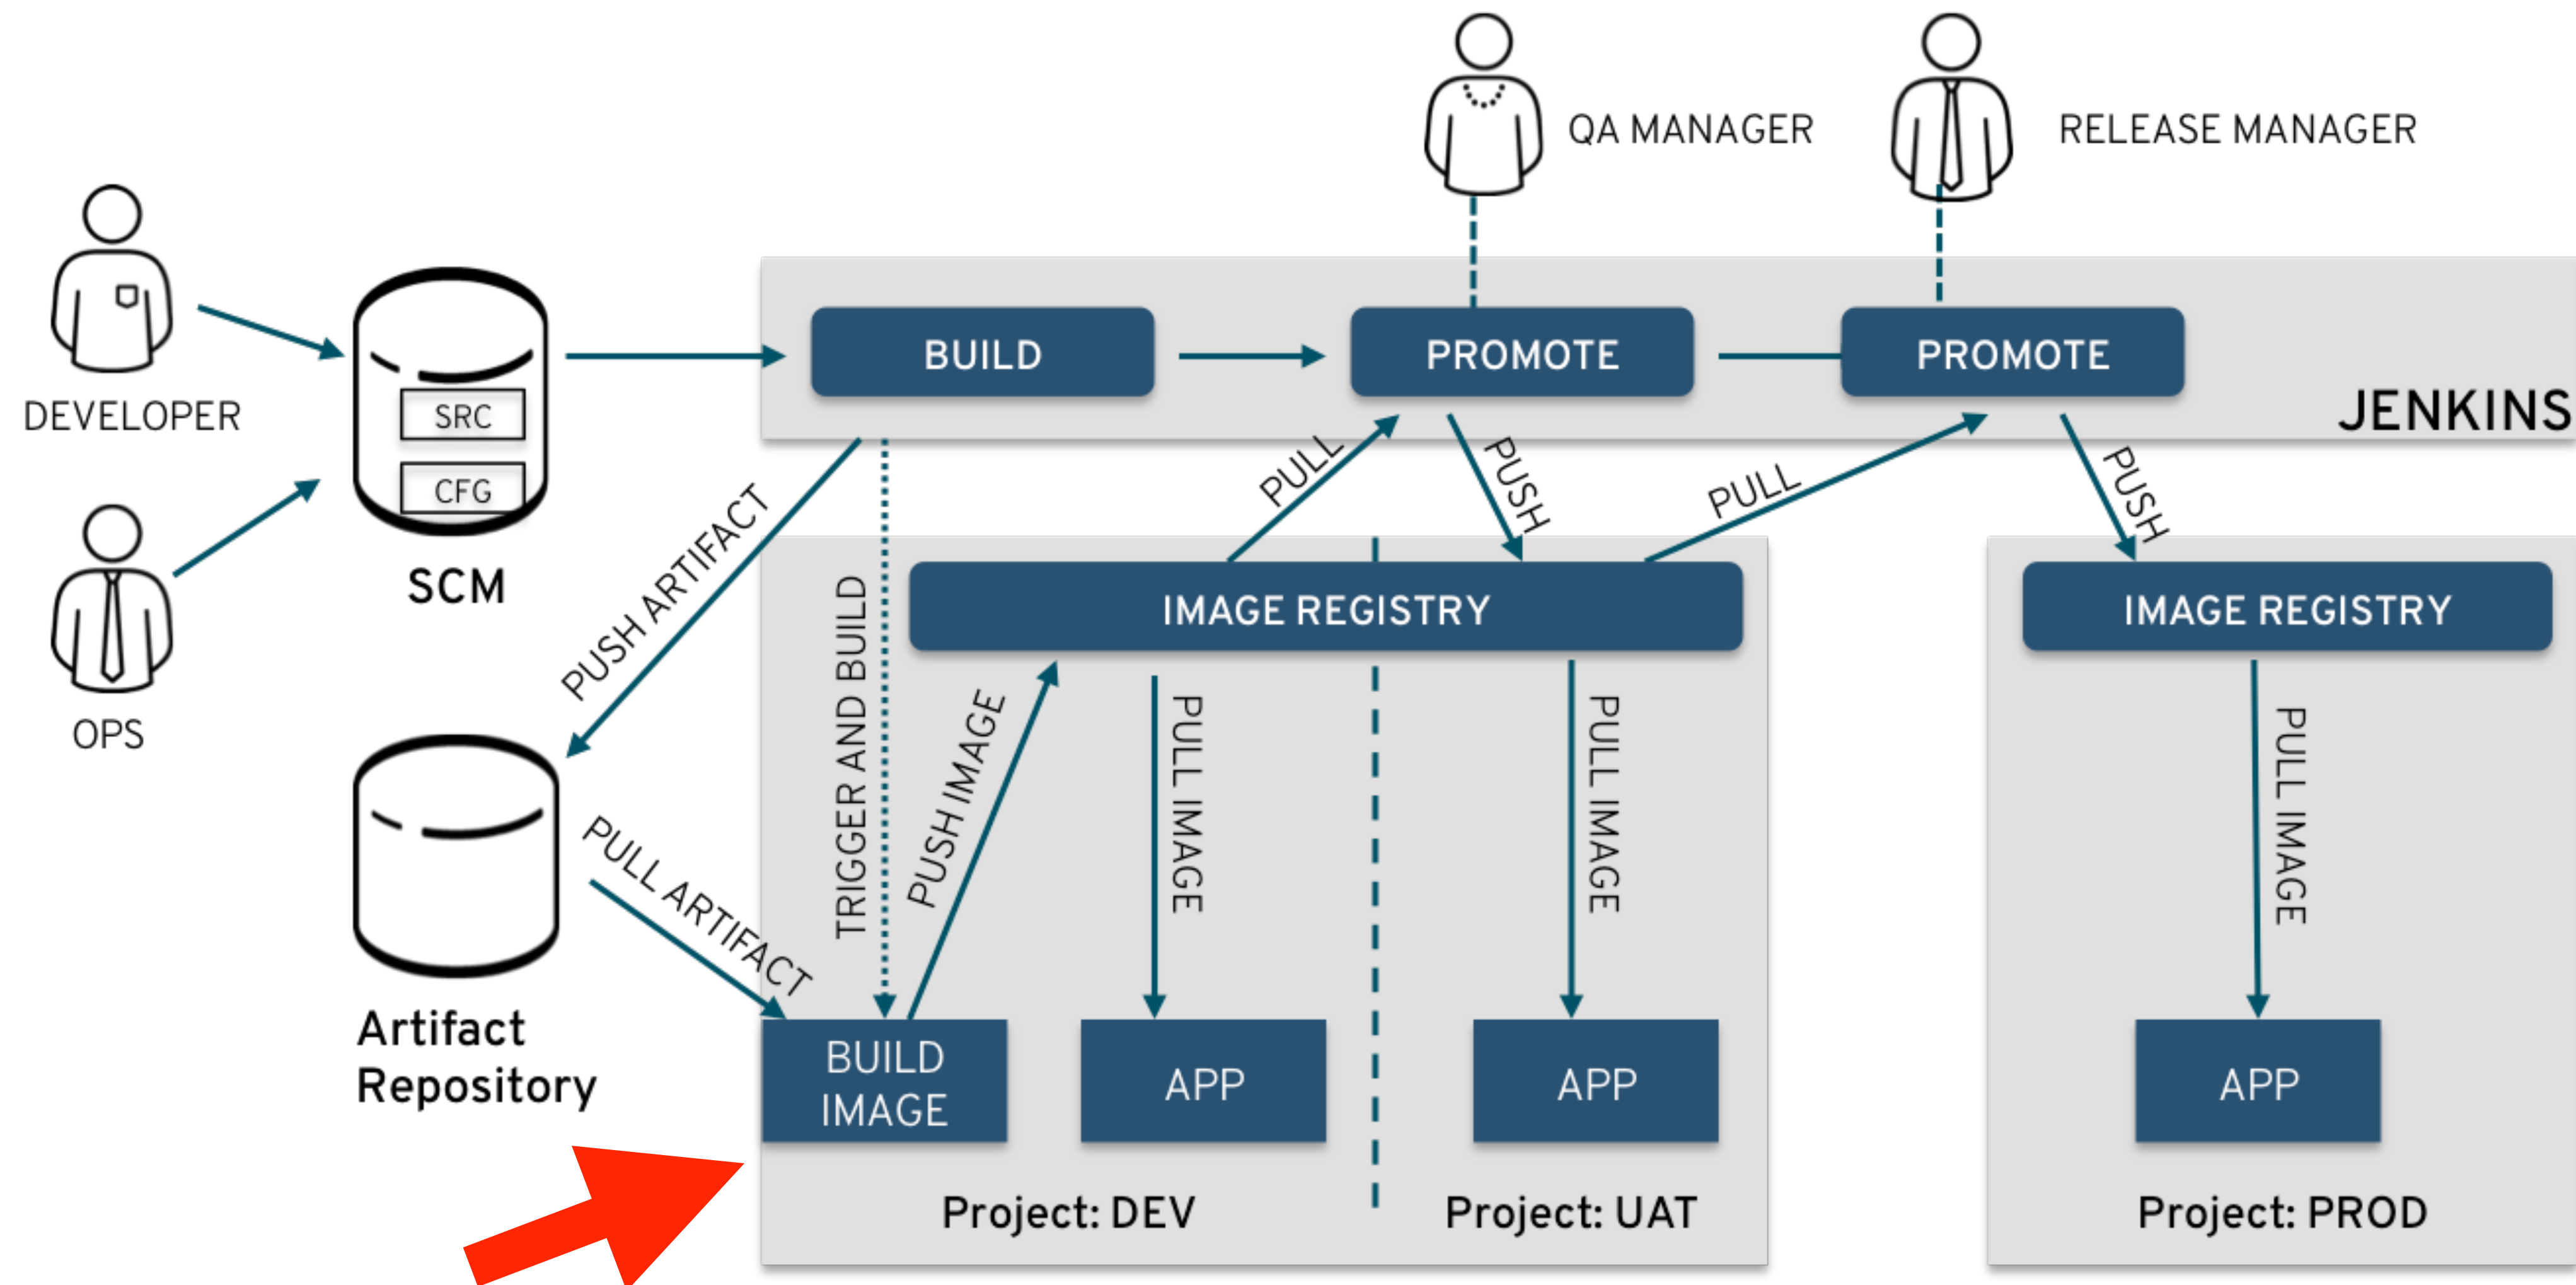
Validate what images and version are running



