

# A DevOps State of Mind: Continuous Security with Kubernetes

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



# “Only the paranoid survive”

- Andy Grove, 1996





A woman with blonde hair in a bun is driving a car. She is looking down at a large map spread across her lap and the center console. The car's interior is visible, including the steering wheel, dashboard, and center console with a gear shifter. Through the windshield, a road with a guardrail and a Volvo car with license plate KWP 031 are visible. The text "THE WORLD IS AUTOMATING" is overlaid in large white letters.

# THE WORLD IS AUTOMATING

Those who succeed in automation will win








# THE CHALLENGE: ENABLE INNOVATION AT SPEED, WHILE EXECUTING AT SCALE WITH EFFICIENCY

**Old**

Innovation

Execution




Static &  
Planned

**New**

Innovation

Execution

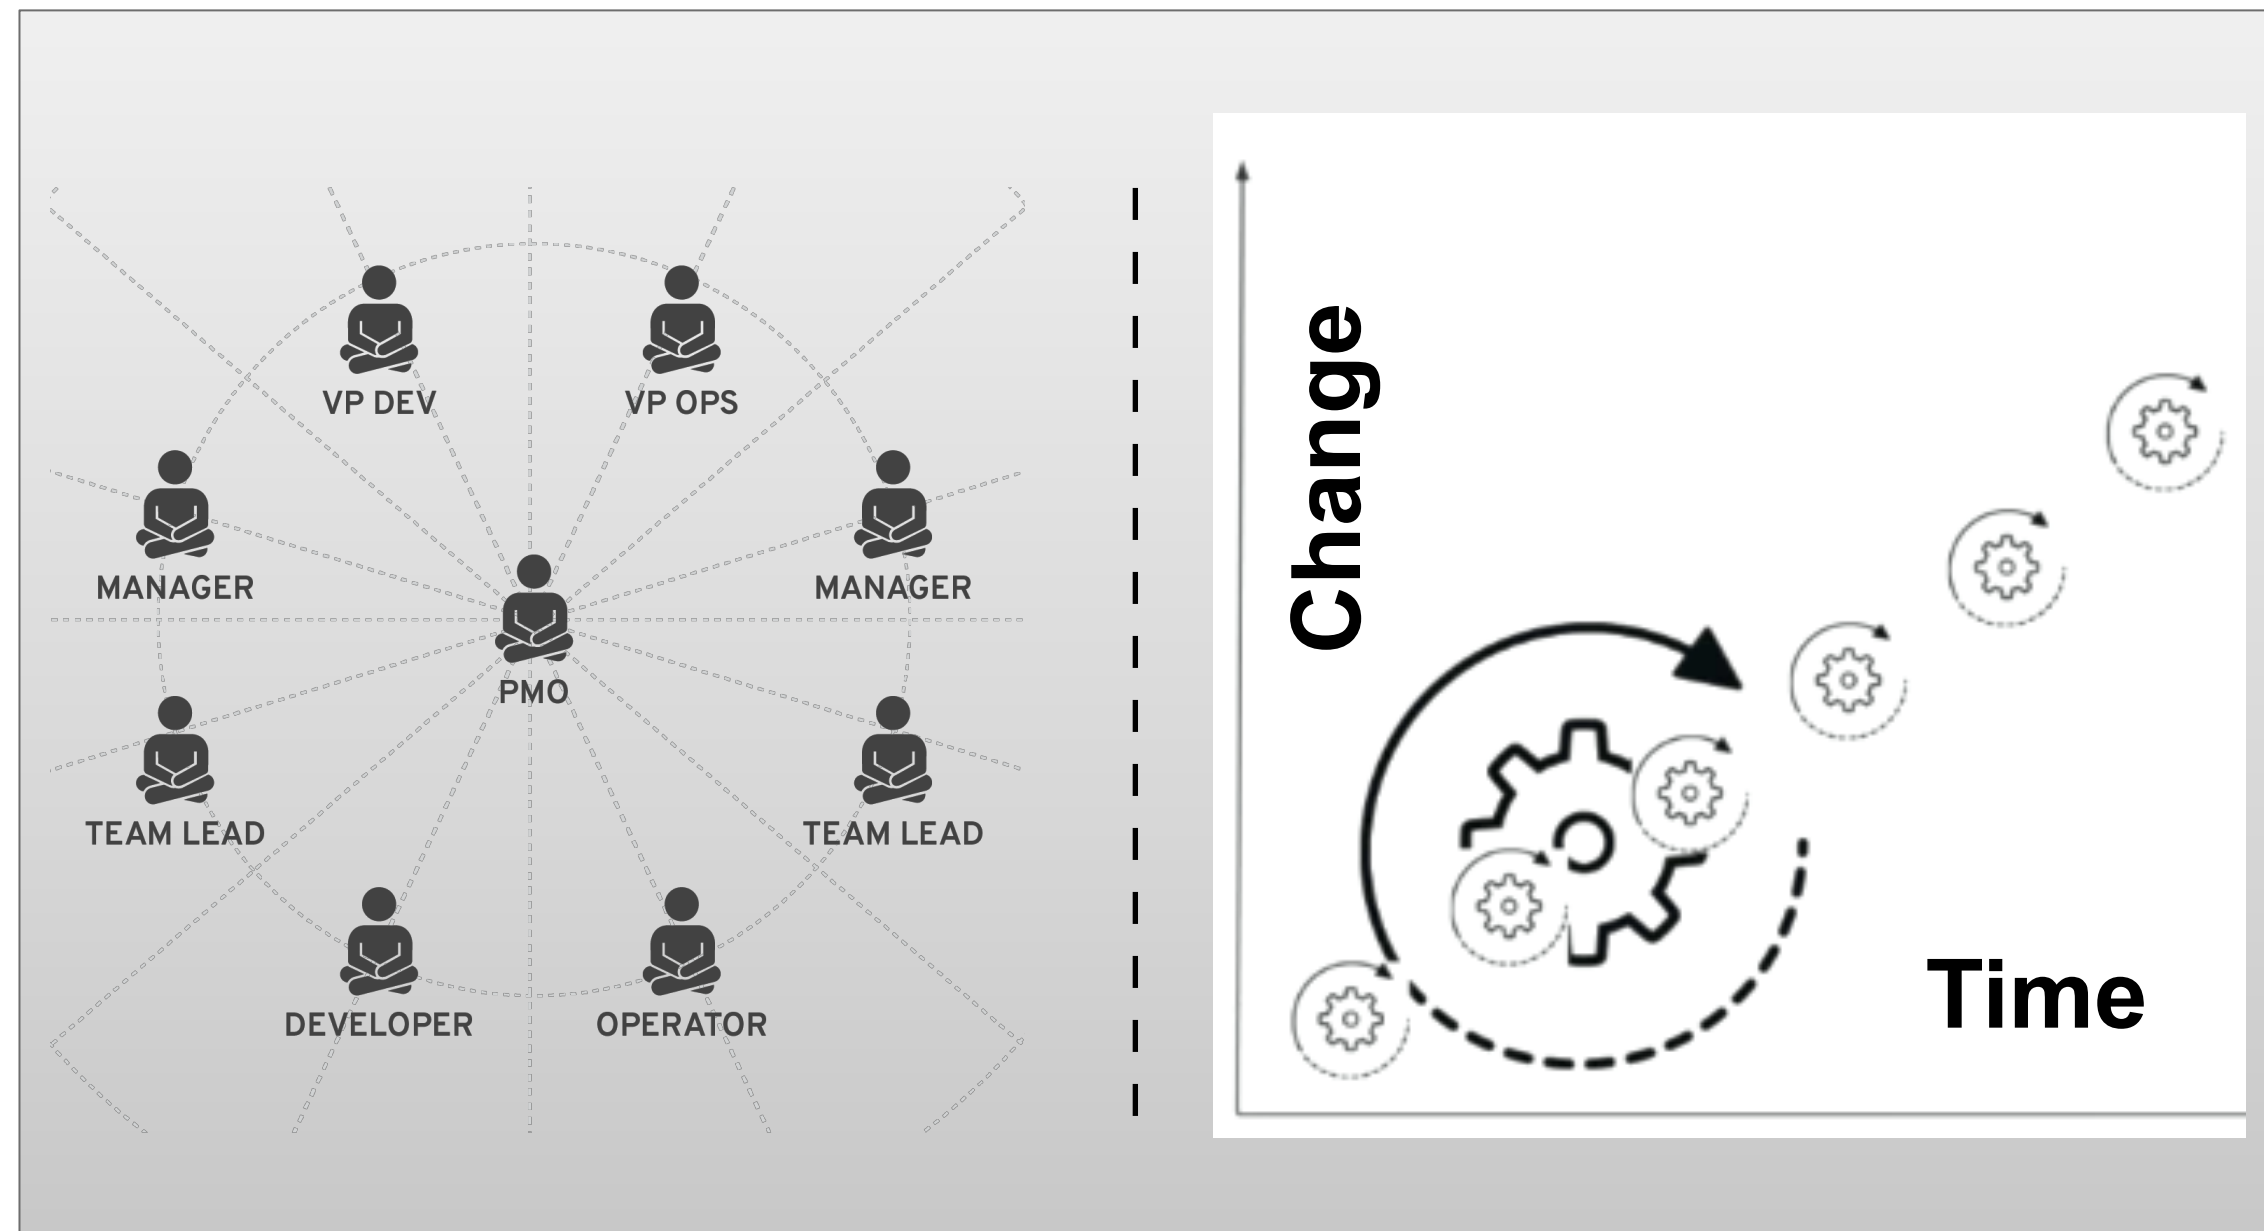


Dynamic &  
Policy Driven



# IT'S NOT JUST SOFTWARE, THE DIGITAL LEADERS =

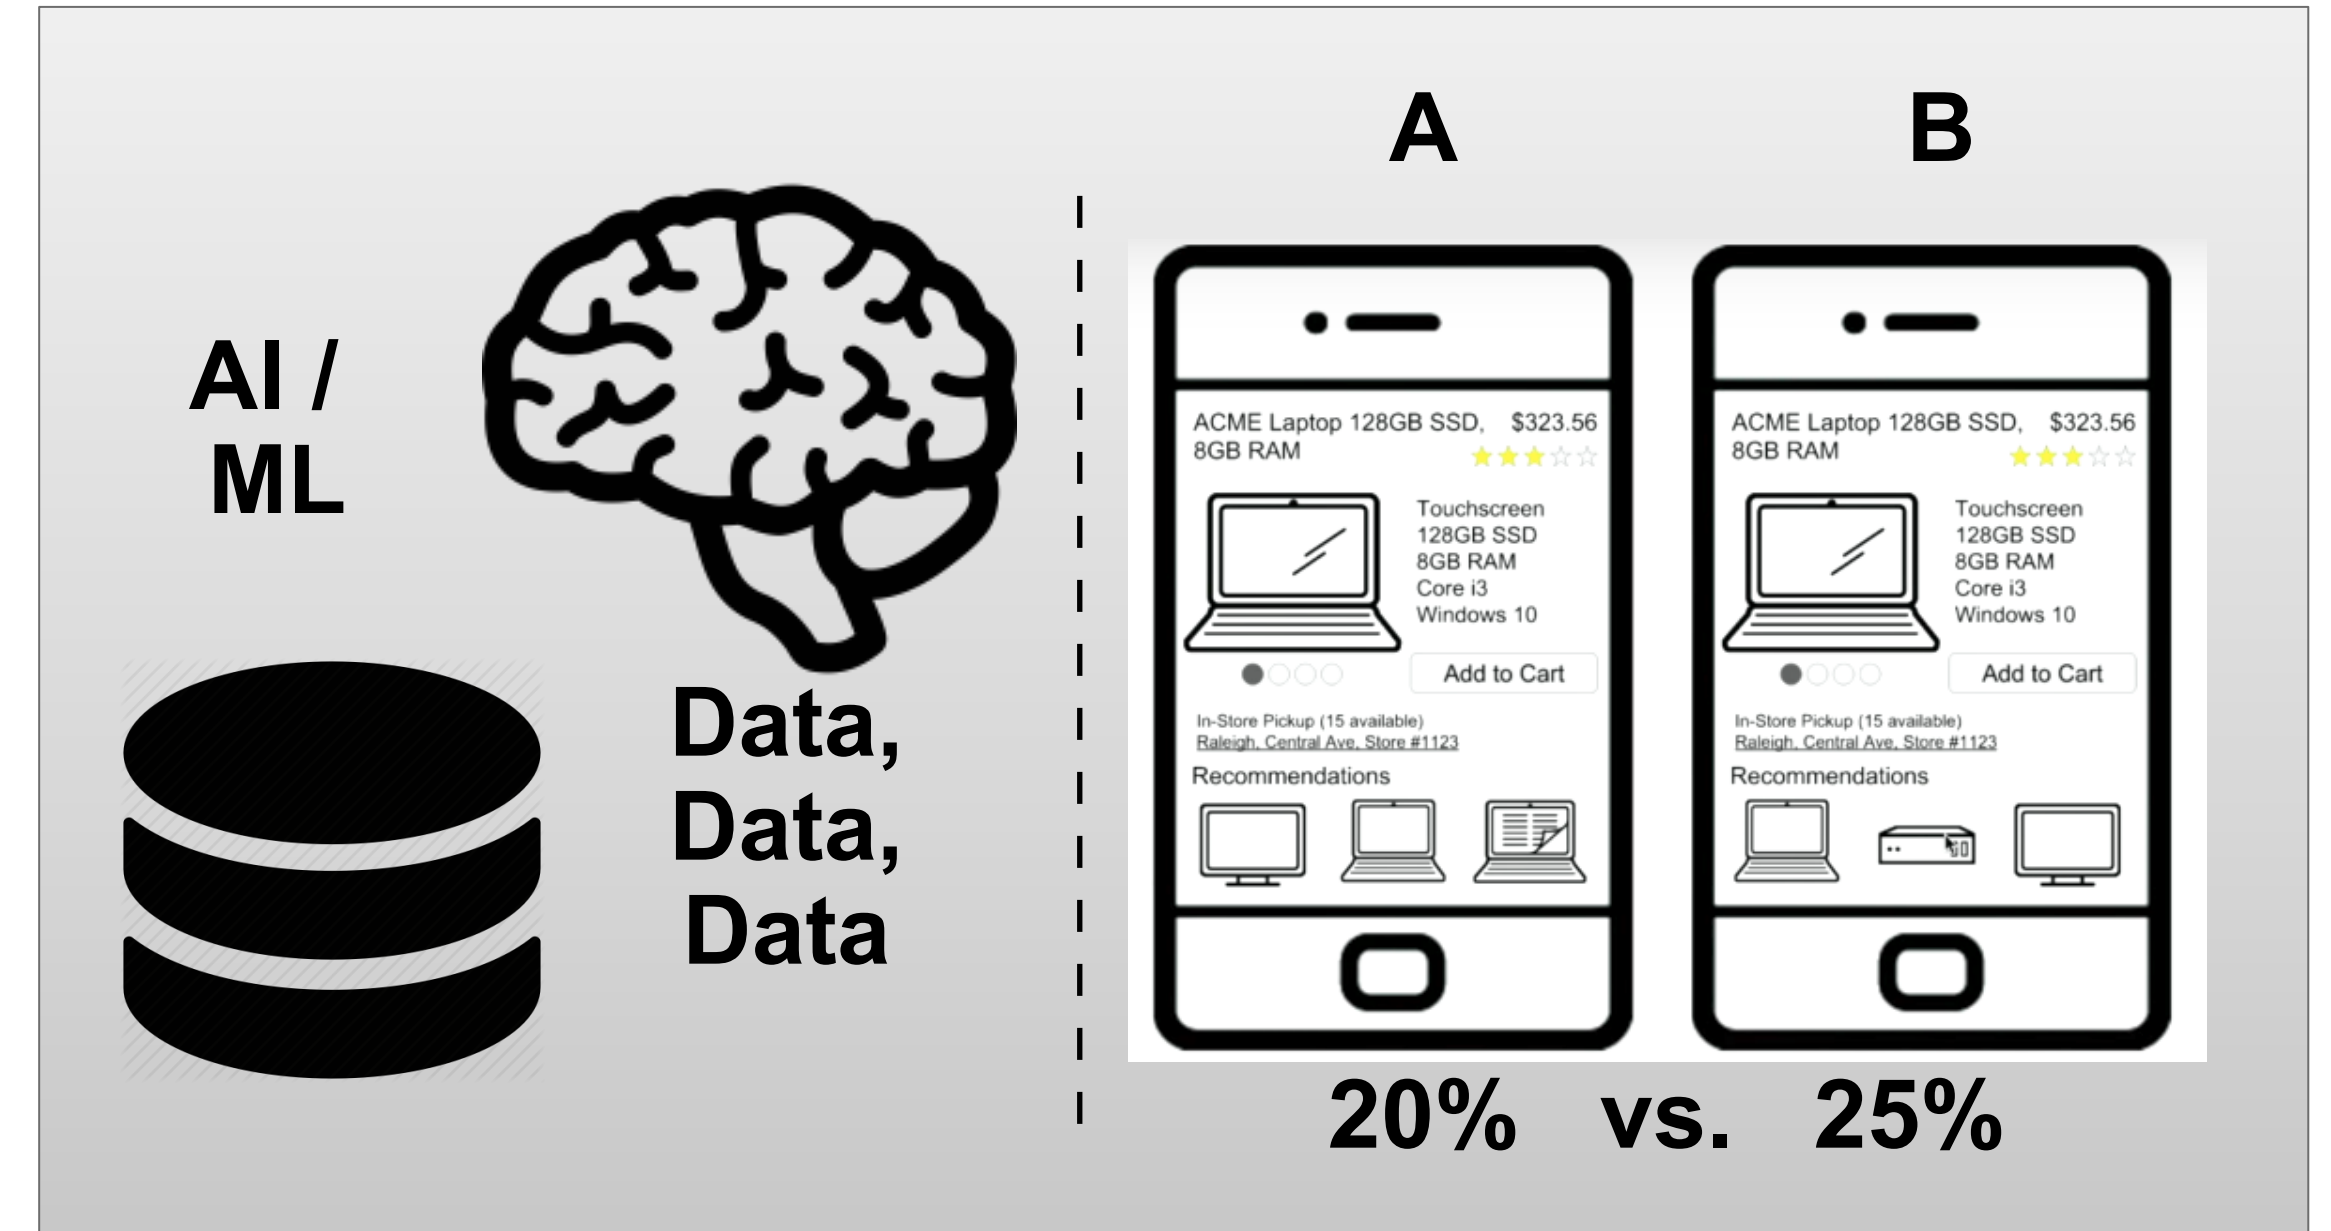
**Speed Up  
Innovation**



**Empowered  
organization**

**Move Fast,  
Break Things**

**Shorten the  
Feedback Loop**

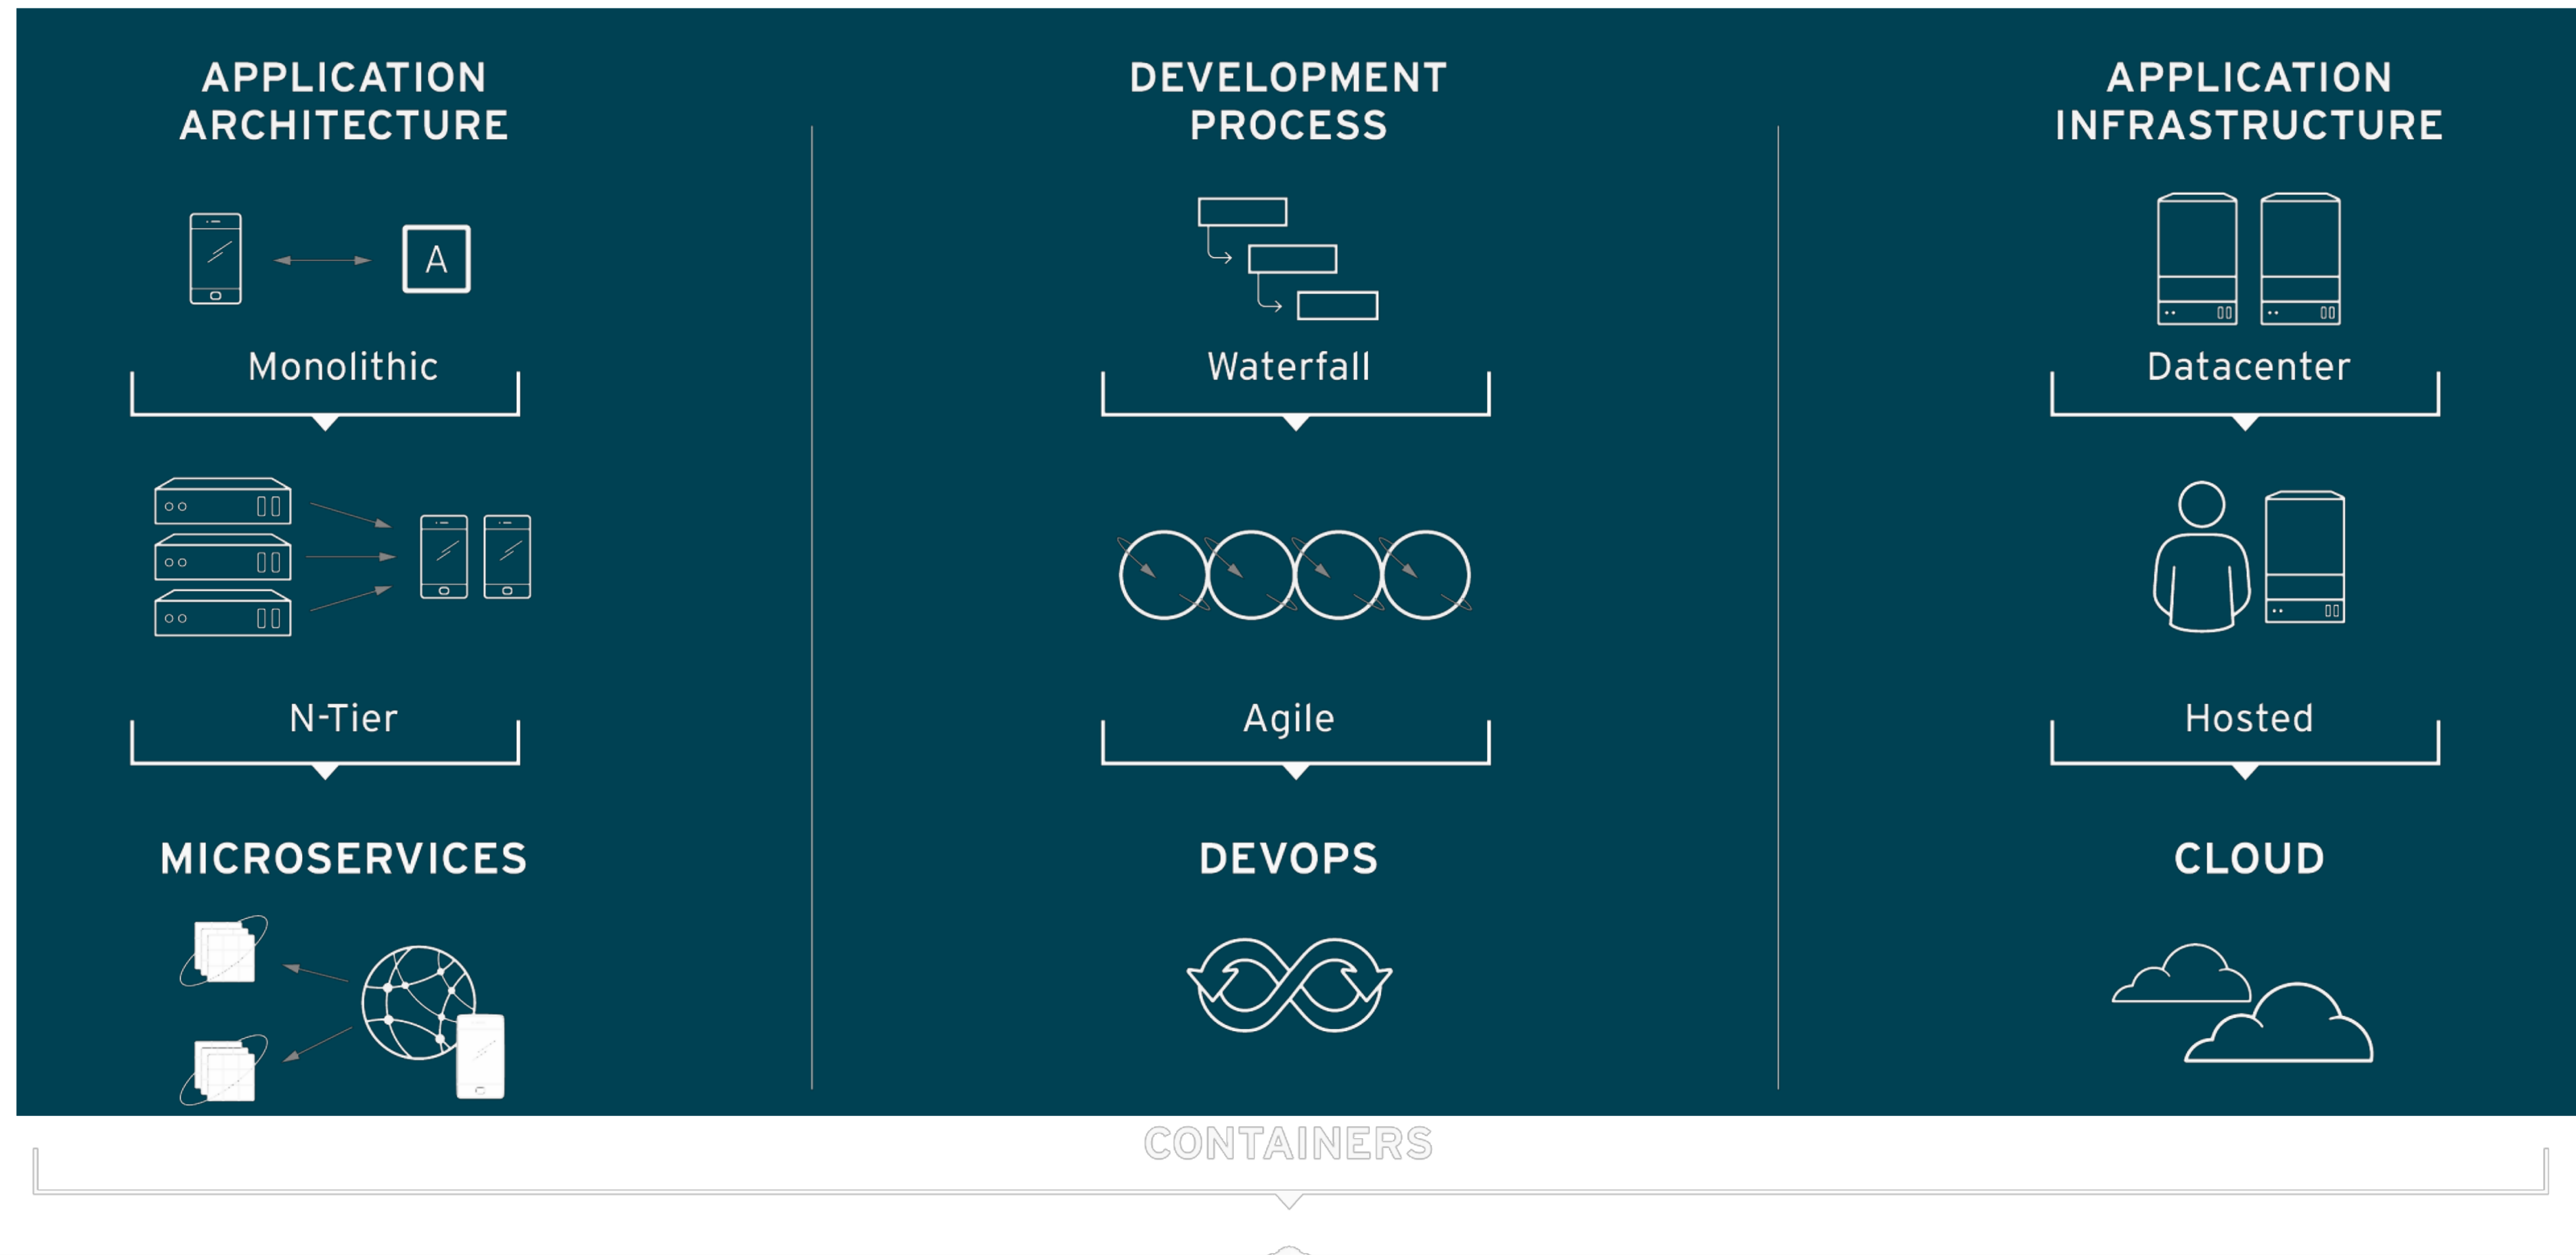


**Real-time  
data-driven  
intelligence &  
personalization**

**Culture of  
experimentation**



# IT MUST EVOLVE & KEEP UP






# SECURING THE ENTERPRISE IS HARDER THAN EVER


*The way we develop, deploy and manage IT is changing dramatically led by DevOps, Cloud Native Applications, and Hybrid Cloud*