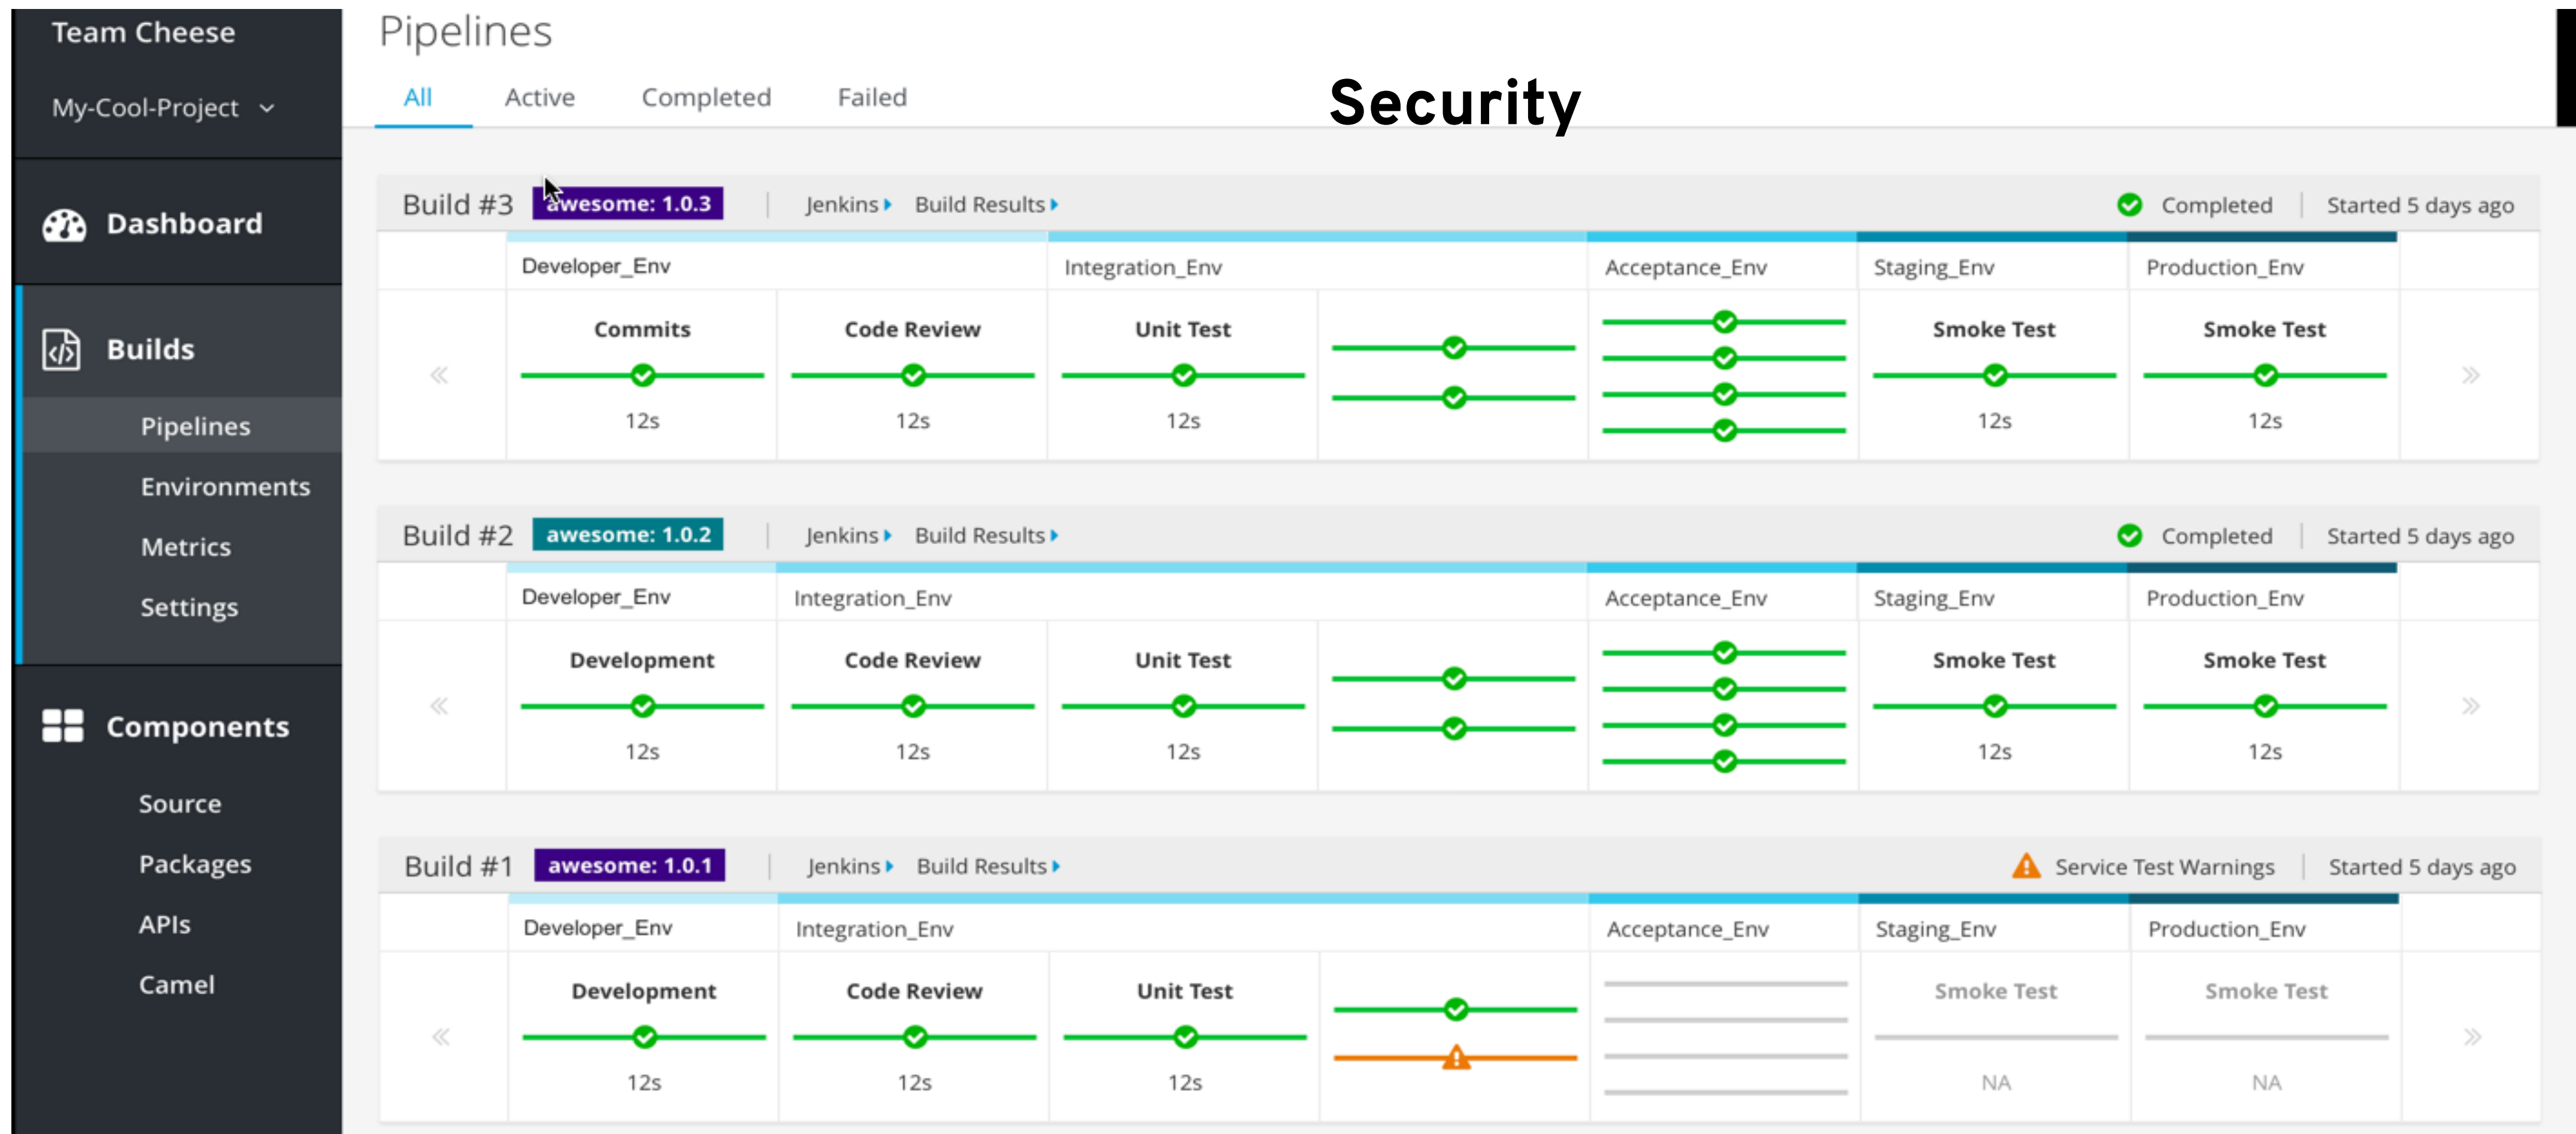


# CONTINUOUS INTEGRATION WITH CONTAINERS

# CONTINUOUS INTEGRATION + SECURITY

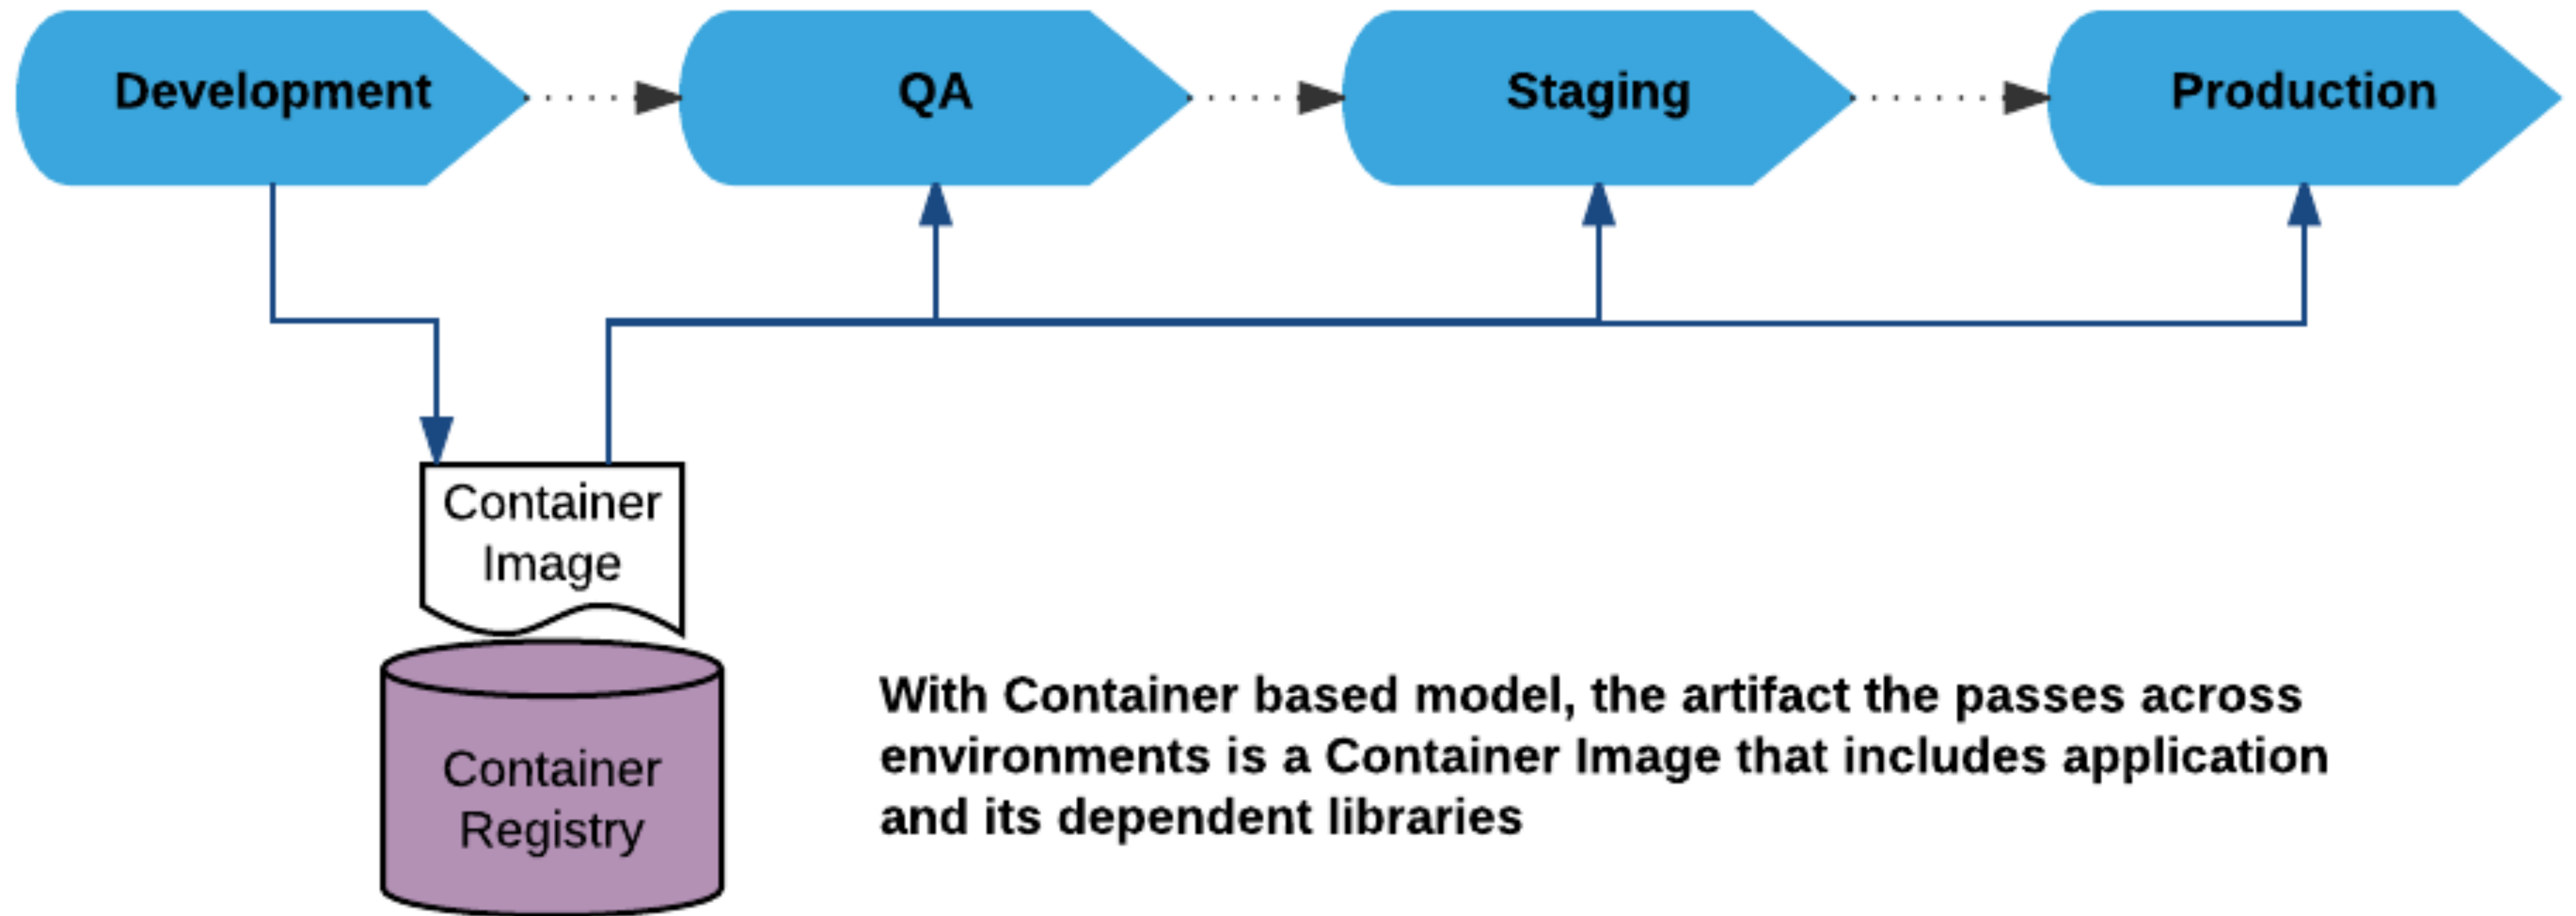


# CONTINUOUS INTEGRATION WITH SECURITY SCAN



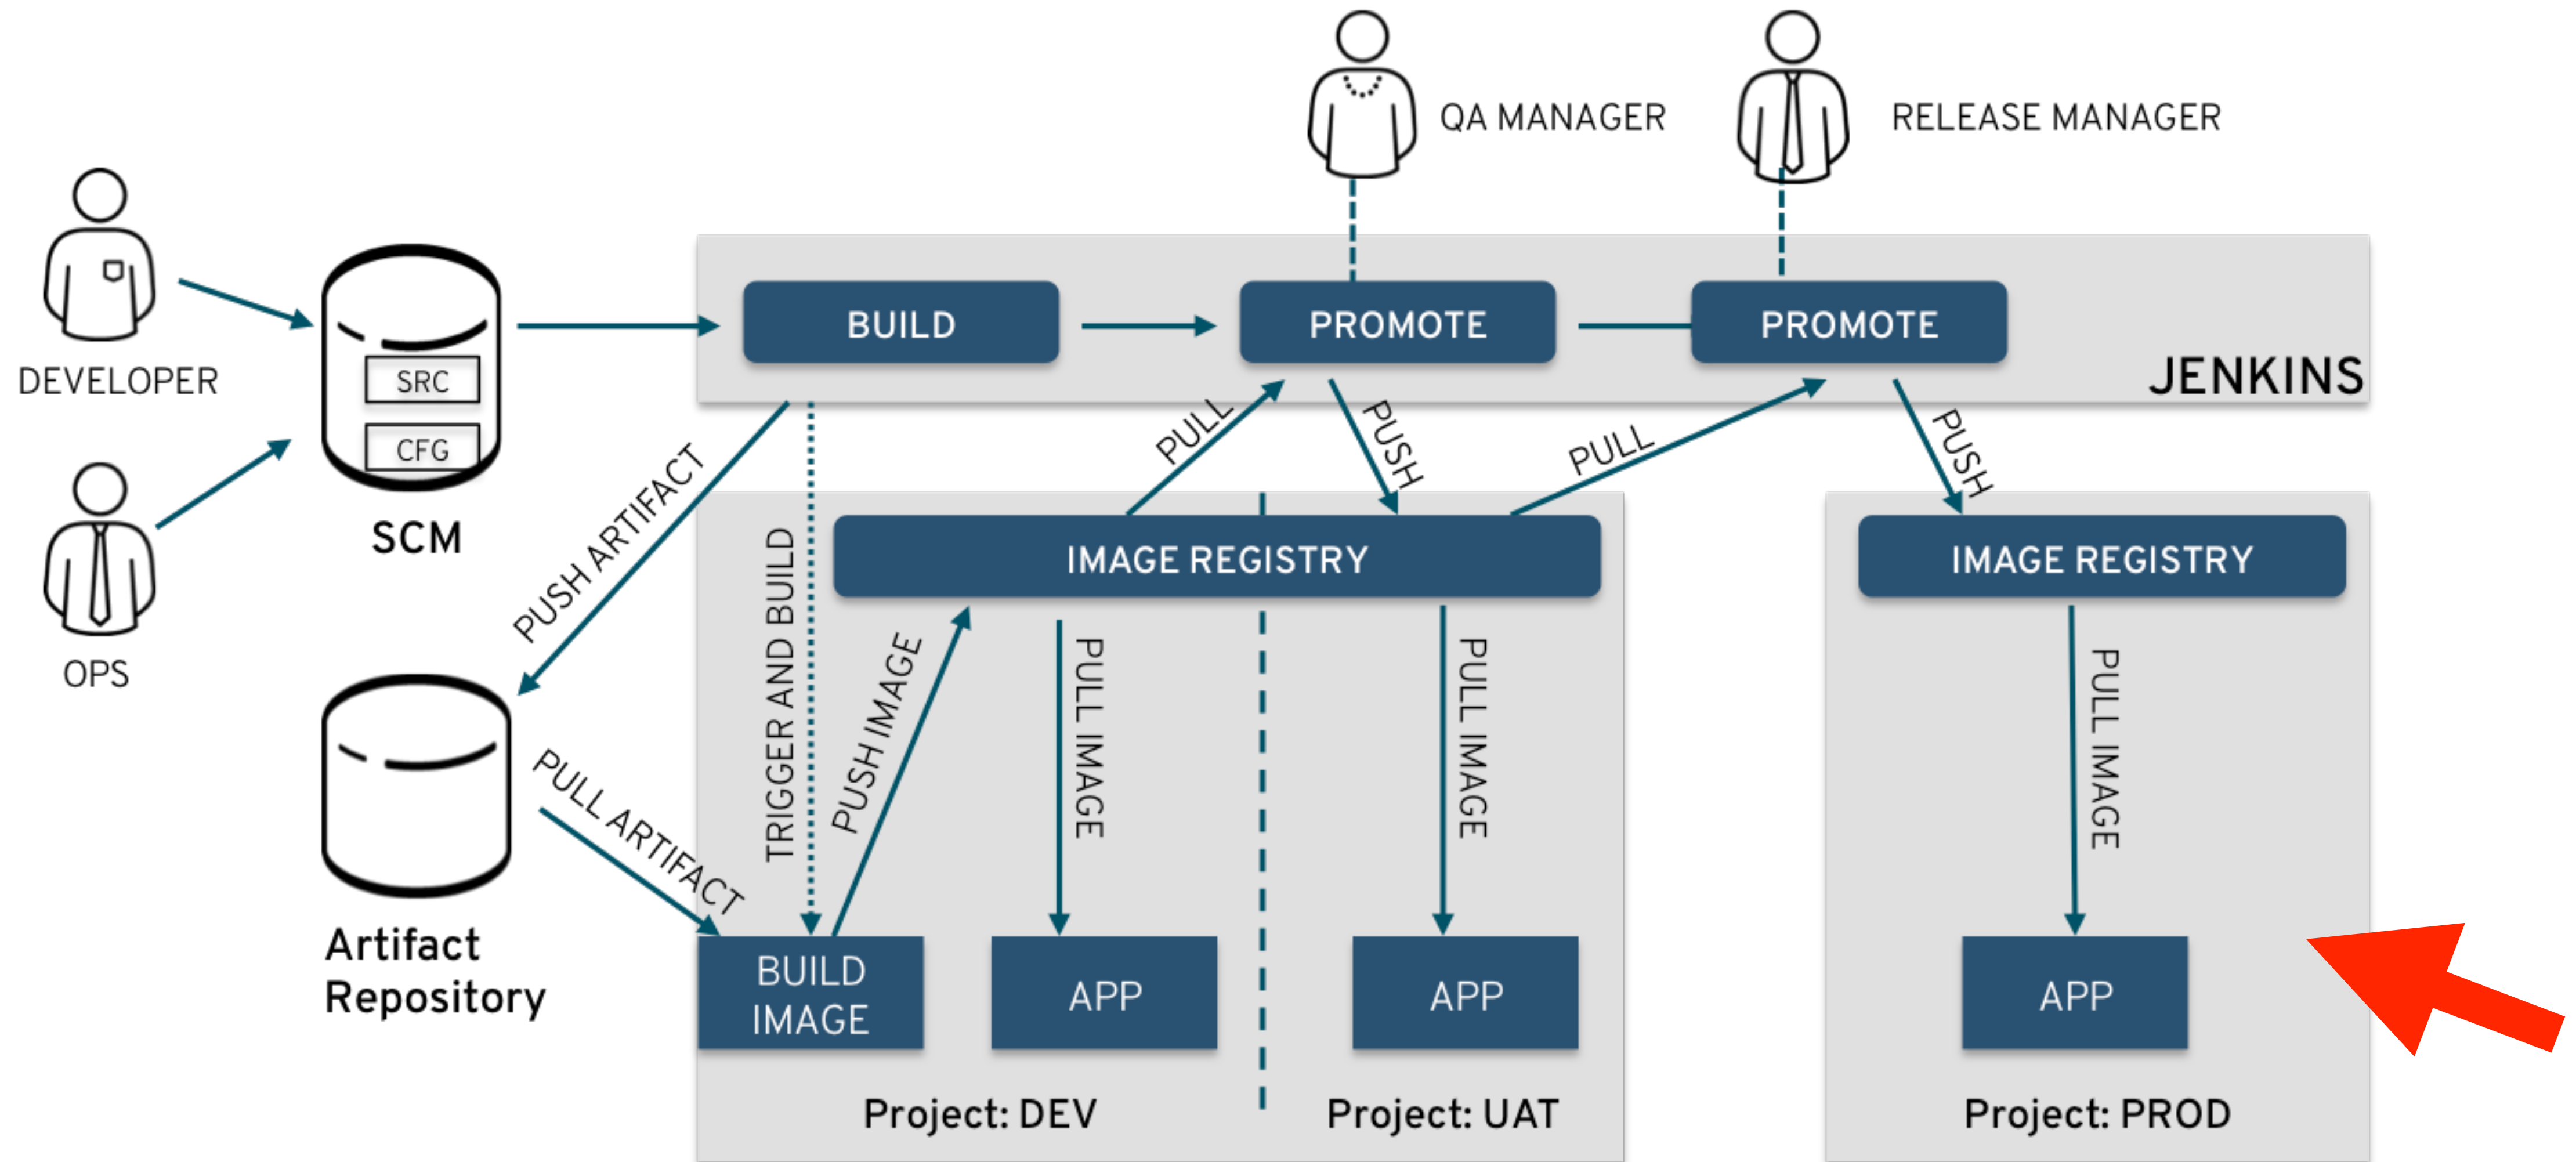
# CONTINUOUS DELIVERY WITH CONTAINERS

# CONTINUOUS DELIVERY WITH CONTAINERS





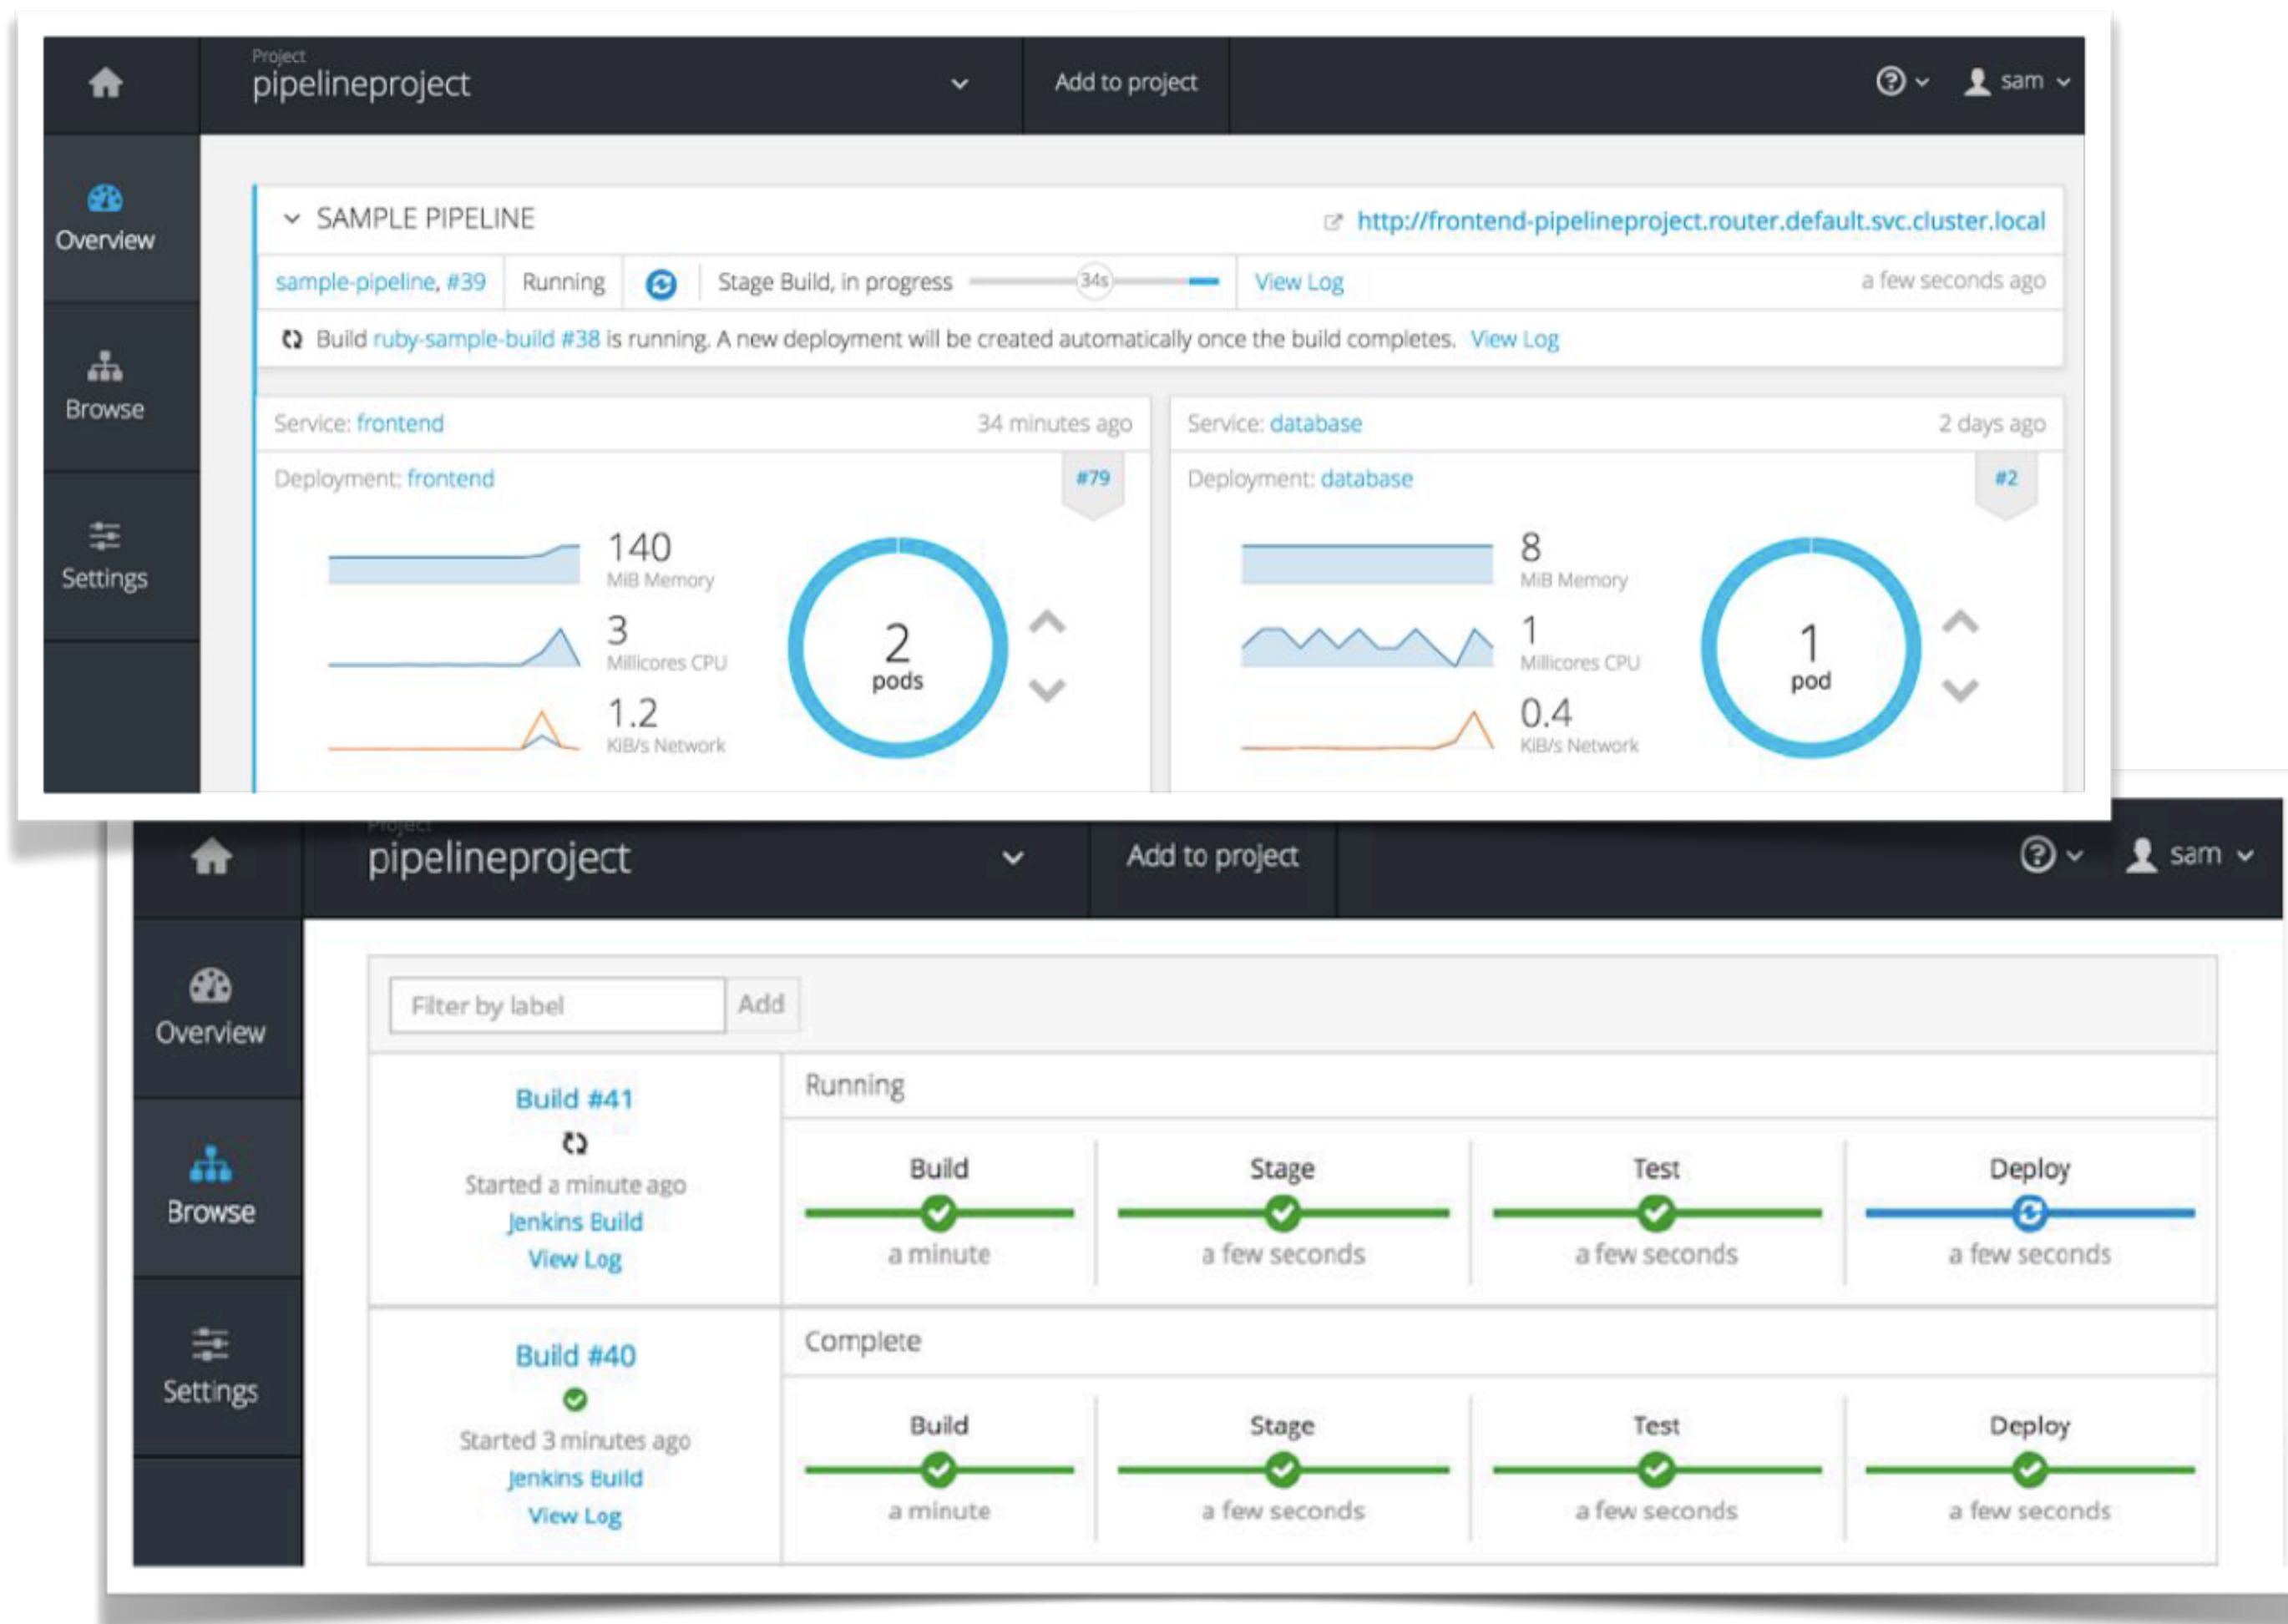
# CONTINUOUS DELIVERY + SECURITY





# CONTINUOUS DELIVERY: DEPLOYMENT STRATEGIES

# CONTINUOUS DELIVERY DEPLOYMENT STRATEGIES




## DEPLOYMENT STRATEGIES


- Recreate
- Rolling updates
- Blue / Green deployment

# Recreate

# RECREATE WITH DOWNTIME



# RECREATE WITH DOWNTIME



# RECREATE WITH DOWNTIME

## Use Case


- Non-mission critical services

## Cons

- Downtime

## Pros


- Simple, clean
- No Schema incompatibilities
- No API versioning



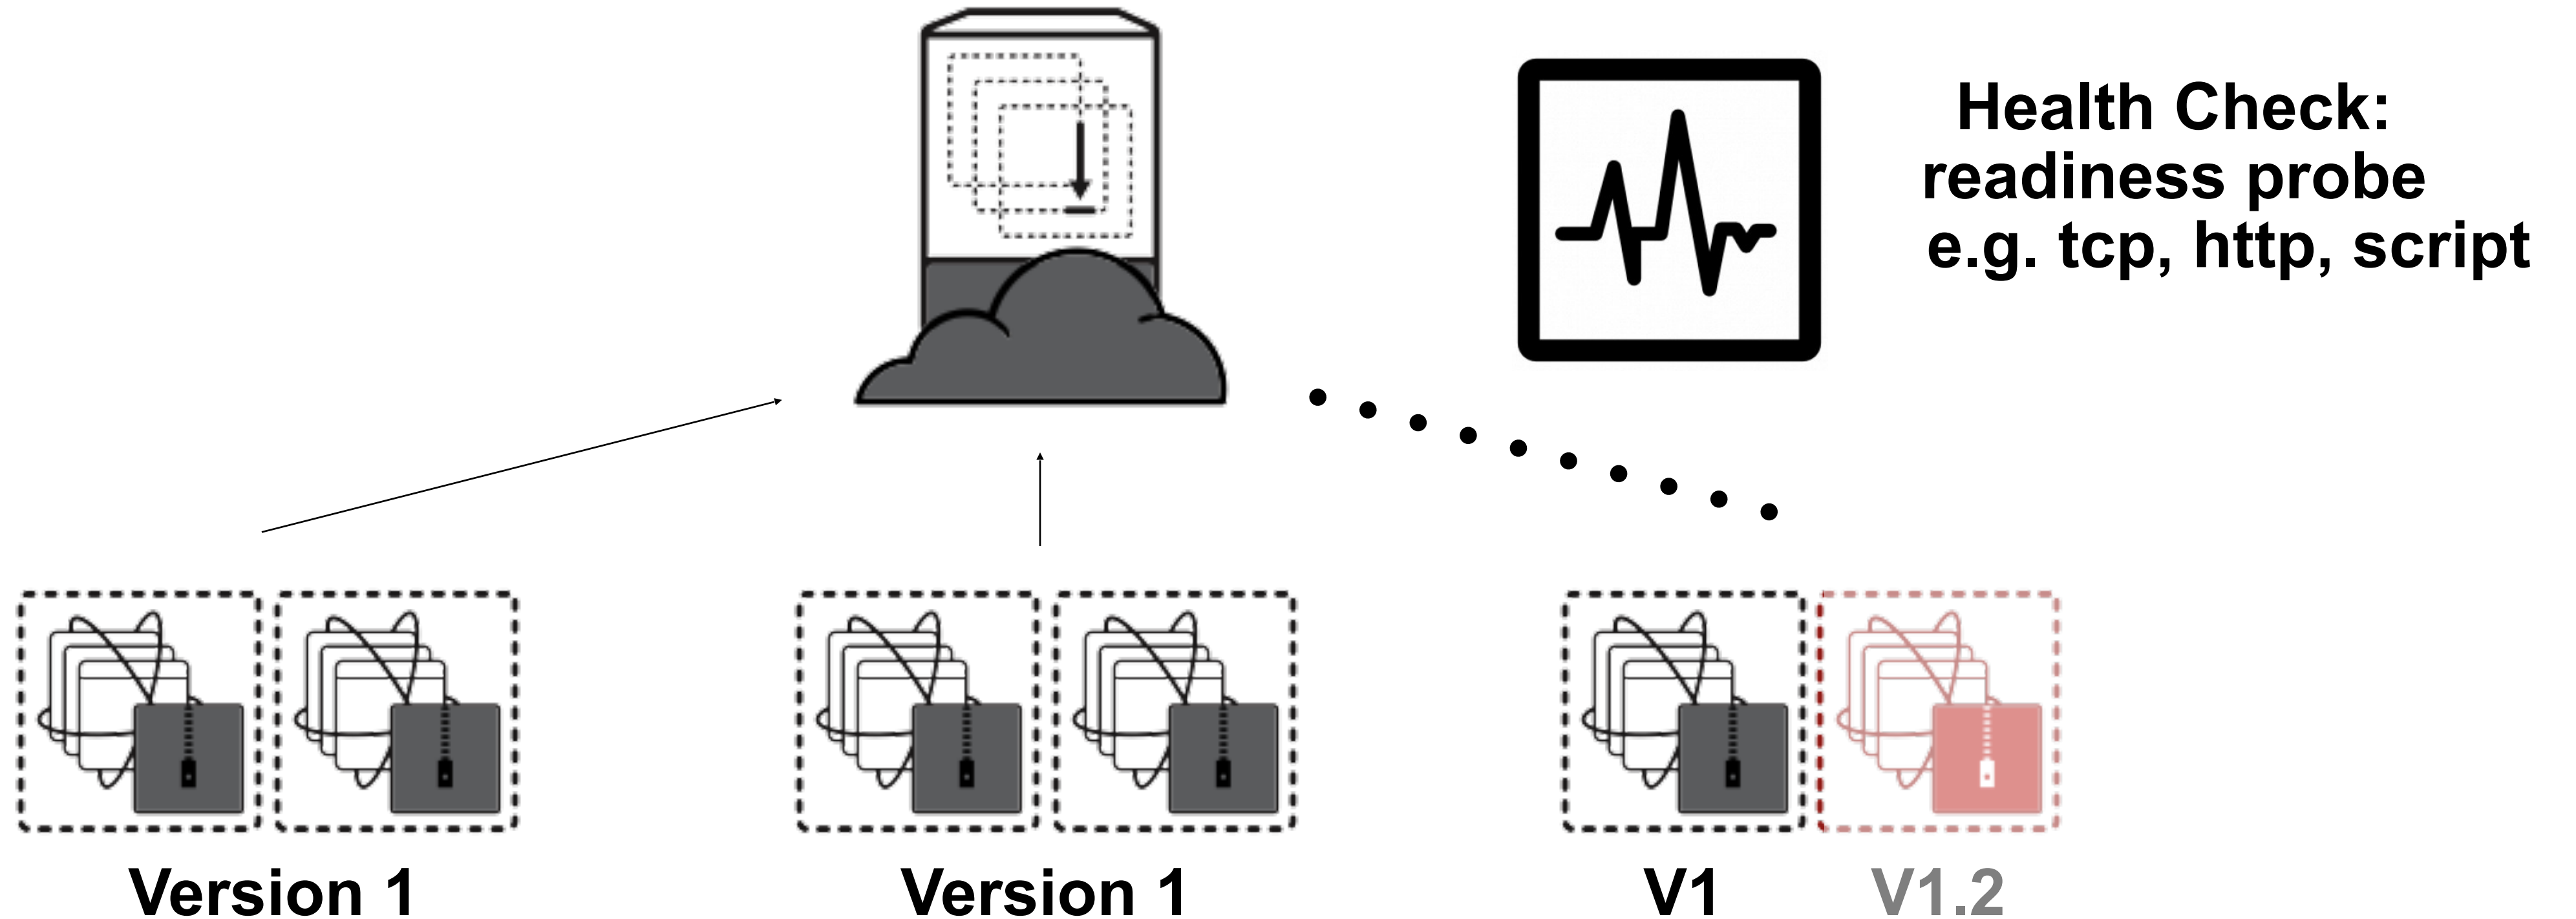