Menacing threat  
landscape




Dissolving  
security  
perimeter



Software-defined  
infrastructure



Cloud  
computing



Applications &  
devices outside of  
IT control

TRADITIONAL NETWORK-BASED DEFENSES ARE NO LONGER ENOUGH




# DEVSECOPS

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



**Security  
Automation**



**Process  
Optimization**




**Continuous  
Security  
Improvement**

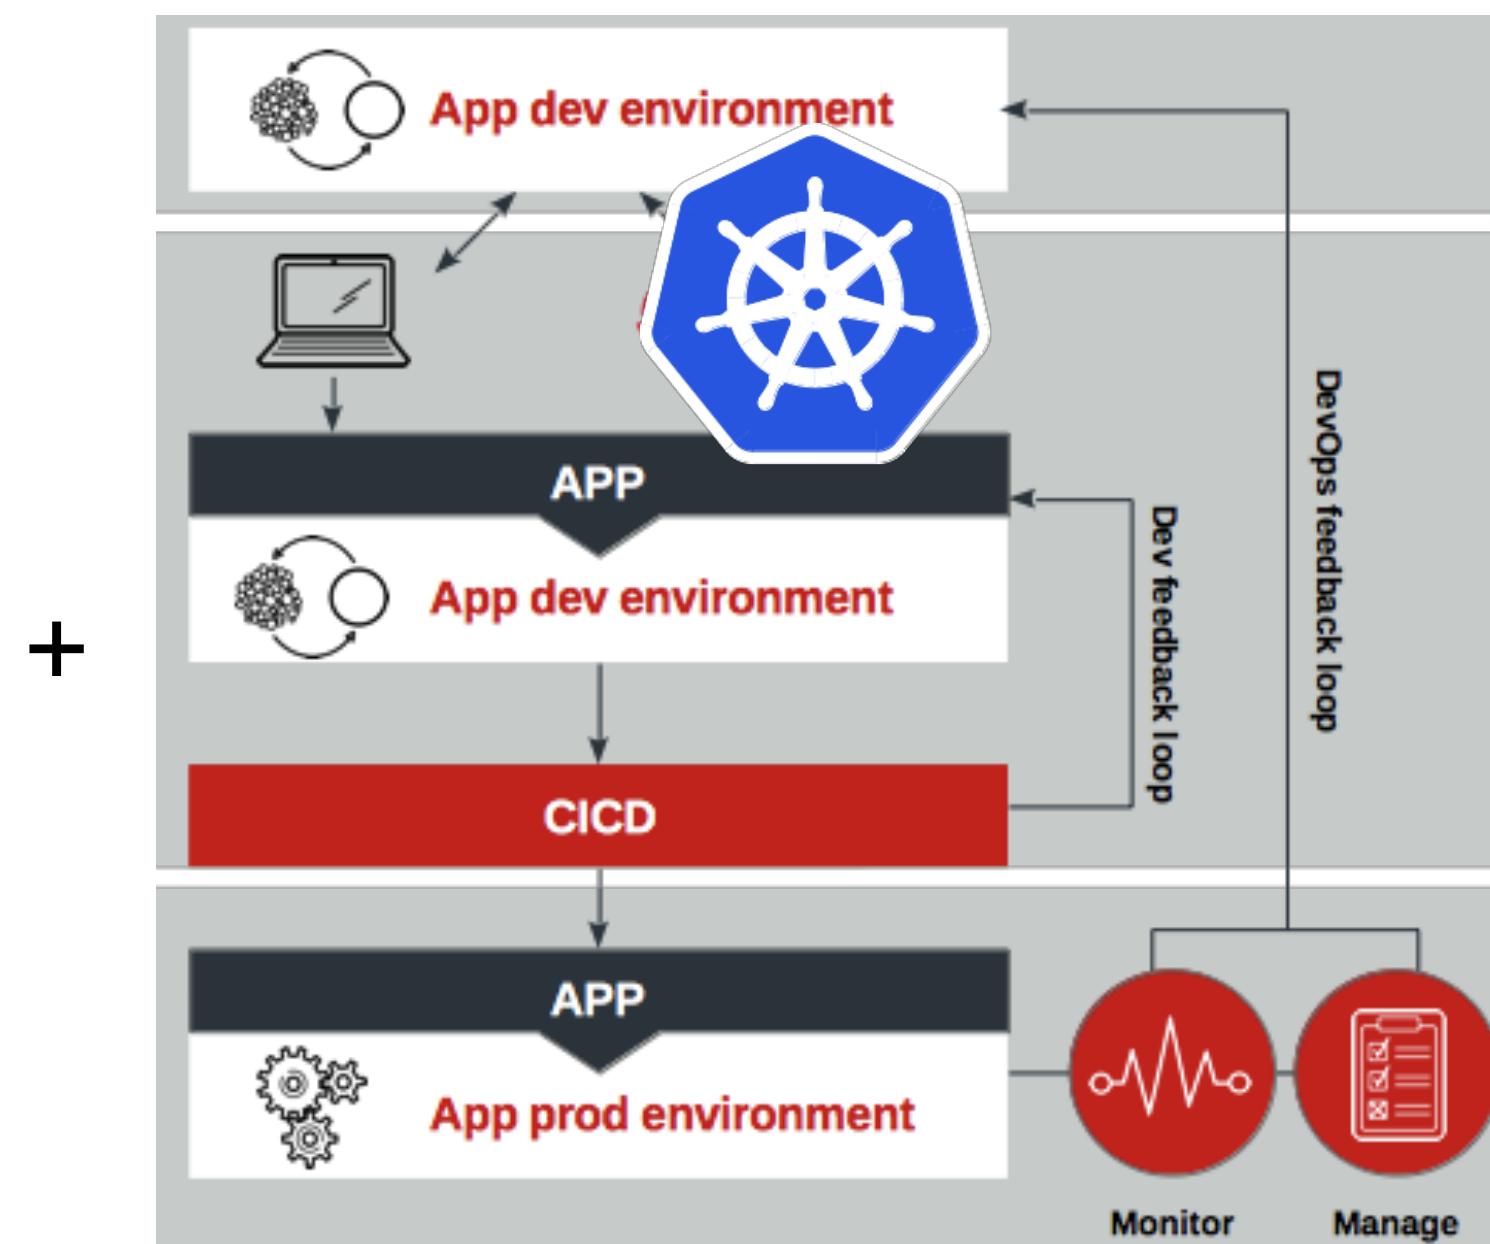


# DEVSECOPS

## Culture

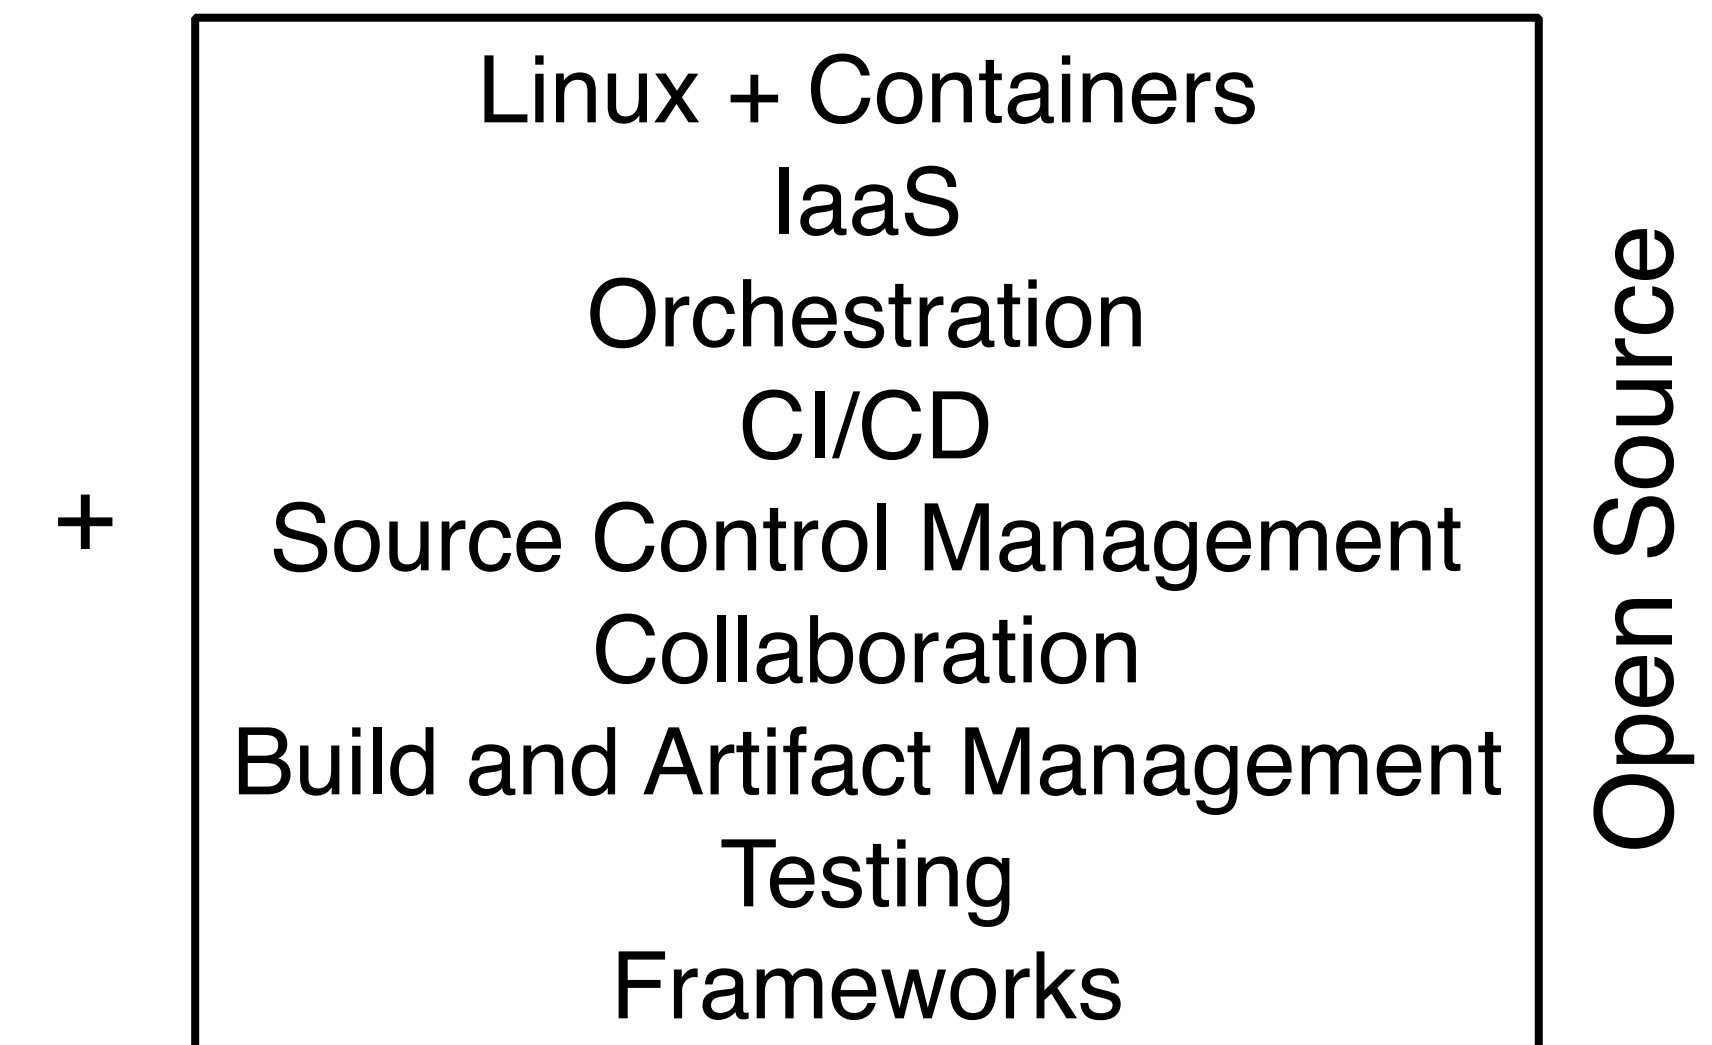


## Process



## Technology

Cloud Native Applications



Hybrid Cloud

## Security





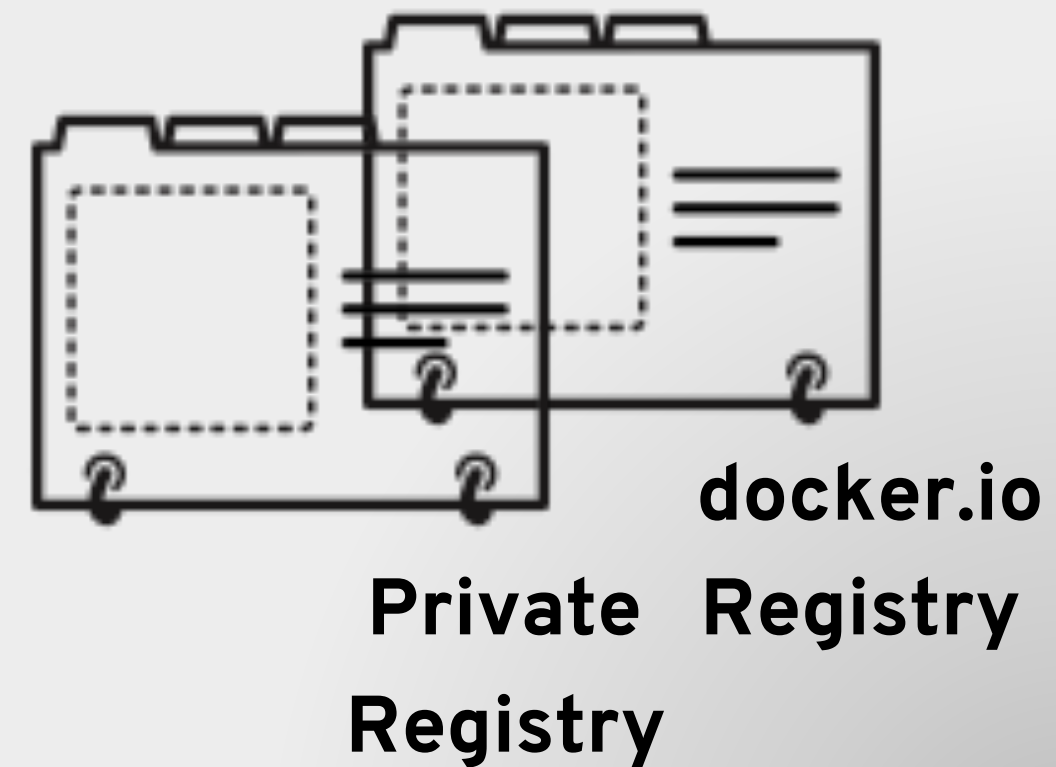
# CONTAINERS ENABLE DEVSECOPS

## Build

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


Build file

## Ship



Container  
Image

## Run




Container  
Instance




# CONTAINERS AT SCALE



Scheduling



Lifecycle & health



Discovery



Monitoring



Scaling



Persistence



Aggregation



Security




BARE METAL



VIRTUAL



PRIVATE CLOUD



PUBLIC CLOUD





# kubernetes

← **Security Platform** →



BARE METAL



VIRTUAL



PRIVATE CLOUD



PUBLIC CLOUD

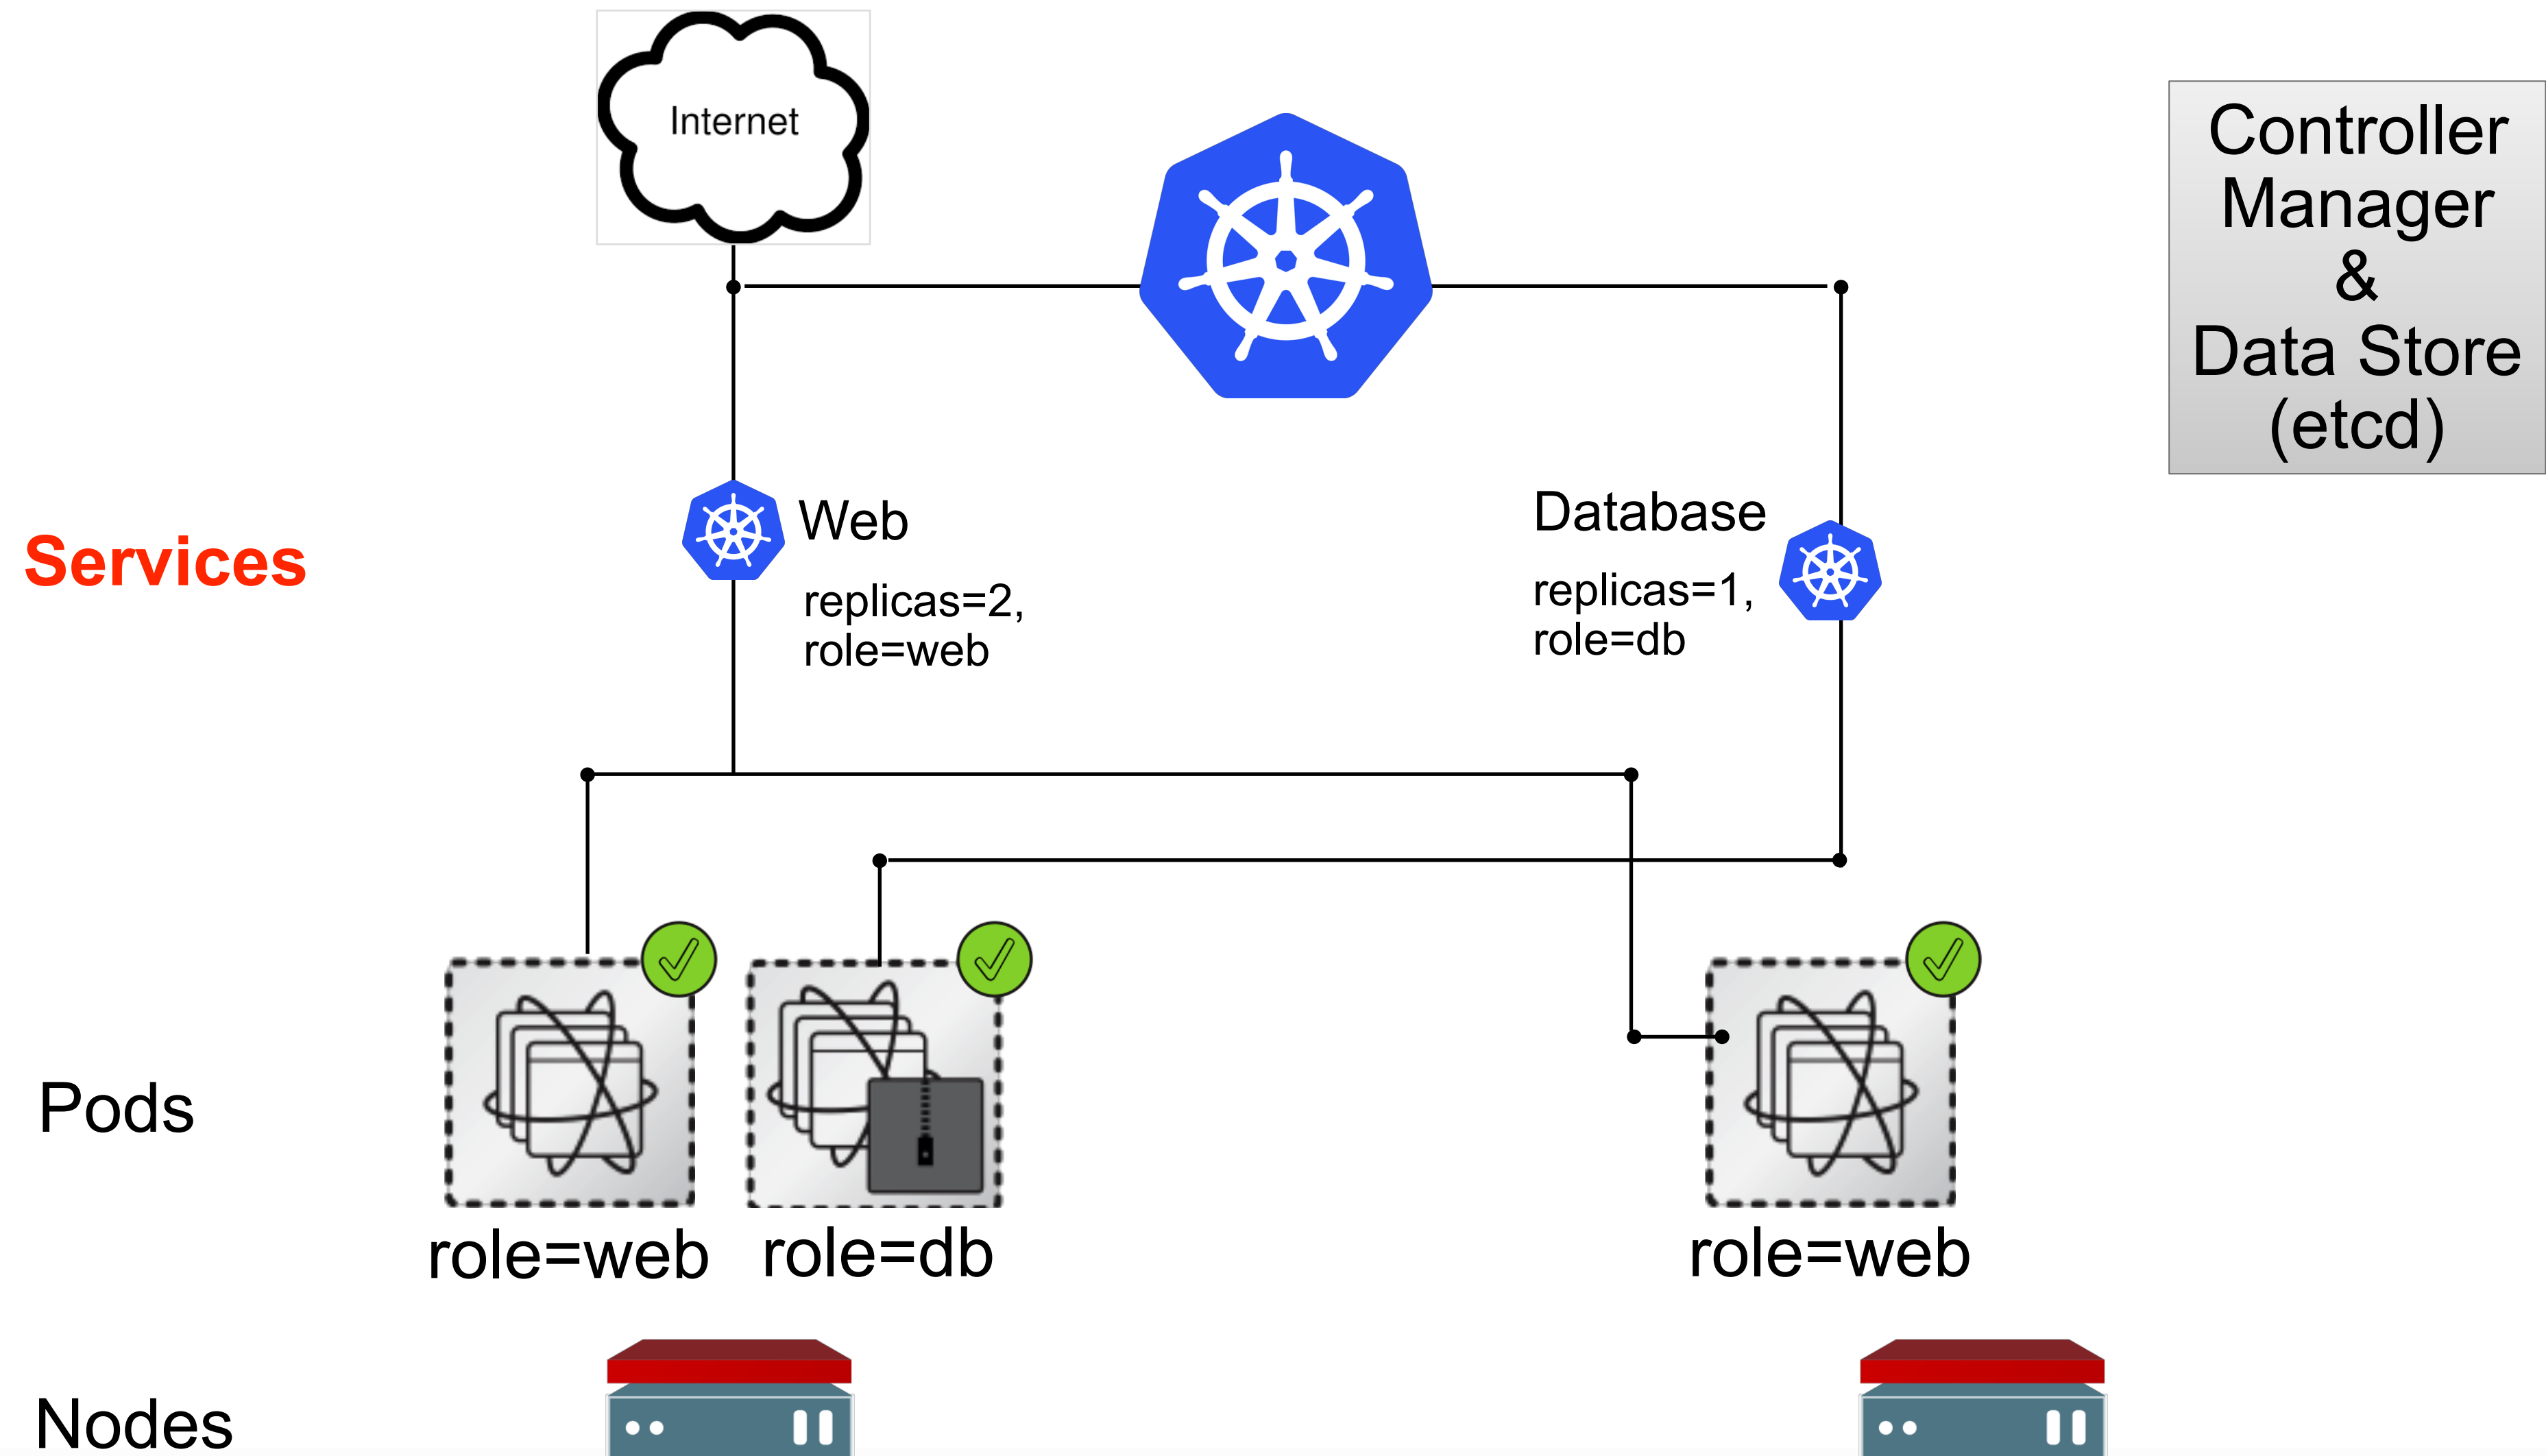


# AUTOMATION



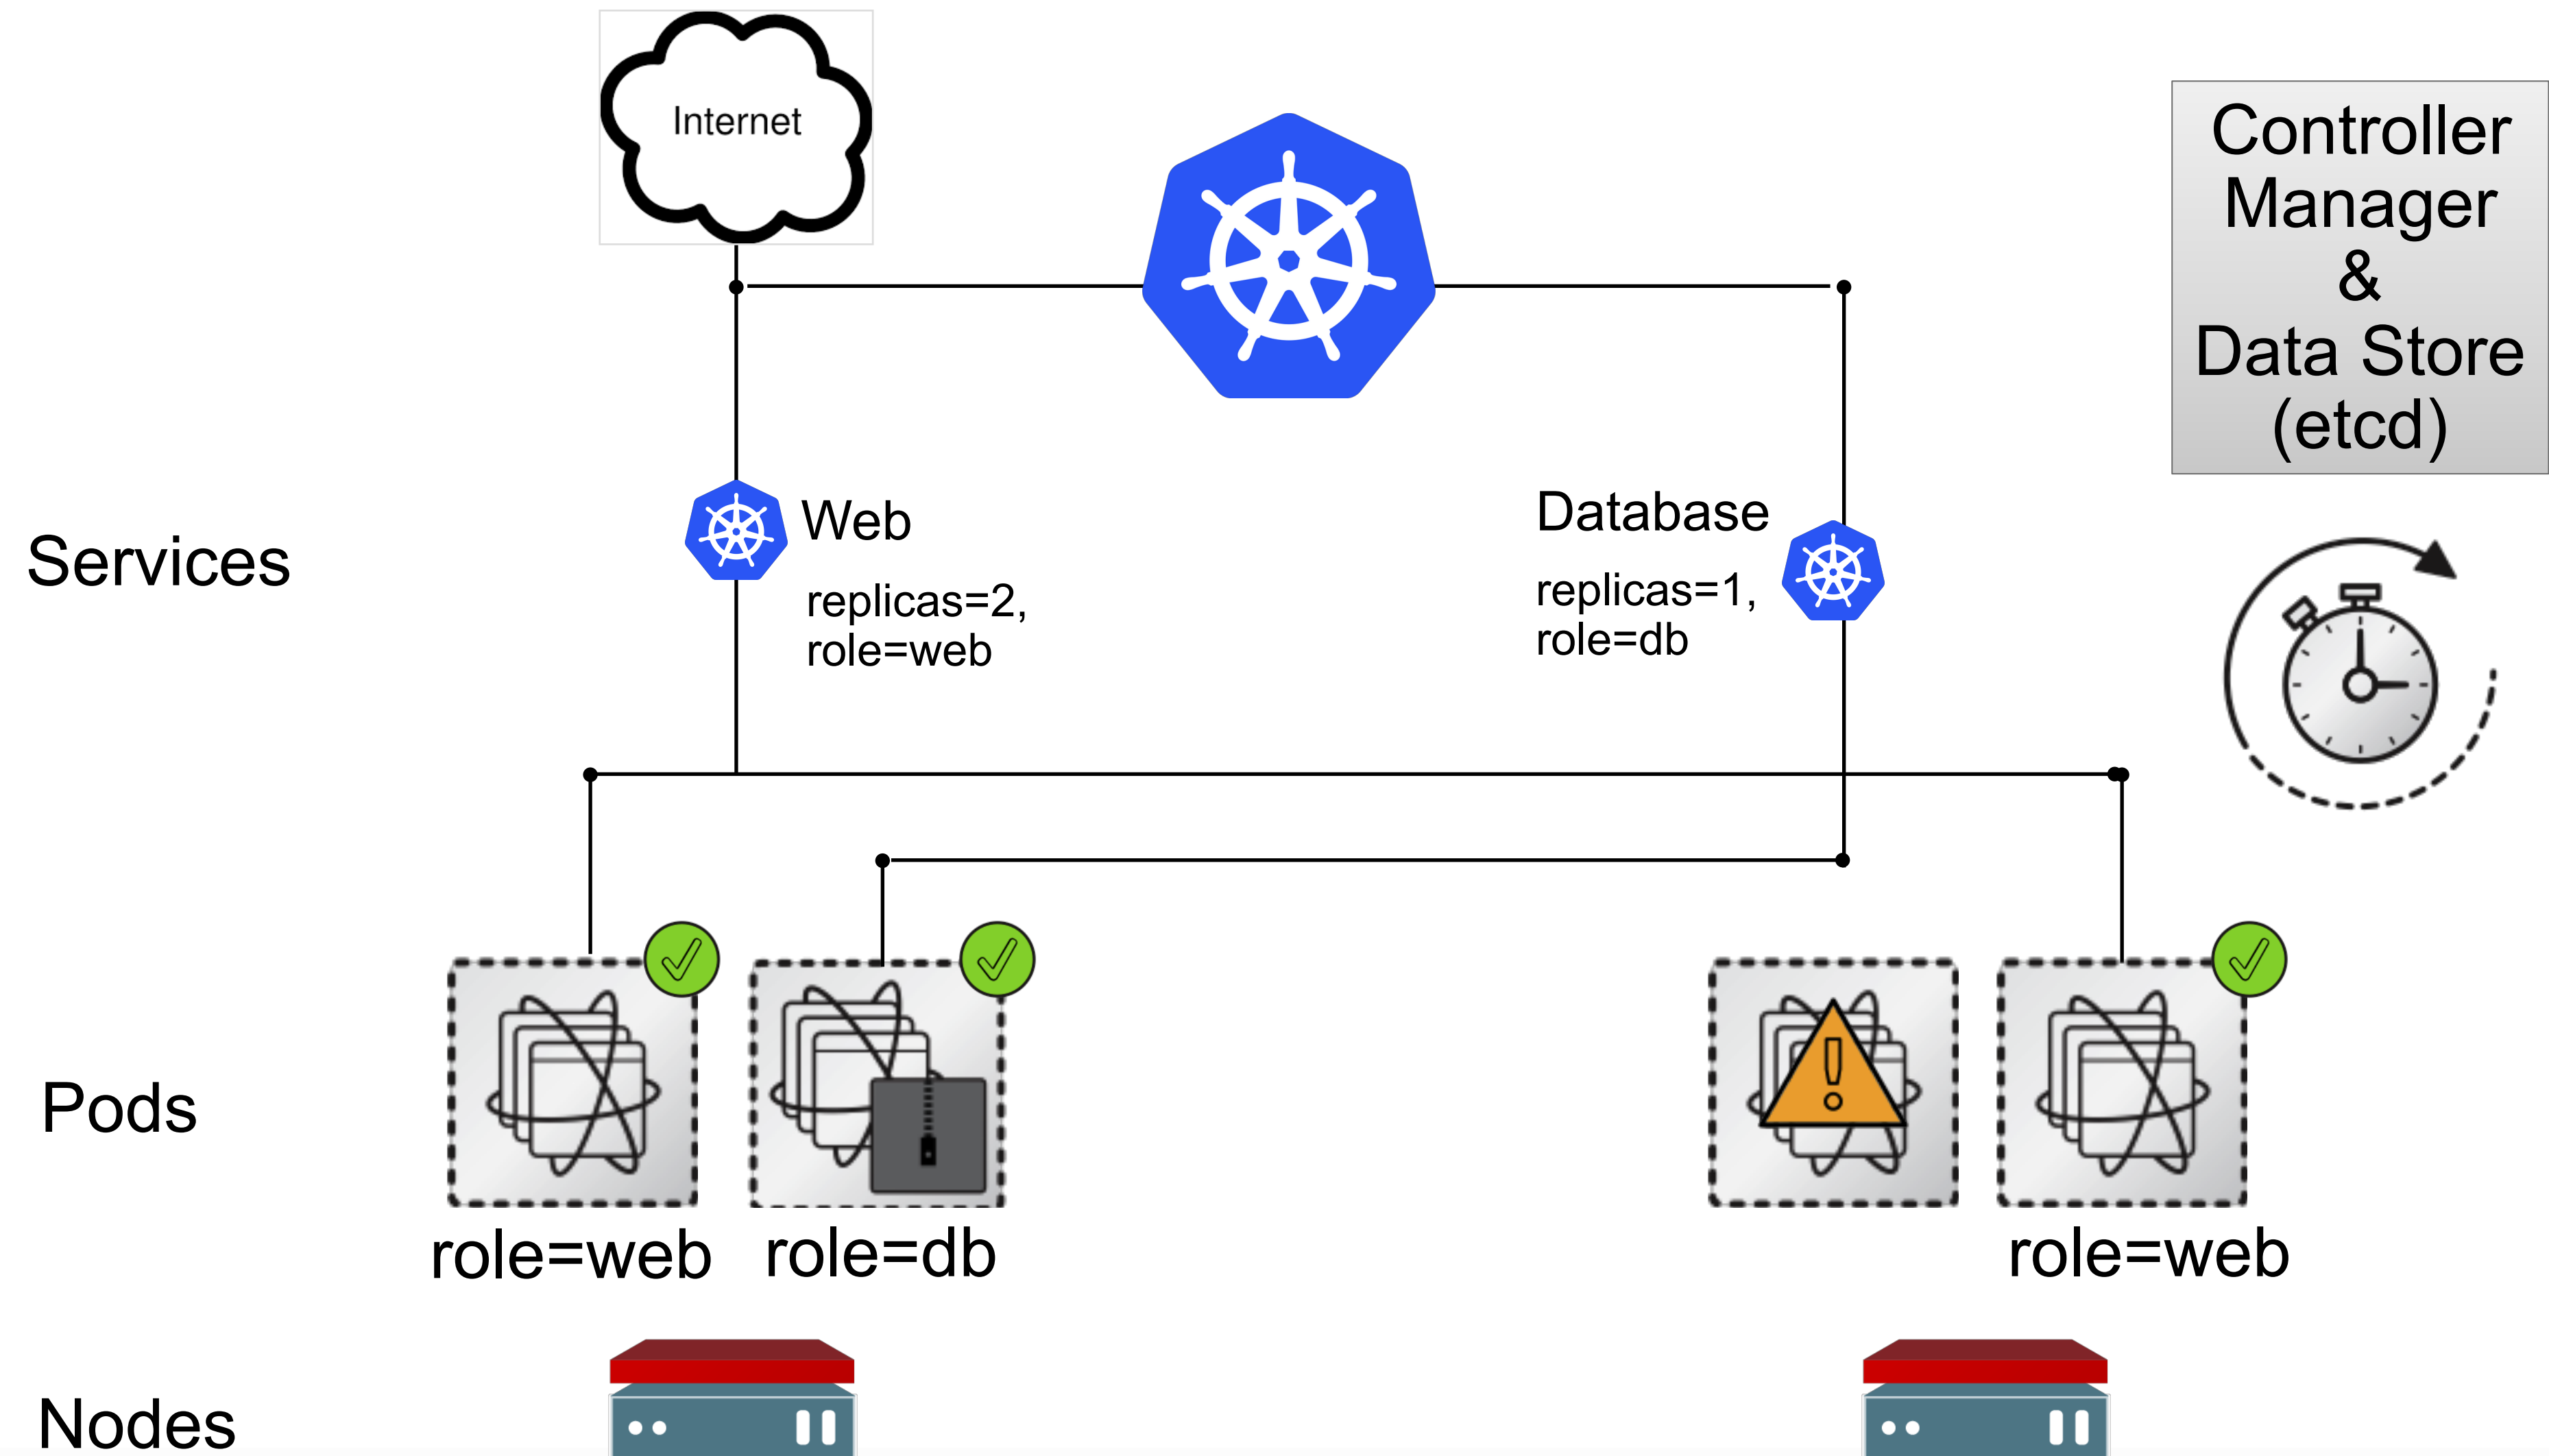
# ORCHESTRATION

Deployment, Declarative



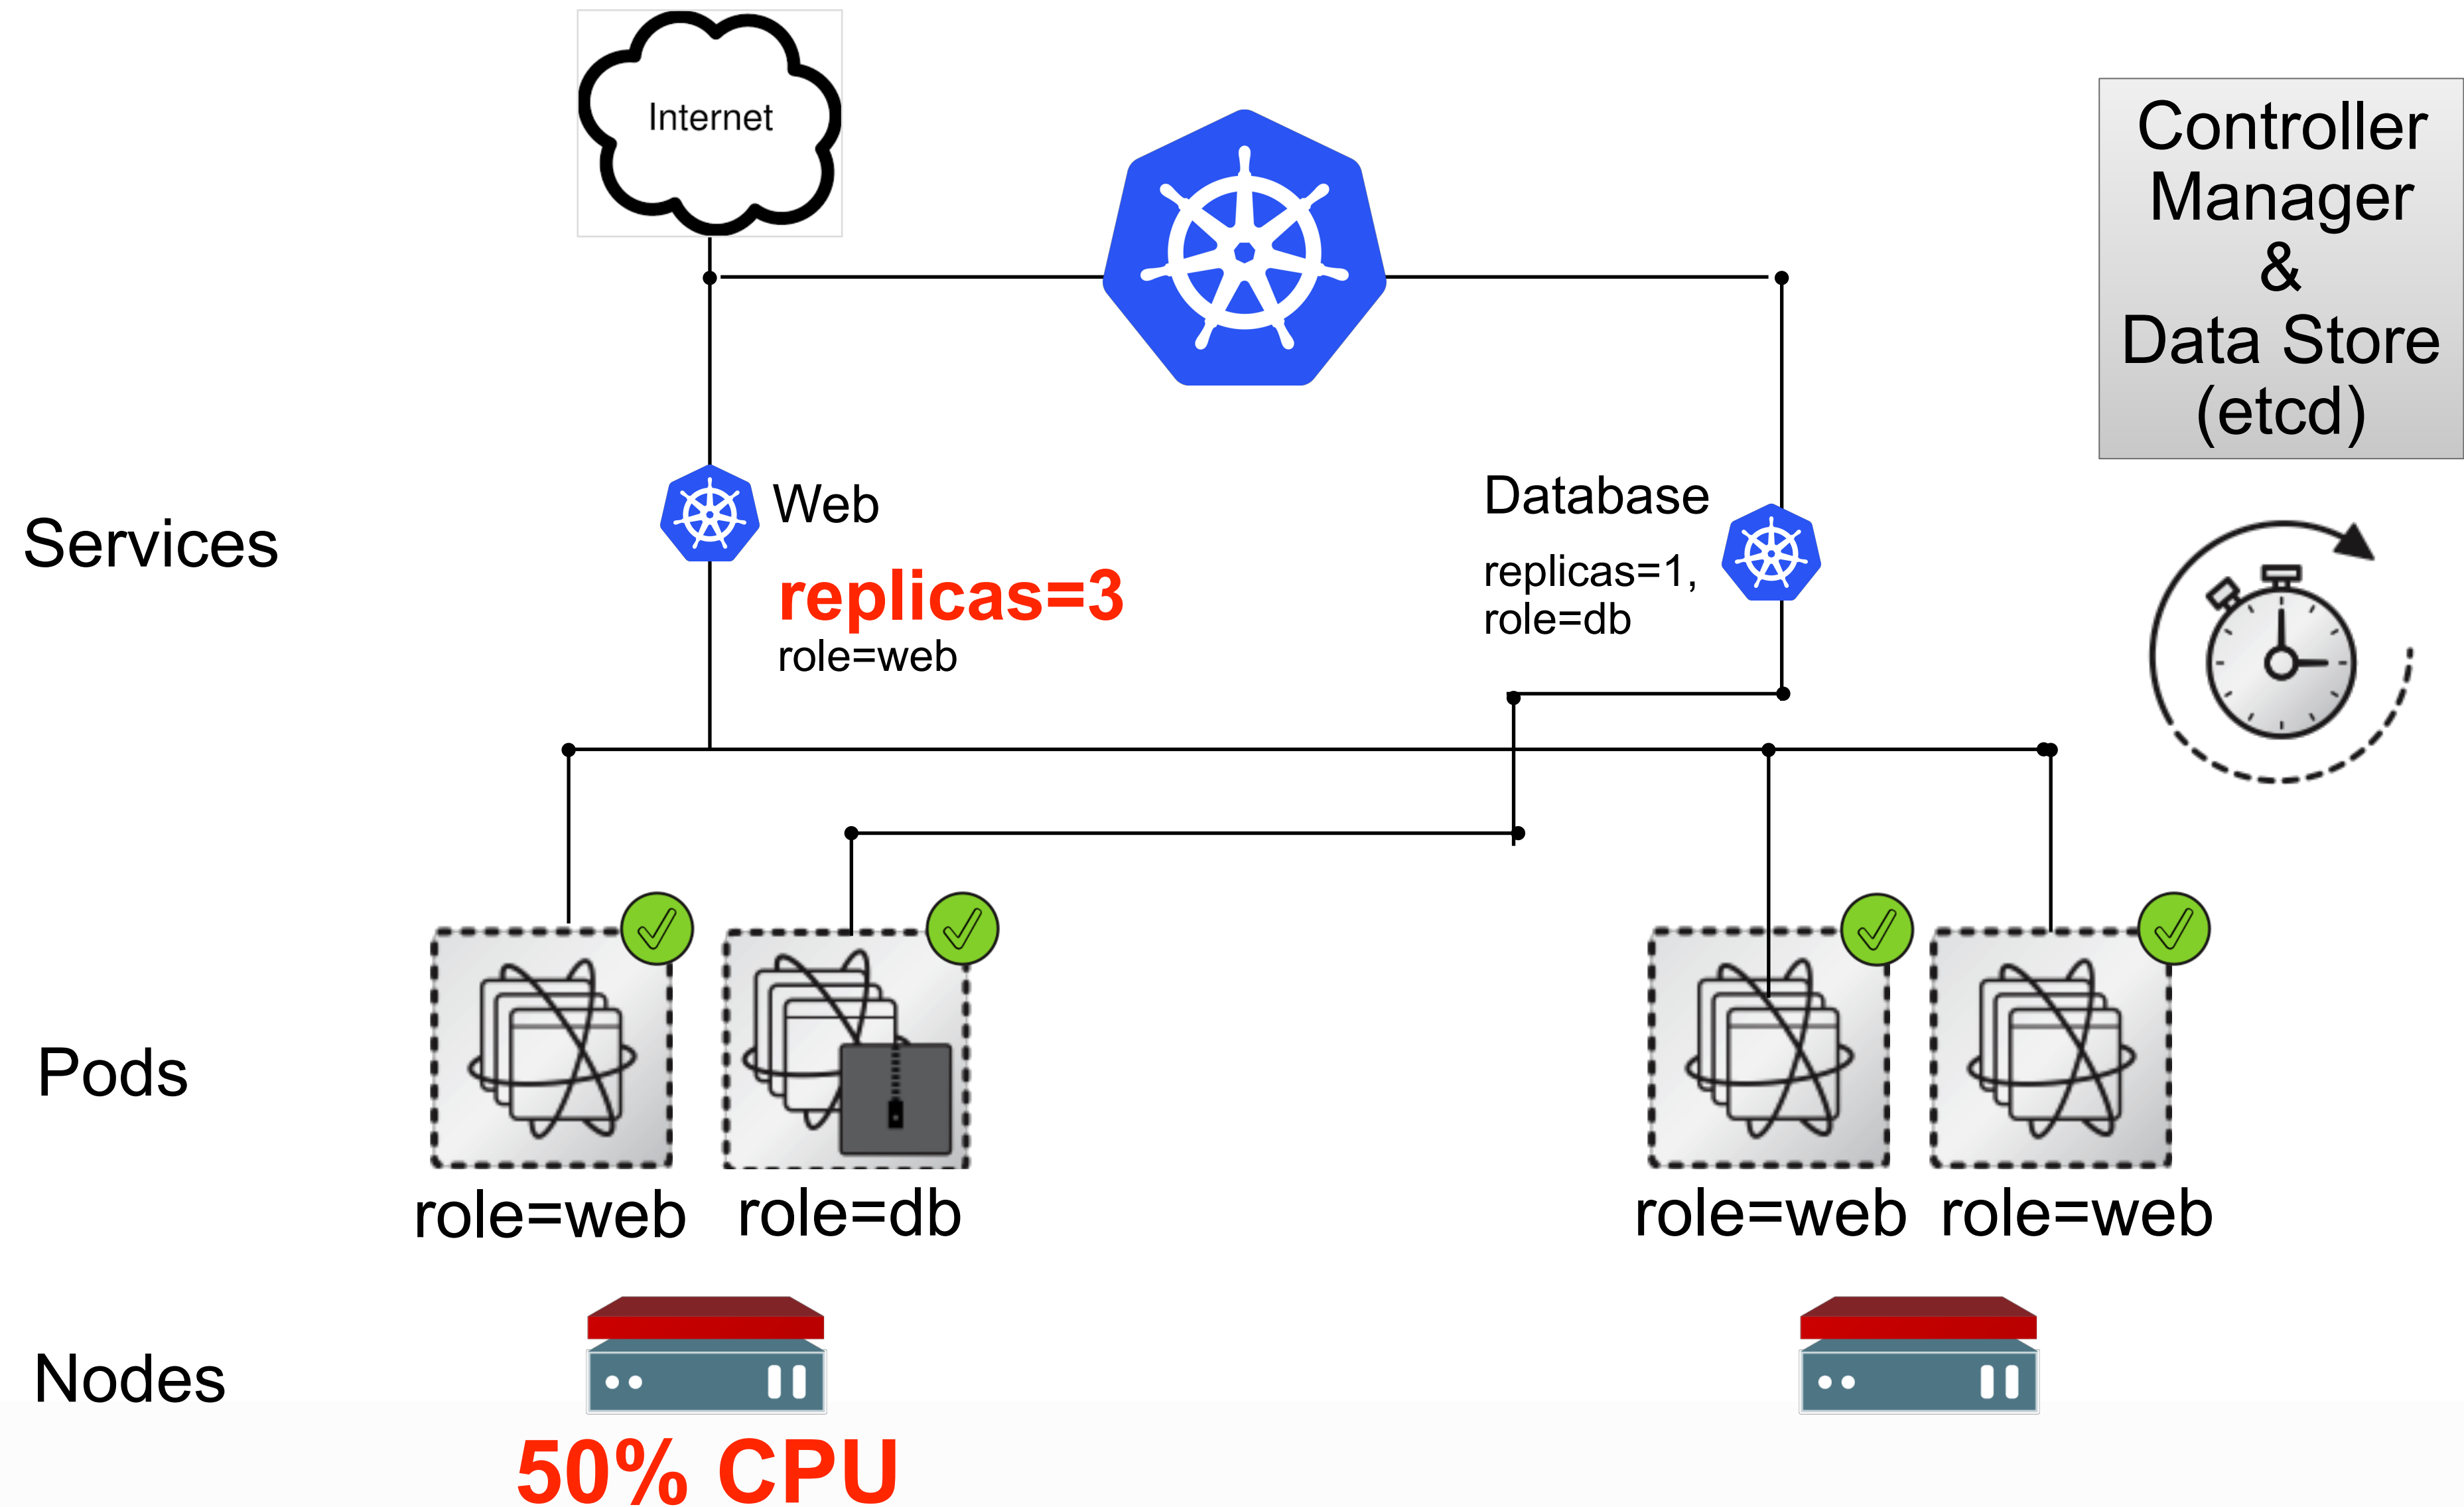


# HEALTH CHECK






# AUTO-SCALE







# SECURING YOUR CONTAINER ENVIRONMENT




**Images**




**Builds**




**Registry**



**Container  
host**



**CI/CD**




**Network  
isolation**




**Monitoring &  
Logging**



**Storage**



**API & Platform  
access**



**Federated  
clusters**

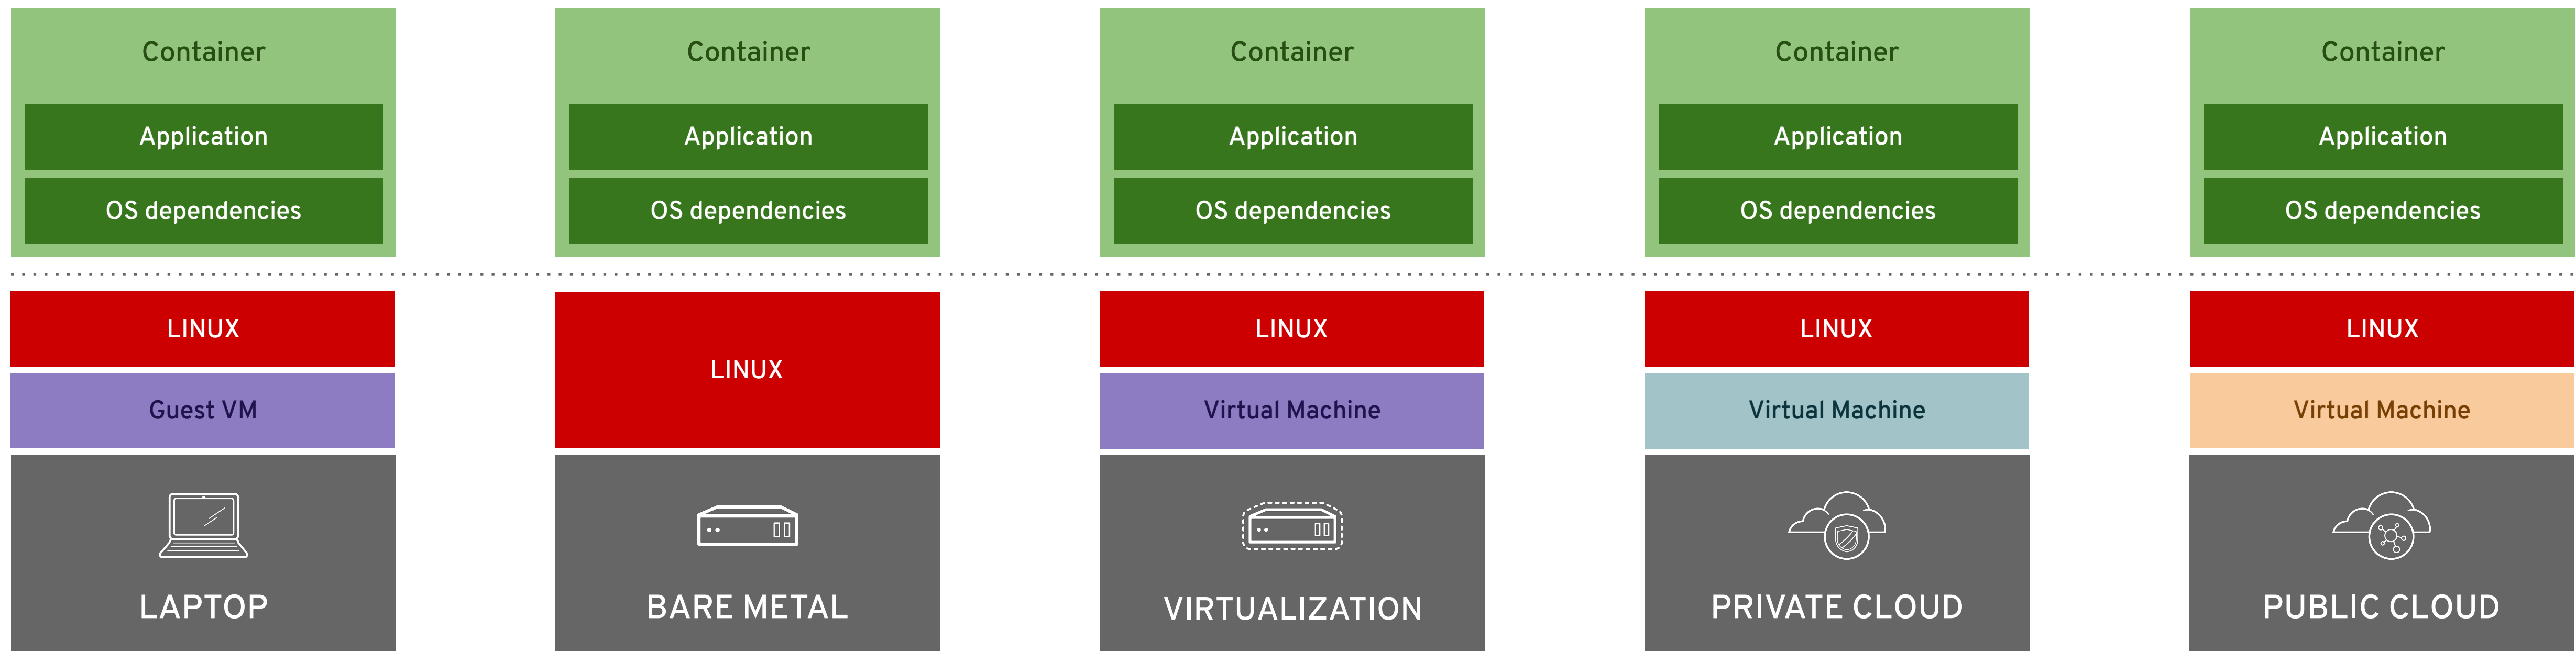


# CONTAINER IMAGES



# CONTAINERS - Build Once, Deploy Anywhere


Reducing Risk and Improving Security with Improved Consistency