# Rolling Updates


# ROLLING UPDATES with ZERO DOWNTIME



# Deploy new version and wait until it's ready...



# Each container/pod is updated one by one



# Each container/pod is updated one by one

## Use Case


- Horizontally scaled
- Backward compatible API/data
- Microservices

## Cons

- Require backward compatible APIs/data
- Resource overhead

## Pros

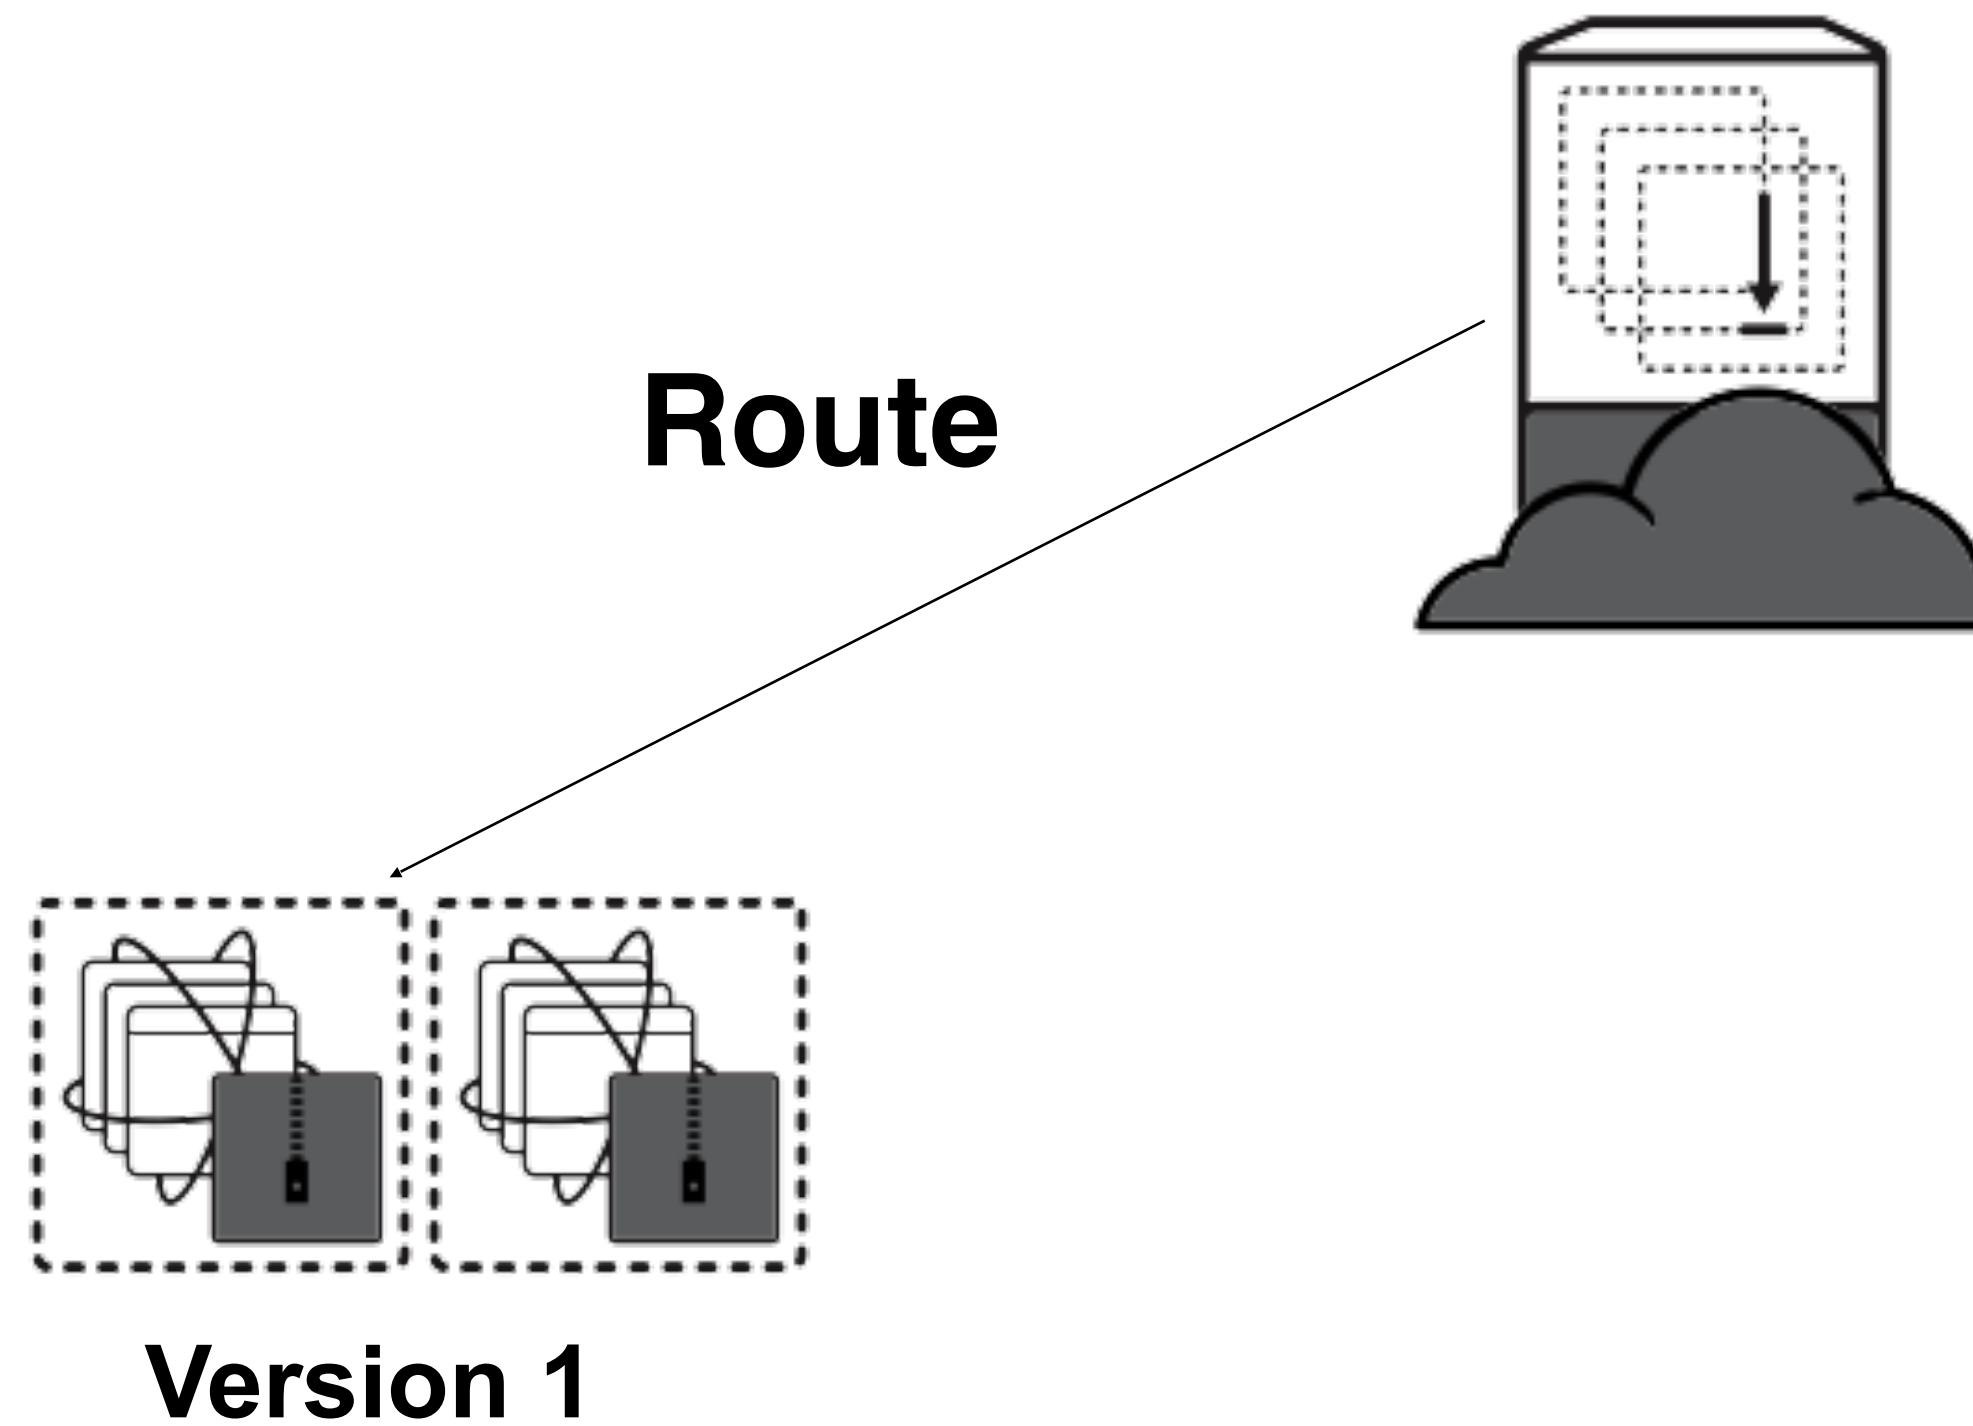
- Zero downtime
- Reduced risk, gradual rollout w/health checks
- Ready for rollback



# Blue / Green Deployment




# BLUE / GREEN DEPLOYMENT




**BLUE**


# BLUE / GREEN DEPLOYMENT



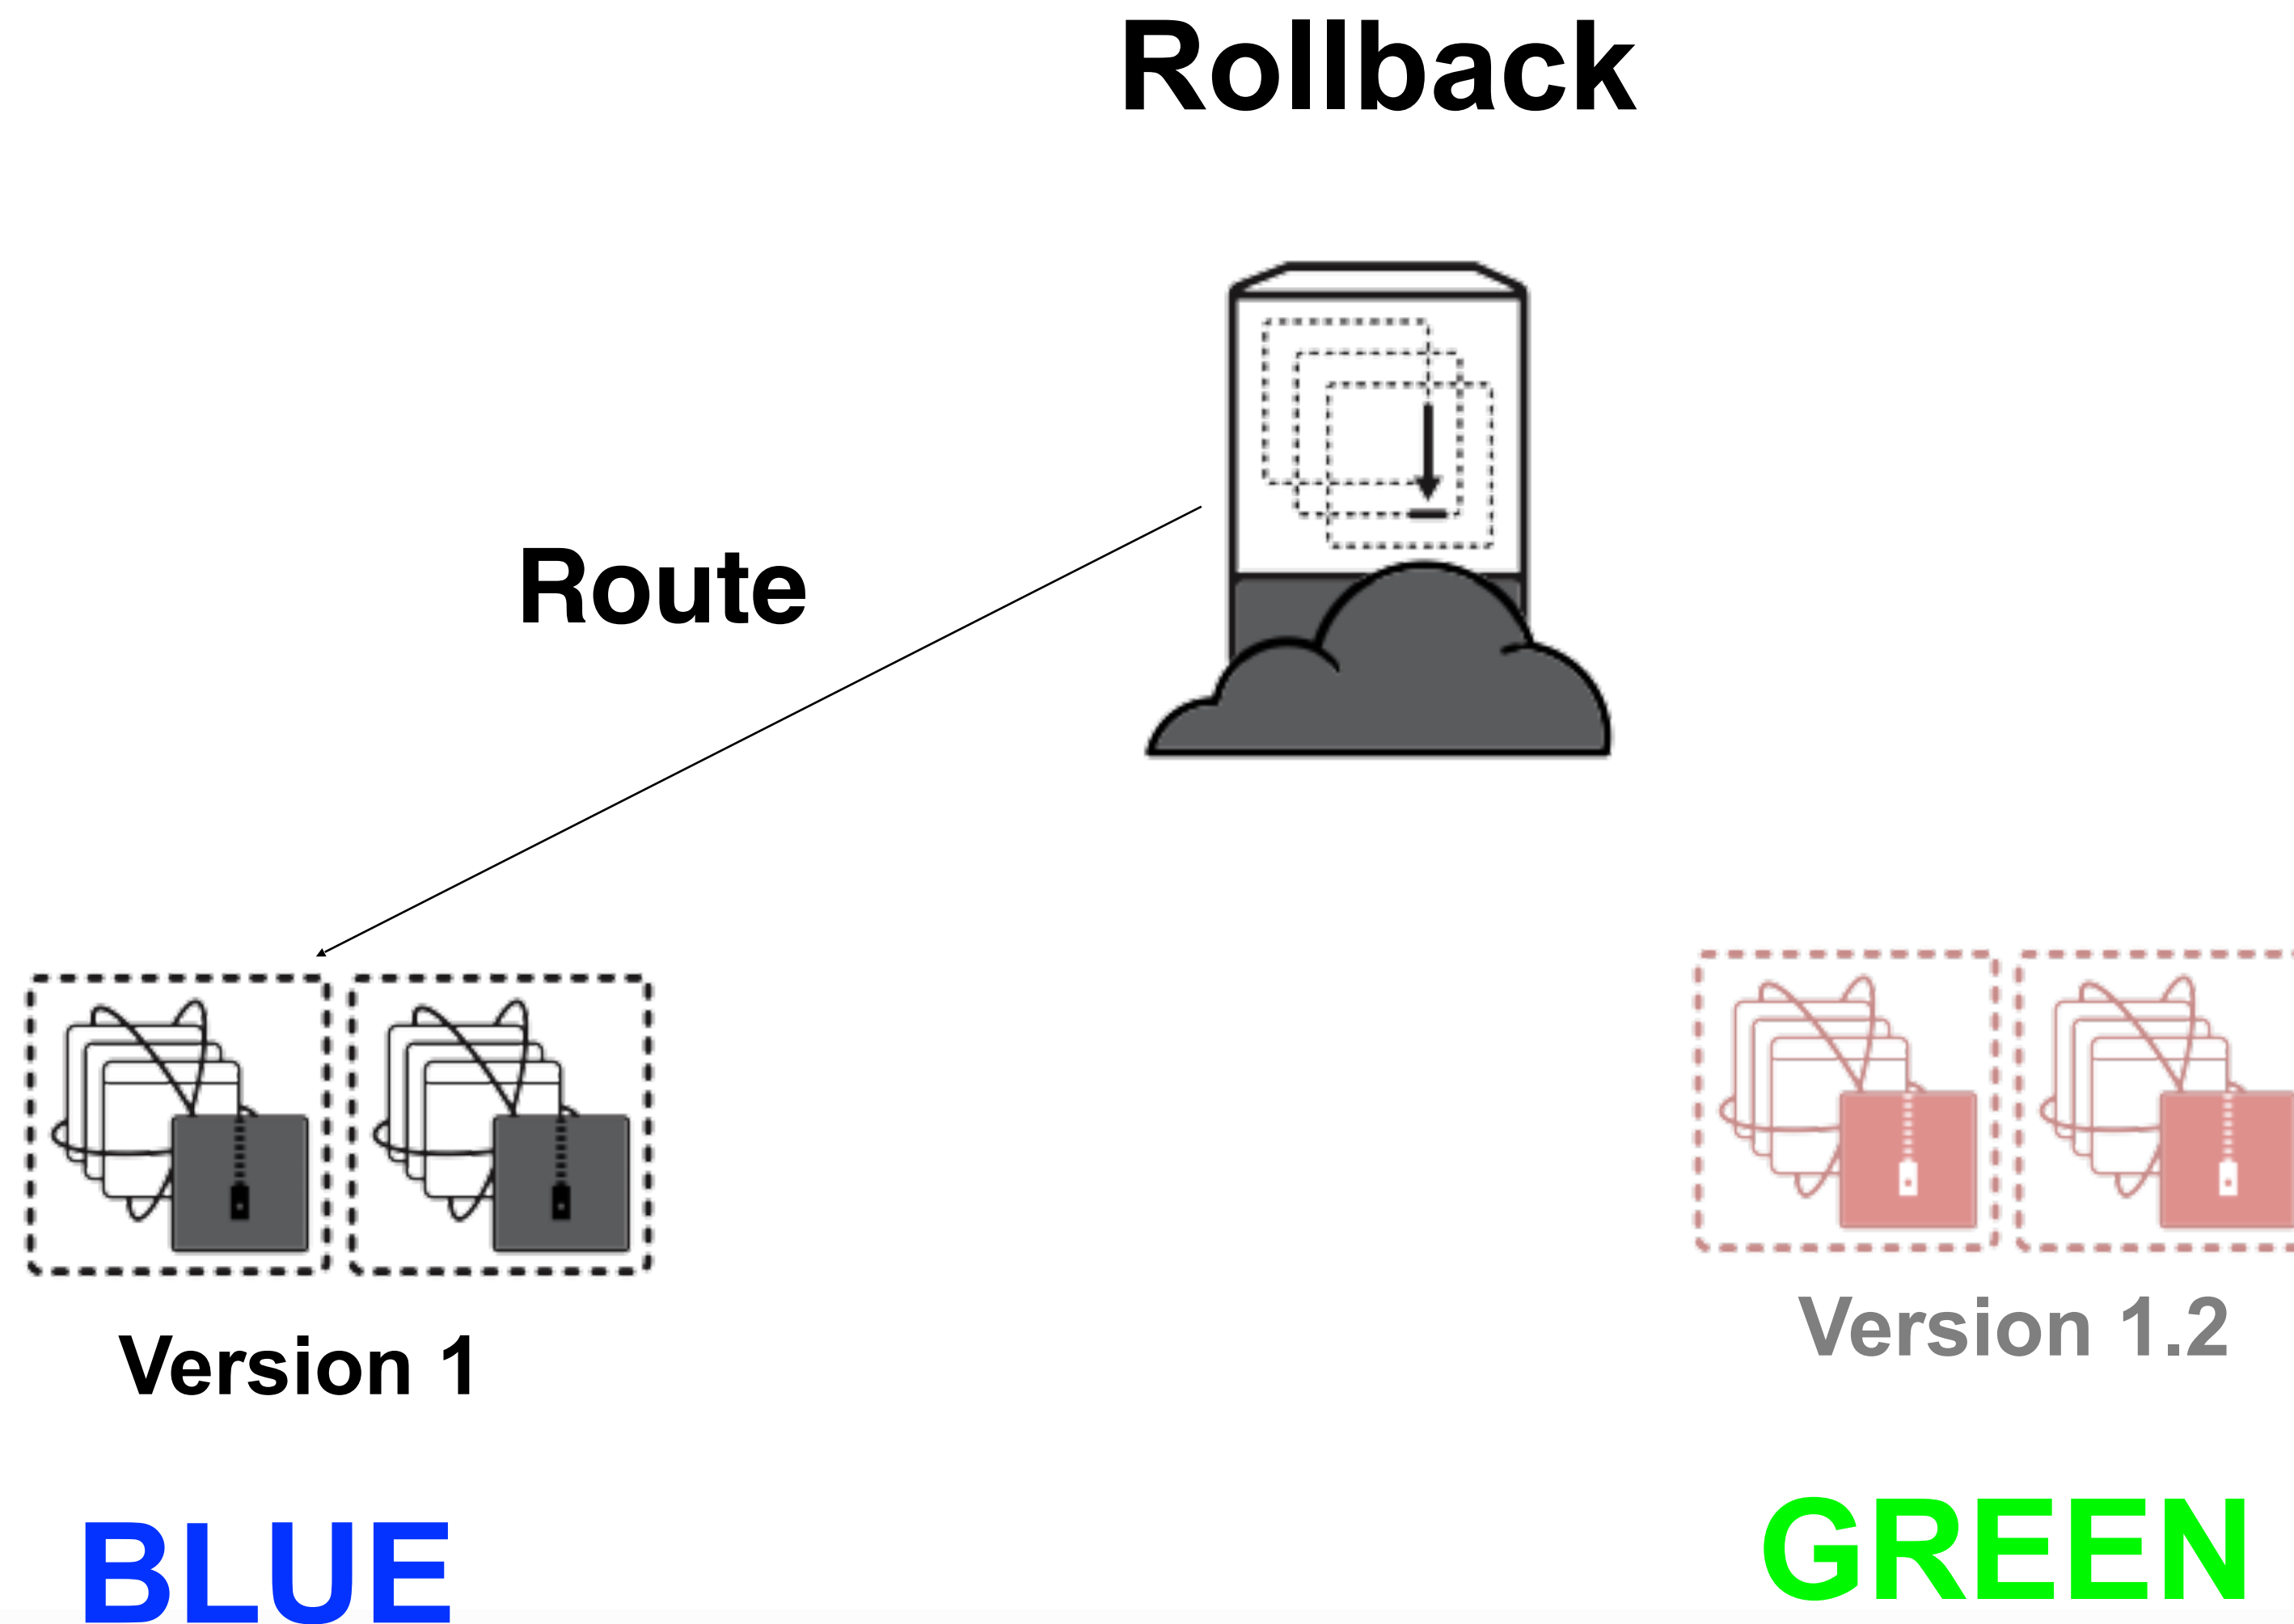
# BLUE / GREEN DEPLOYMENT



# BLUE / GREEN DEPLOYMENT



# BLUE / GREEN DEPLOYMENT



## Use Case

- Self-contained micro services (data)

## Cons

- Resource overhead
- Data synchronization

## Pros

- Low risk, never change production
- No downtime
- Production like testing
- Rollback

# RAPID INNOVATION & EXPERIMENTATION



# MICROSERVICES

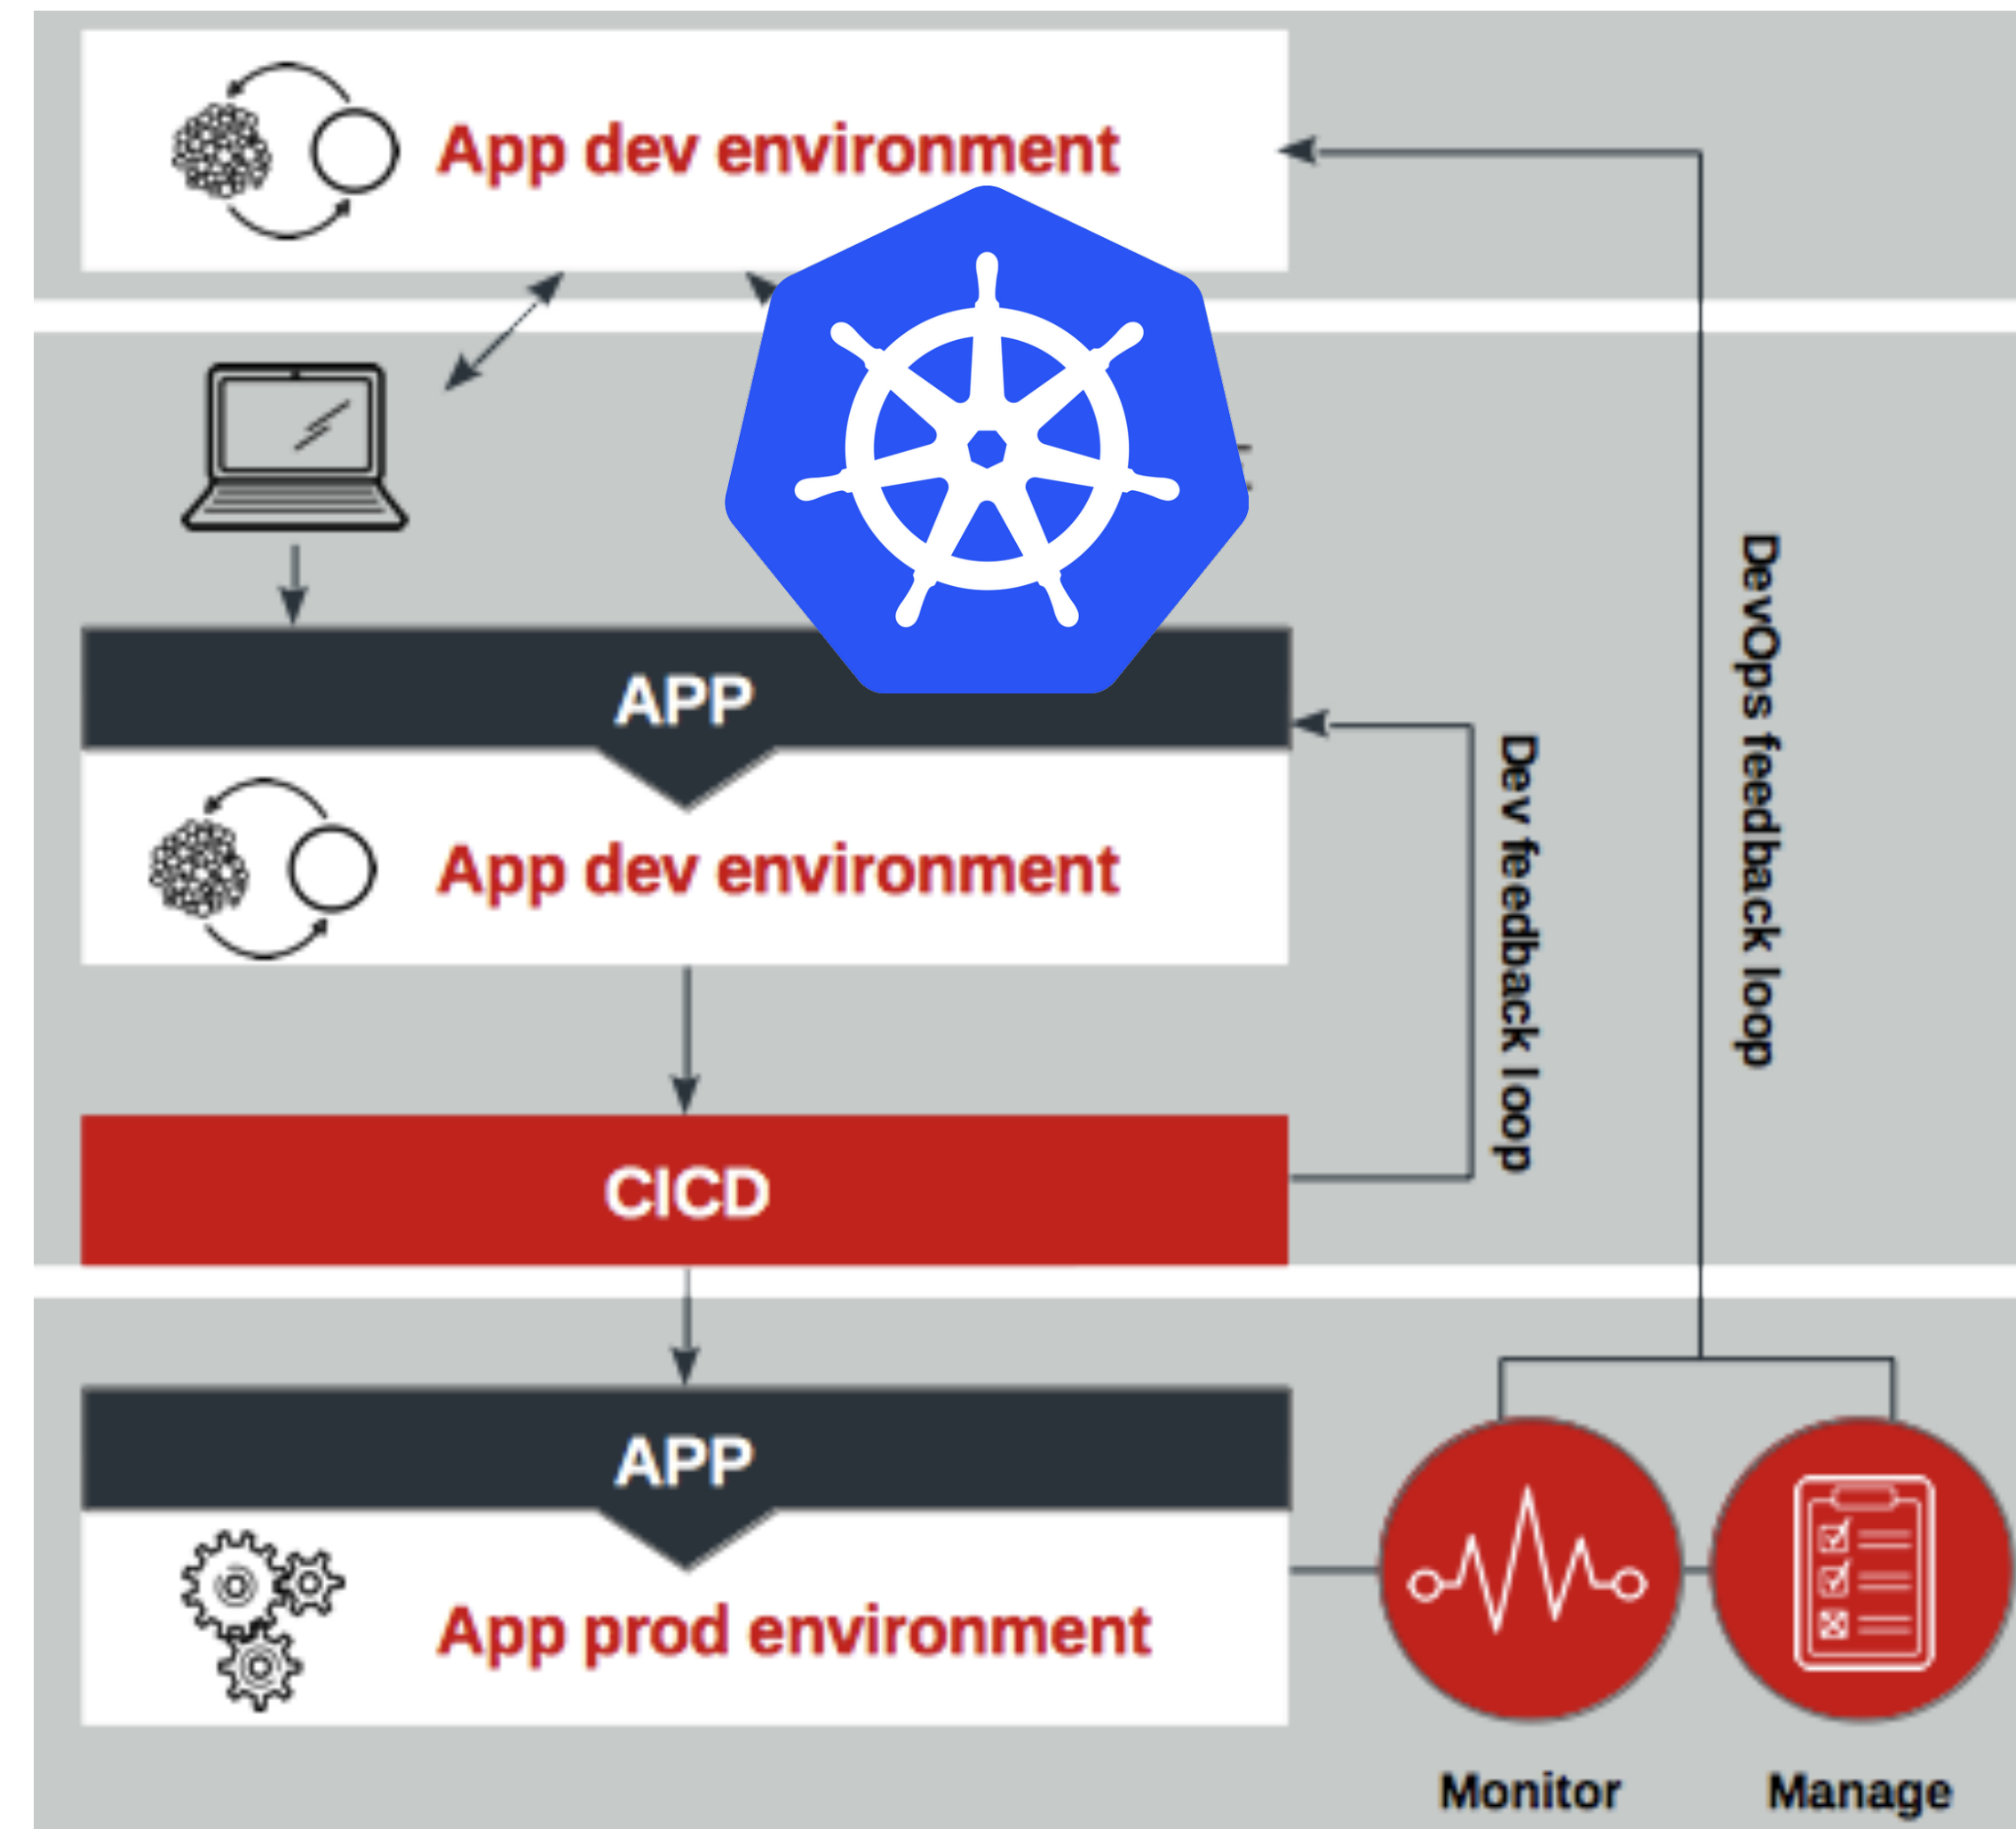
## RAPID INNOVATION & EXPERIMENTATION



”only about 1/3 of ideas improve the metrics they were designed to improve.”

Ronny Kohavi, Microsoft (Amazon)