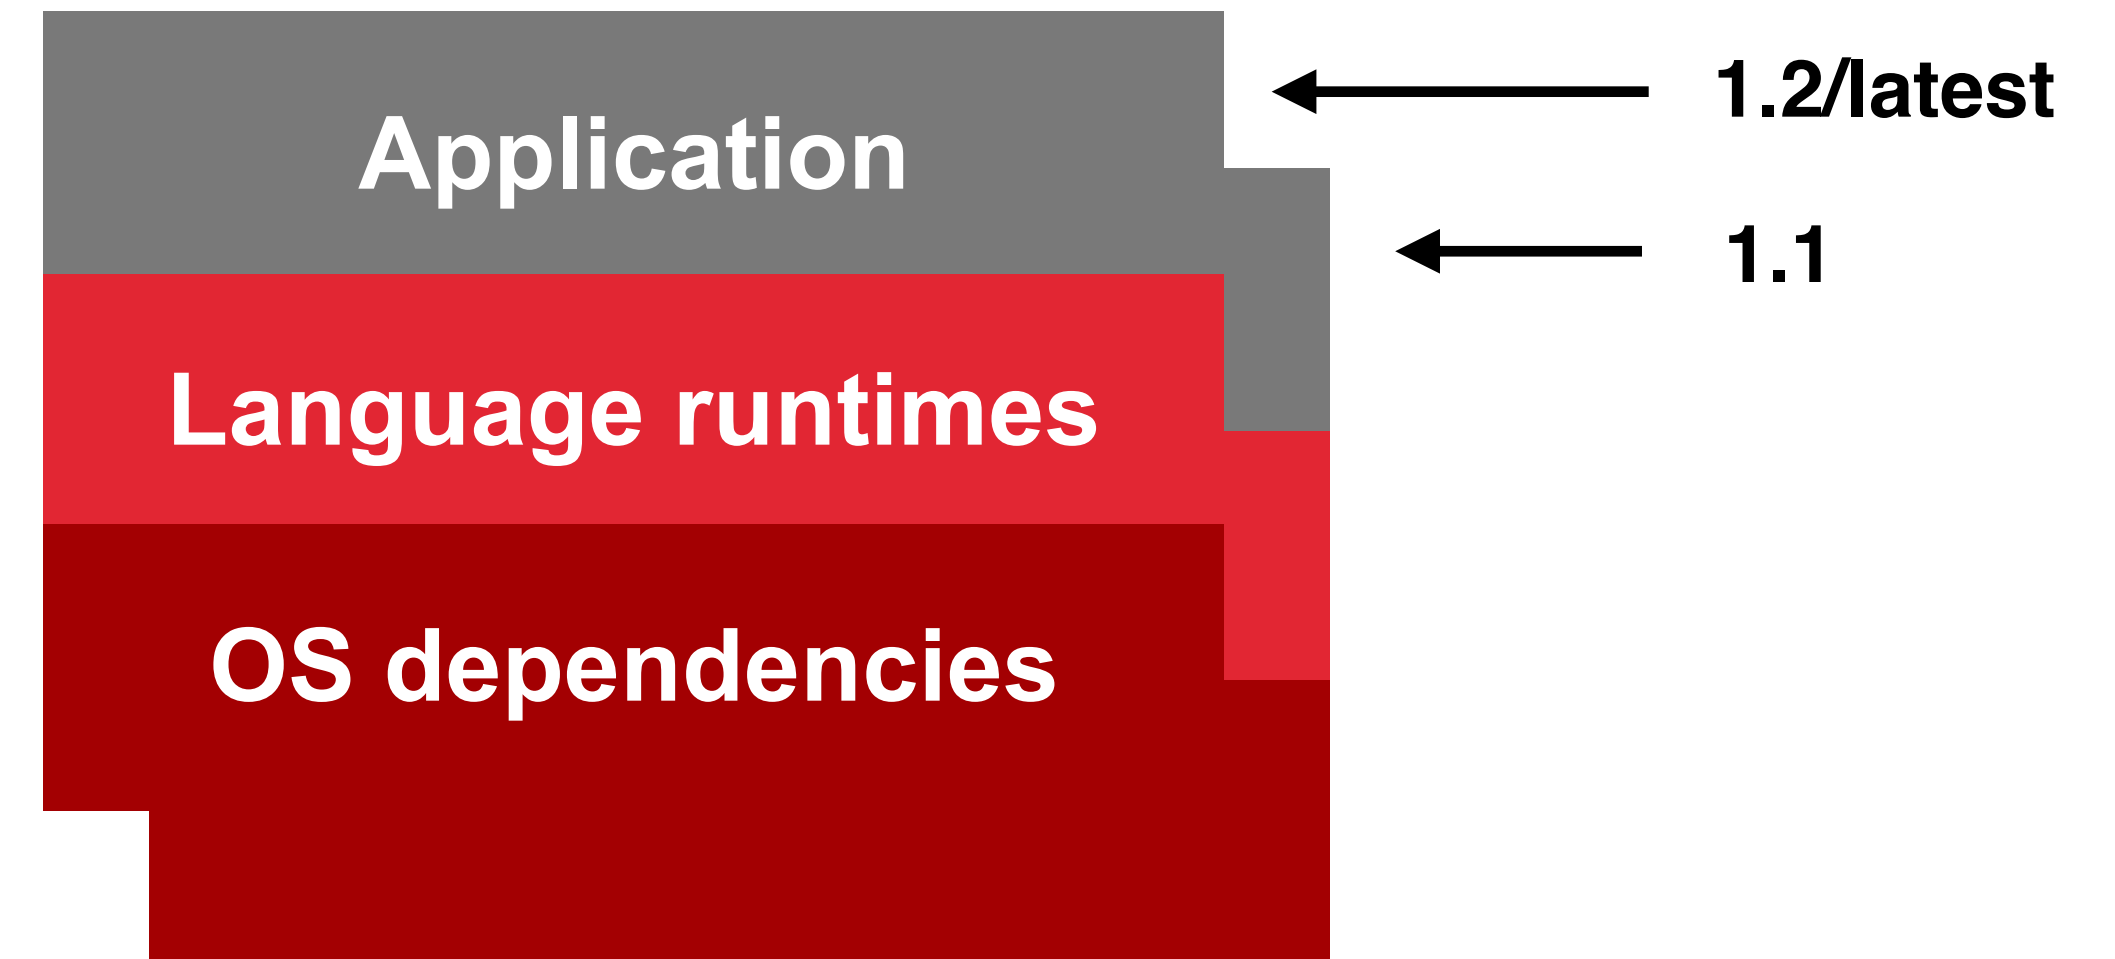


# CONTAINER IMAGE

JAR




CONTAINER IMAGE





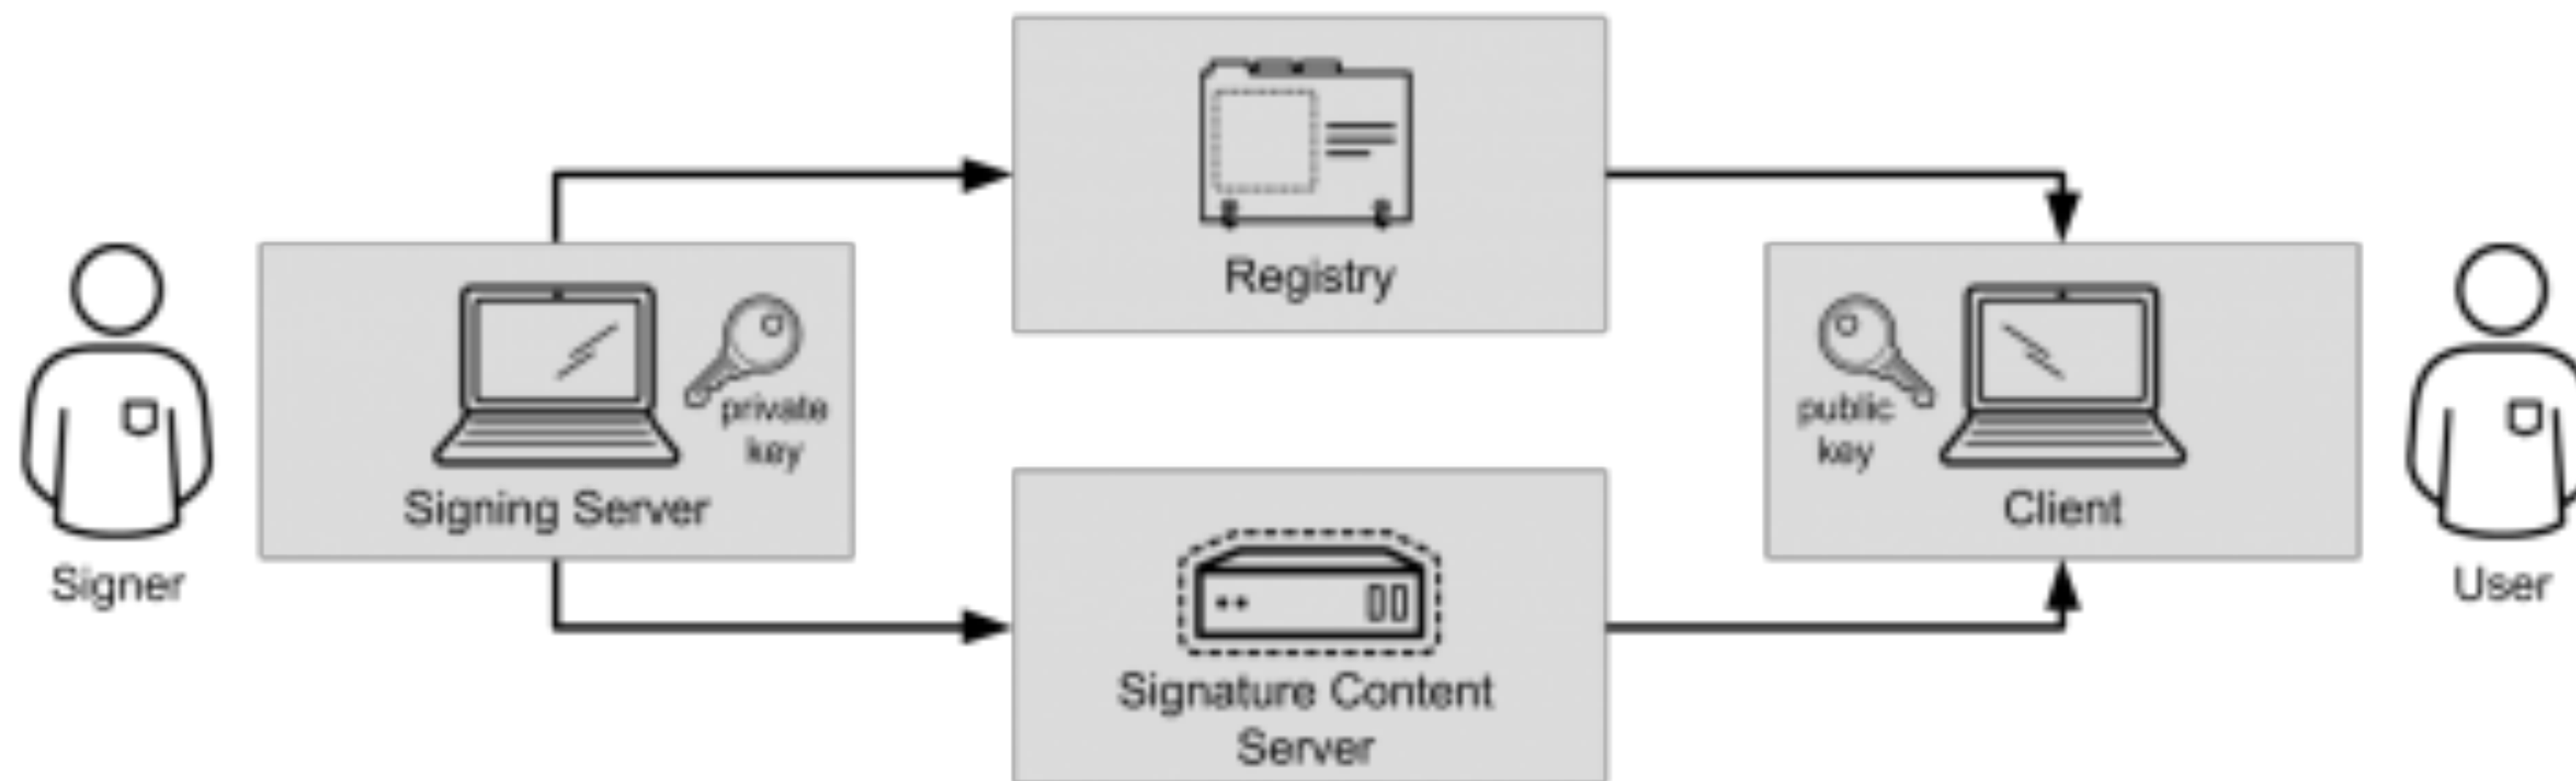
# TREAT CONTAINERS AS IMMUTABLE





# CONTAINER IMAGE SIGNING

Validate what images and version are running



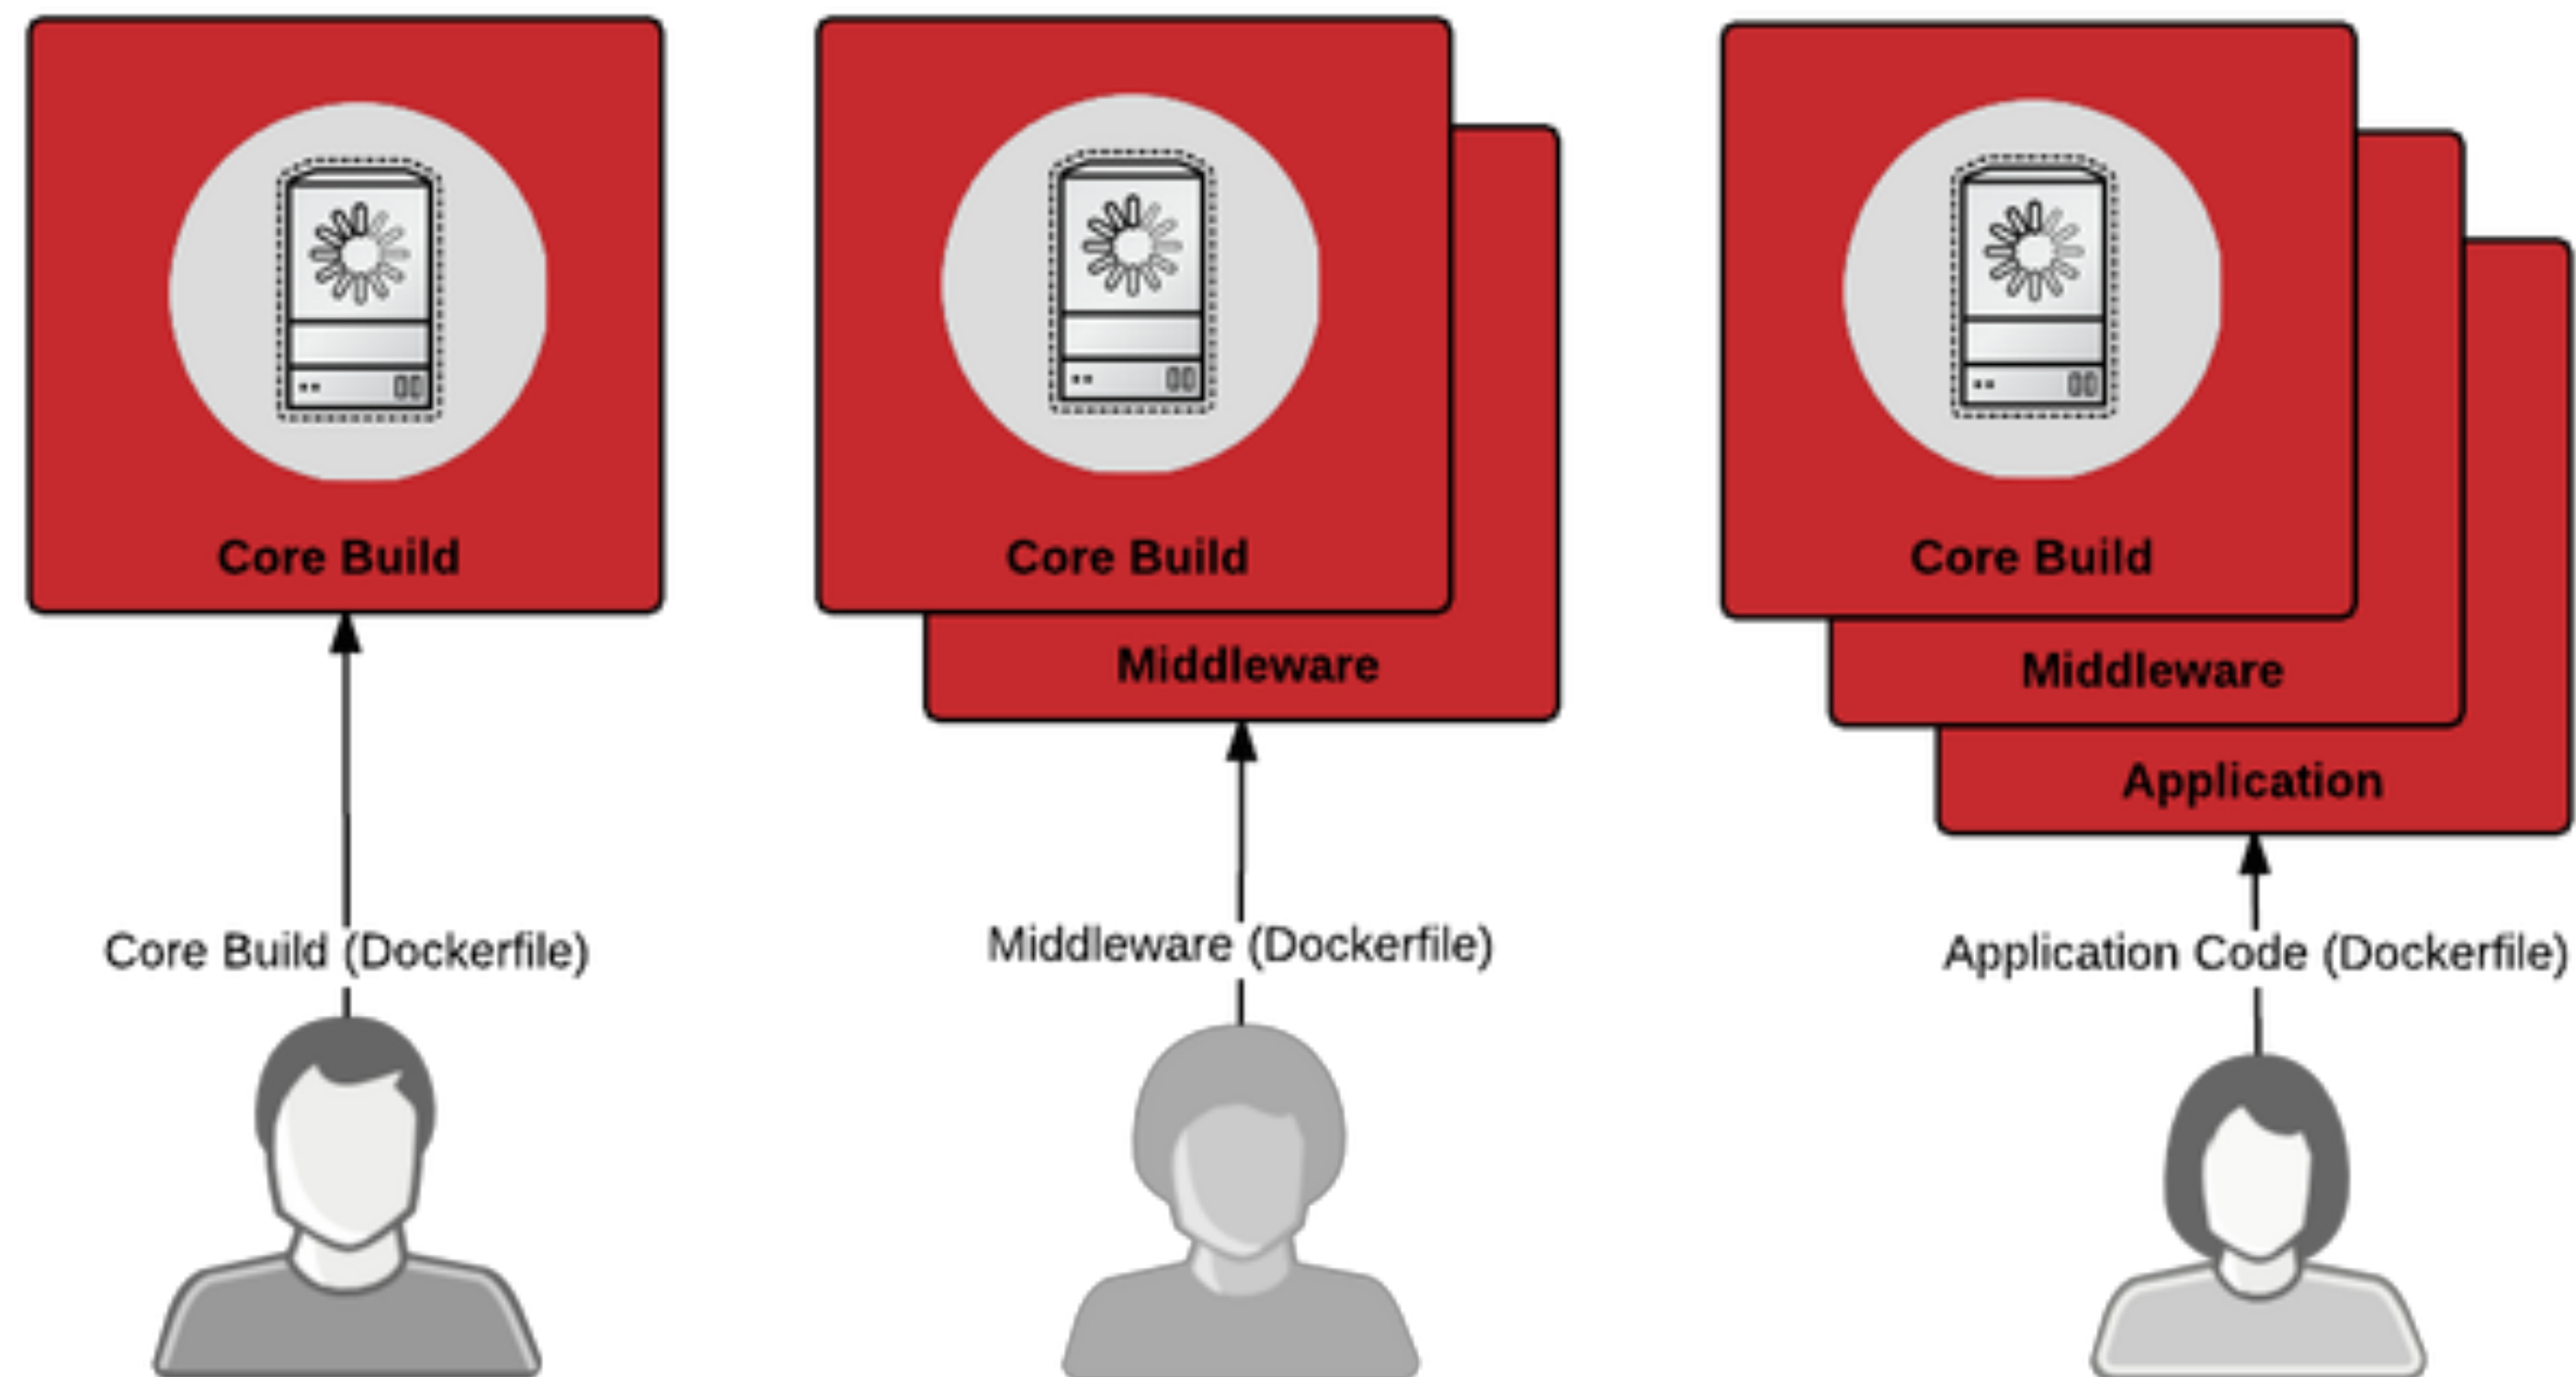
- Authenticating authorship
- Non-repudiation
- Ensuring image integrity



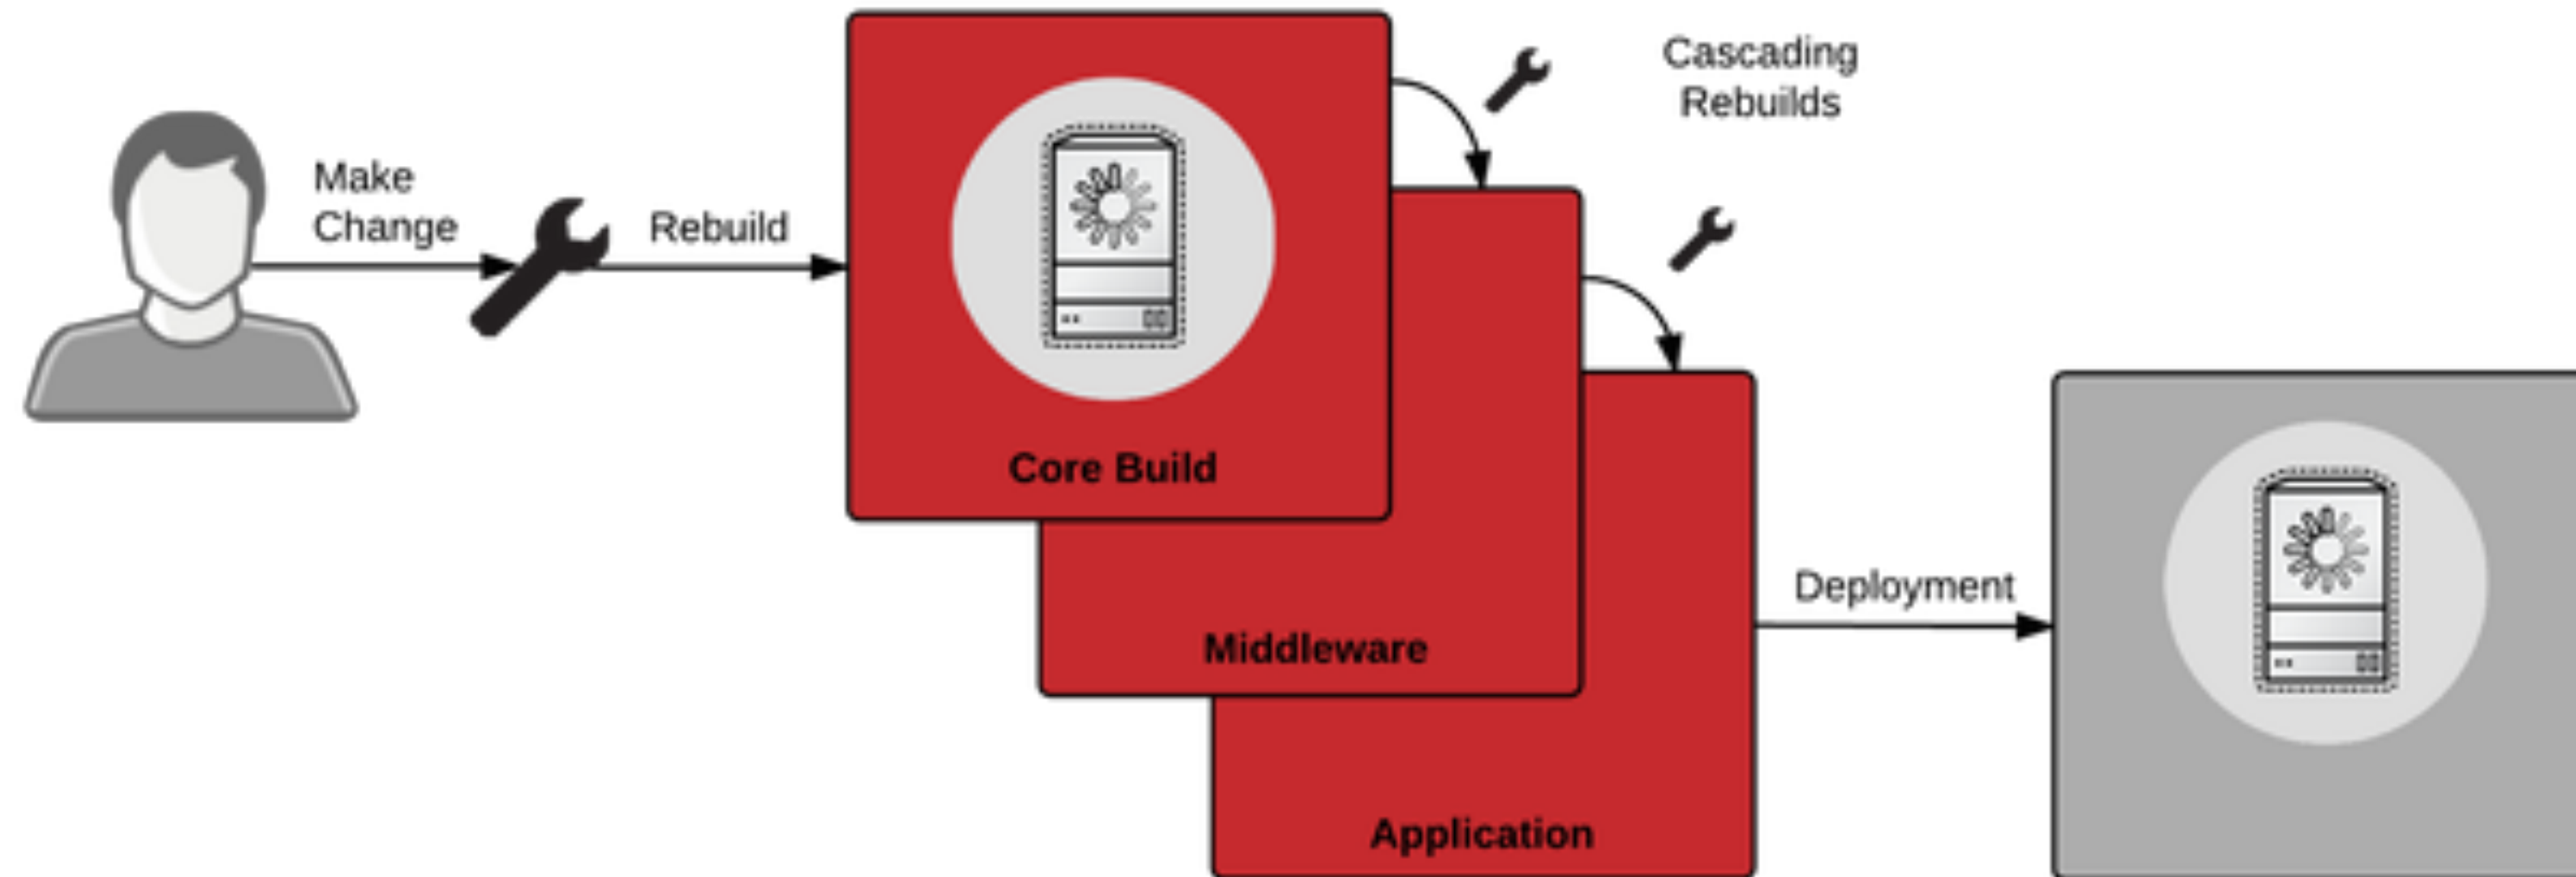
# CONTAINER BUILDS



# A CONVERGED SOFTWARE SUPPLY CHAIN



# CUSTOM SUPPLY CHAIN





# BUILD FILE BEST PRACTICES

## Build file

```
FROM registry.redhat.com/rhel7  
RUN groupadd -g 999 appuser && \  
    useradd -r -u 999 -g appuser appuser  
USER appuser  
CMD echo "Hello"
```


- Treat build file as a Blueprint
- Version control build file
- Don't login to build/configure
- Be explicit with versions, not latest
- Always list registry pulling FROM
- Specify USER, default is root
- Each Run creates a new layer