# CONTINUOUS FEEDBACK LOOP



# A/B TESTING USING CANARY DEPLOYMENTS

A




B




# CANARY DEPLOYMENTS

100%



Route




Version 1

25% Conversion Rate



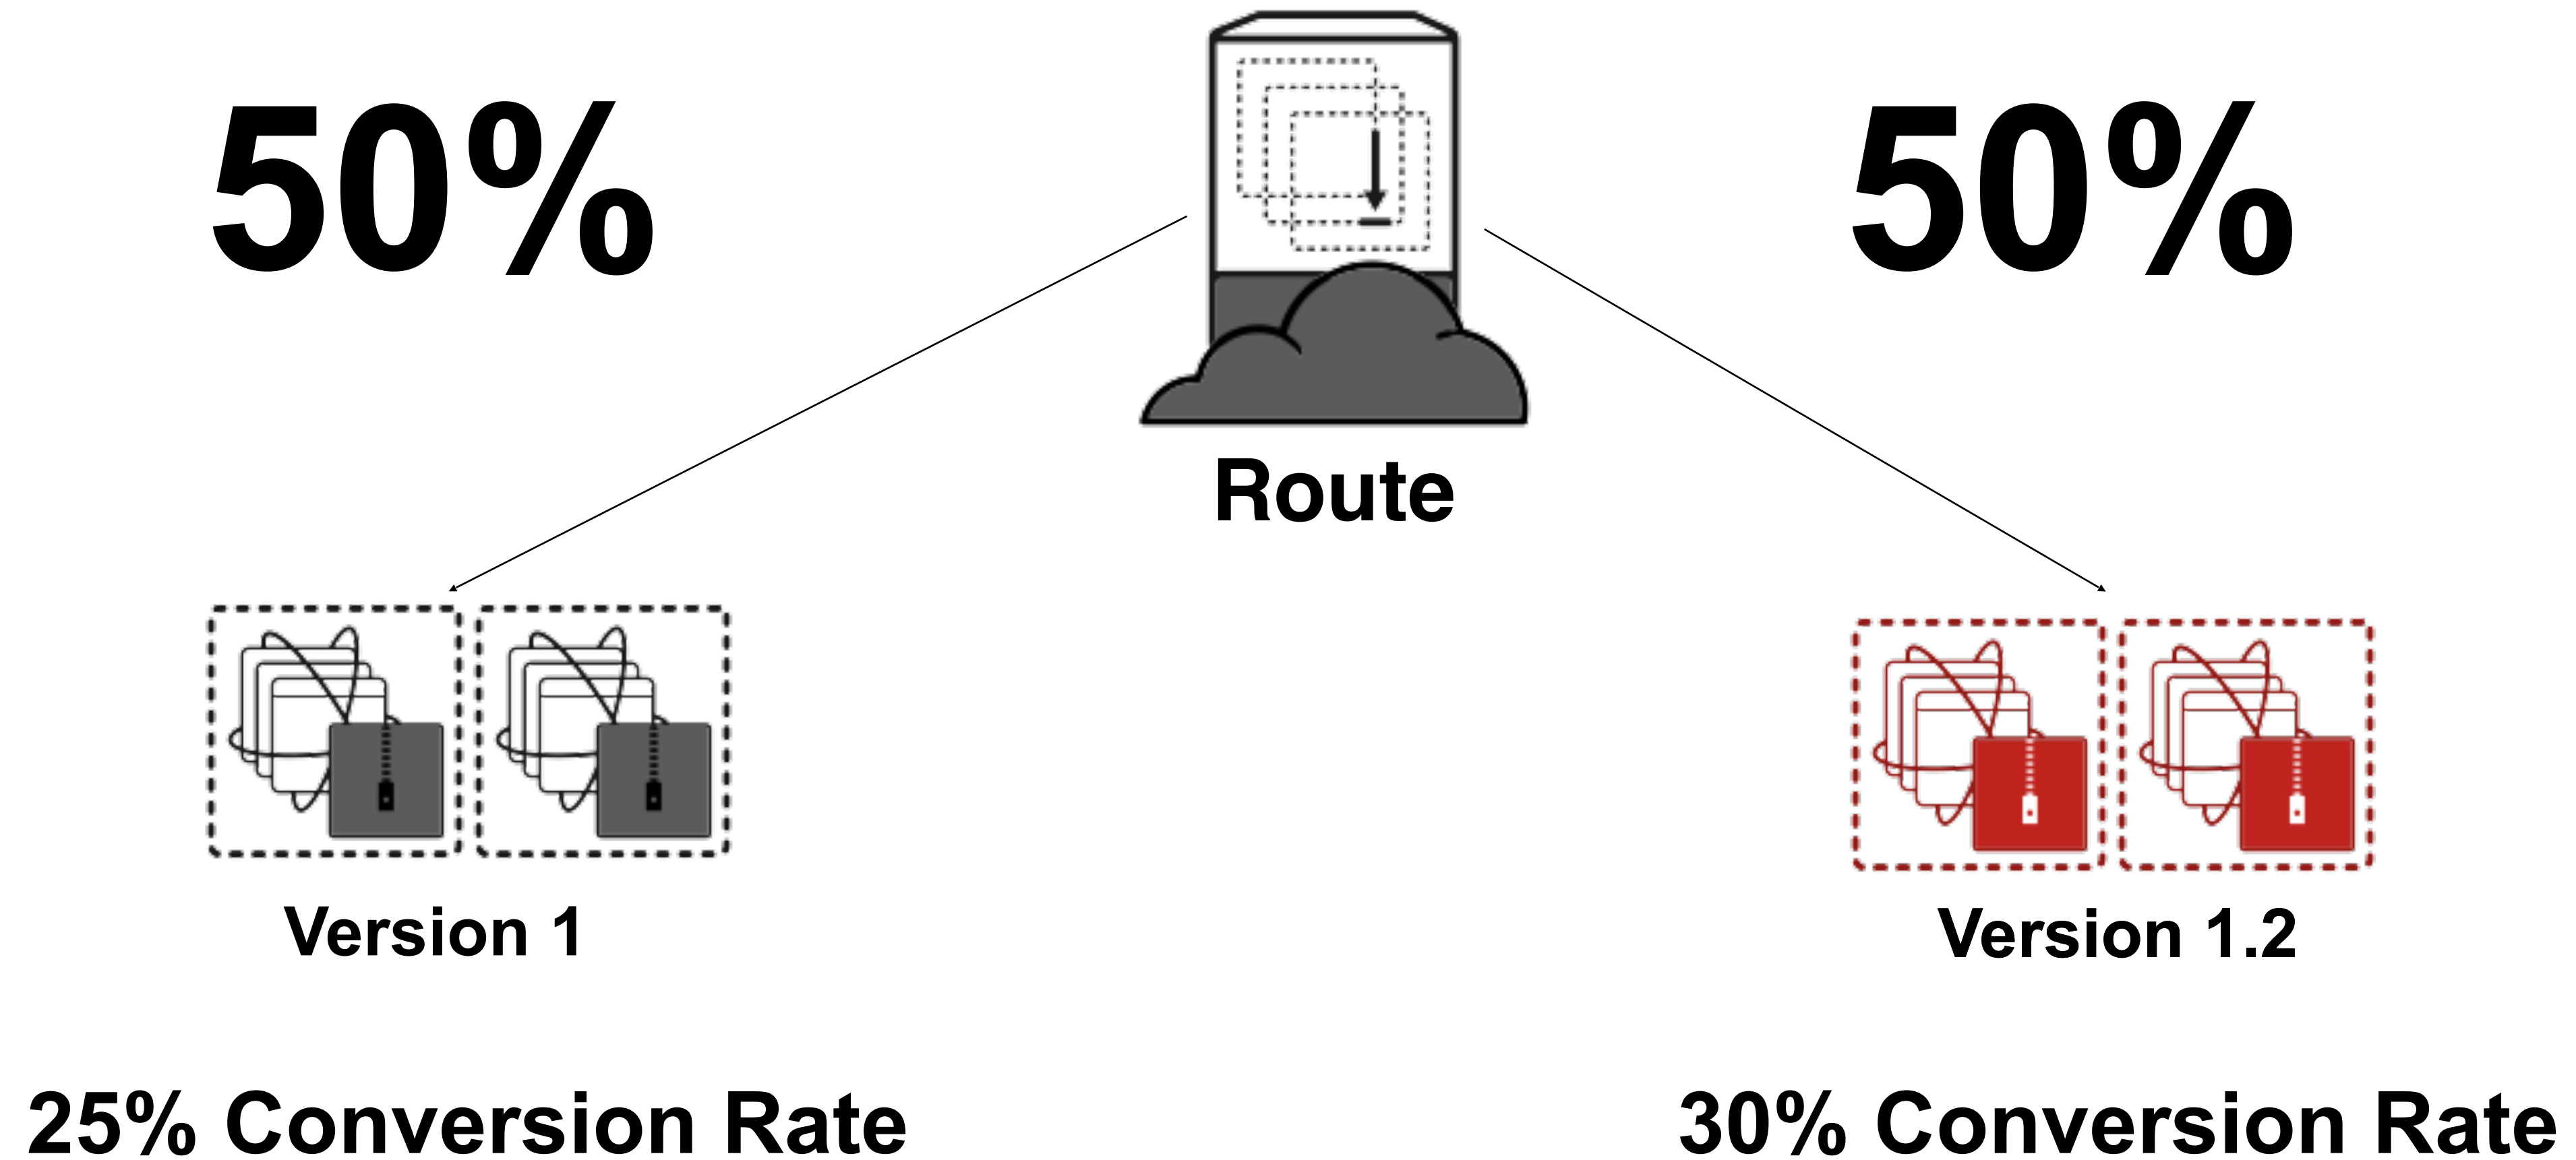
Tests / CI



Version 1.2

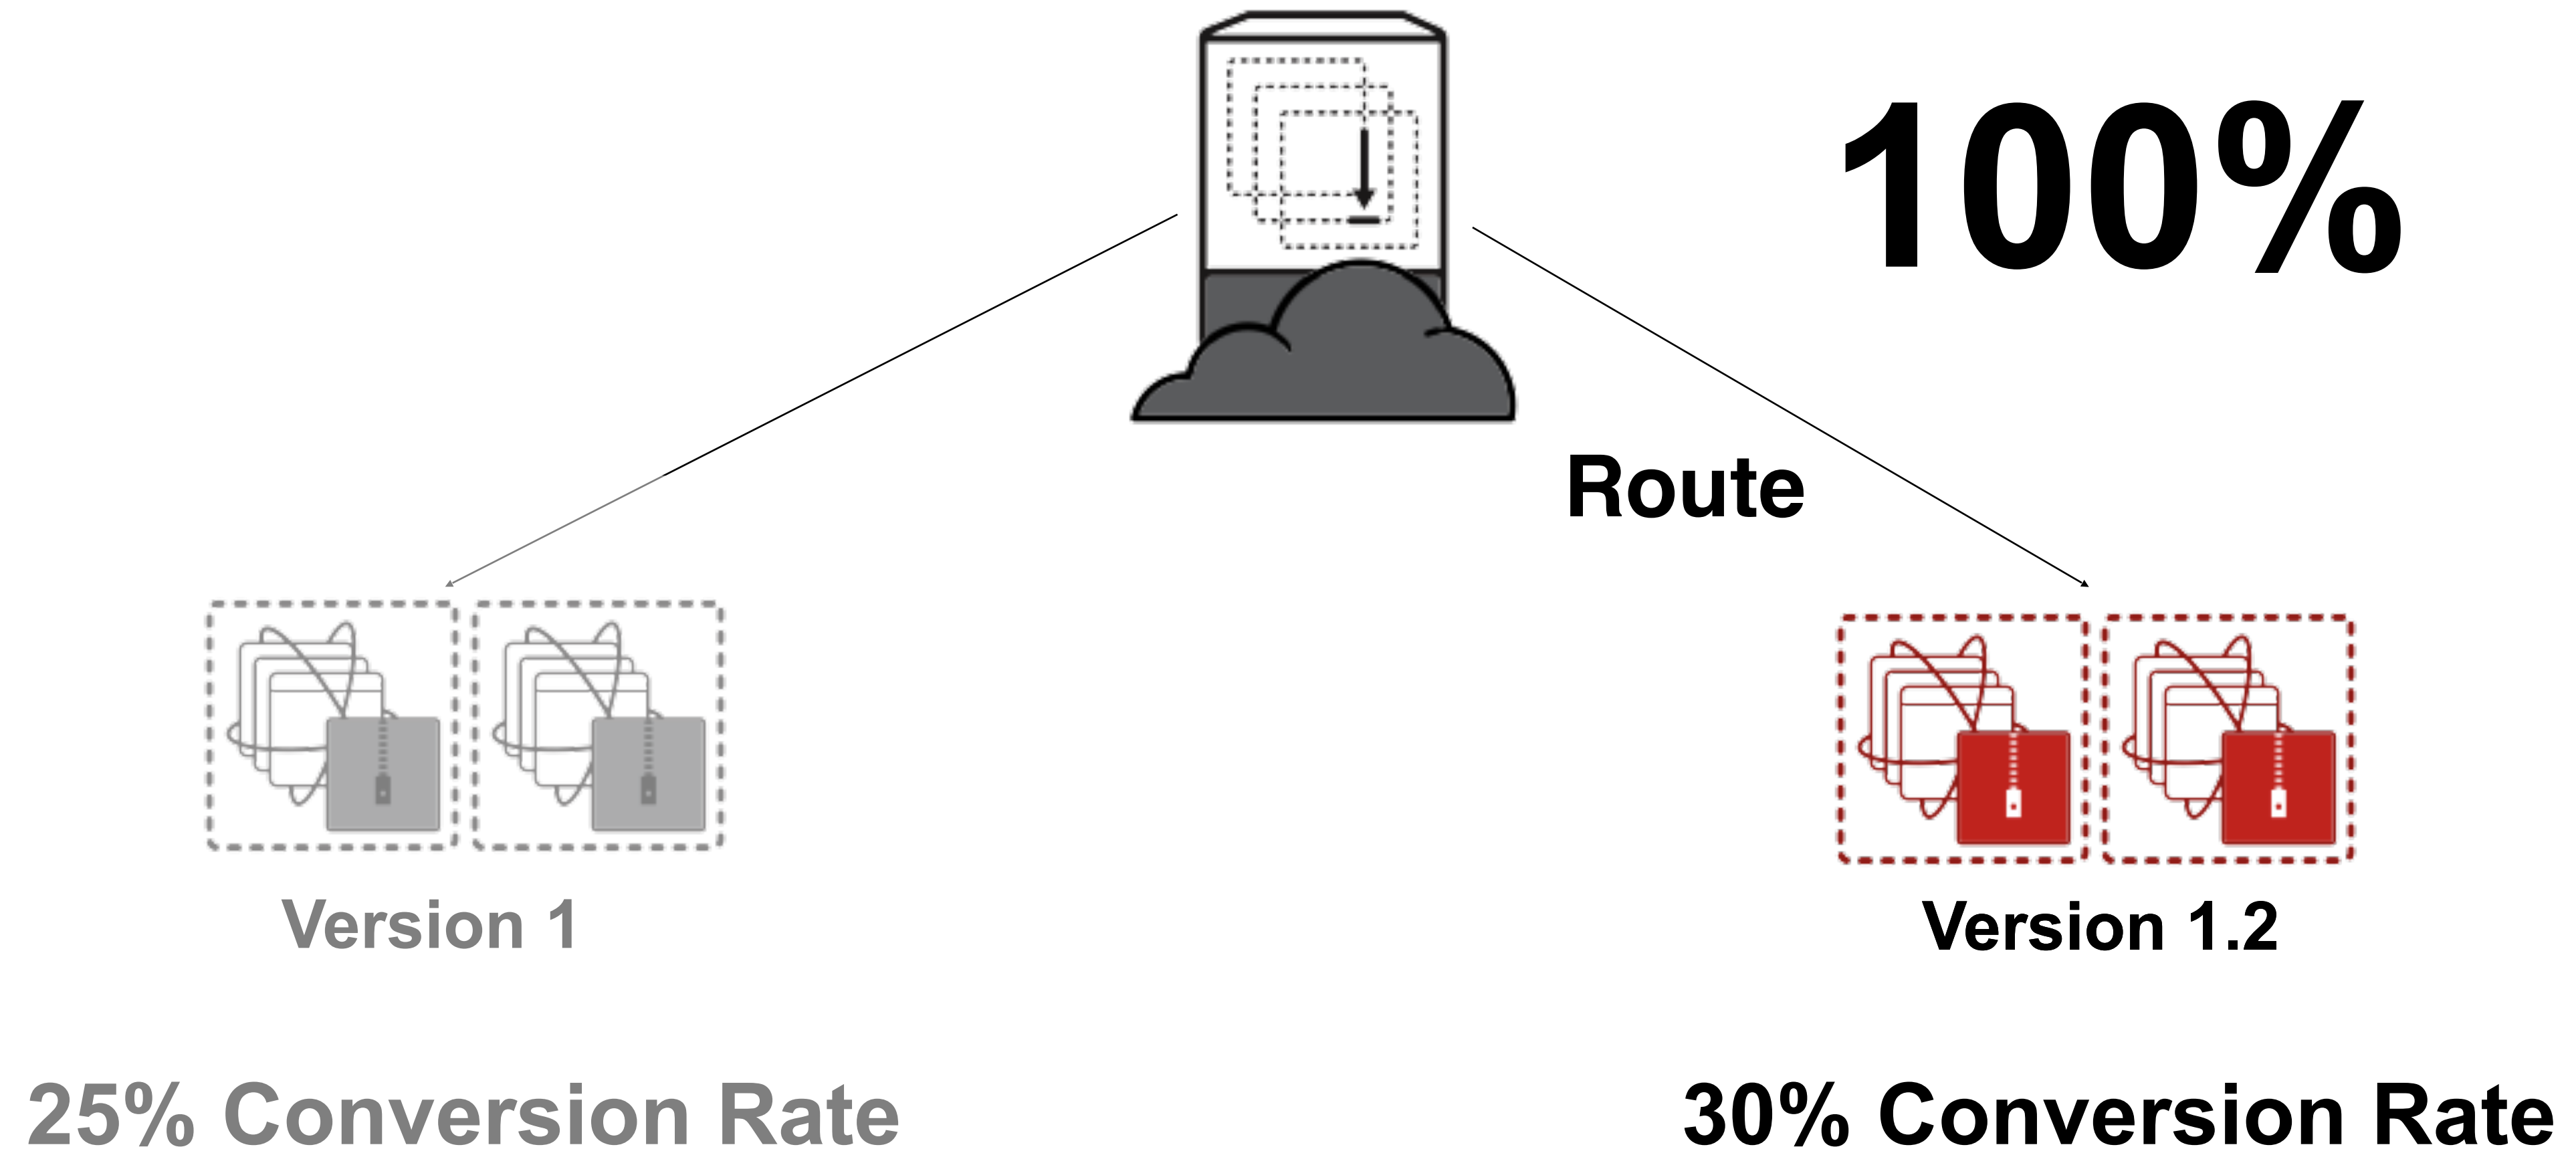
?! Conversion Rate

# CANARY DEPLOYMENTS



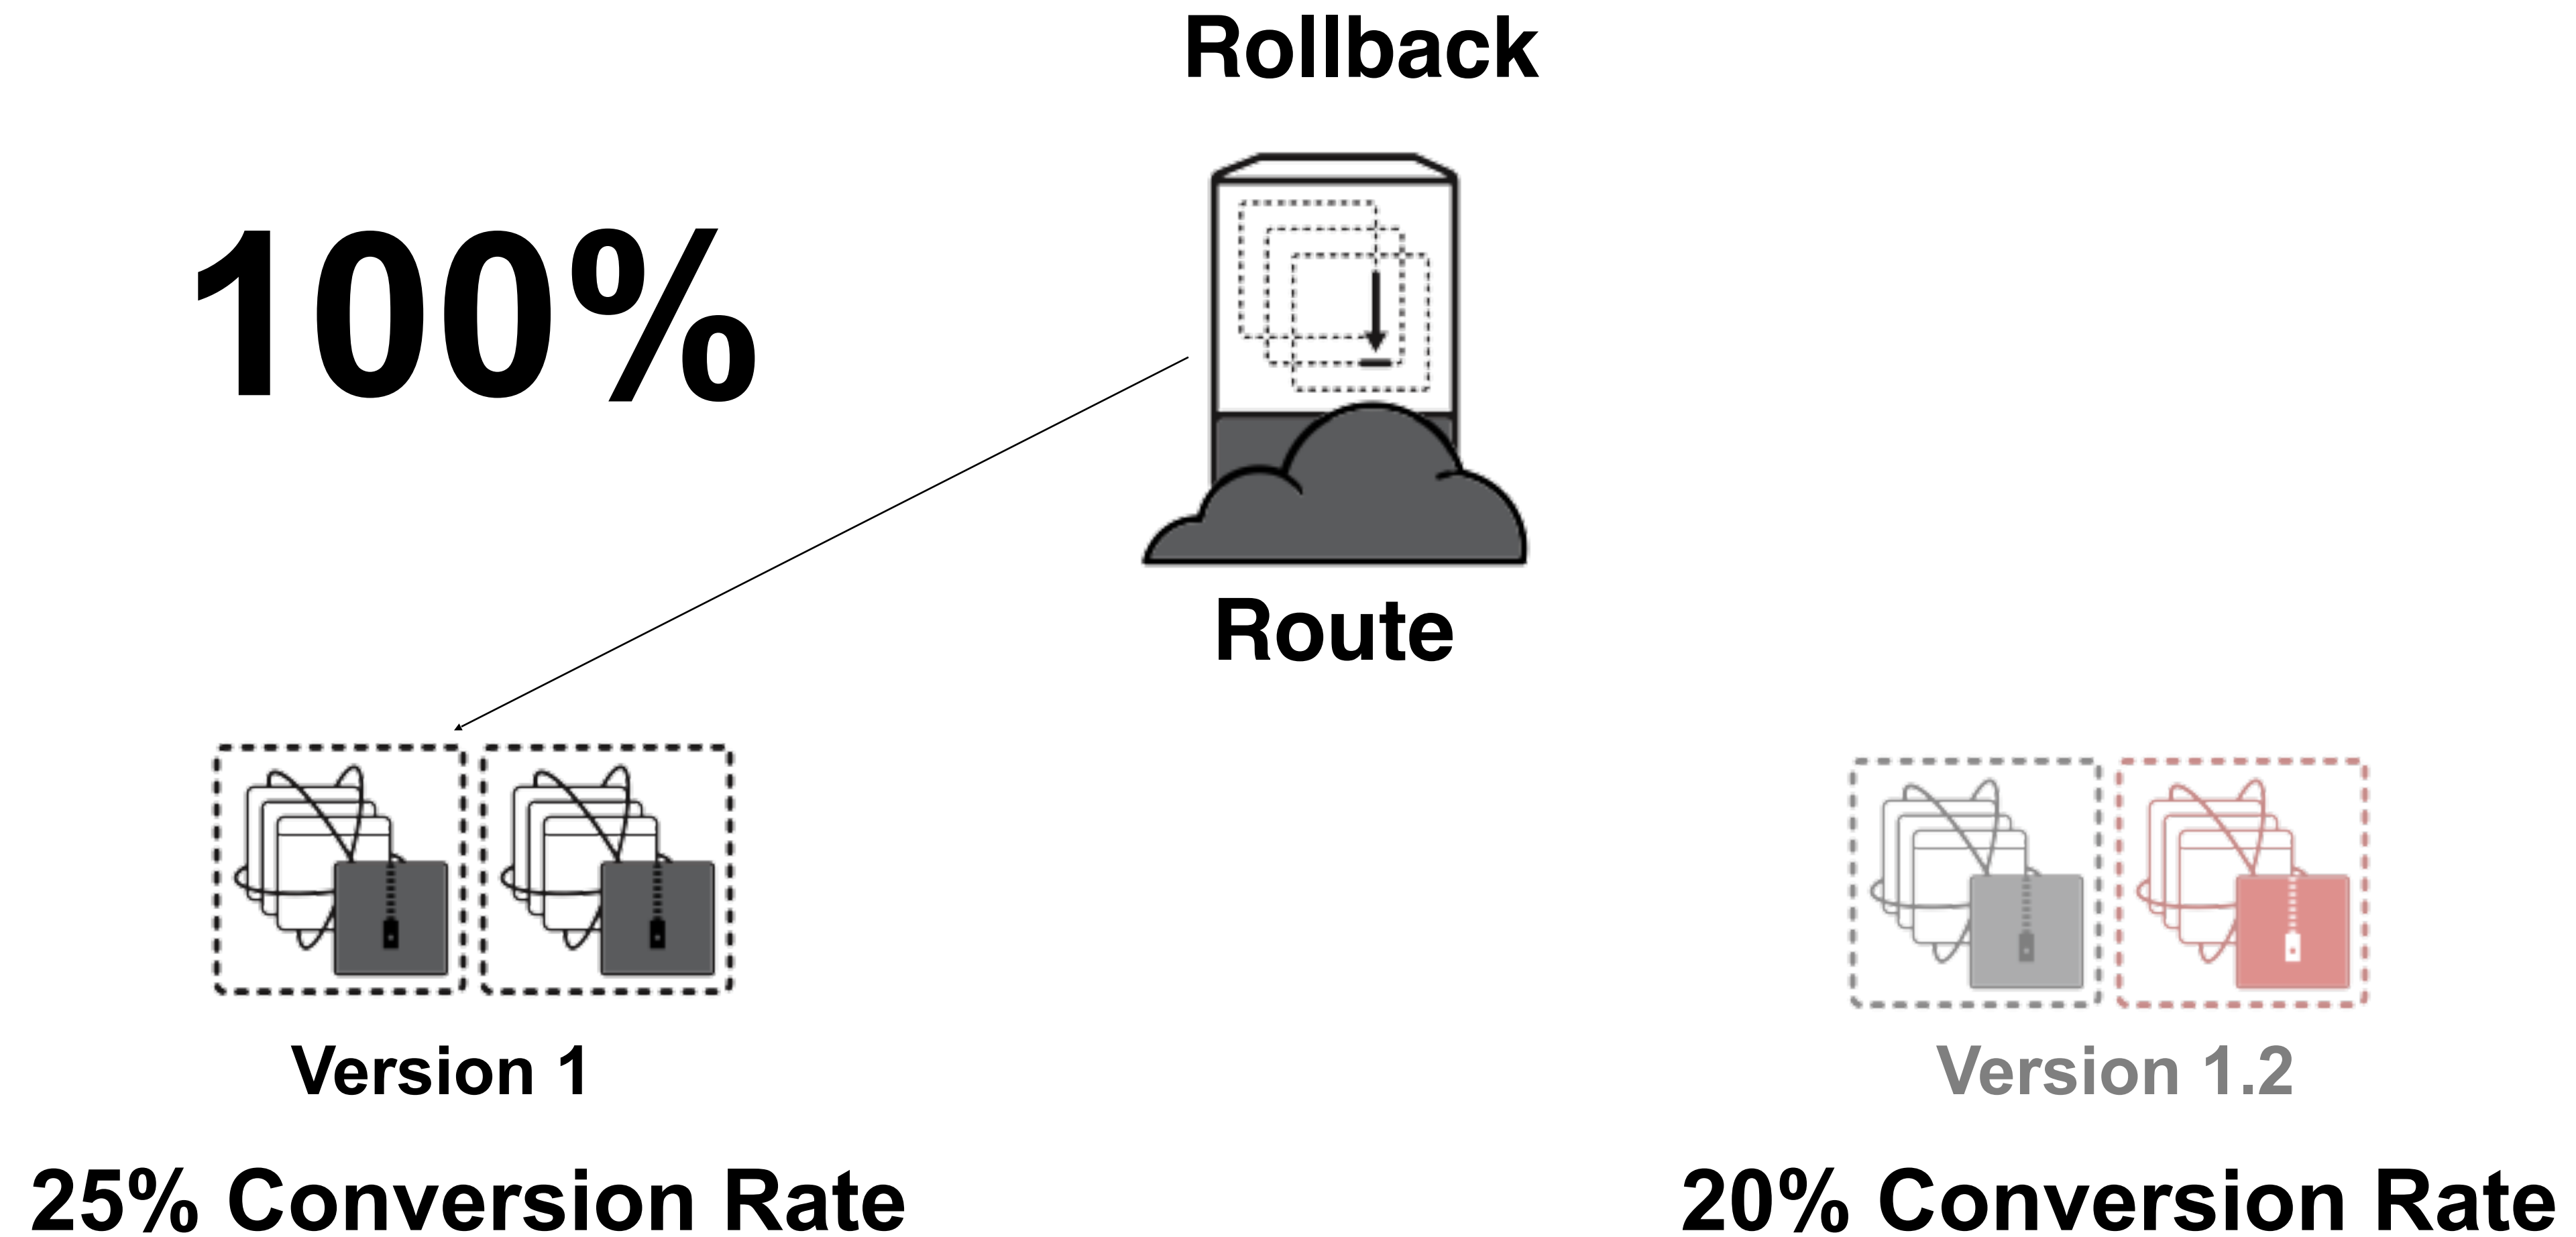


# CANARY DEPLOYMENTS



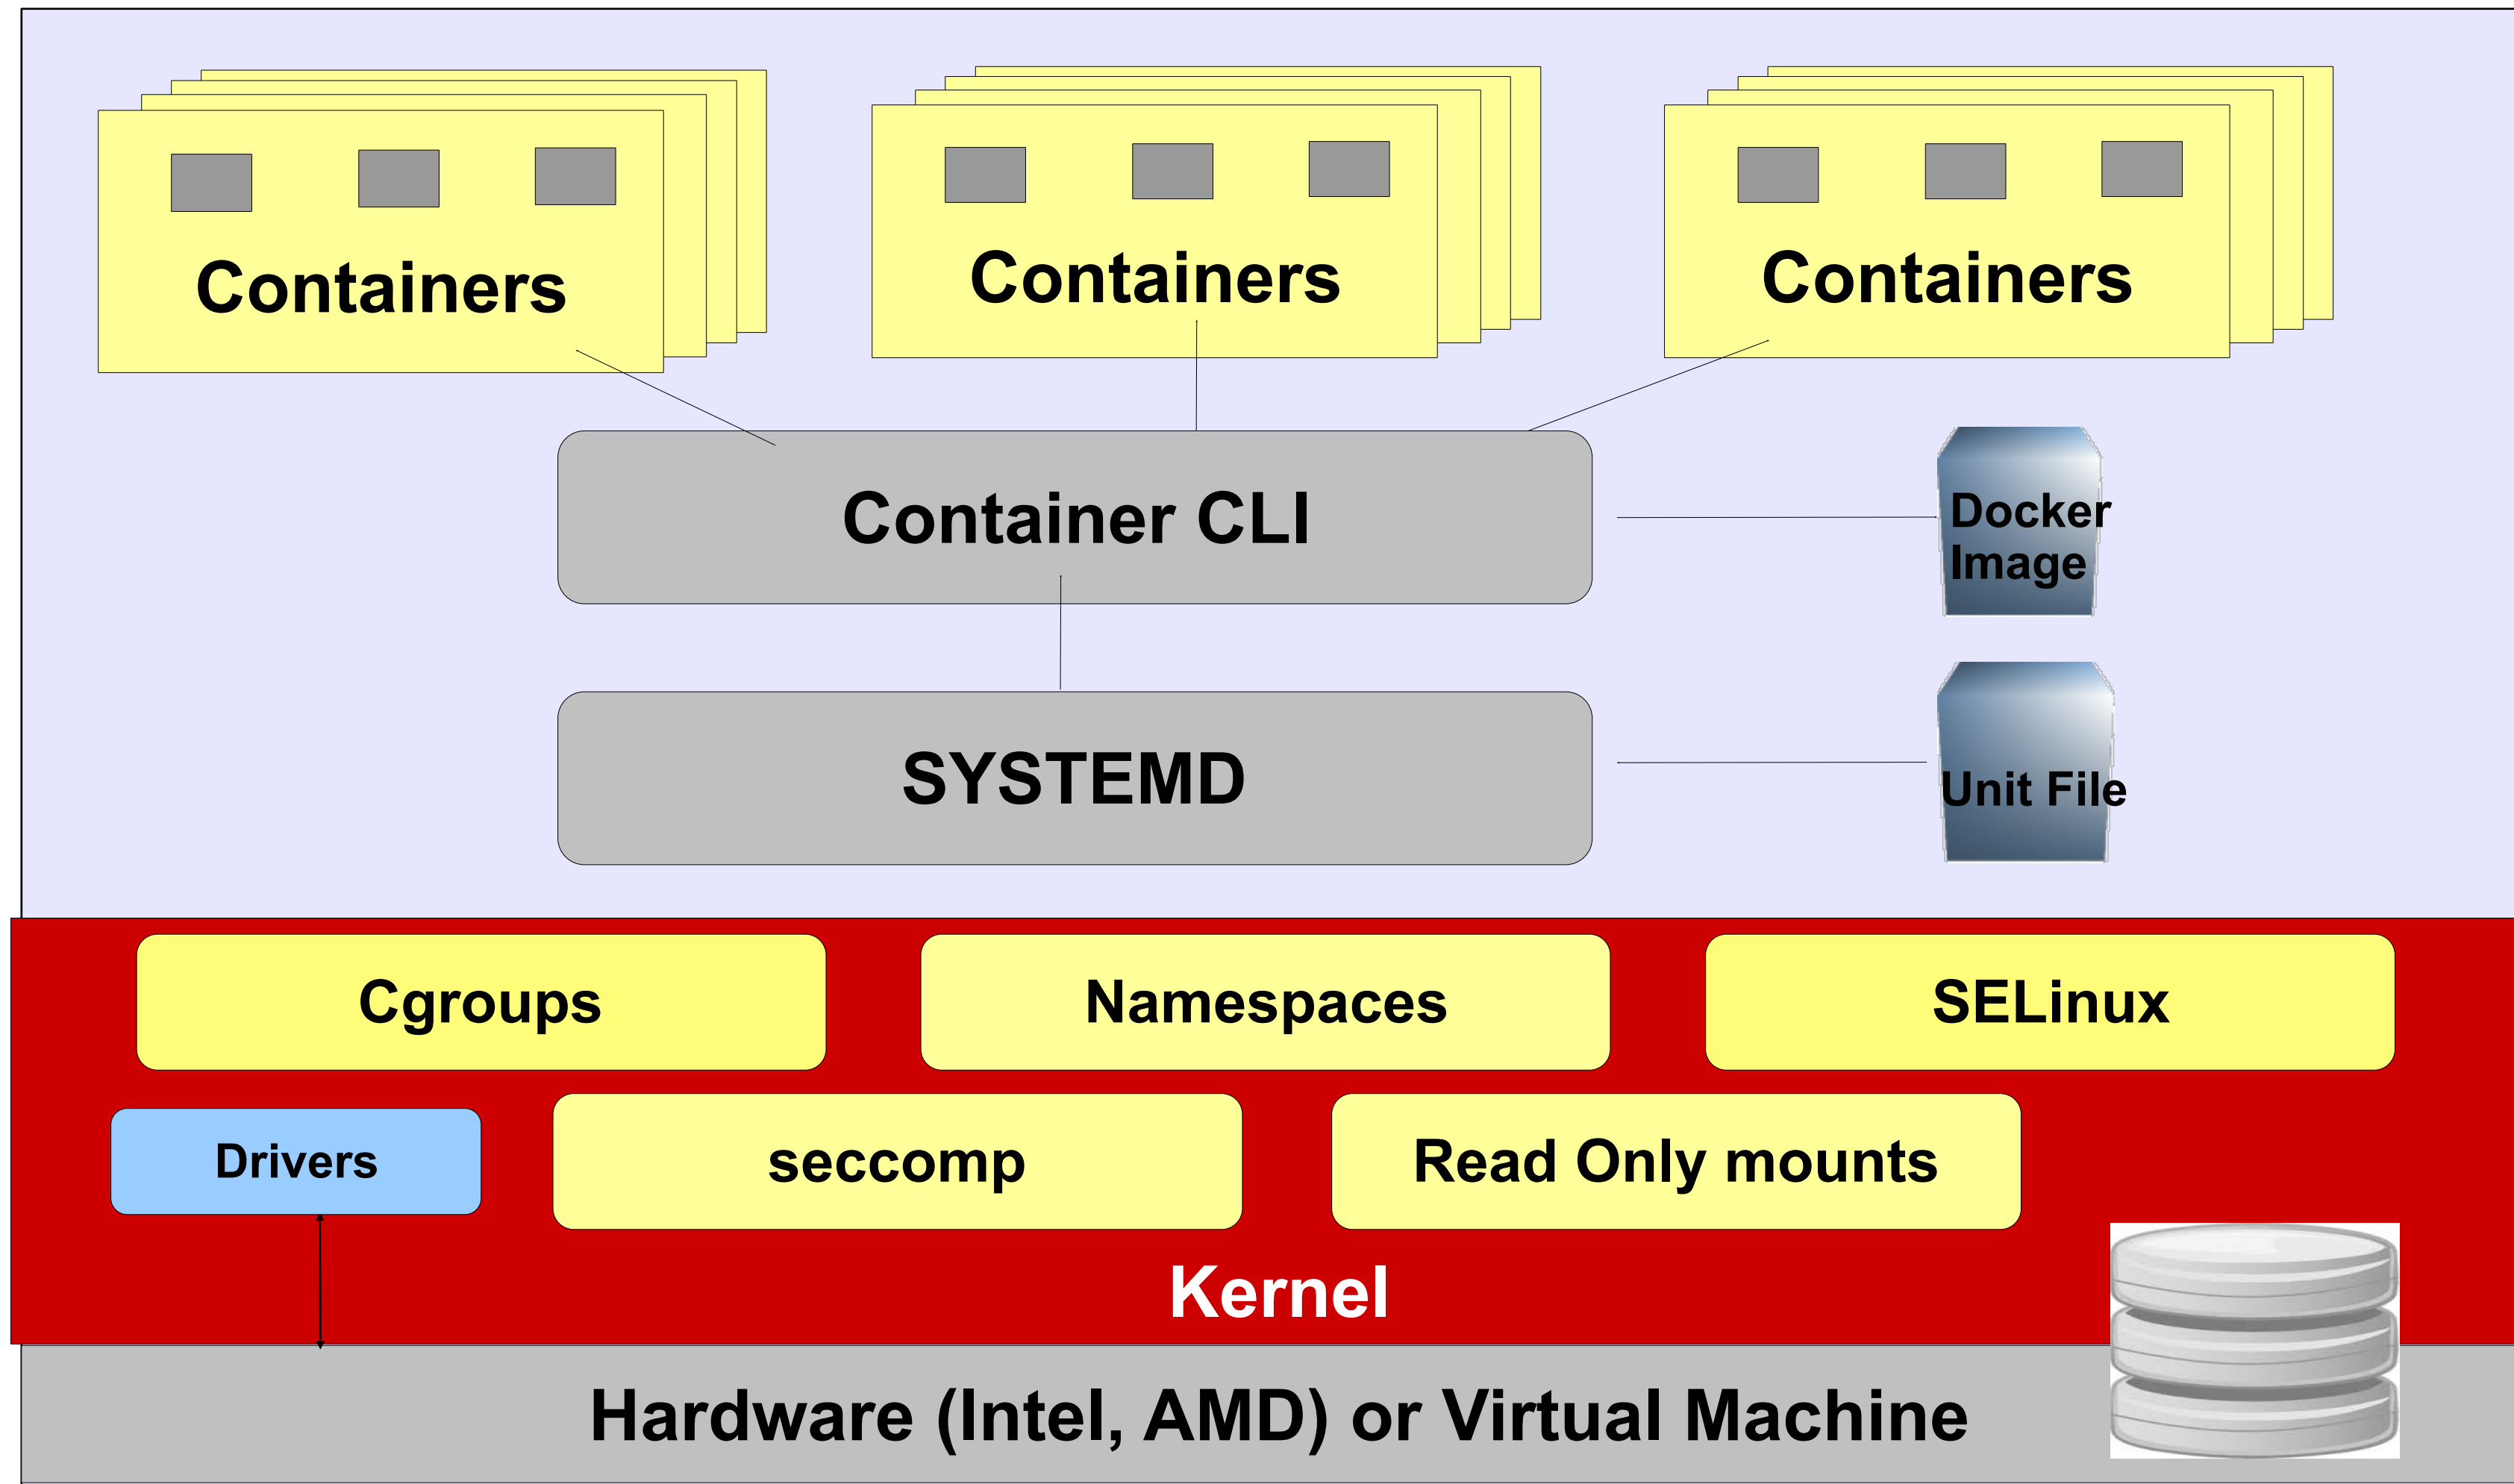


# CANARY DEPLOYMENTS

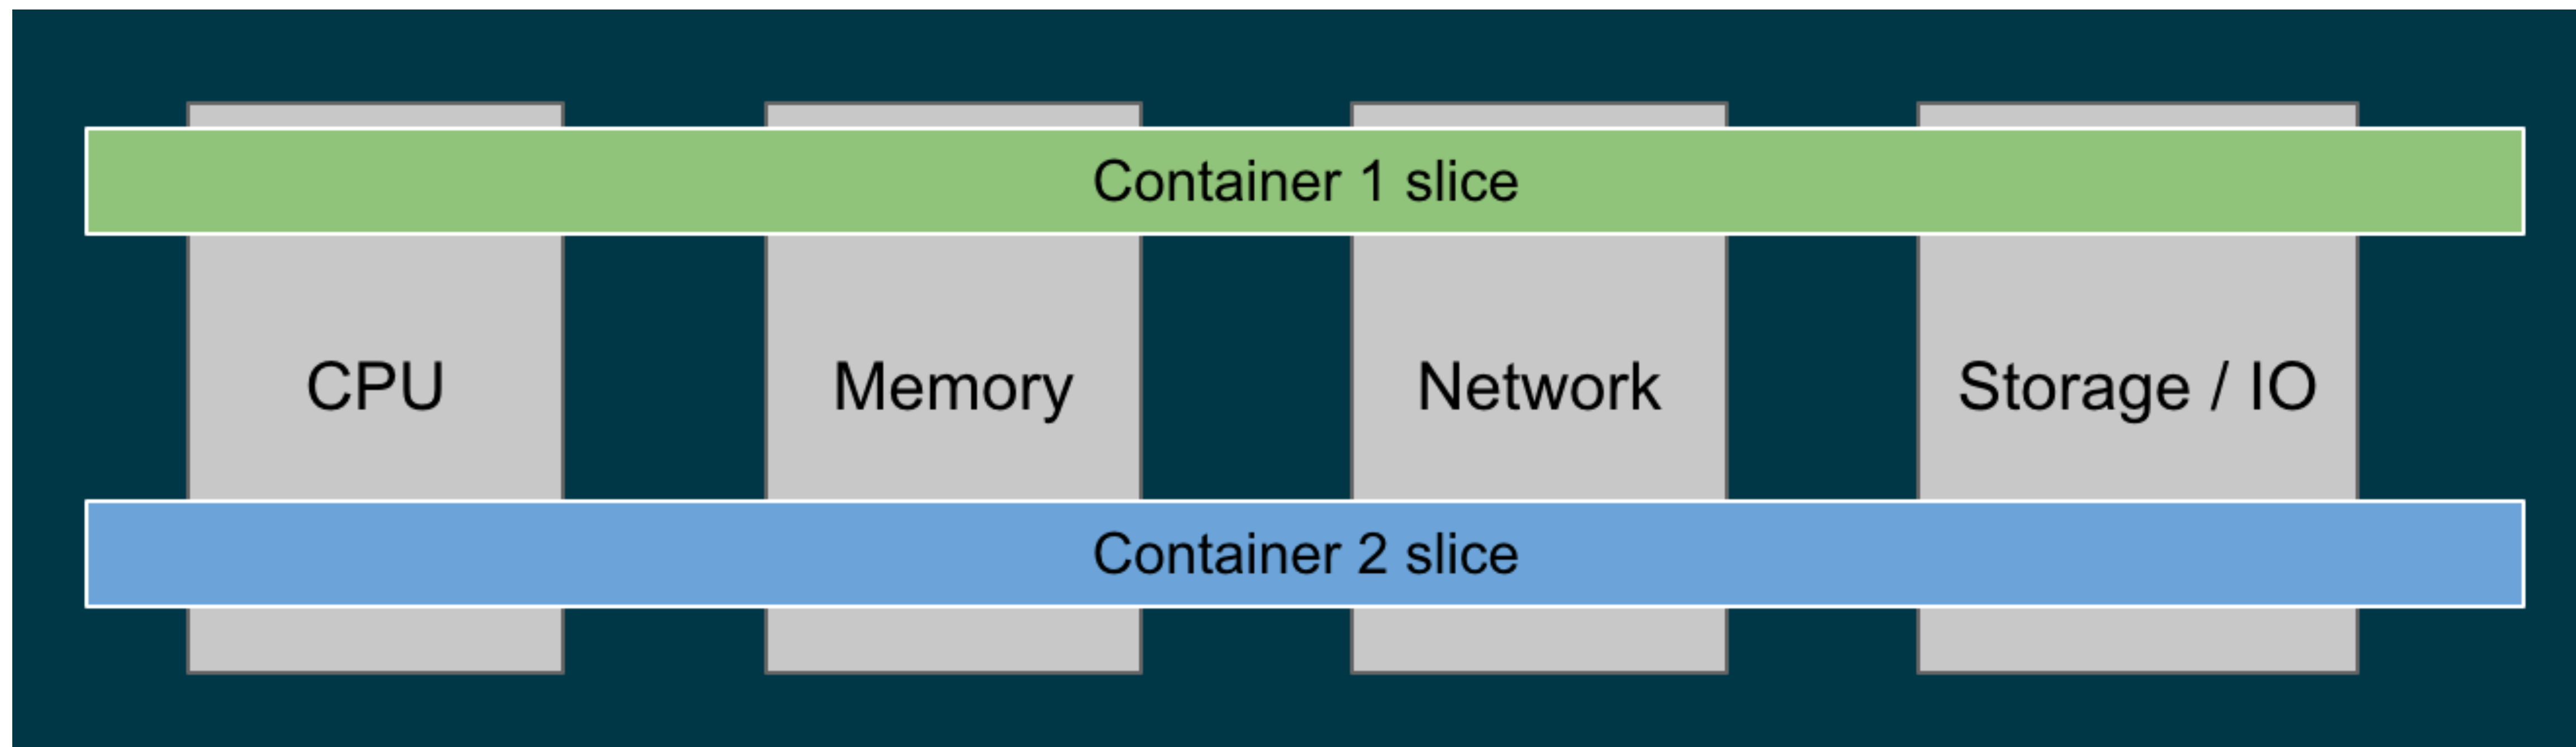


# CONTAINER HOST SECURITY

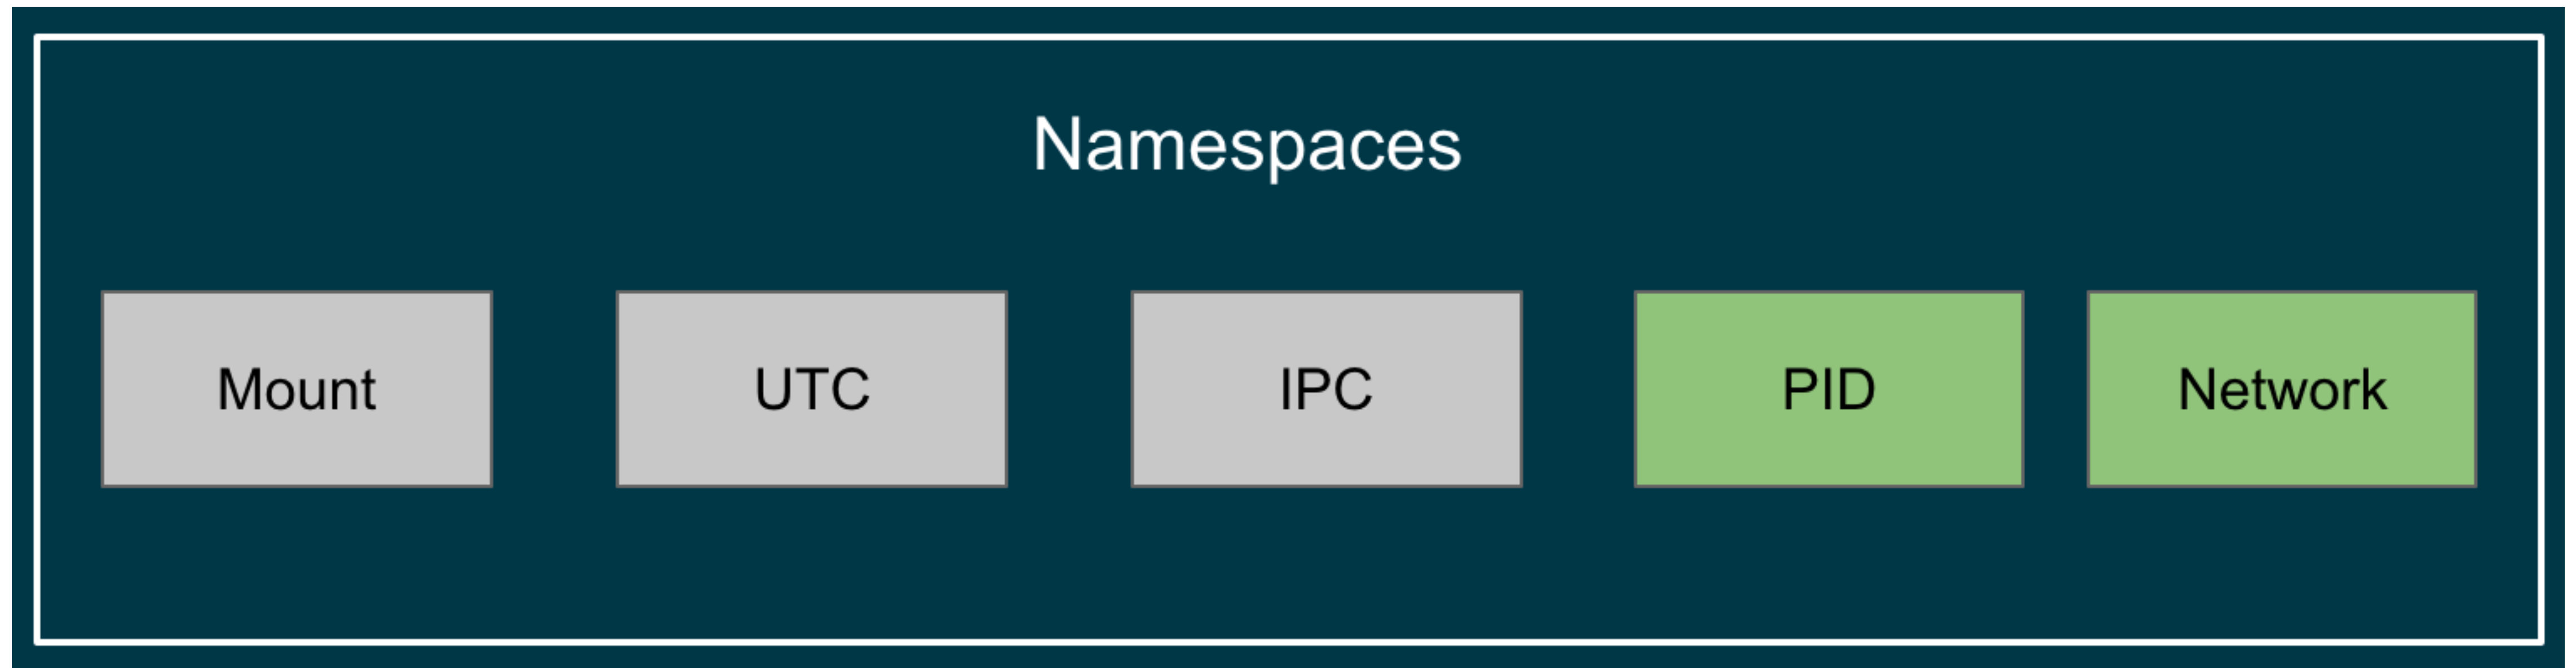
# CONTAINERS ARE LINUX



# CGROUPS - RESOURCE ISOLATION




# NAMESPACES - PROCESS ISOLATION



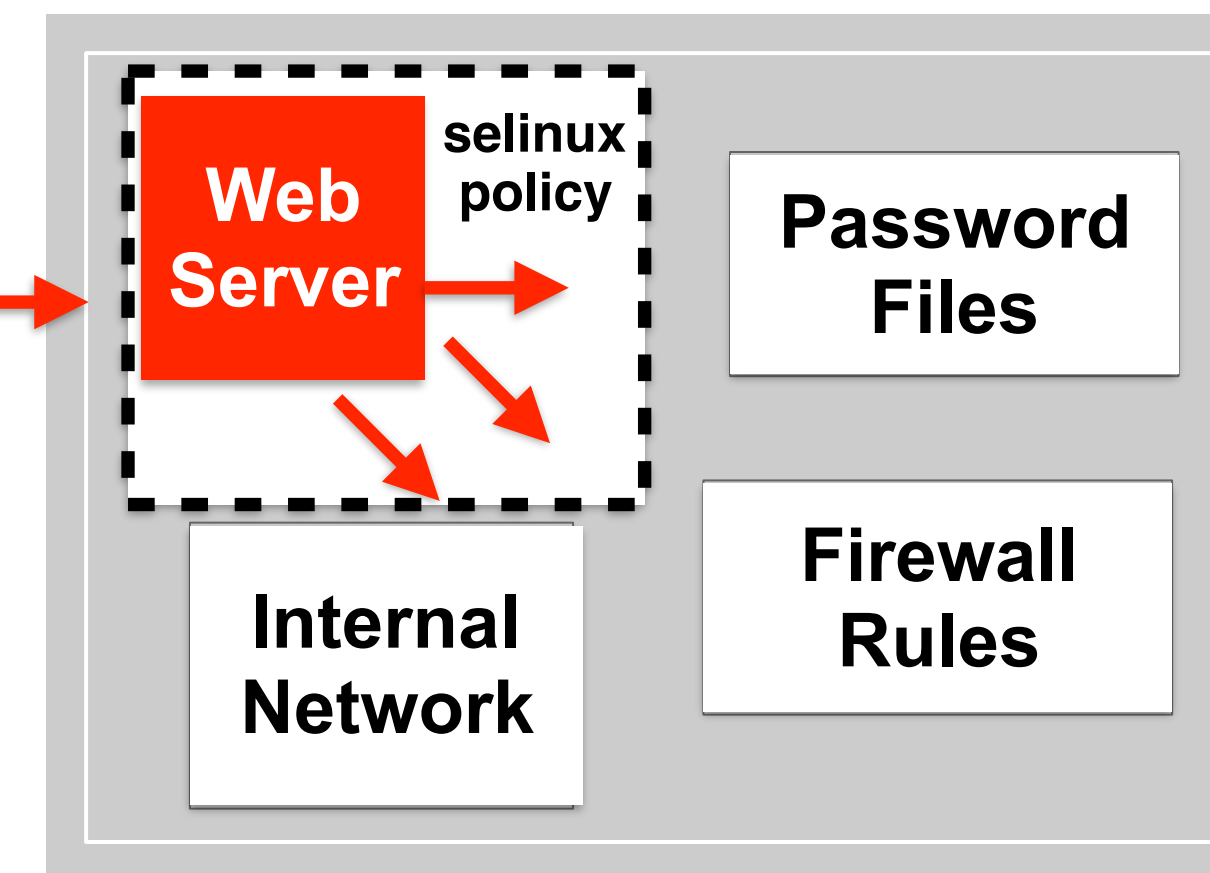
# SELINUX - MANDATORY ACCESS CONTROLS

## Discretionary Access Controls (file permissions)



## Mandatory Access Controls (selinux)

Attacker





# SECCOMP - DROPPING PRIVILEGES

CAP_SETPCAP	Modify process capabilities
CAP_SYS_MODULE	Insert/Remove kernel modules