# CONTAINER REGISTRY 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>)

# 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

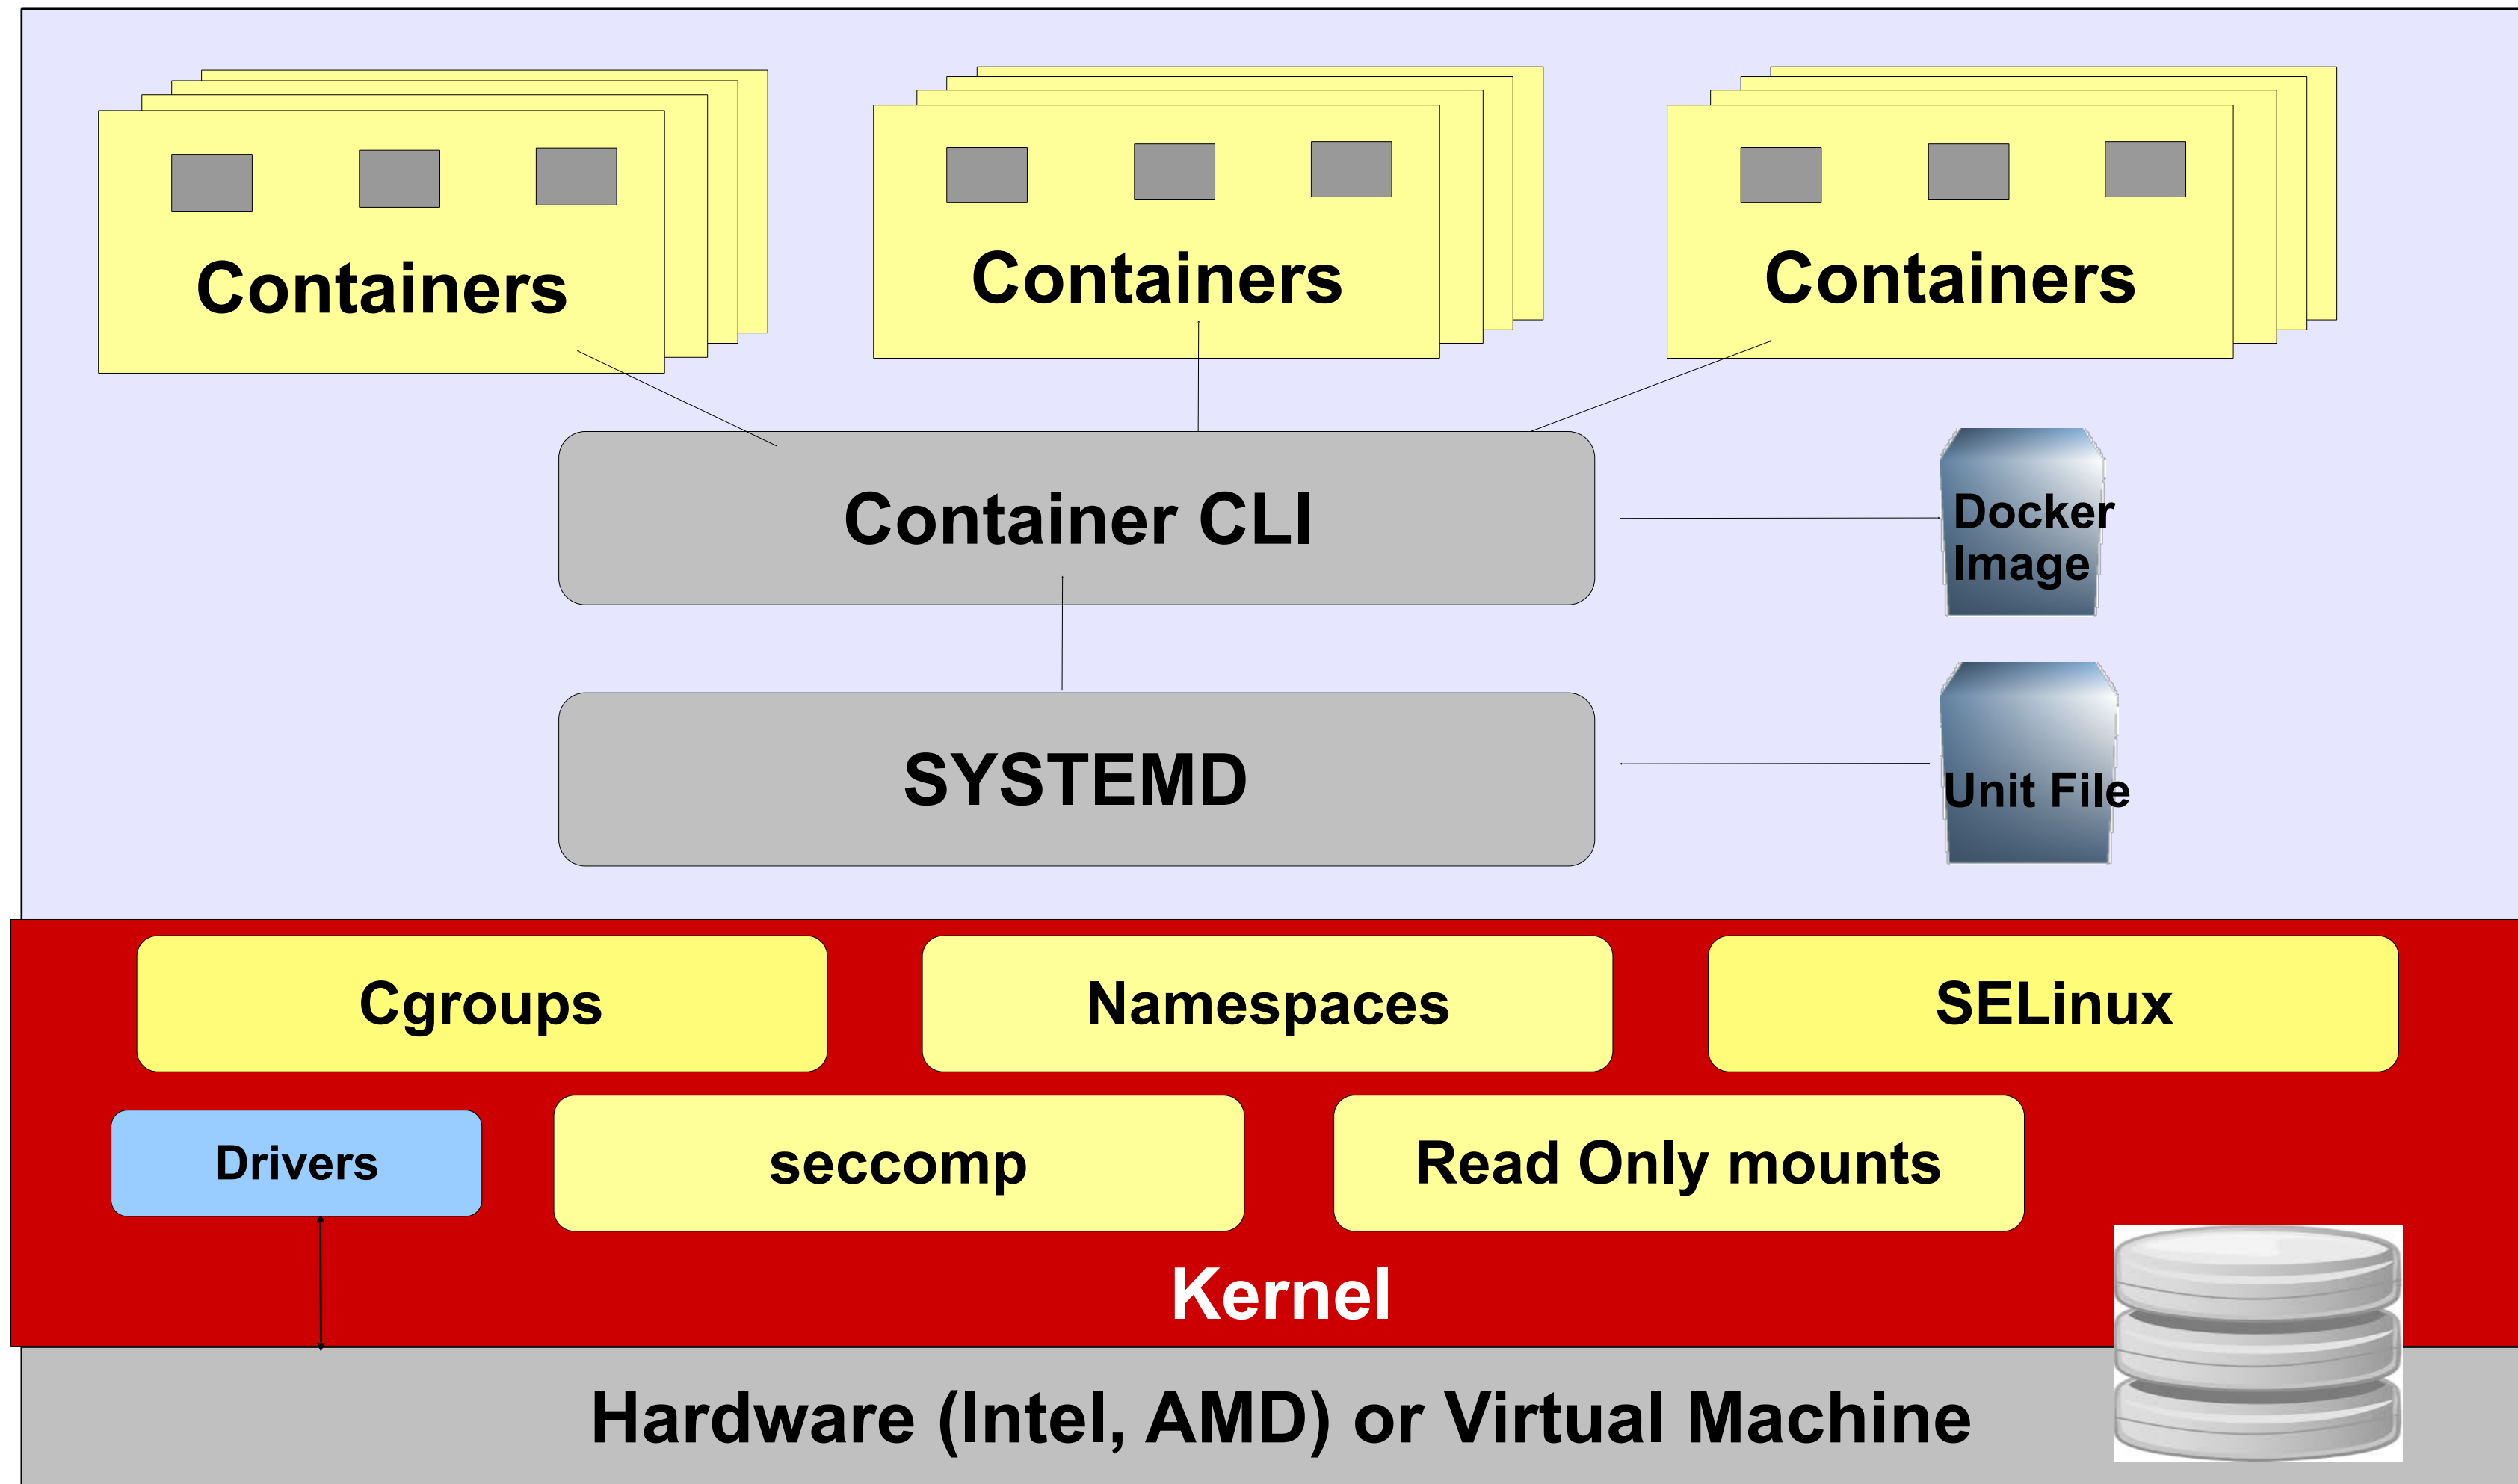
Docker Version

1.7.4



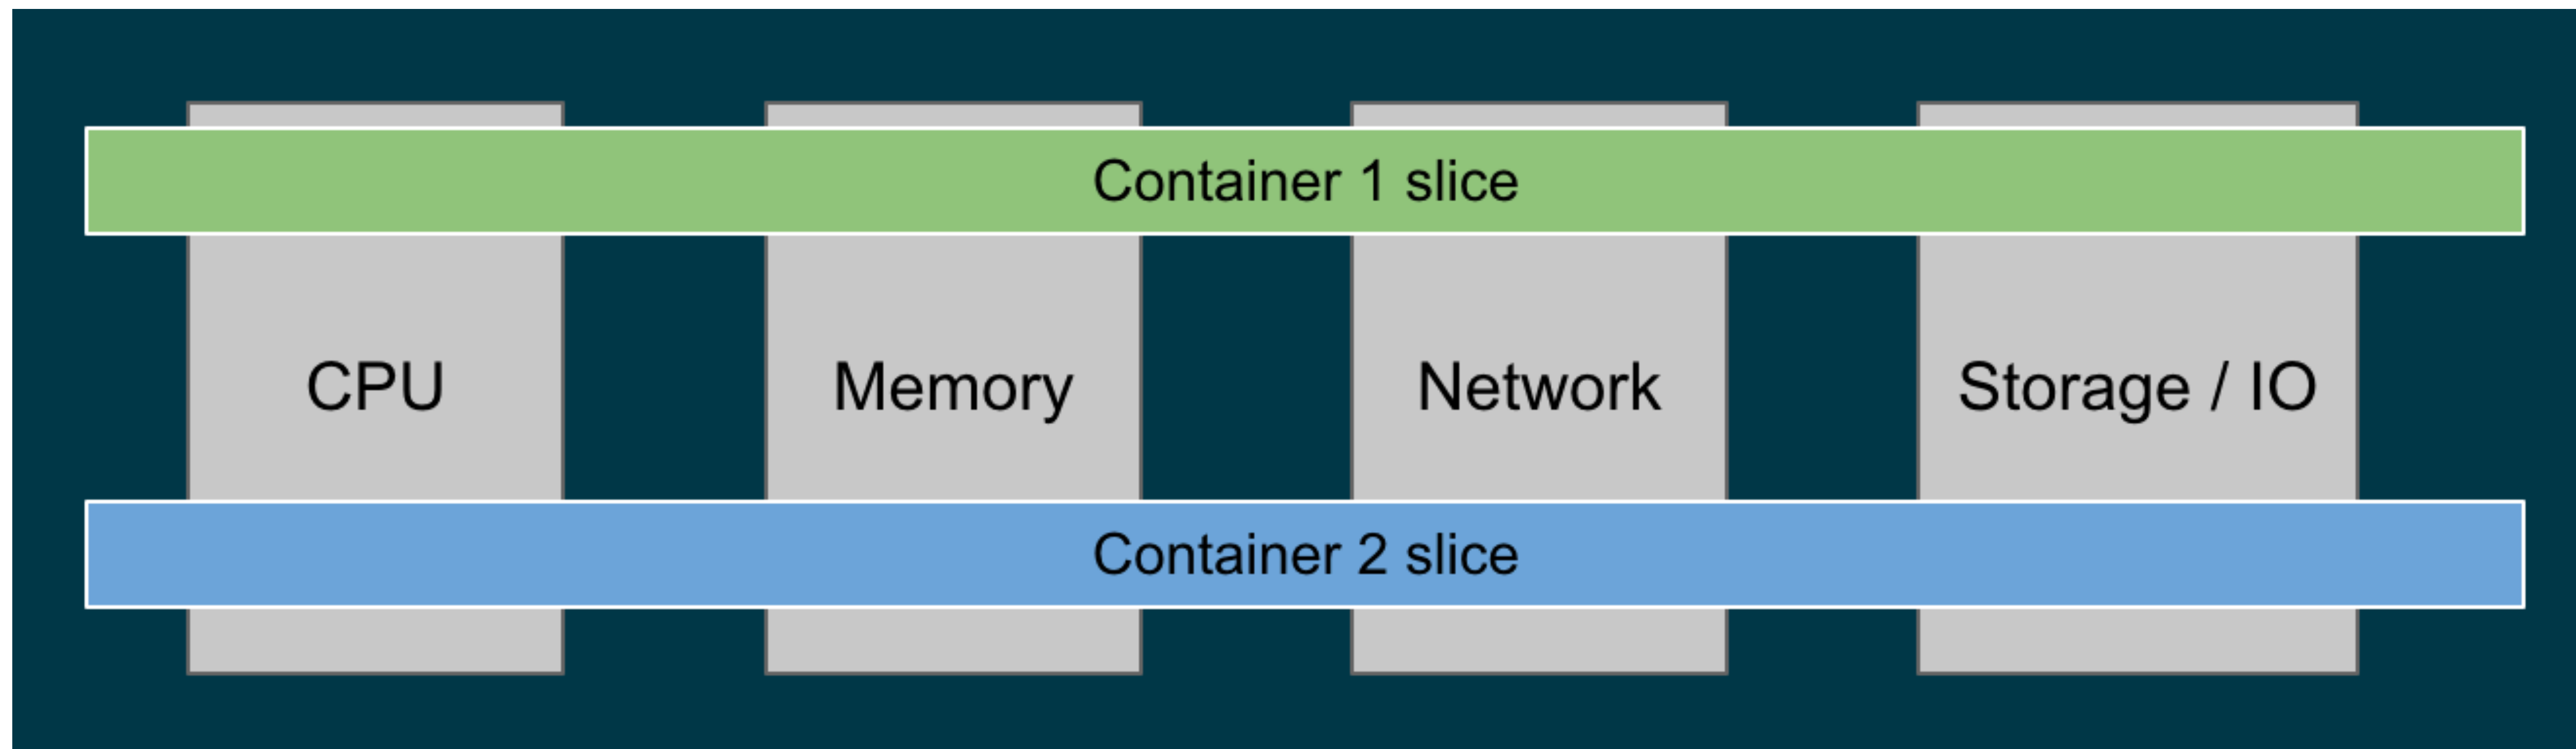
# CONTAINER HOST SECURITY

# CONTAINERS ARE LINUX

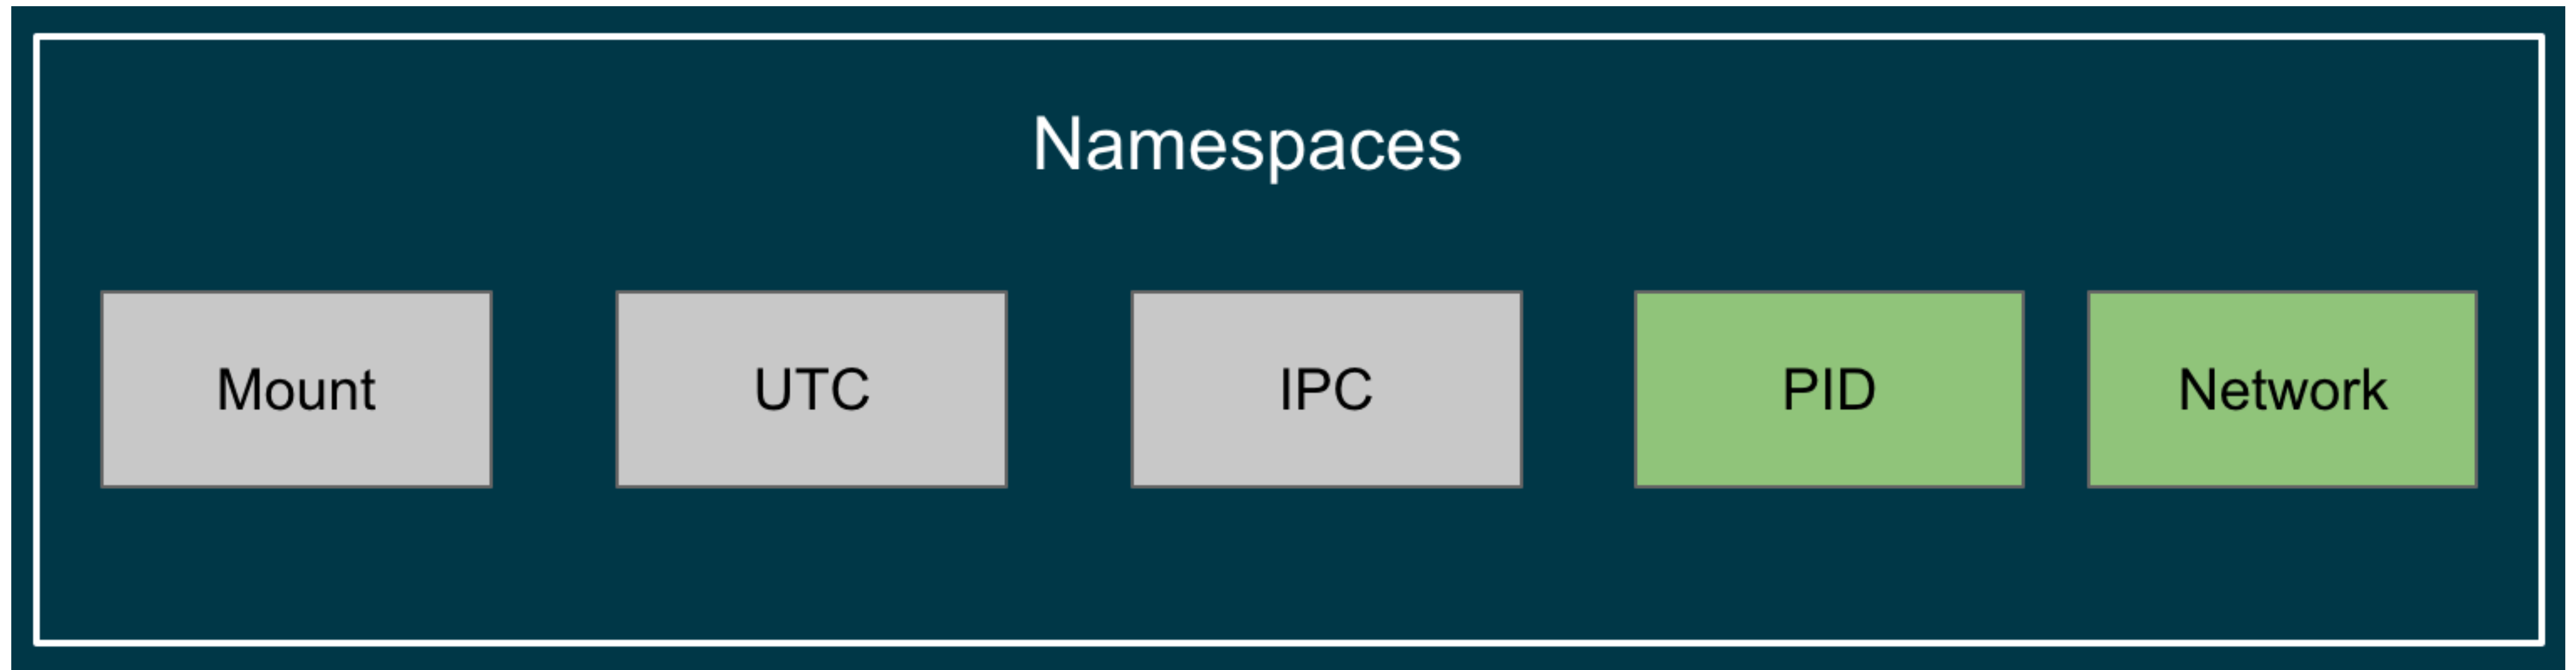




# CGROUPS - RESOURCE ISOLATION



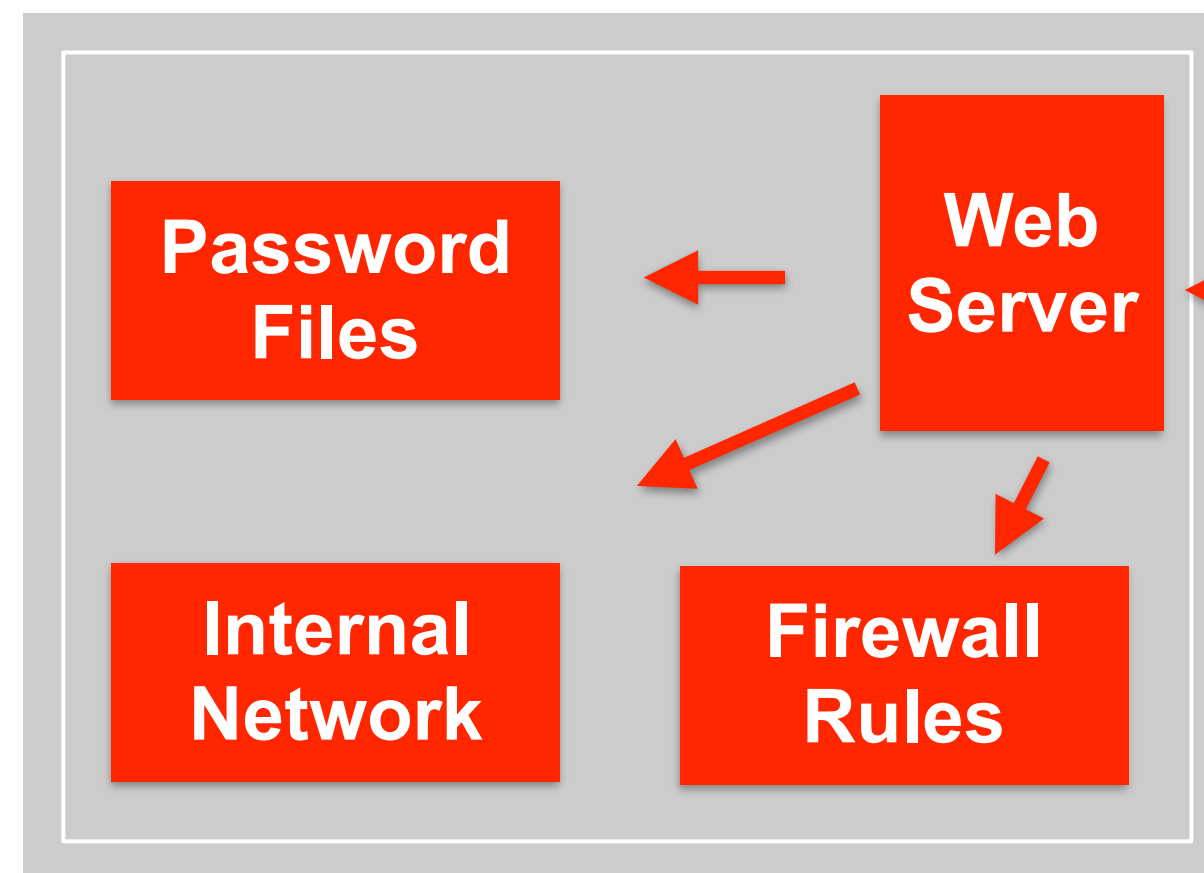
# NAMESPACES - PROCESS ISOLATION





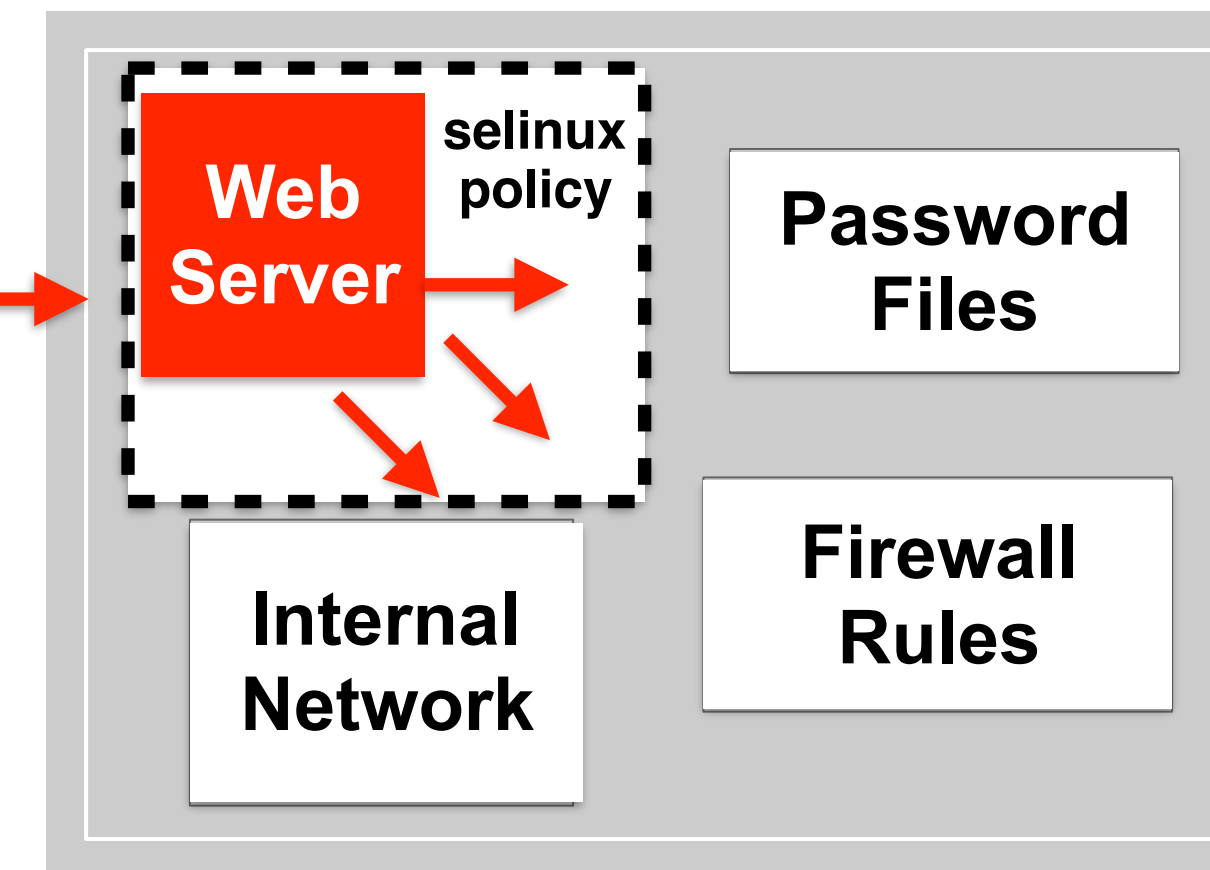
# SELINUX - MANDATORY ACCESS CONTROLS

## Discretionary Access Controls (file permissions)



**Attacker**

## Mandatory Access Controls (selinux)



# SECCOMP AND LINUX CAPABILITIES

## FILTERING SYSTEM CALLS and 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

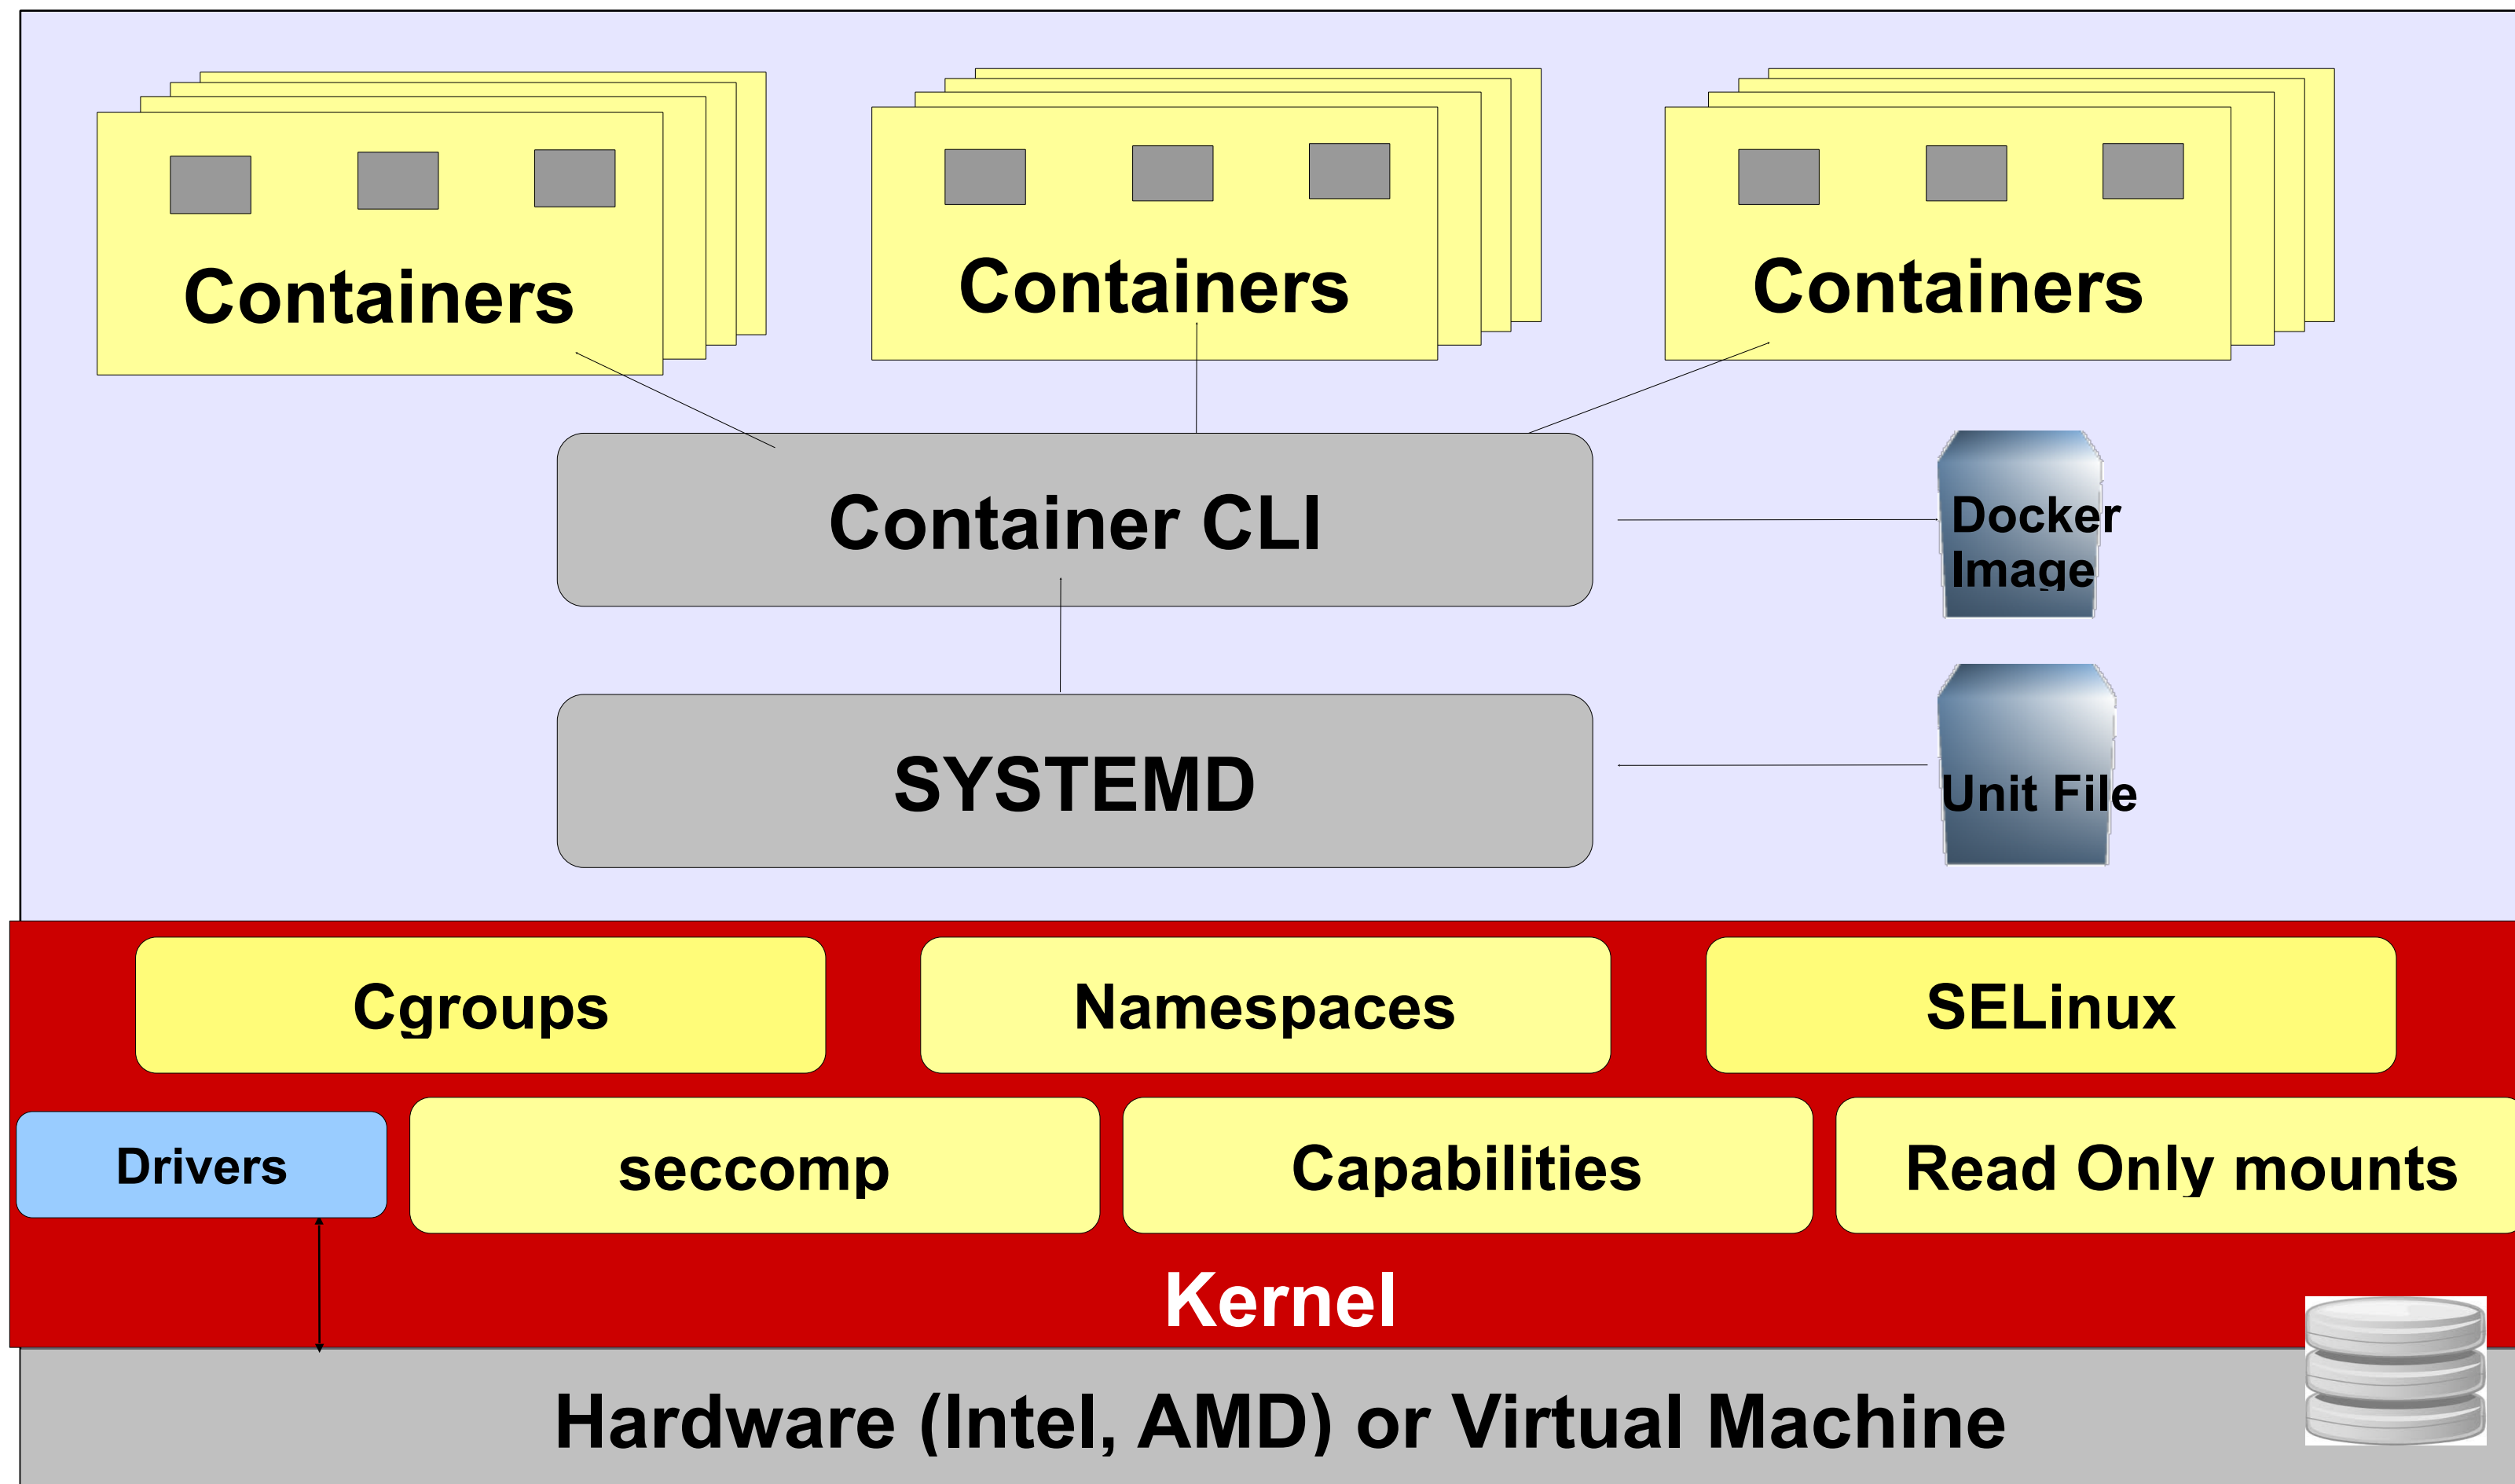
/proc/sysrg-trigger

/proc/irq

/proc/bus

R/O

# CONTAINER HOST SECURITY



## Best Practices

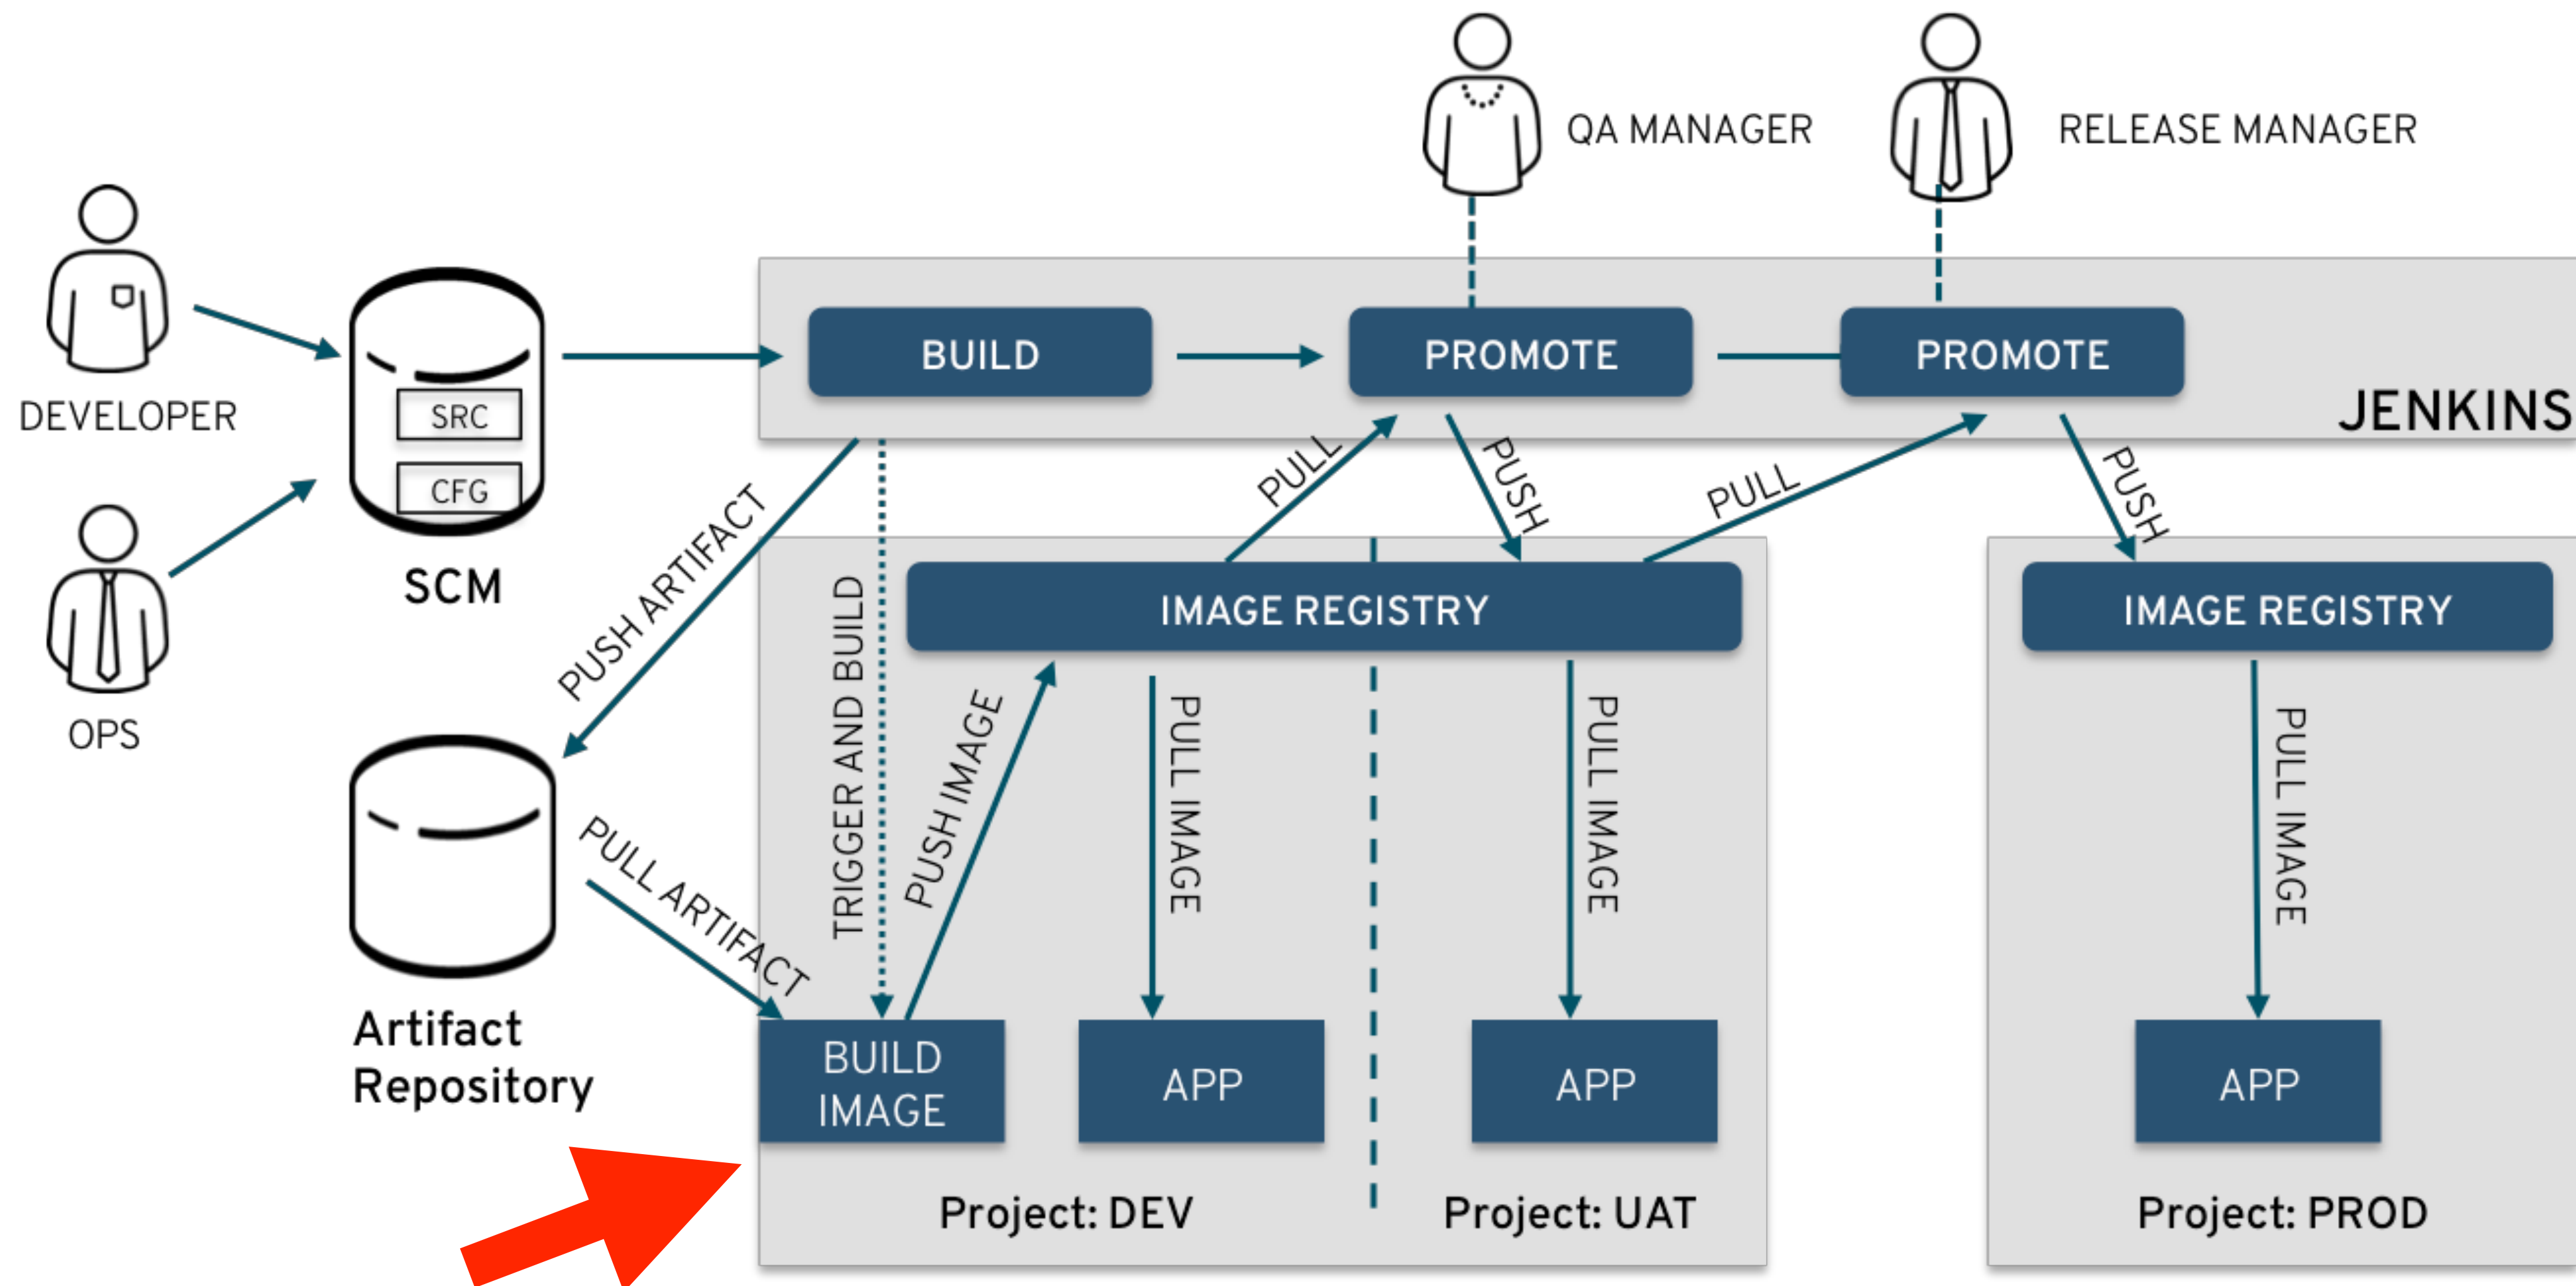
- Don't run as root
- If you must, limit Linux Capabilities
- Limit SSH Access
- Use namespaces
- Define resource quotas
- Enable logging
- Apply Security Errata
- Apply Security Context and seccomp filters
- Run production unprivileged containers as read-only

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



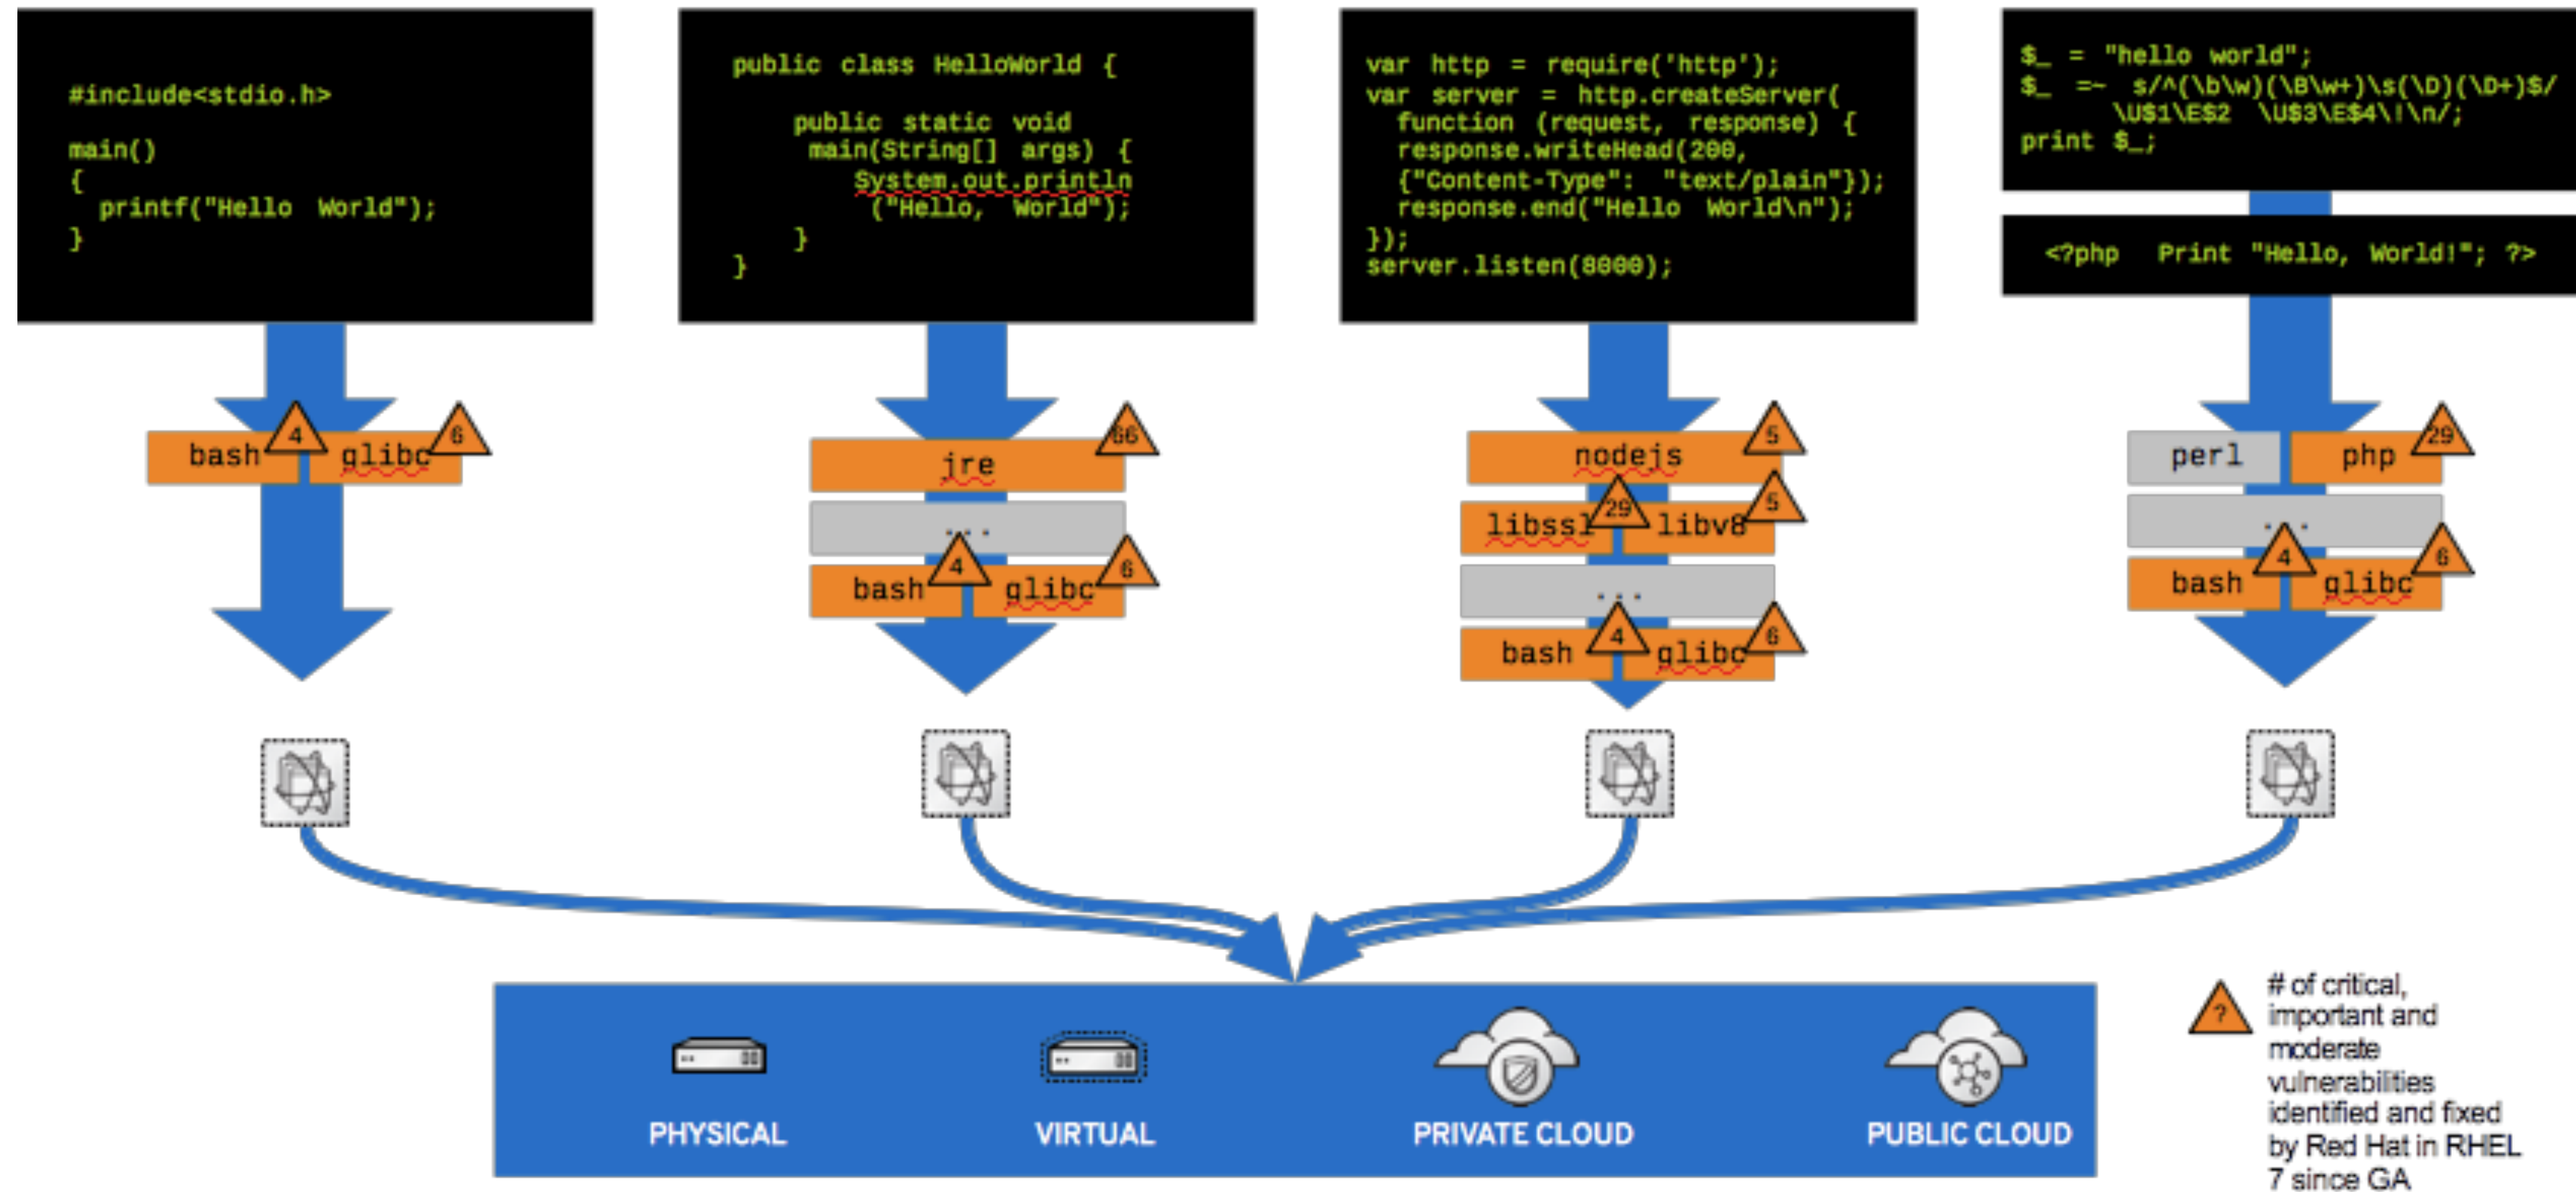
# CONTINUOUS INTEGRATION WITH CONTAINERS

# CONTINUOUS INTEGRATION + BUILDS





# WHAT'S INSIDE MATTERS...



# CONTINUOUS INTEGRATION WITH SECURITY SCAN

Team Cheese

My-Cool-Project ▾

Dashboard

Builds

Pipelines

Environments

Metrics

Settings

Components

Source

Packages

APIs

Camel

Pipelines

AllActiveCompletedFailed

Security

Build #3

awesome: 1.0.3

Jenkins ▶ Build Results ▶

Completed

Started 5 days ago

	Developer_Env	Integration_Env		Acceptance_Env	Staging_Env	Production_Env	
«	Commits	Code Review	Unit Test		Smoke Test	Smoke Test	»
	12s	12s	12s		12s	12s	

Build #2

awesome: 1.0.2

Jenkins ▶ Build Results ▶

Completed

Started 5 days ago

	Developer_Env	Integration_Env		Acceptance_Env	Staging_Env	Production_Env	
«	Development	Code Review	Unit Test		Smoke Test	Smoke Test	»
	12s	12s	12s		12s	12s	

Build #1

awesome: 1.0.1

Jenkins ▶ Build Results ▶

Service Test Warnings

Started 5 days ago

	Developer_Env	Integration_Env		Acceptance_Env	Staging_Env	Production_Env	
«	Development	Code Review	Unit Test		Smoke Test	Smoke Test	»
	12s	12s	12s		NA	NA	



# AUTOMATED SECURITY SCANNING with OpenSCAP

Scan physical servers, virtual machines, docker images and containers for Security Policy Compliance (CCEs) and known Security Vulnerabilities (CVEs)

## Content



SCAP Security  
Guide  
for RHEL

**CCE-27002-5**

Set Password Minimum  
Length

CVE DATABASE

**CVE-2015-5477**

Impact: **Important**  
Public: 2015-07-28


## Scan



OpenSCAP



**FOREMAN**



## Reports

### Compliance and Scoring

The target system did not satisfy conditions of 33 rules! Please review rule results and consider applying remediation.

#### Rule result breakdown

34 passed 33 failed 1

#### Failed rules by severity breakdown

3 high 16 medium 14 low

#### Score

Scoring system	Score	Maximum	%
um:xccdf:scoring:default	48.935184	100.000000	48.94%

### Existence of Password Hashes (1x fail)

Accounts With Empty Password	high	fail
Password Hashes are Shadowed	medium	pass

### Parameters (2x fail)

Length in login.defs	medium	fail
Age	medium	fail
Age	low	pass

### Using PAM (10x fail)

### Requirements (5x fail)

Prompts Permitted Per-Session	low	pass
Minimum Digit Characters	low	fail
Minimum Uppercase Characters	low	fail
Minimum Special Characters	low	fail
Minimum Lowercase Characters	low	fail

# Security Policies in SCAP Security Guide (partial)

Standard Docker Host Security Profile

Java Runtime Environment (JRE)

Upstream Firefox STIG

RHEL OSP STIG

Red Hat Corporate Profile for Certified Cloud Providers (RH CCP)


STIG for Red Hat Enterprise Linux 6, 7 Server

STIG for Red Hat Virtualization Hypervisor

Common Profile for General-Purpose Debian Systems

Common Profile for General-Purpose Fedora Systems

Common Profile for General-Purpose Ubuntu Systems



Payment Card Industry – Data Security Standard (PCI-DSS) v3

U.S. Government Commercial Cloud Services (C2S)

CNSSI 1253 Low/Low/Low Control Baseline for Red Hat Enterprise Linux 7

Criminal Justice Information Services (CJIS) Security Policy

Unclassified Information in Non-federal Information Systems and Organizations (NIST 800-171)

U.S. Government Configuration Baseline (NIAP OSPP v4.0, USGCB, STIG)



# SECURITY POLICY REPORT

▼ Verify Proper Storage and Existence of Password Hashes 1x fail		
Prevent Log In to Accounts With Empty Password	high	fail
Verify All Account Password Hashes are Shadowed	medium	pass
▼ Set Password Expiration Parameters 2x fail		
Set Password Minimum Length in login.defs	medium	fail
Set Password Minimum Age	medium	fail
Set Password Warning Age	low	pass
▼ Protect Accounts by Configuring PAM 10x fail		
▼ Set Password Quality Requirements 5x fail		
▼ Set Password Quality Requirements, if using pam_pwquality 5x fail		
Set Password Retry Prompts Permitted Per-Session	low	pass
Set Password Strength Minimum Digit Characters	low	fail
Set Password Strength Minimum Uppercase Characters	low	fail
Set Password Strength Minimum Special Characters	low	fail
Set Password Strength Minimum Lowercase Characters	low	fail

# SECURITY POLICY REMEDIATION

## Set Password Strength Minimum Digit Characters

Rule ID	accounts_password_pam_dcredit
Result	<div>fail</div>
Time	2015-07-31T14:57:17
Severity	low
Identifiers and References	<div>identifiers: CCE-27163-5</div> <div>references: IA-5(b), IA-5(c), 194, 194, 71,</div>

The pam\_pwquality module's dcredit parameter controls requirements for usage of digits in a password. When set to a negative number, any password will be required to contain that many digits. When set to a positive number, pam\_pwquality will grant +1 additional length credit for each digit. Add dcredit=-1 after pam\_pwquality.so to require use of a digit in passwords.

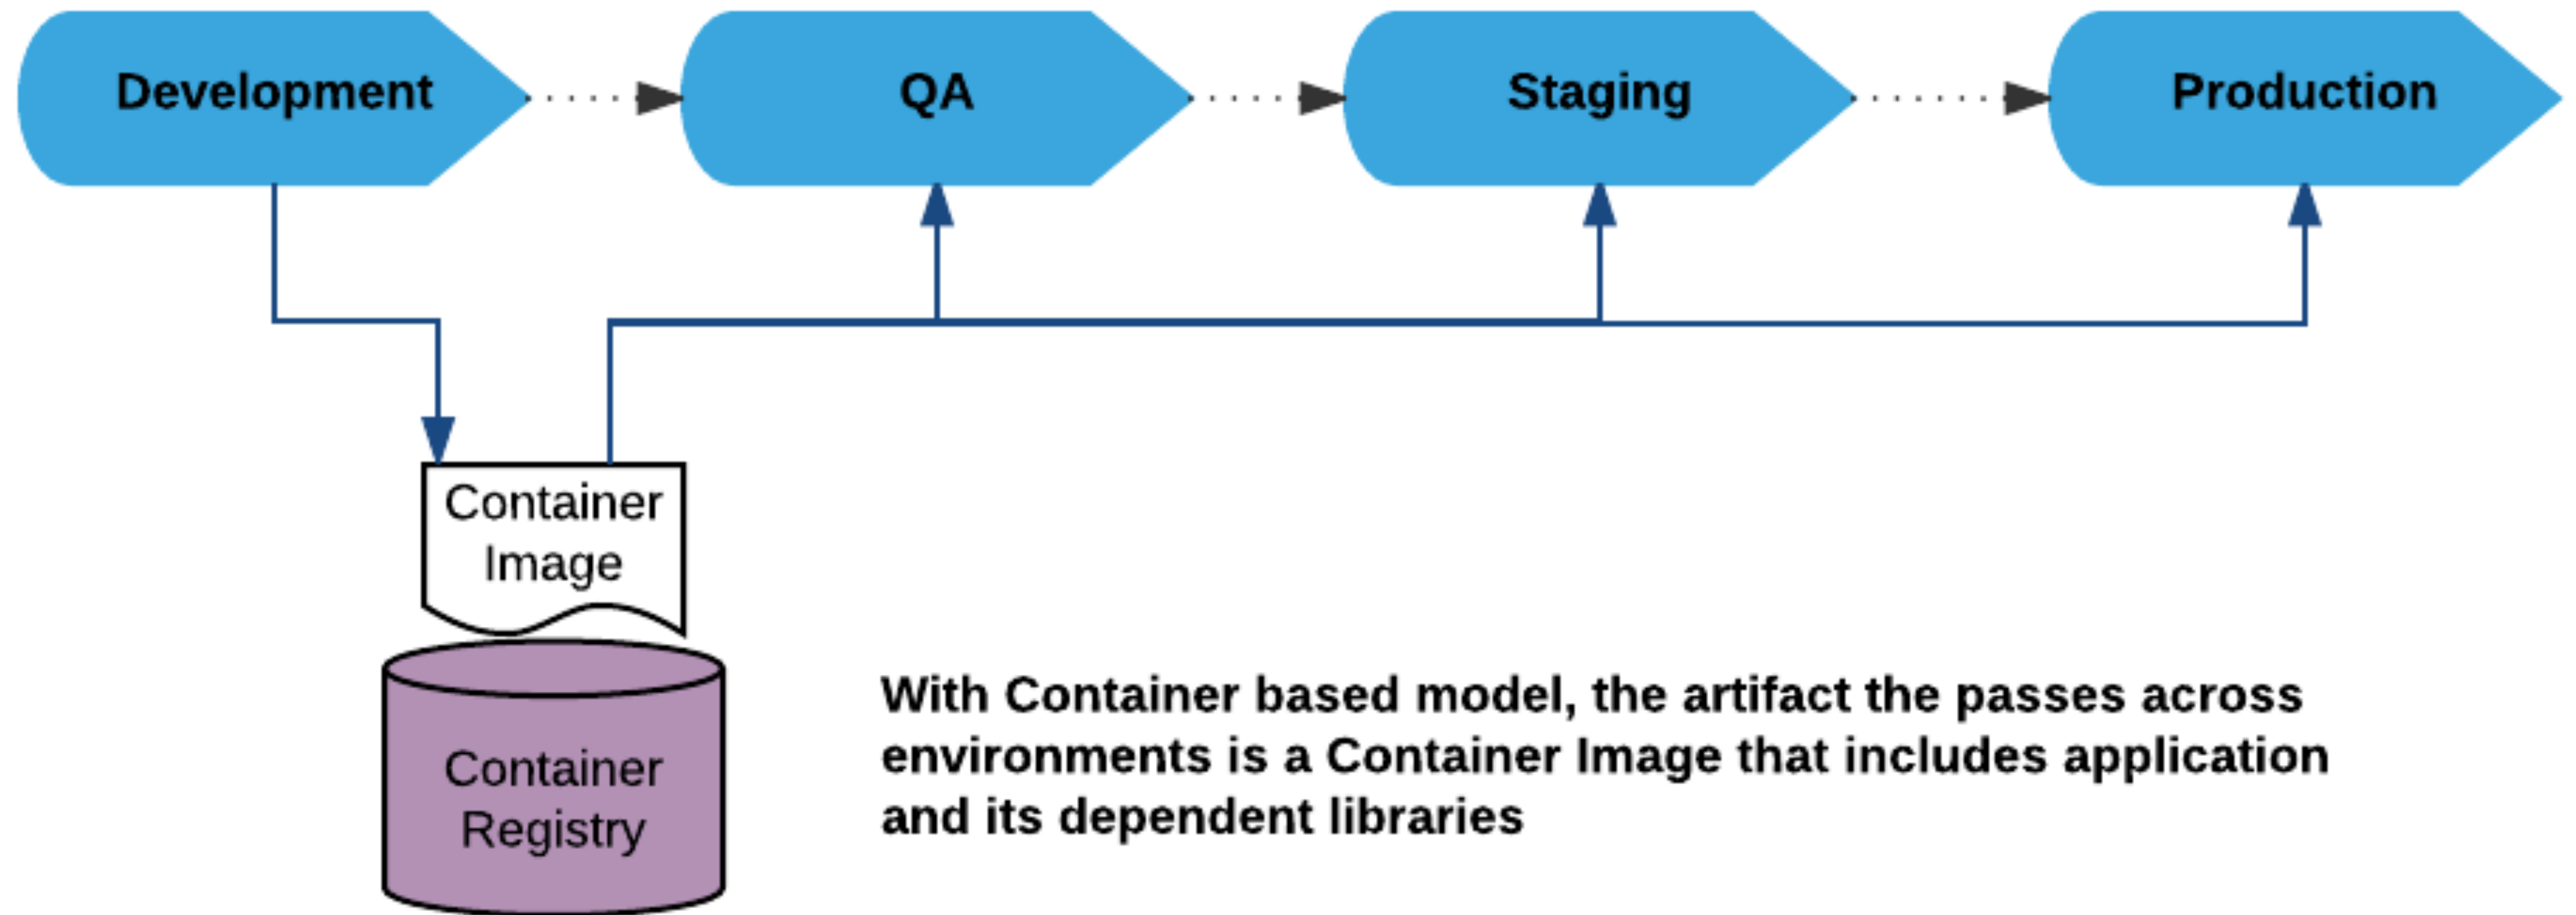
### Remediation script:

```
var_password_pam_dcredit="-1"
if grep -q "dcredit=" /etc/pam.d/system-auth; then
    sed -i --follow-symlink "s/\(dcredit *= *\).*$/\1$var_password_pam_dcredit/" /etc/pam.d/system-auth
else
    sed -i --follow-symlink "/pam_pwquality.so/ s/$/ dcredit=$var_password_pam_dcredit/" /etc/pam.d/system-auth
fi
```