CAP_SYS_RAWIO	Modify Kernel Memory
CAP_SYS_PACCT	Configure process accounting
CAP_SYS_NICE	Modify Priority of processes
CAP_SYS_RESOURCE	Override Resource Limits
CAP_SYS_TIME	Modify the system clock
CAP_SYS_TTY_CONFIG	Configure tty devices
CAP_AUDIT_WRITE	Write the audit log
CAP_AUDIT_CONTROL	Configure Audit Subsystem
CAP_MAC_OVERRIDE	Ignore Kernel MAC Policy
CAP_MAC_ADMIN	Configure MAC Configuration
CAP_SYSLOG	Modify Kernel printk behaviour
CAP_NET_ADMIN	Configure the network:
CAP_SYS_ADMIN	<ul style="list-style-type: none"><li>- Setting the hostname/domainname</li><li>- mount(),unmount()</li><li>- nfsservctl</li><li>- ....</li></ul>

# READ ONLY MOUNTS

/sys

/proc/sys

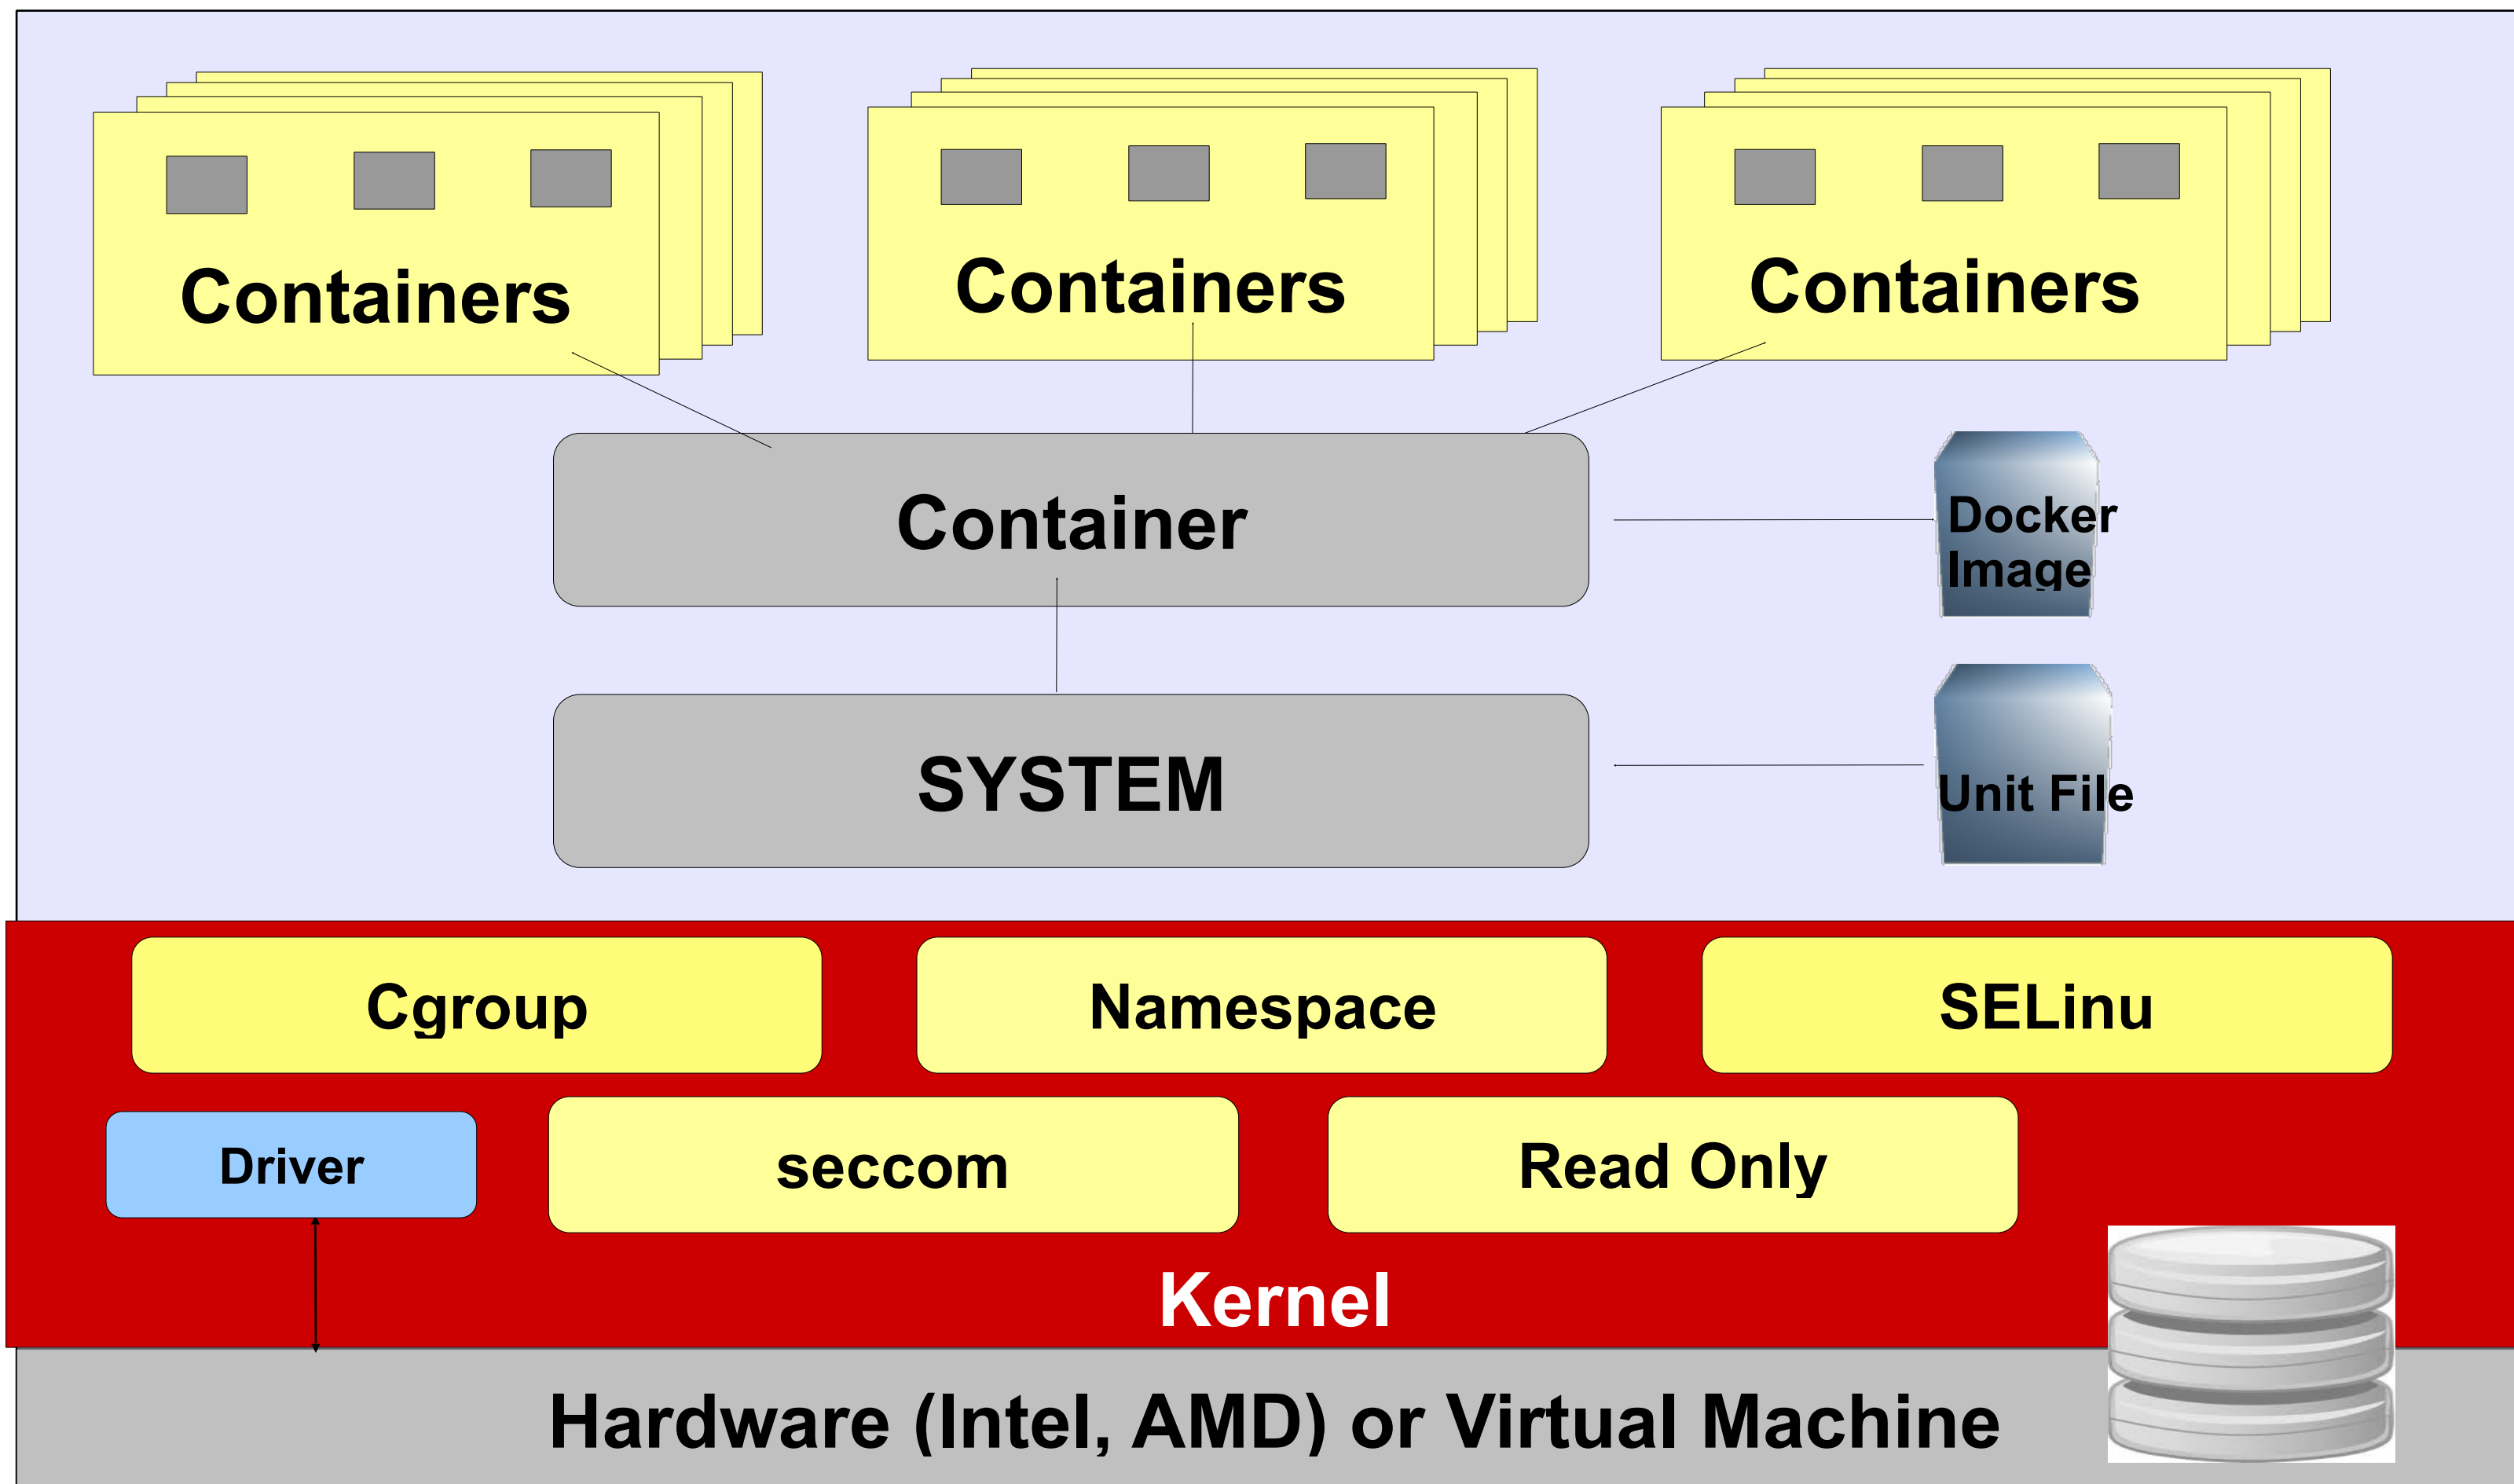
/proc/sysrg-trigger

/proc/irq

/proc/bus

R/O

# CONTAINER HOST SECURITY




## Best Practices


- Don't run as root
- Limit SSH Access
- Use namespaces
- Define resource quotas
- Enable logging
- Apply Security Errata
- Apply Security Context and seccomp filters