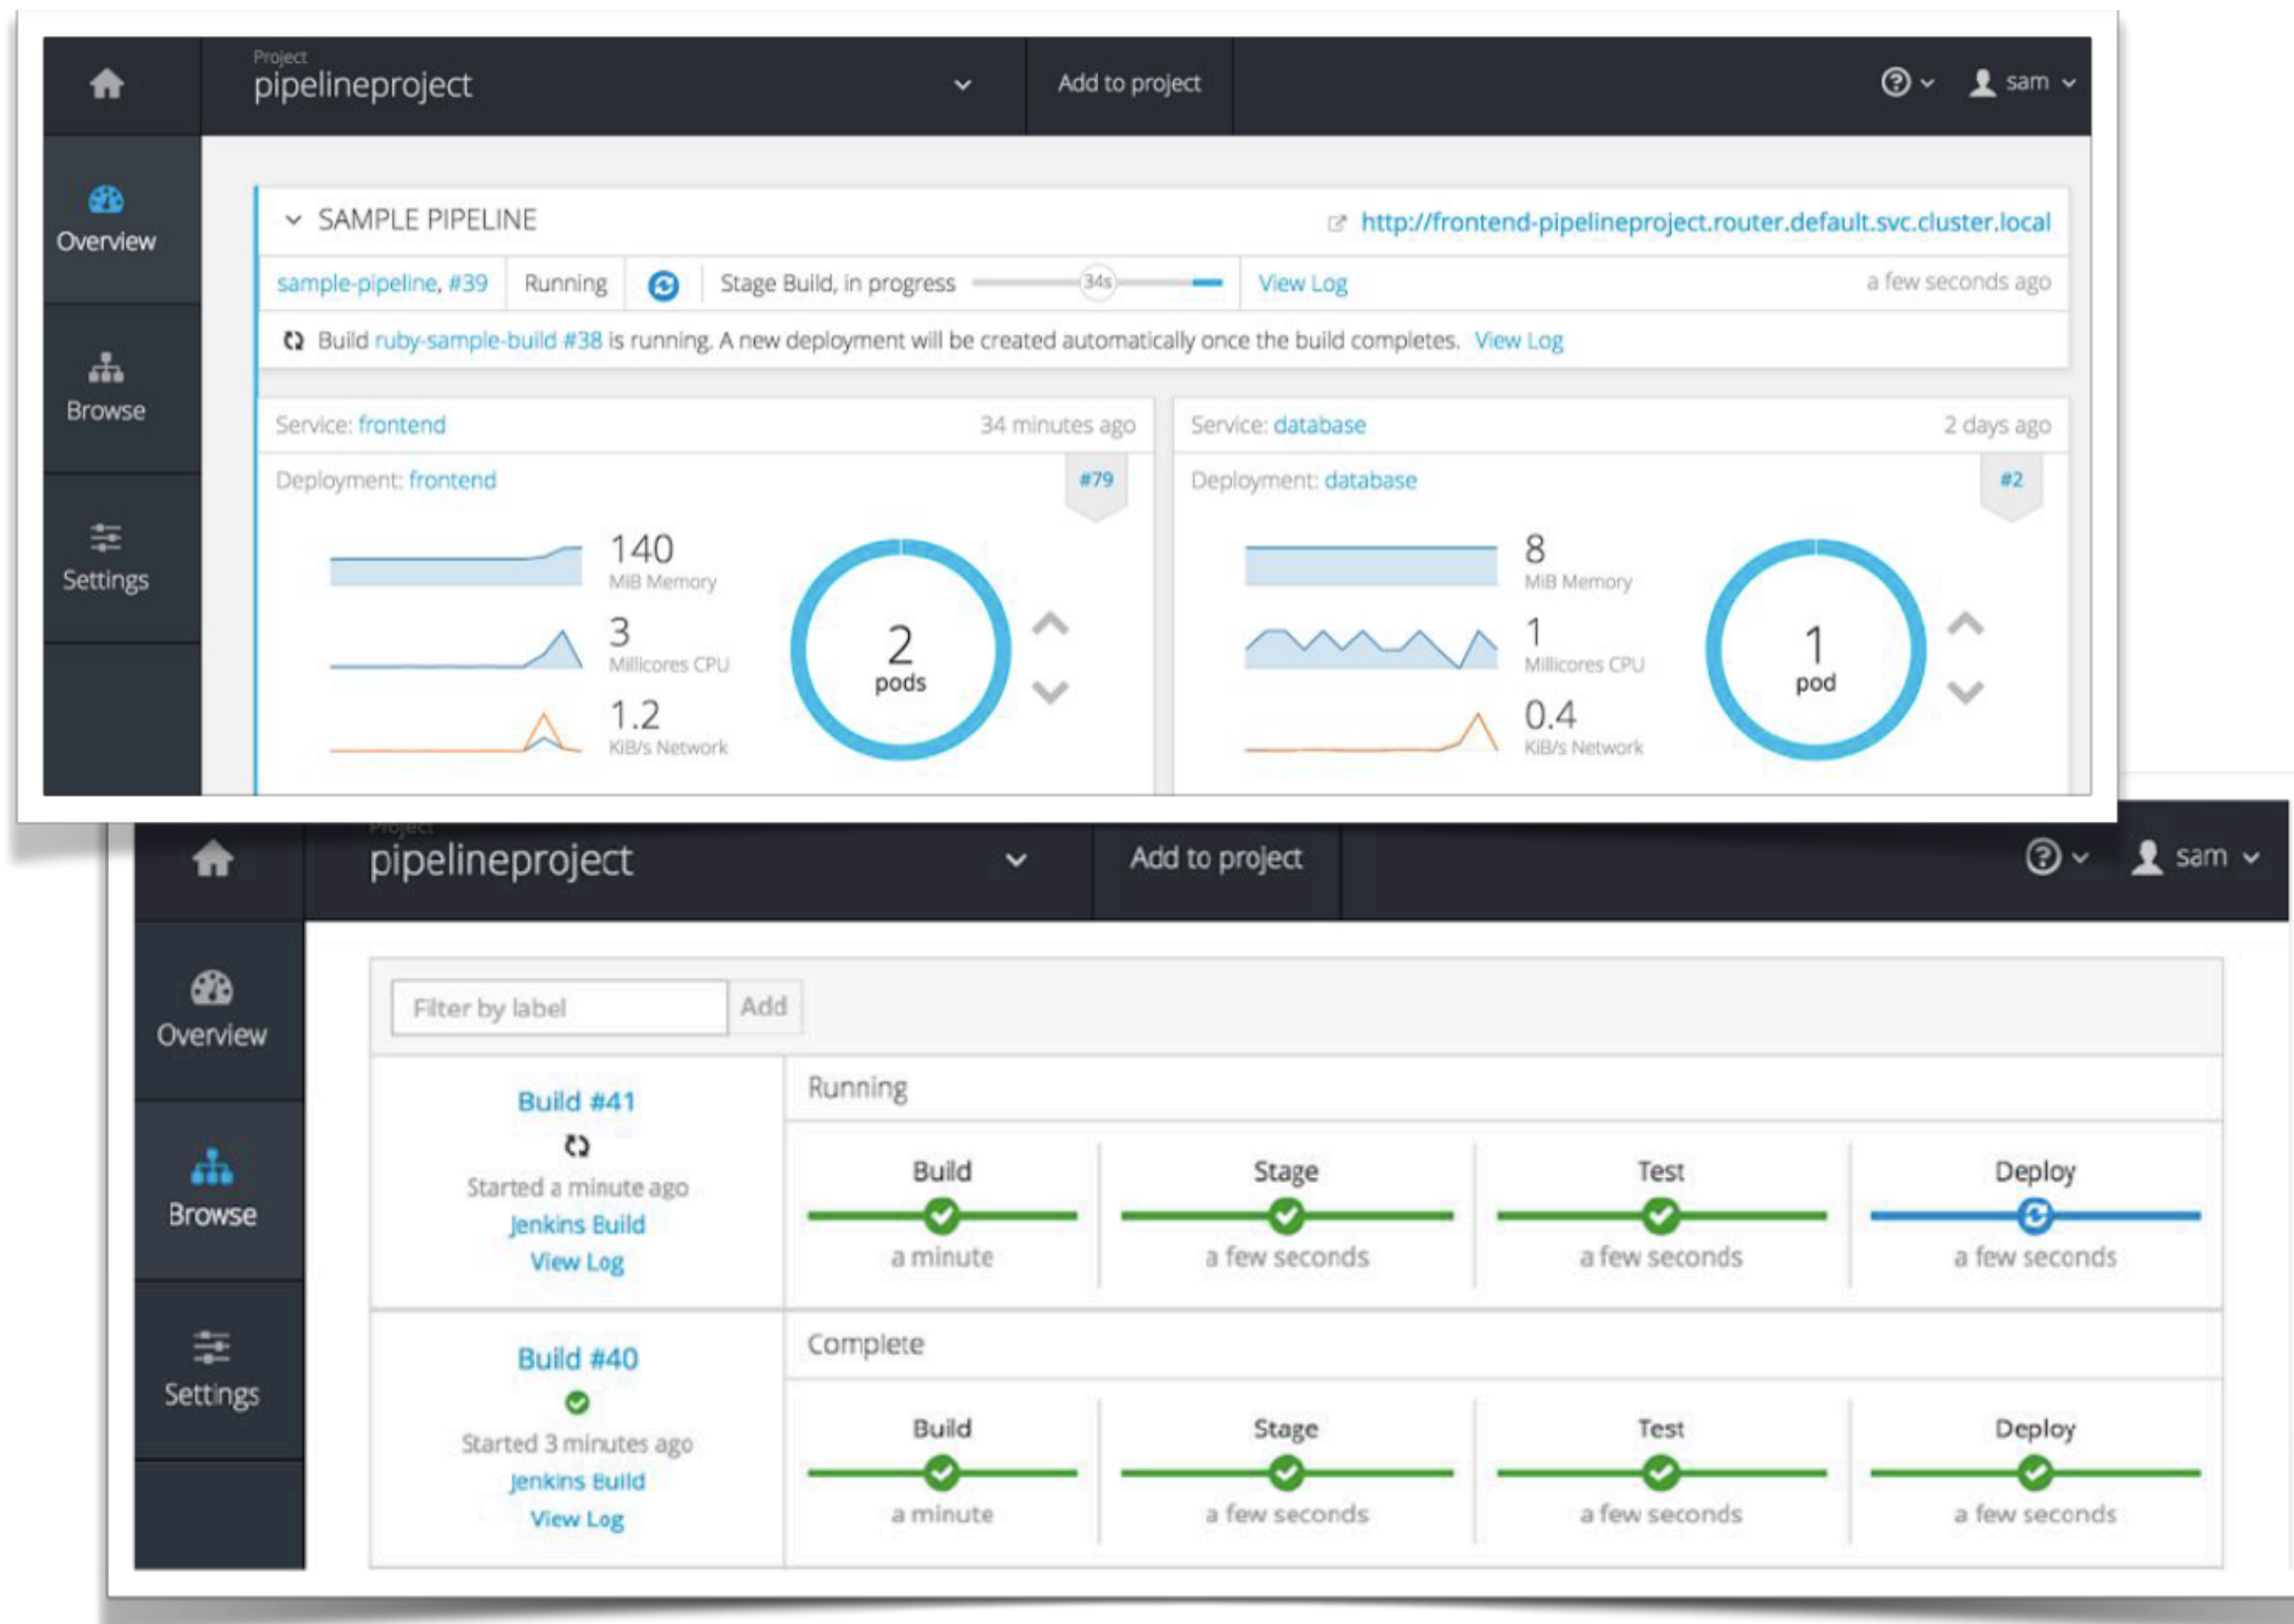
# CONTINUOUS DELIVERY WITH CONTAINERS

# CONTINUOUS DELIVERY WITH CONTAINERS





# CONTINUOUS DELIVERY DEPLOYMENT STRATEGIES




## DEPLOYMENT STRATEGIES

- Recreate
- Rolling updates
- Blue / Green deployment
- Canary with A/B testing


# 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



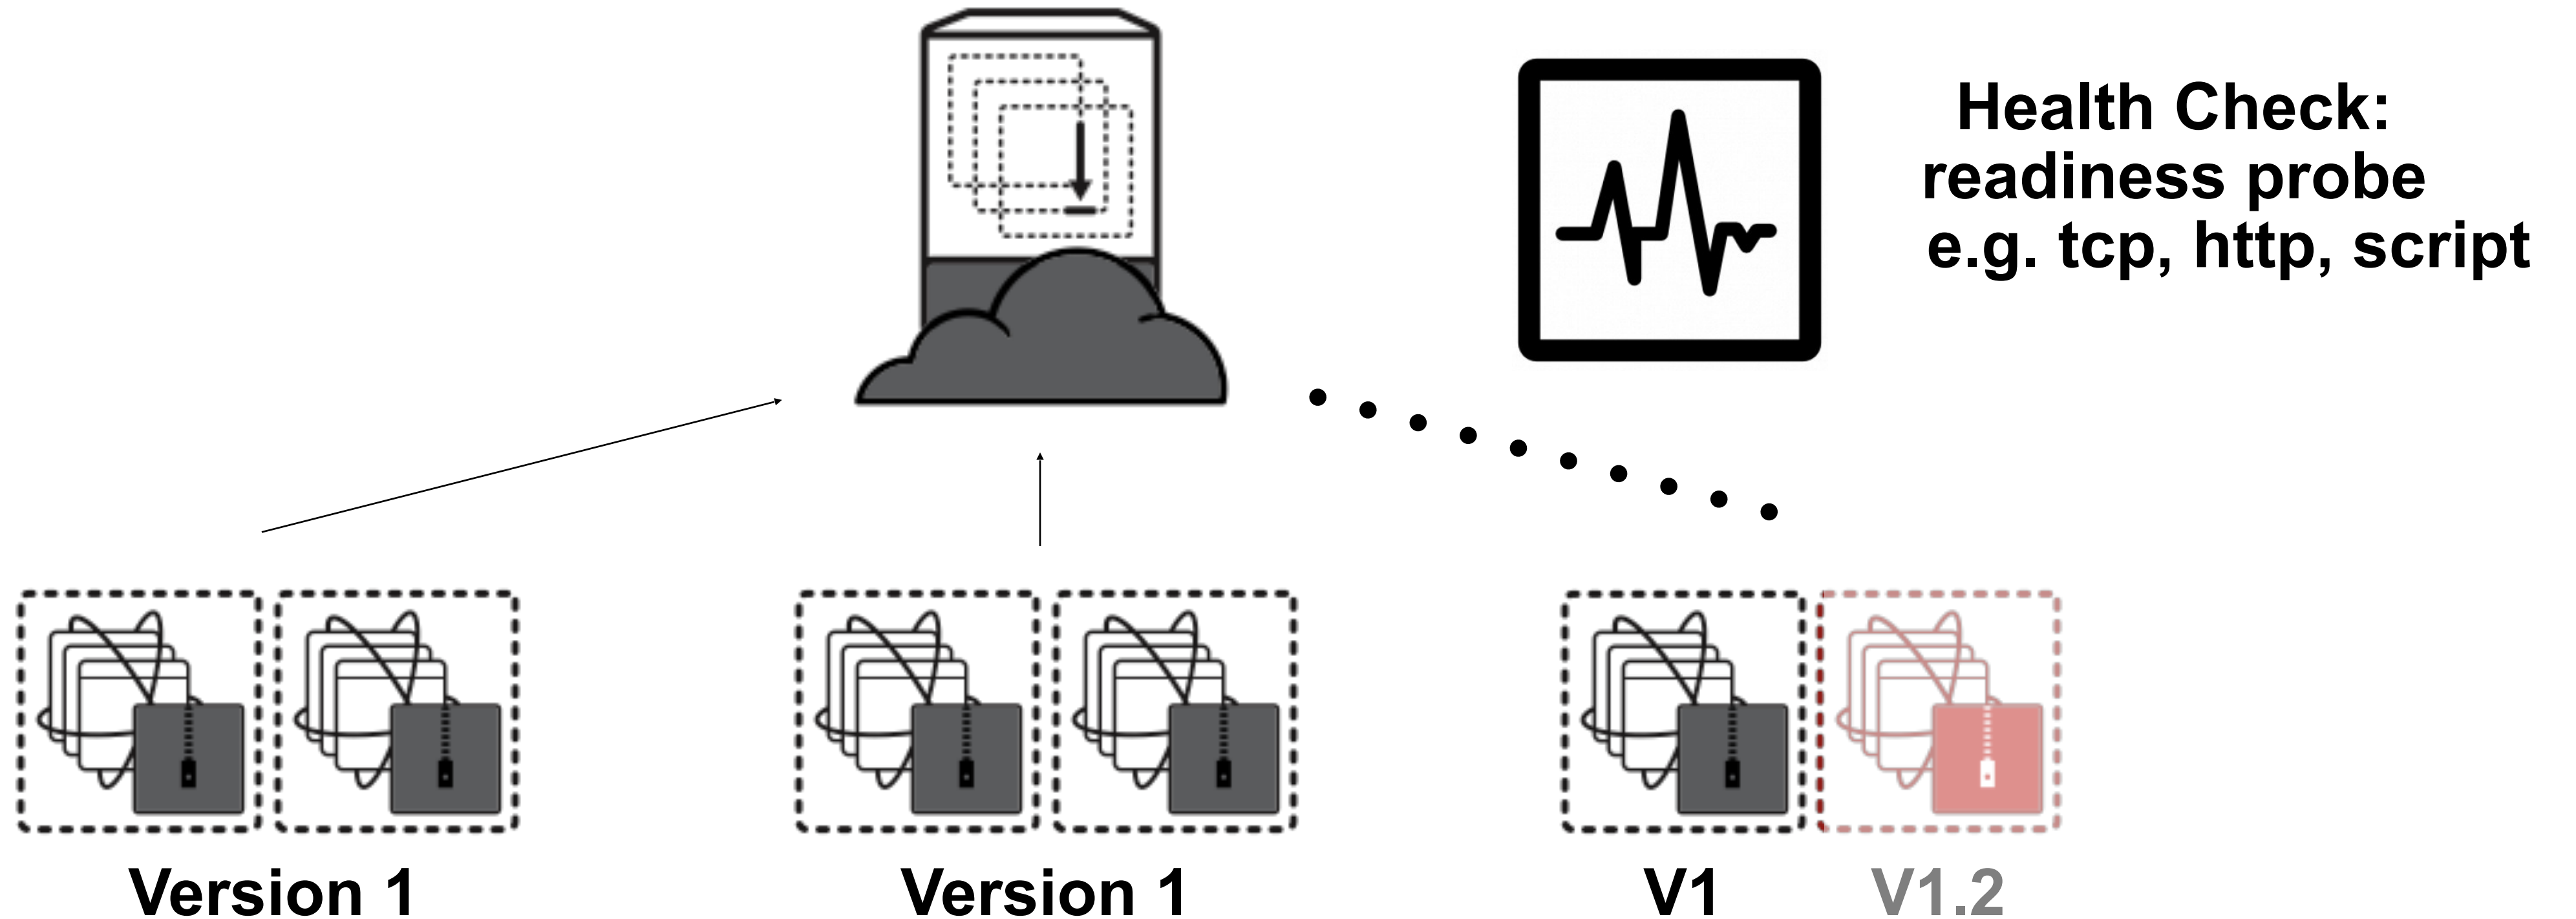
# Rolling Updates




# 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

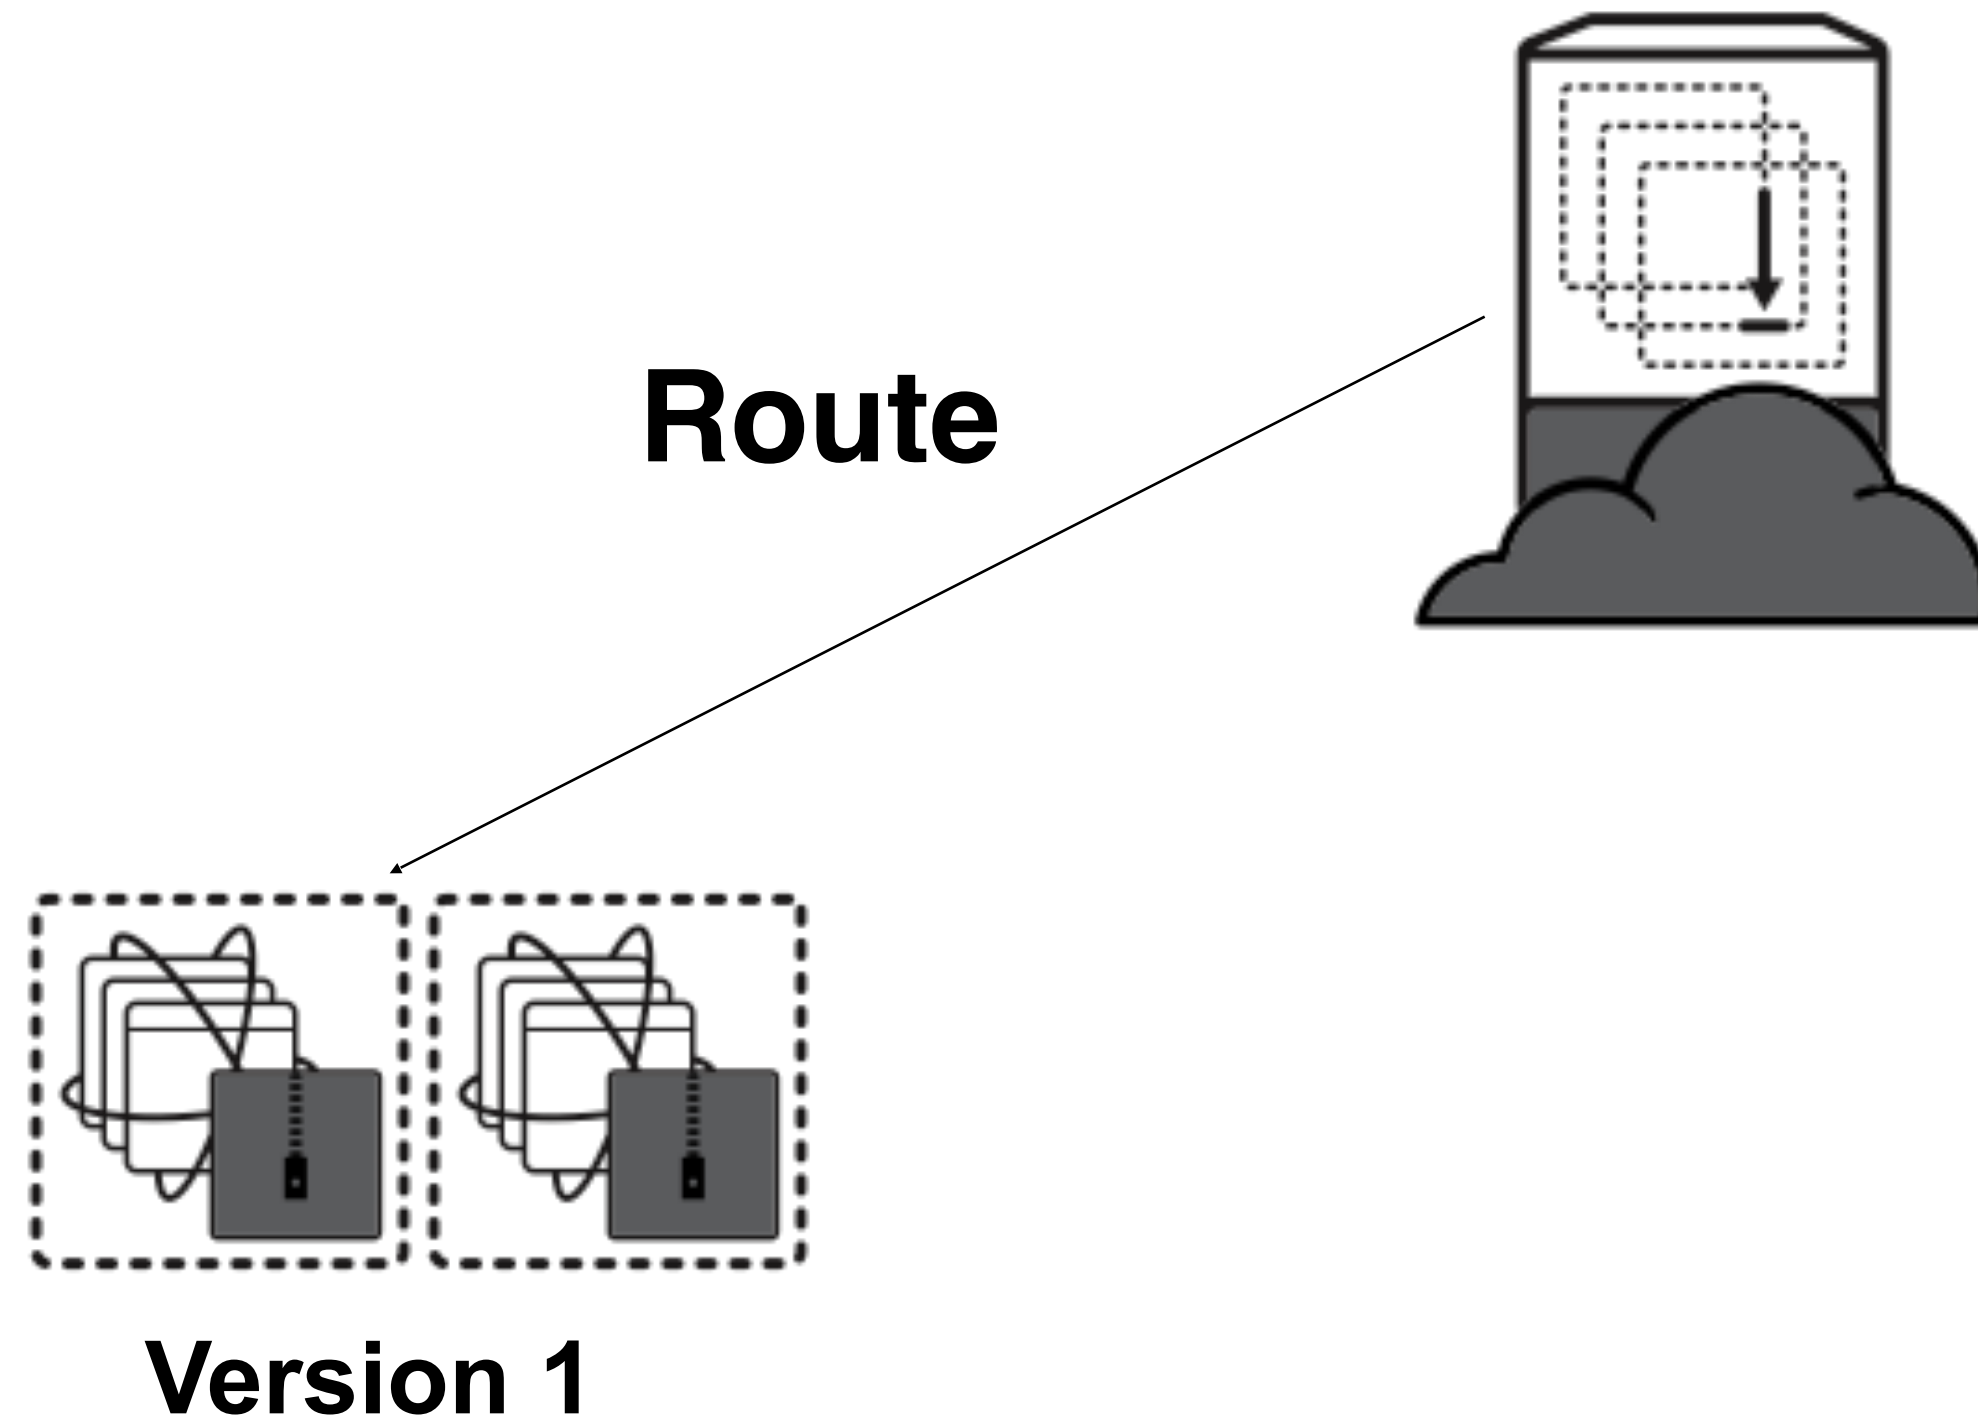
## Pros

- 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