<http://blog.kubernetes.io/2016/08/security-best-practices-kubernetes-deployment.html>


# SECURING CONTAINERS




**Builds**




**Images**




**Registry**



**CI/CD**




**Container  
host**



**Network  
isolation**




**Storage**



**API & Platform  
access**



**Monitoring &  
Logging**



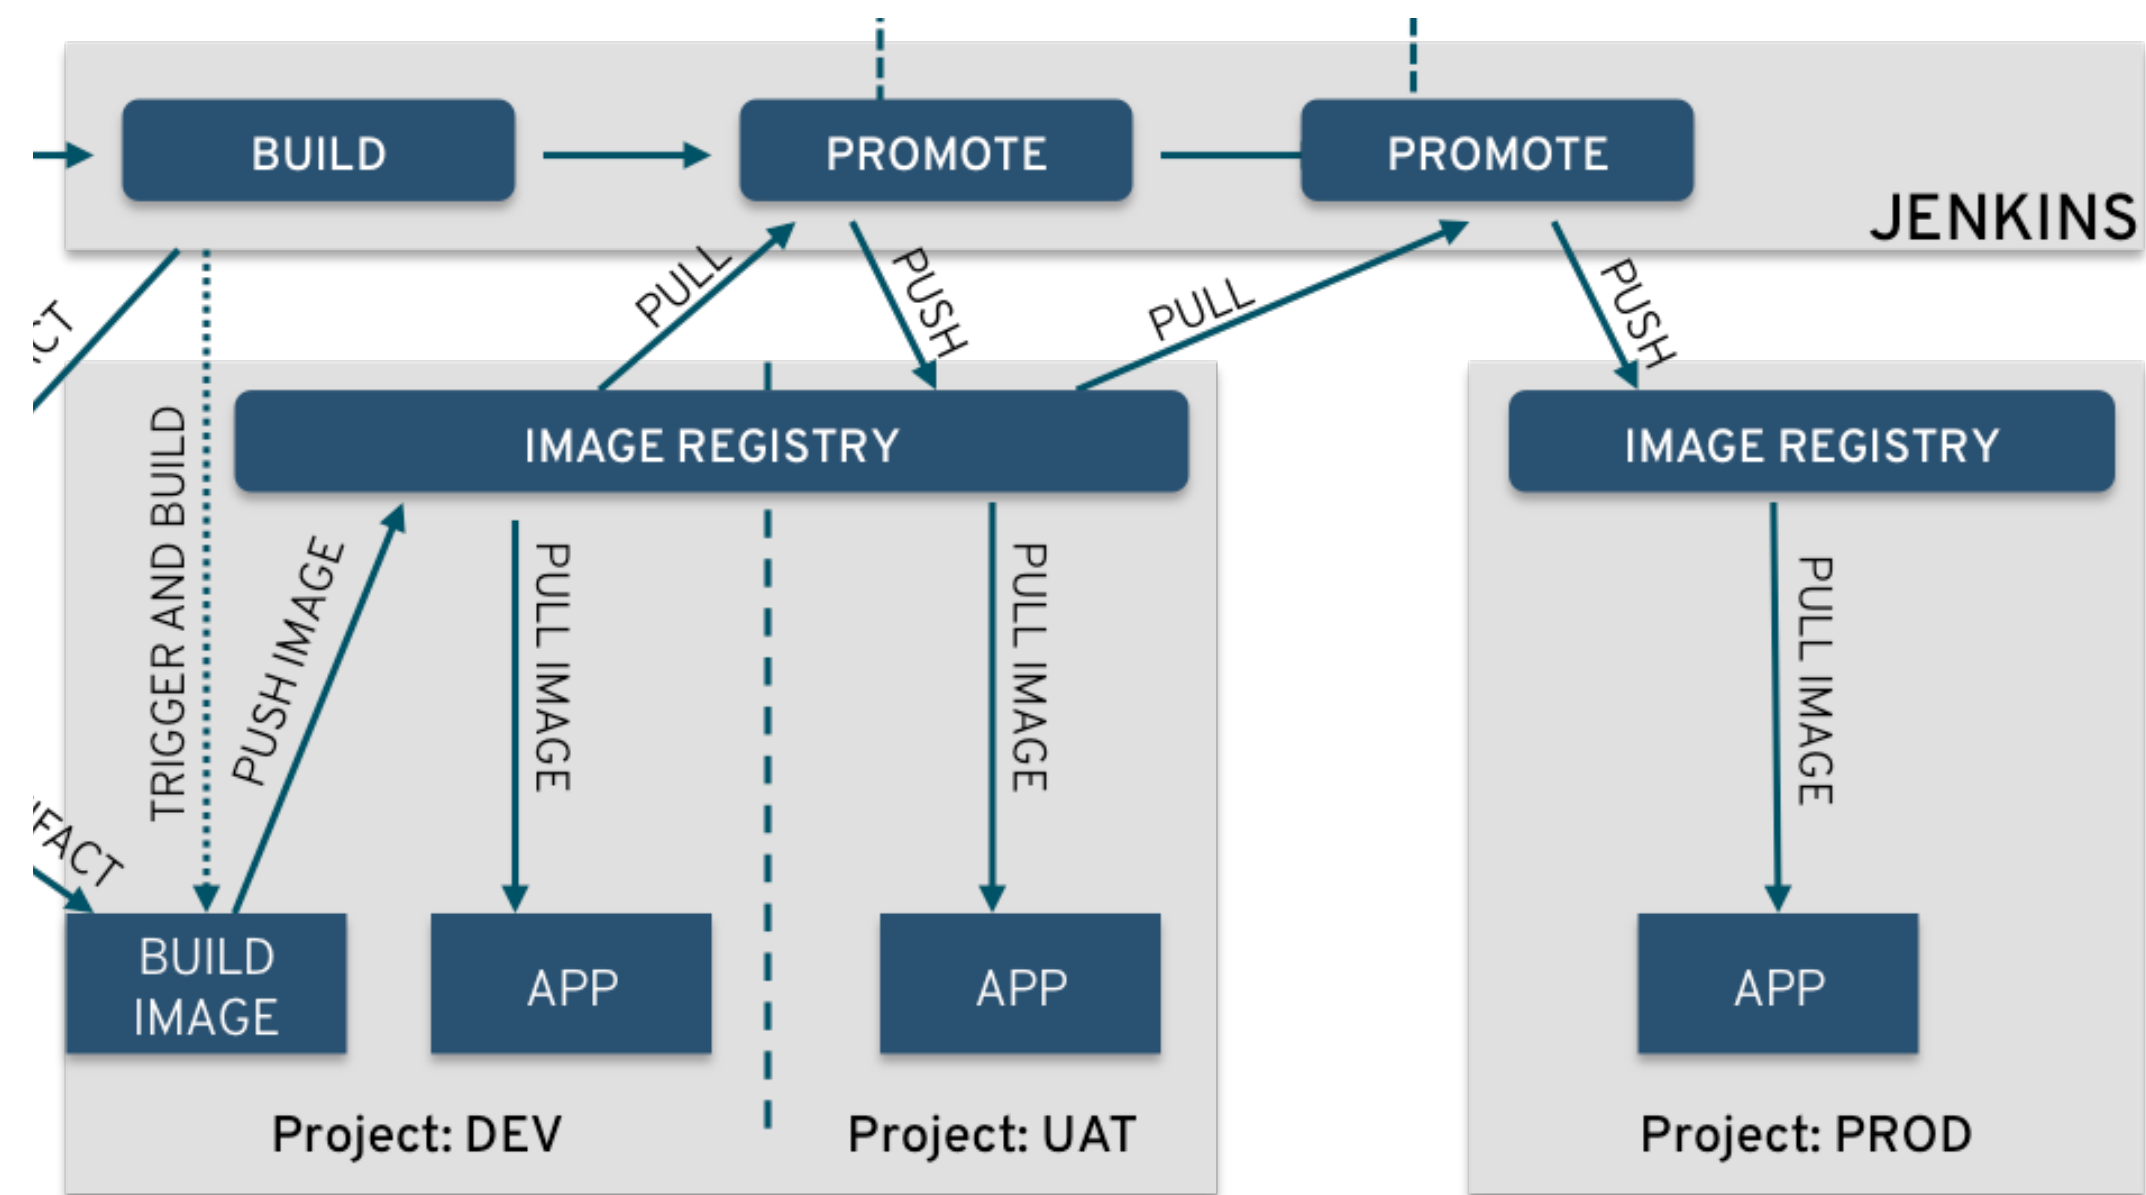
**Federated  
clusters**

# NETWORK ISOLATION

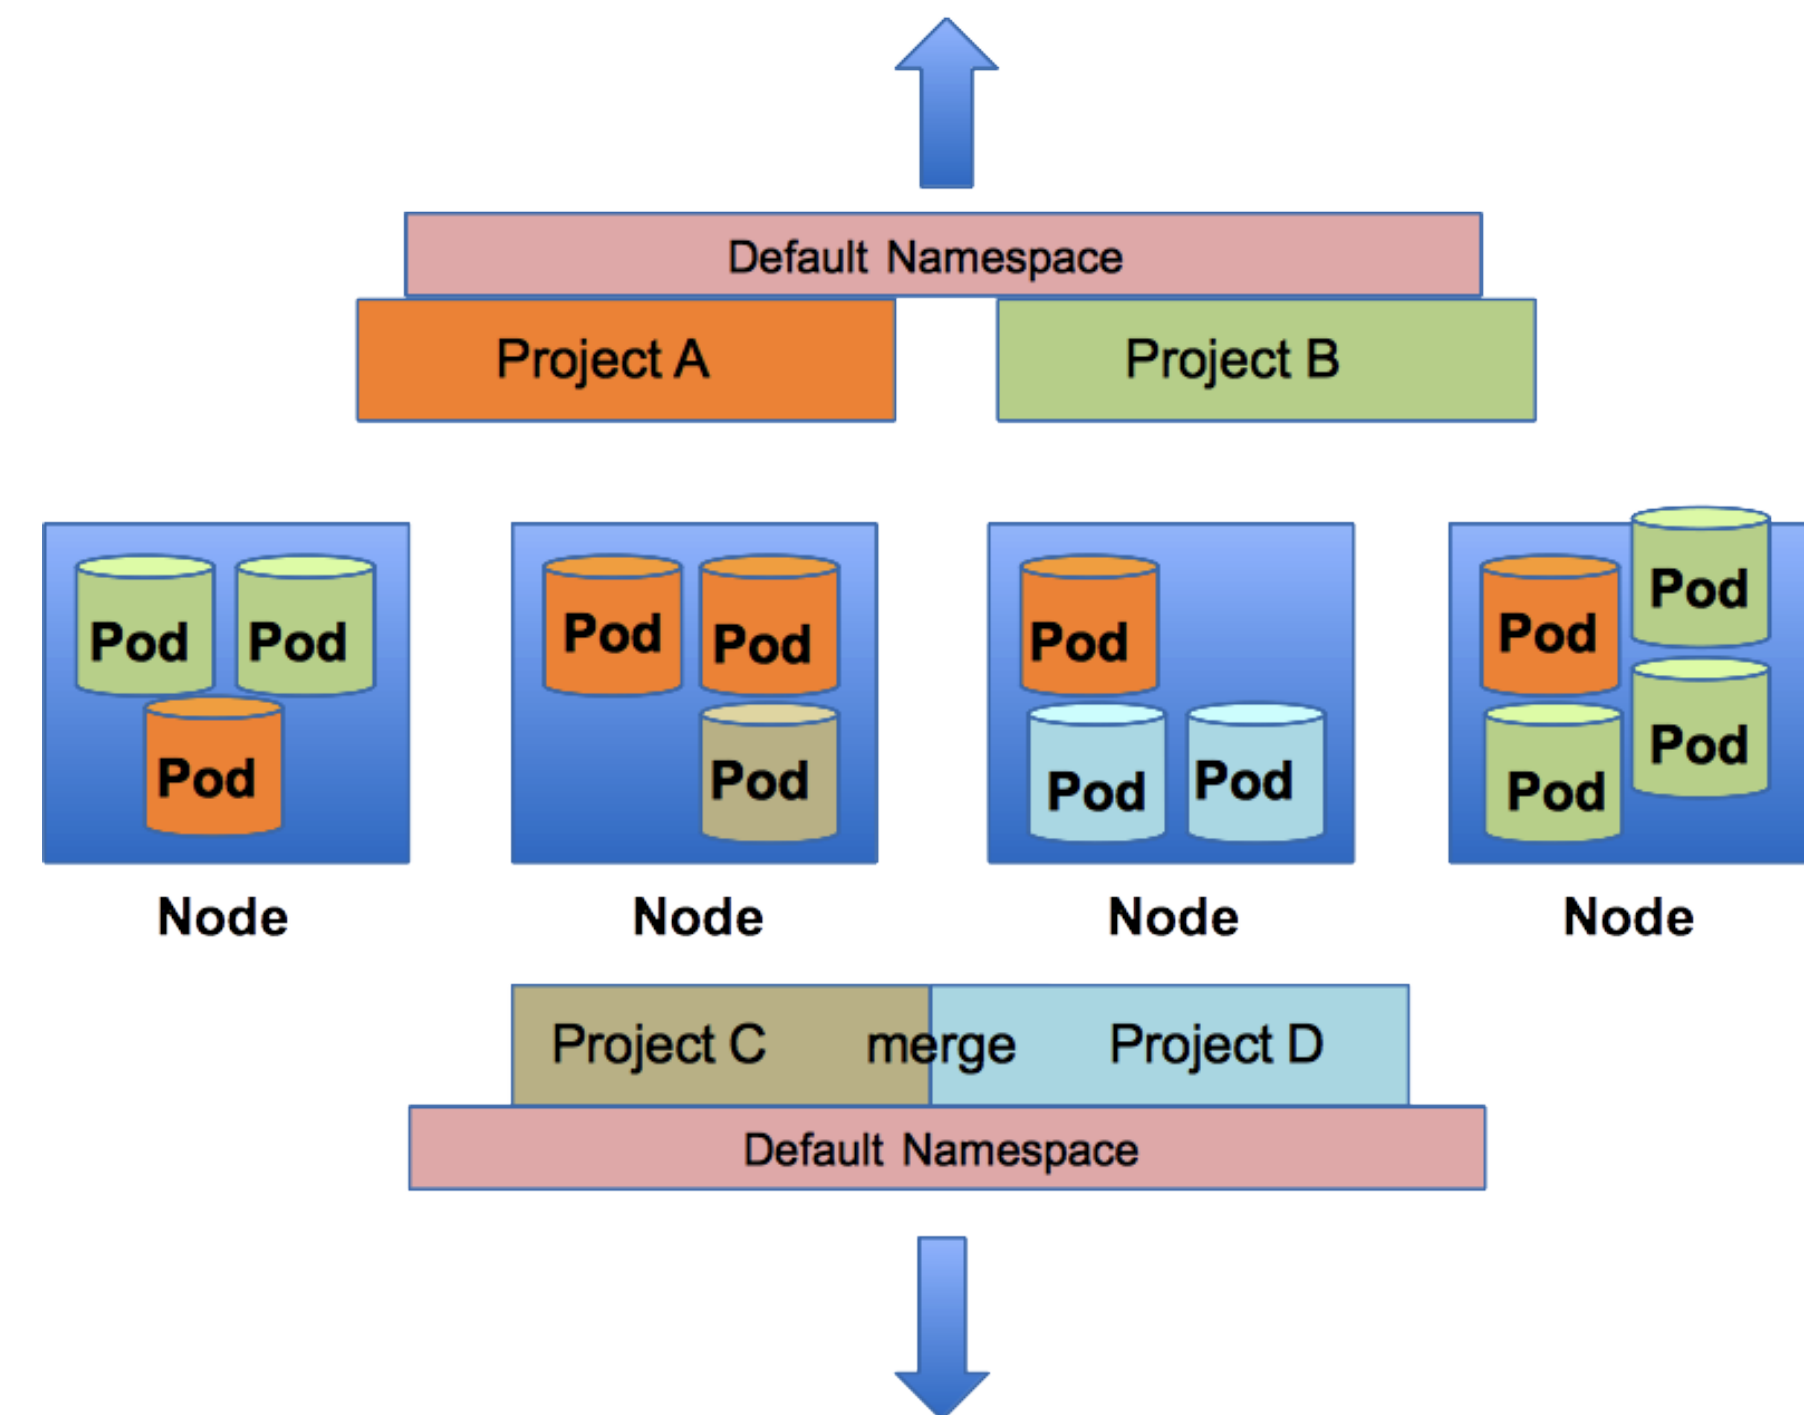


# NETWORK ISOLATION

## Multi-Environment



## Multi-Tenant



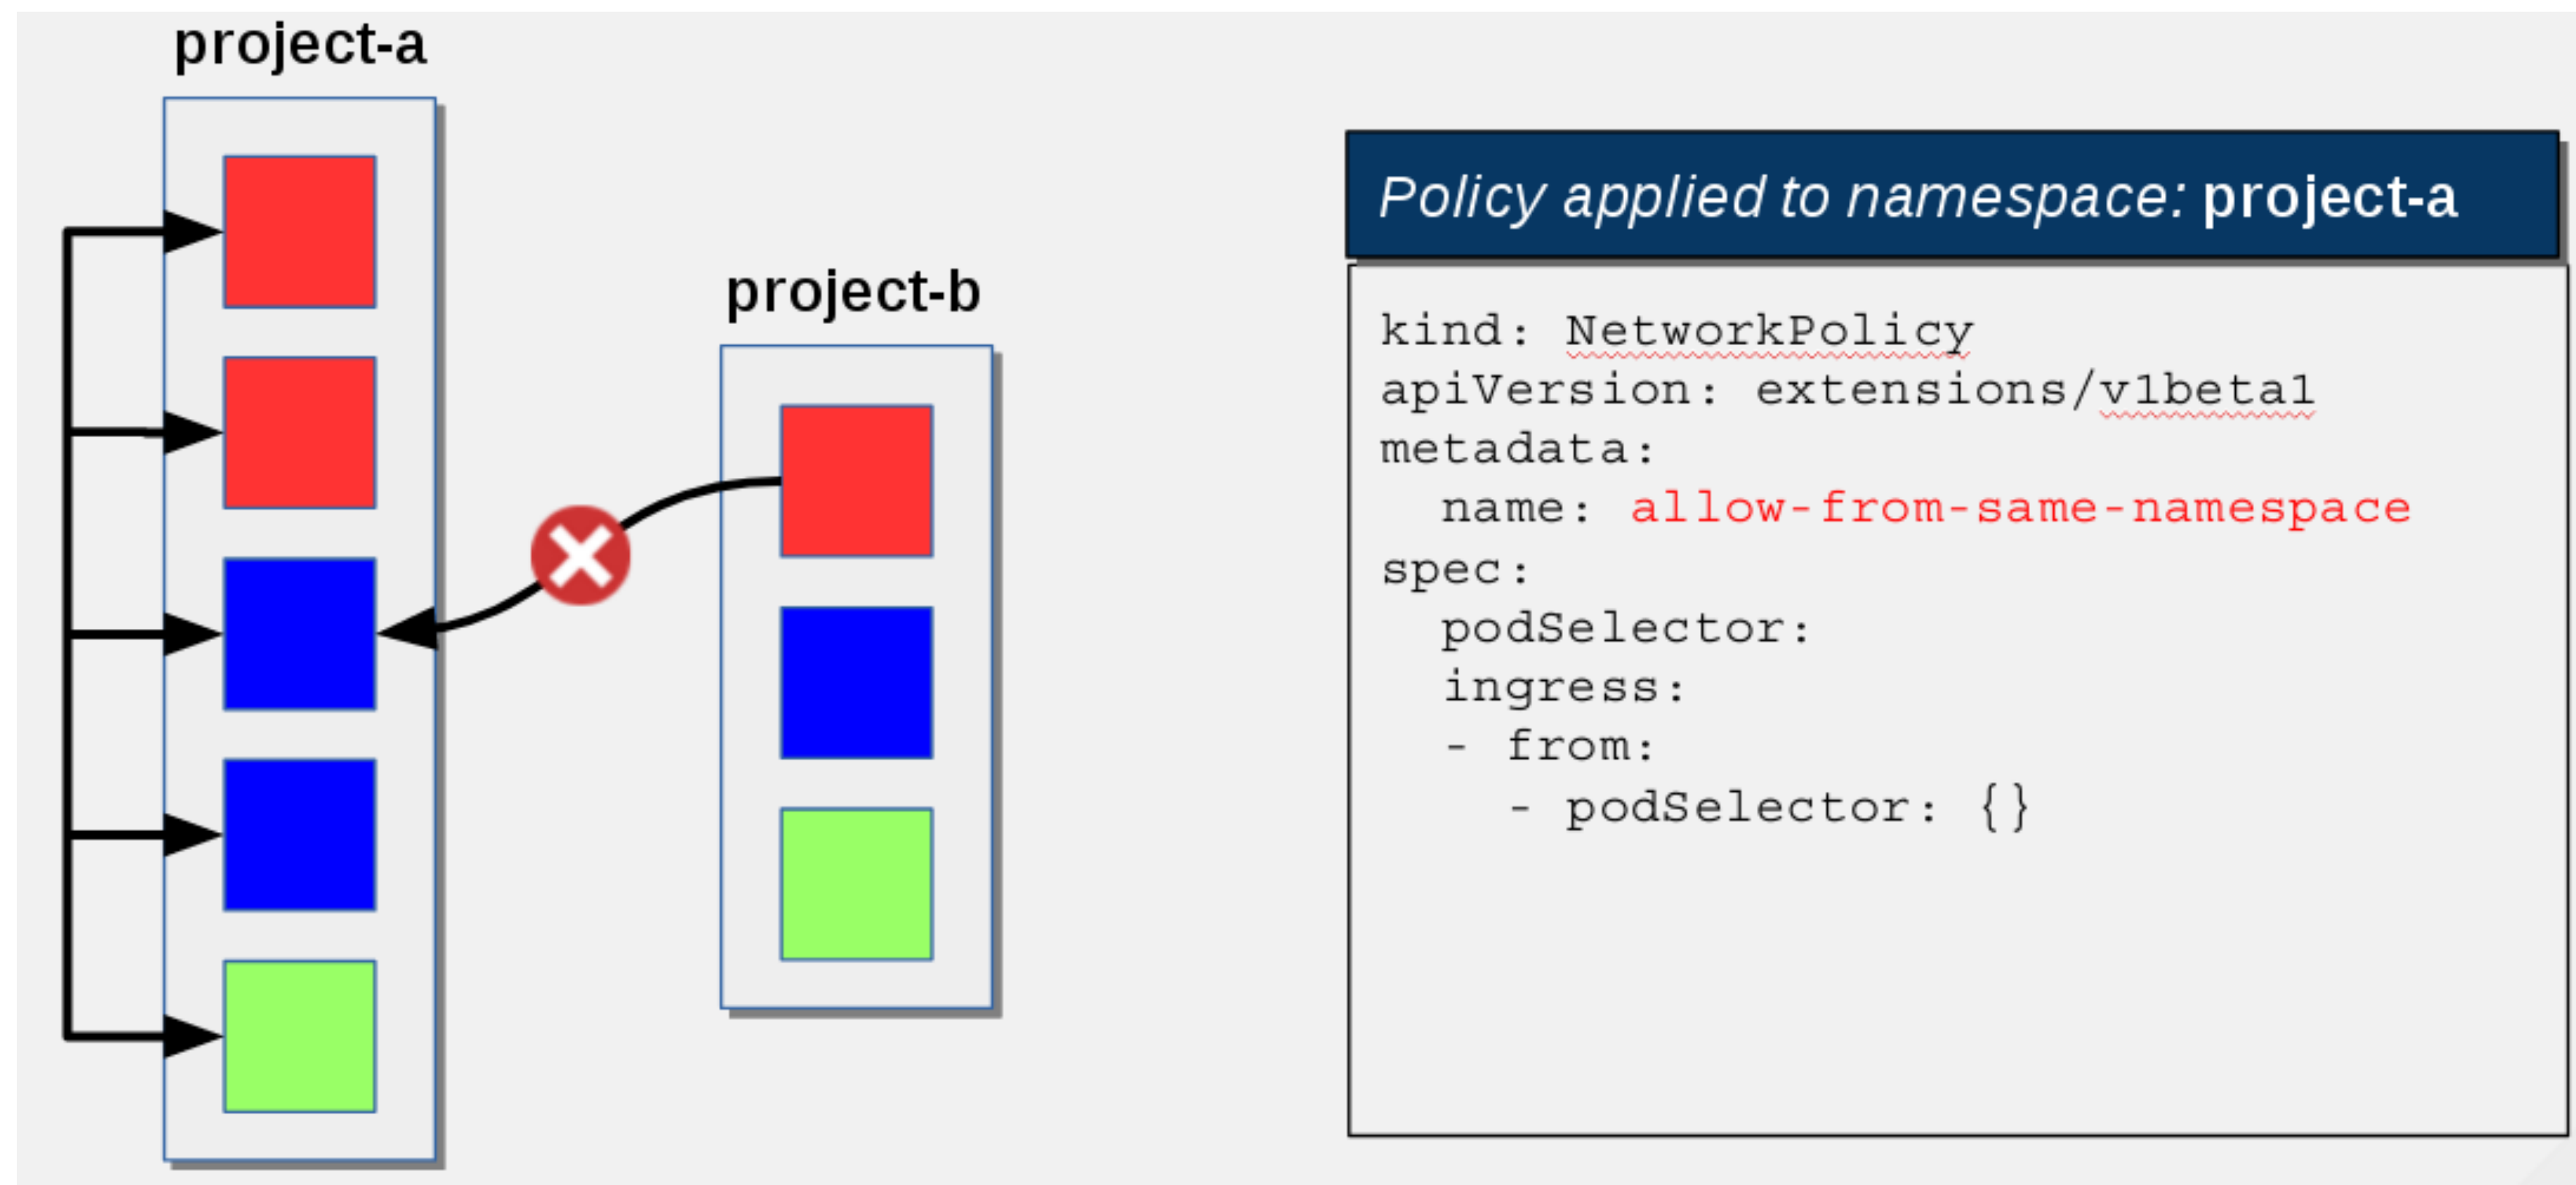
**Network Namespace  
provides resource isolation**



# NETWORK POLICY

example:

all pods in namespace 'project-a' allow traffic from any other pods in the same namespace."




# STORAGE SECURITY

# STORAGE SECURITY

Local Storage Quota

Security Context Constraints



Node's /var/lib/origin

# API & PLATFORM ACCESS

# API & PLATFORM ACCESS

Authentication  
via  
OAuth tokens and  
SSL certificate

Authorization  
via  
Policy Engine  
checks  
User/Group  
Defined Roles

# MONITORING & LOGGING




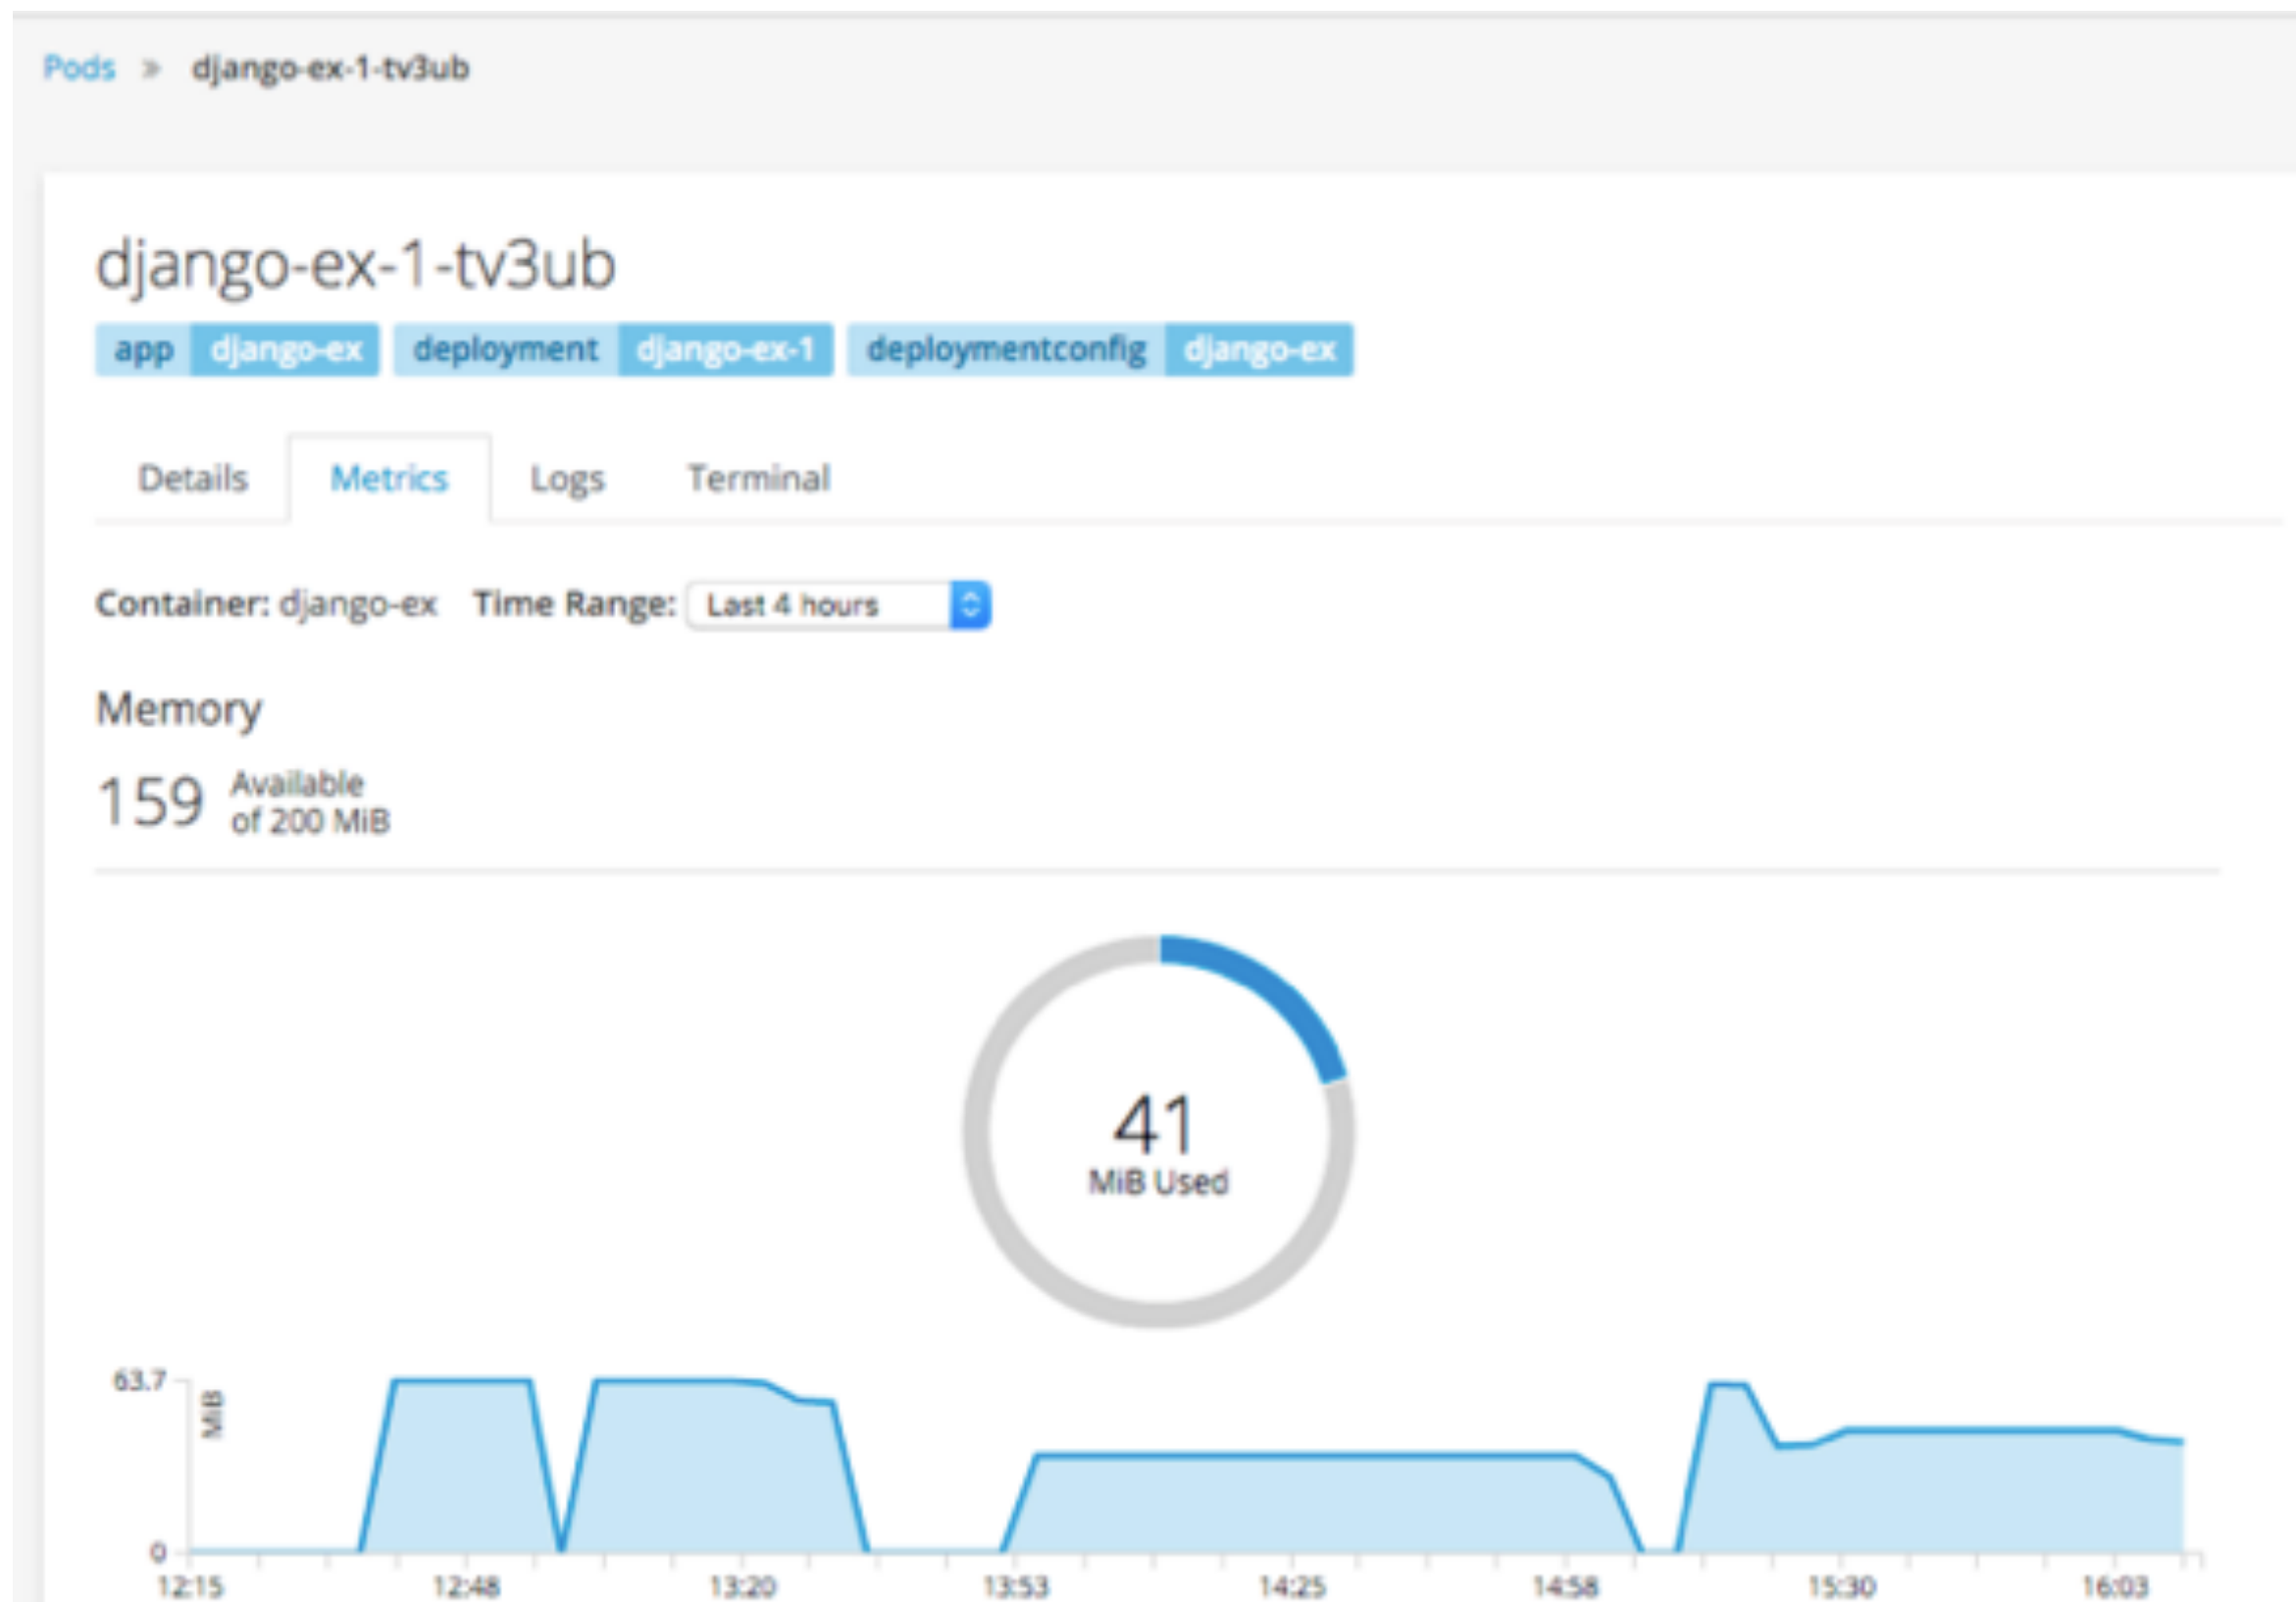
# LOGGING

## Aggregate platform and application log access via Kibana + Elasticsearch



# MONITORING

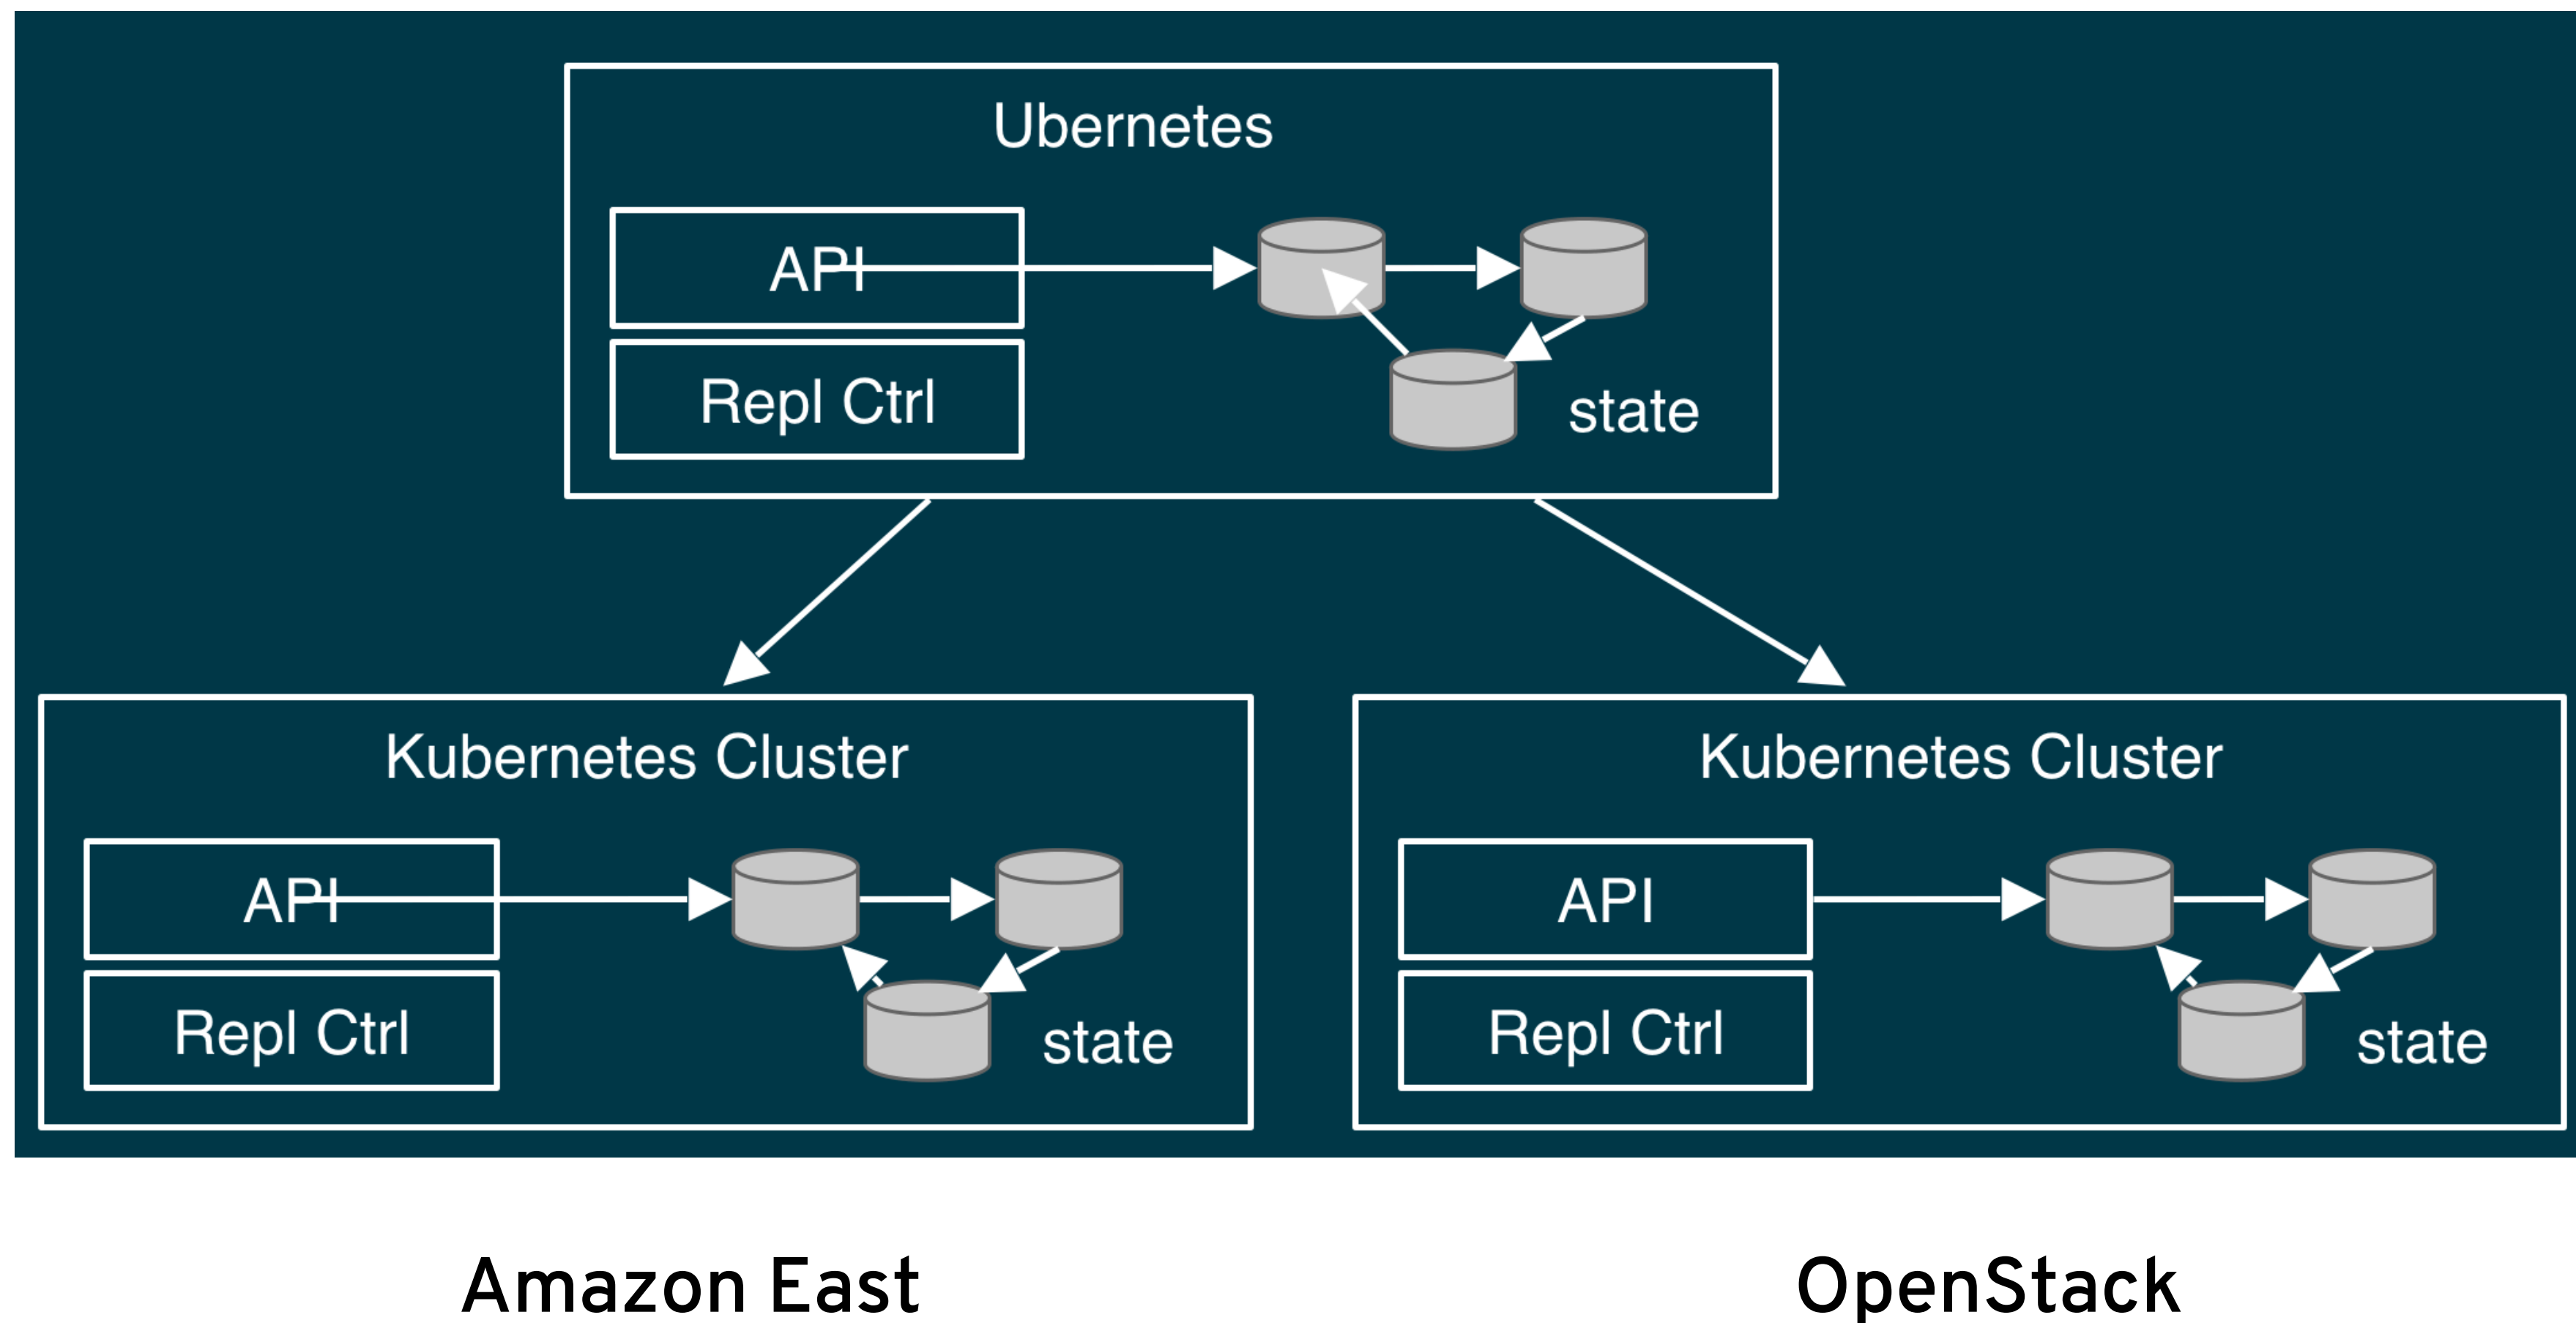
## Historical CPU and Memory usage



# FEDERATION

# FEDERATED CLUSTERS

Roles & access management (in-dev)



# MICROSERVICES



# SERVICE MESH




## Monitoring & Metrics

- prometheus (logs)
- grafana (visual)



**Istio**



## Traffic routing

- pilot
- circuit breaker
- a/b testing
- traffic mirroring



## Access Control & usage policies

- mixr (policy decisions)



## Encryption & Auth

- citadel
- service 2 service
- user auth



## Fault injections


- envoy
- corner cases: abort & delays



# DEVSECOPS METRICS



**Compliance  
Score**



**Deployment  
Frequency**




**Lead  
Time**

**404**

Page not found

**Deployment  
Failure Rate**



**Mean Time  
to Recover**

**99.999**

**Service  
Availability**

# THANK YOU

**linkedin: Chris Van Tuin**

**email: [cvantuin@redhat.com](mailto:cvantuin@redhat.com)**

**twitter: [@chrisvantuin](https://twitter.com/chrisvantuin)**