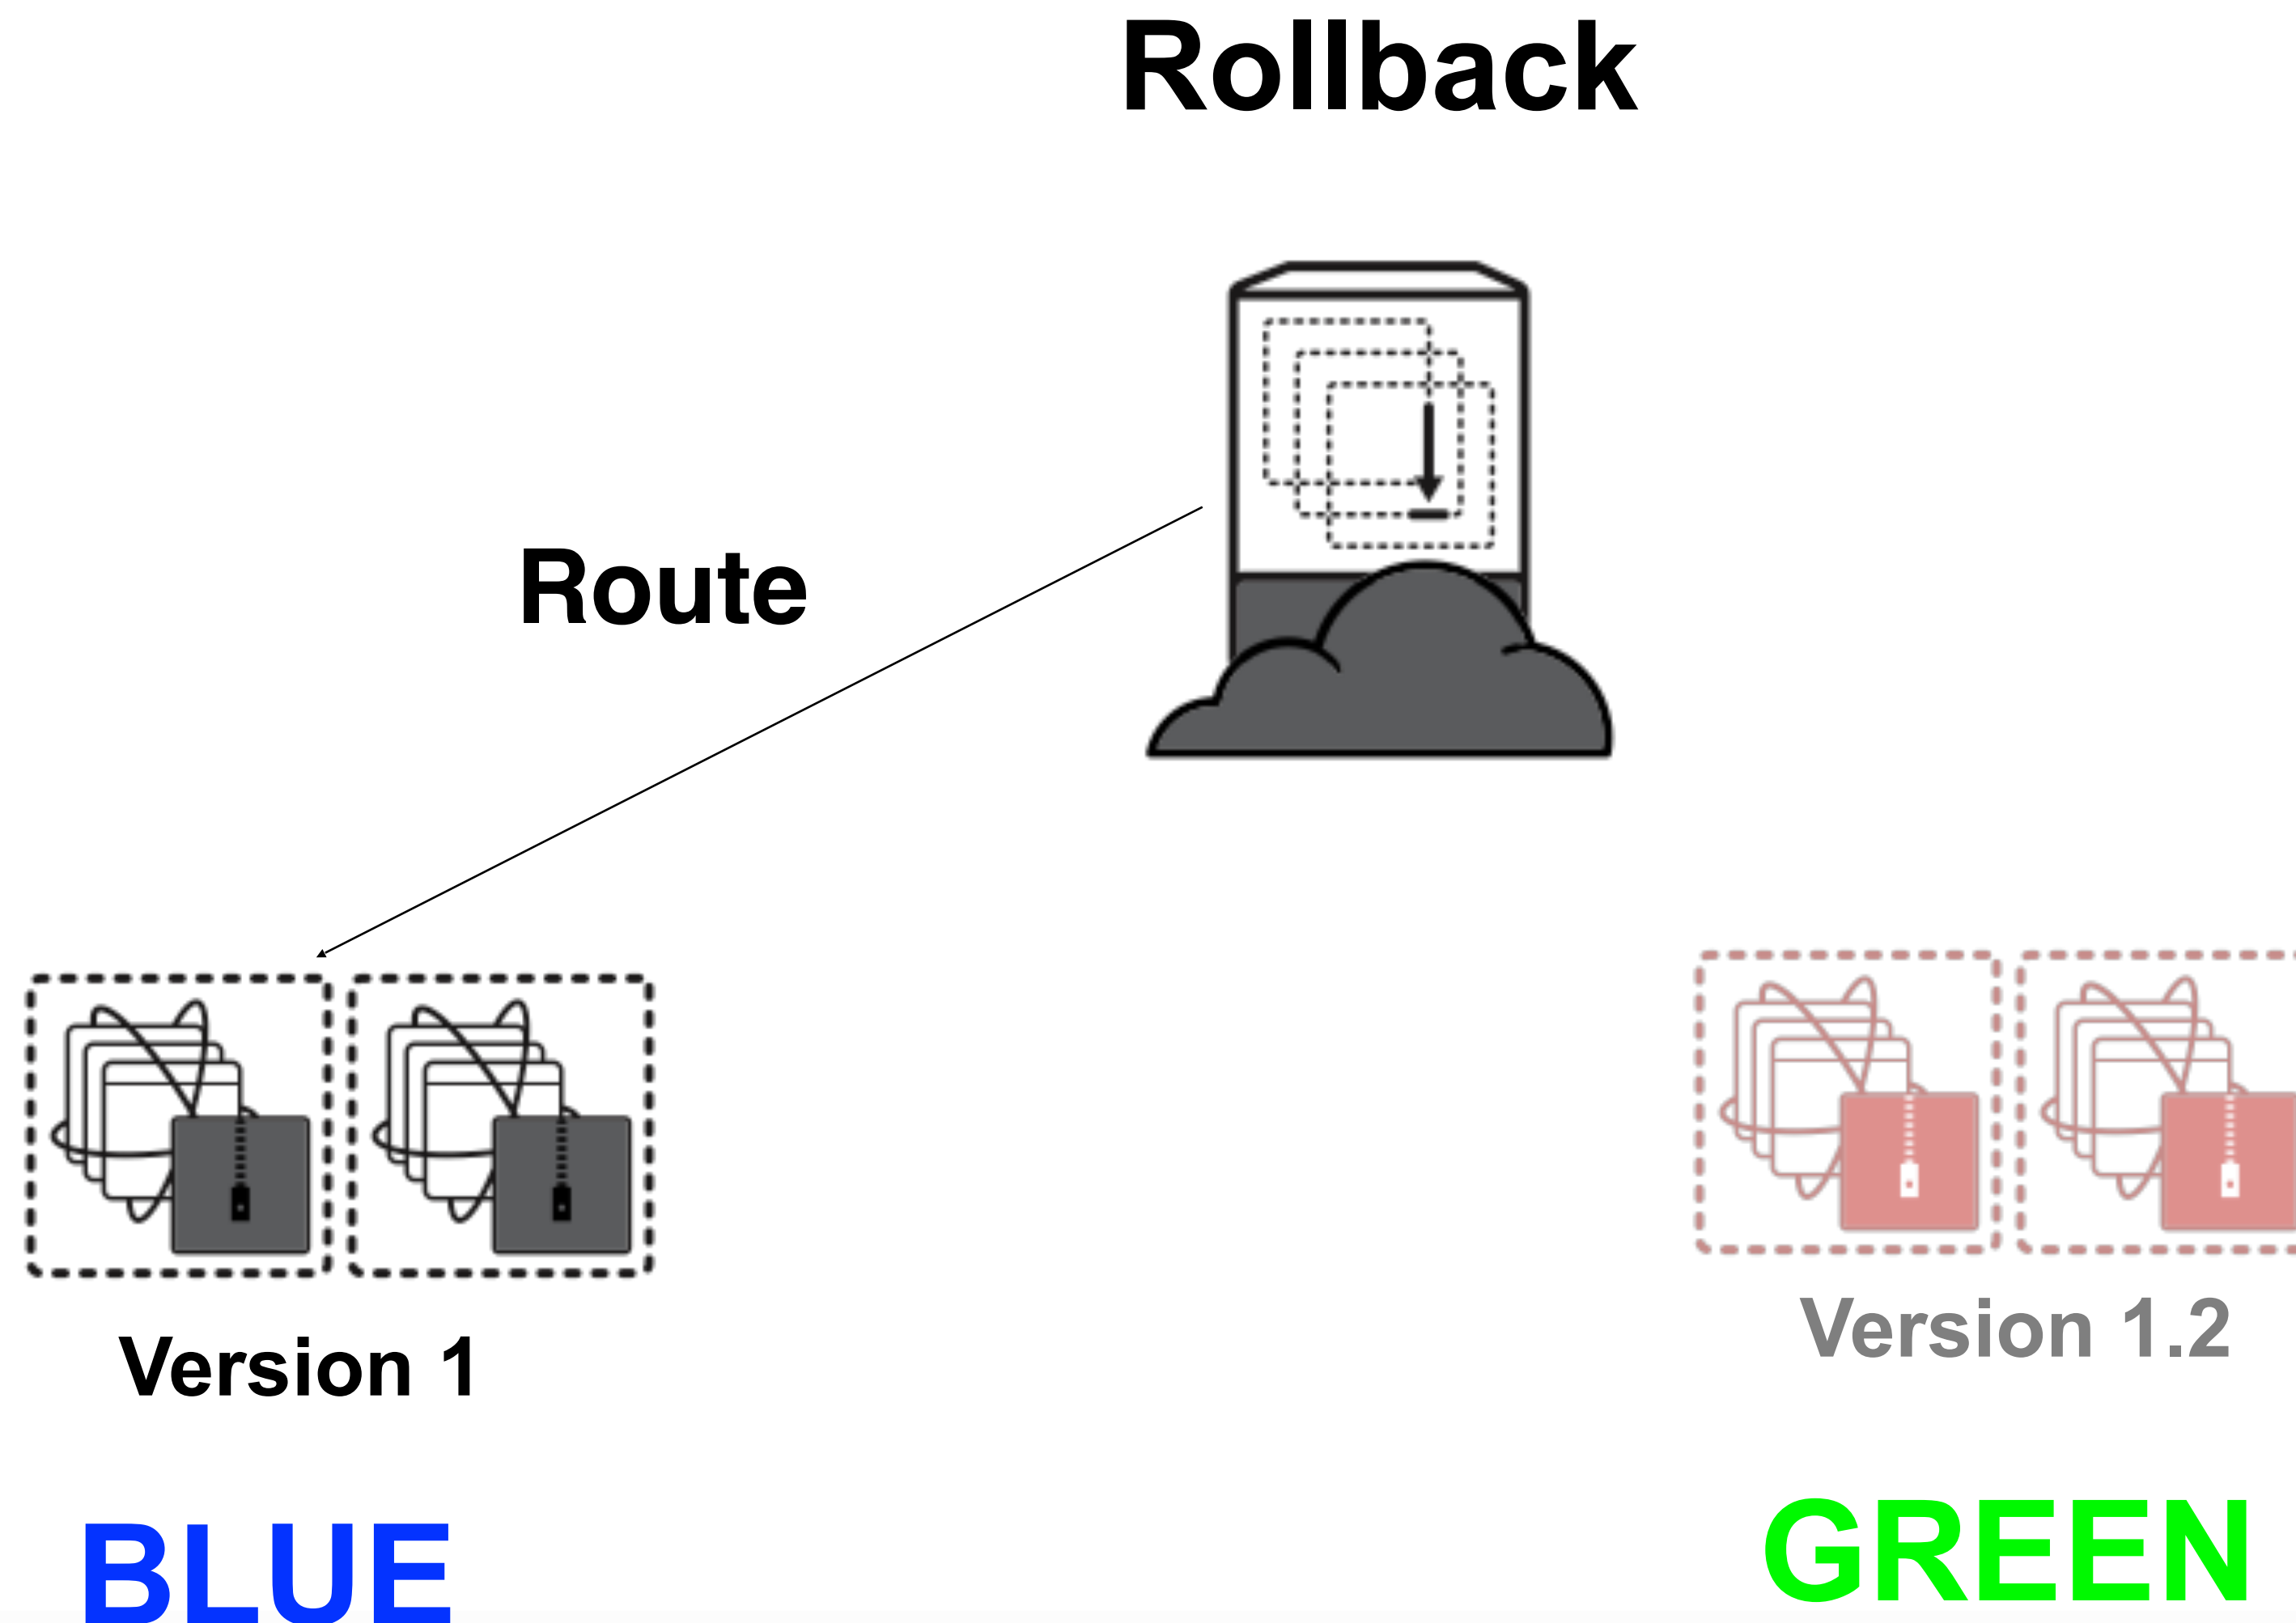


# BLUE / GREEN DEPLOYMENT





# 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

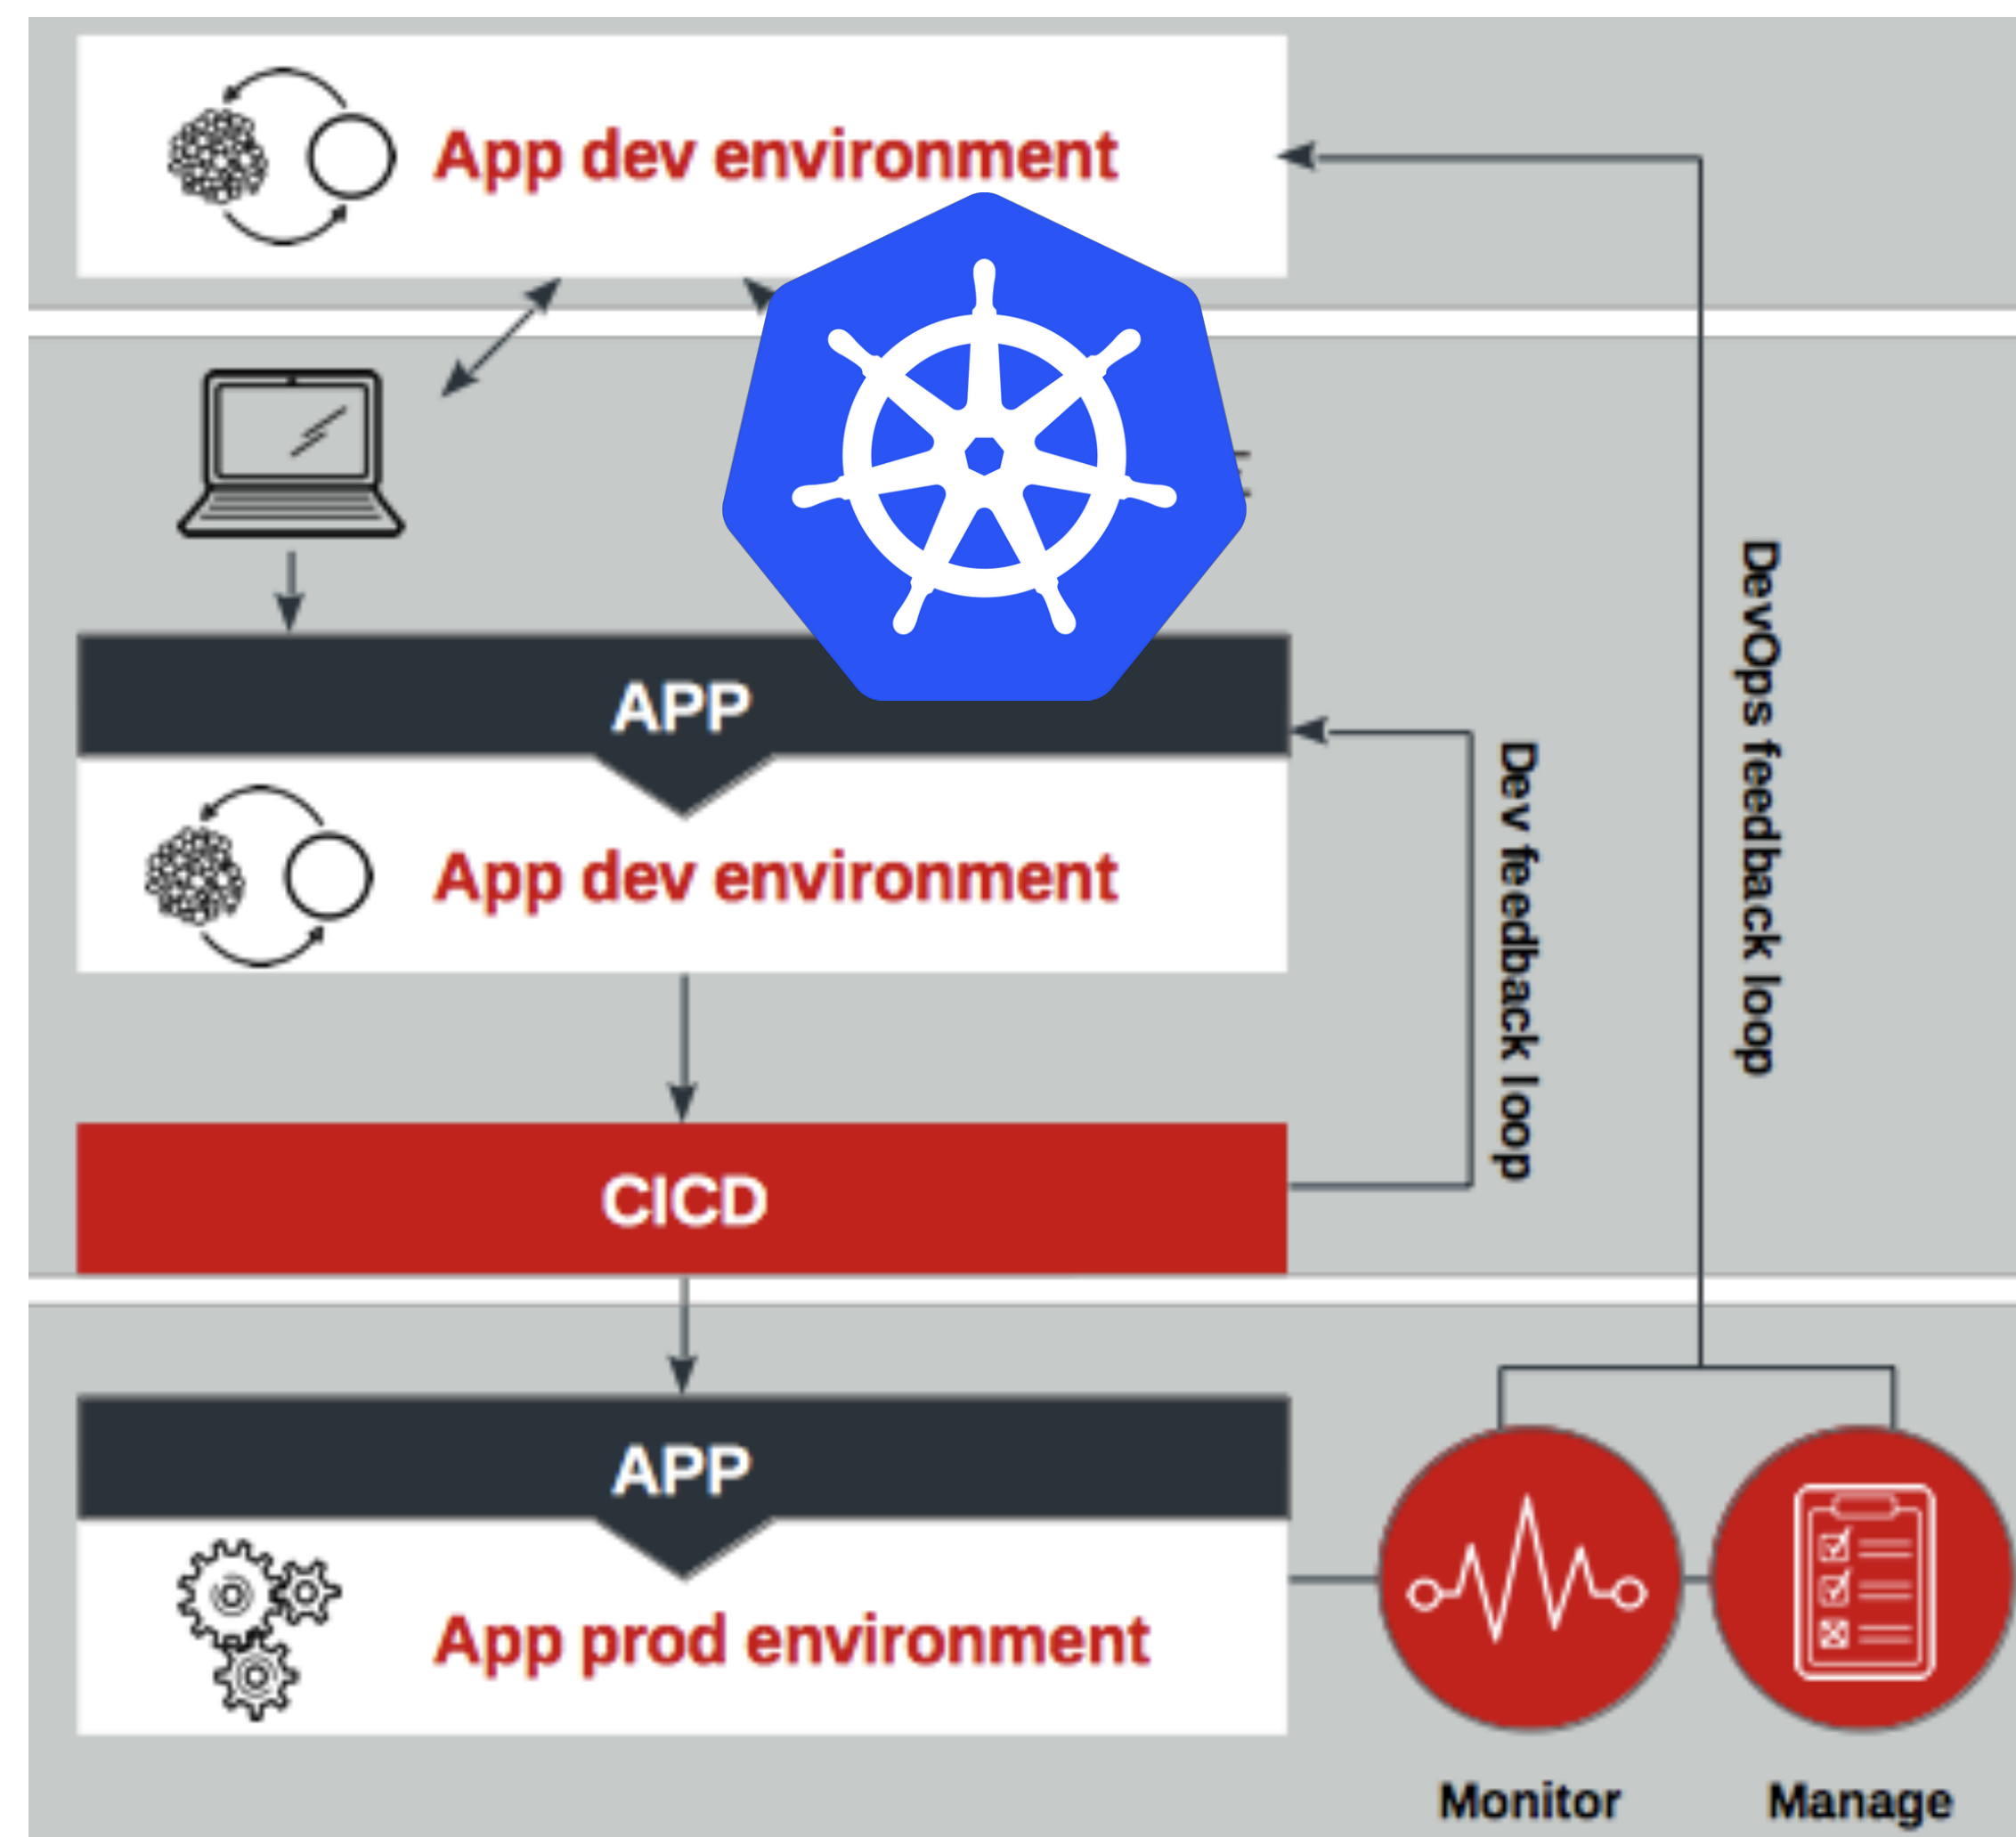


”only about 1/3 of ideas improve the metrics they were designed to improve.”

Ronny Kohavi, Microsoft (Amazon)



# CONTINUOUS FEEDBACK LOOP



# A/B TESTING USING CANARY DEPLOYMENTS

A




B




# CANARY DEPLOYMENTS

100%



Route



Version A

25% Conversion Rate



Tests / CI

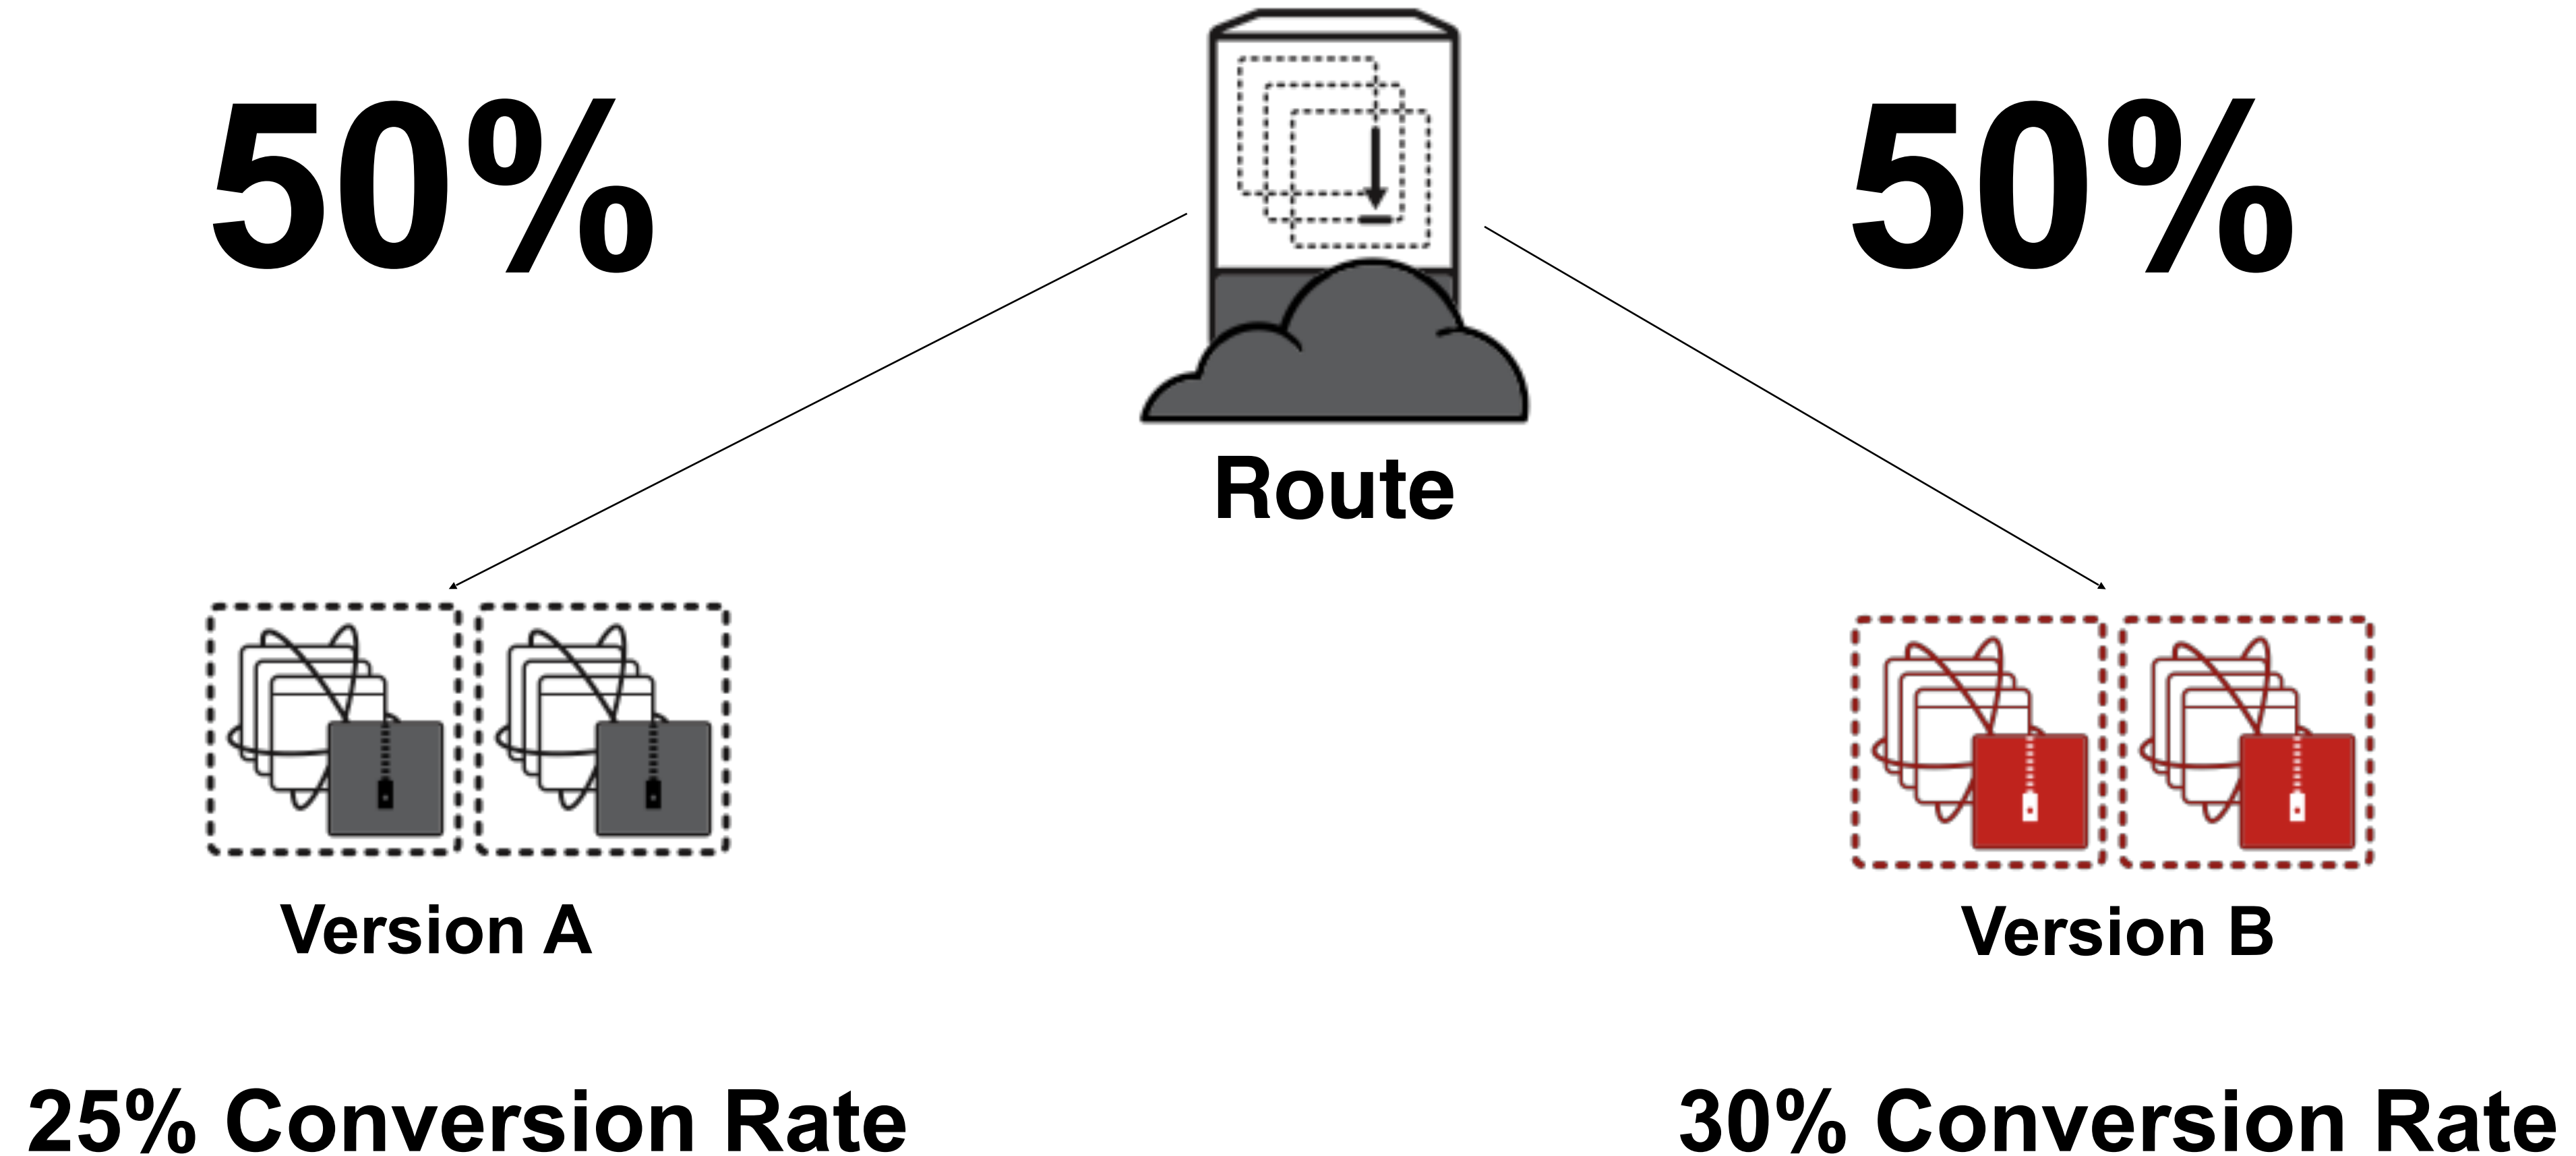


Version B

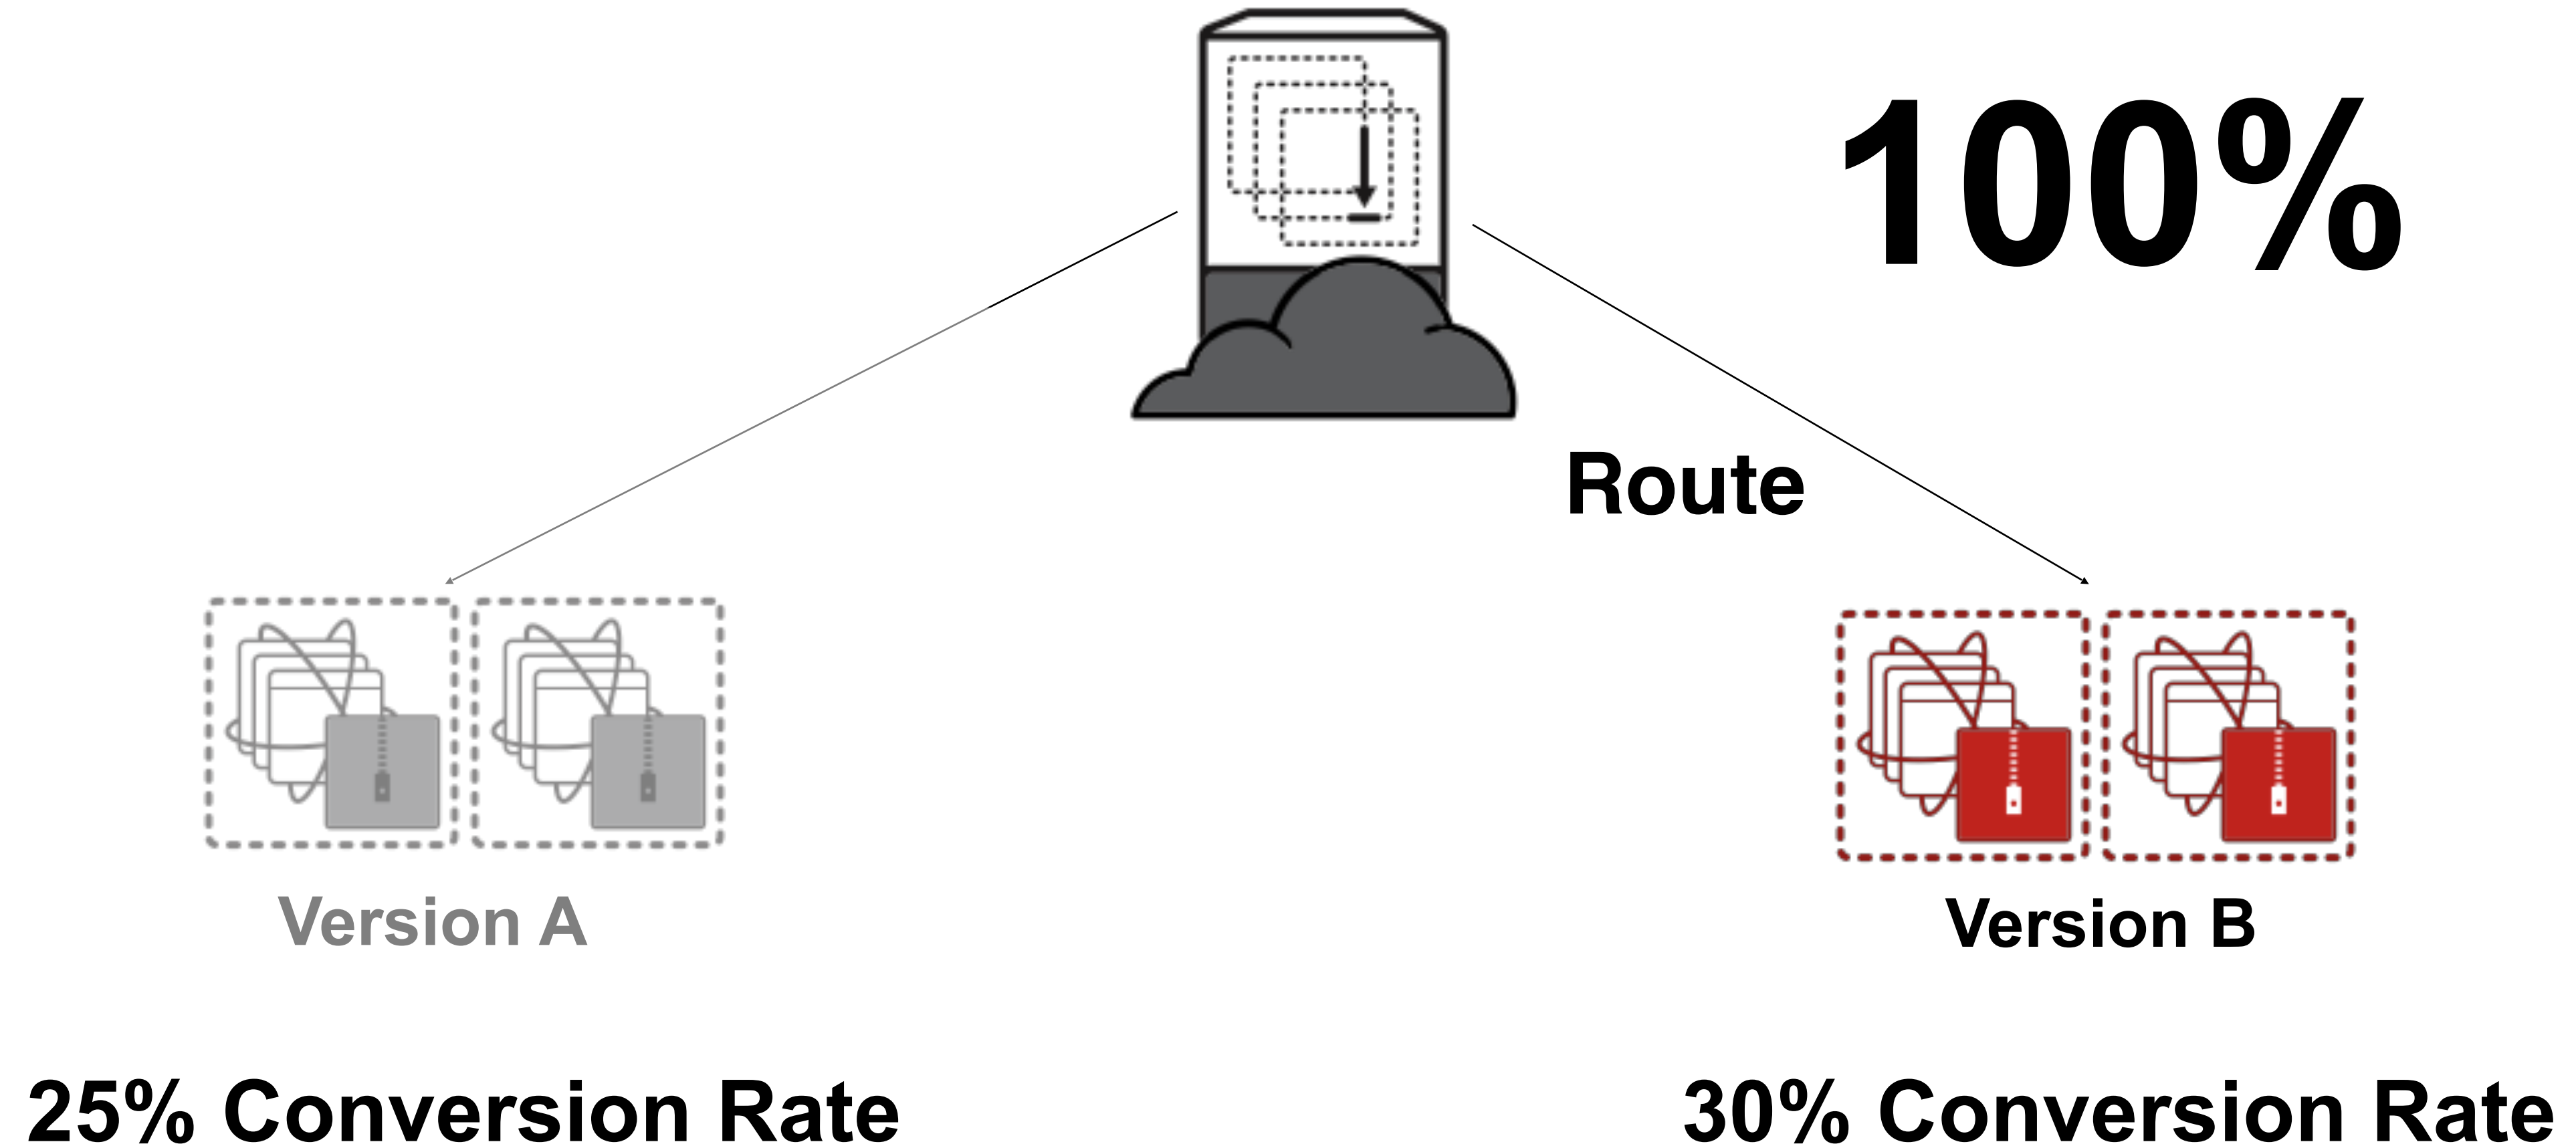
?! Conversion Rate



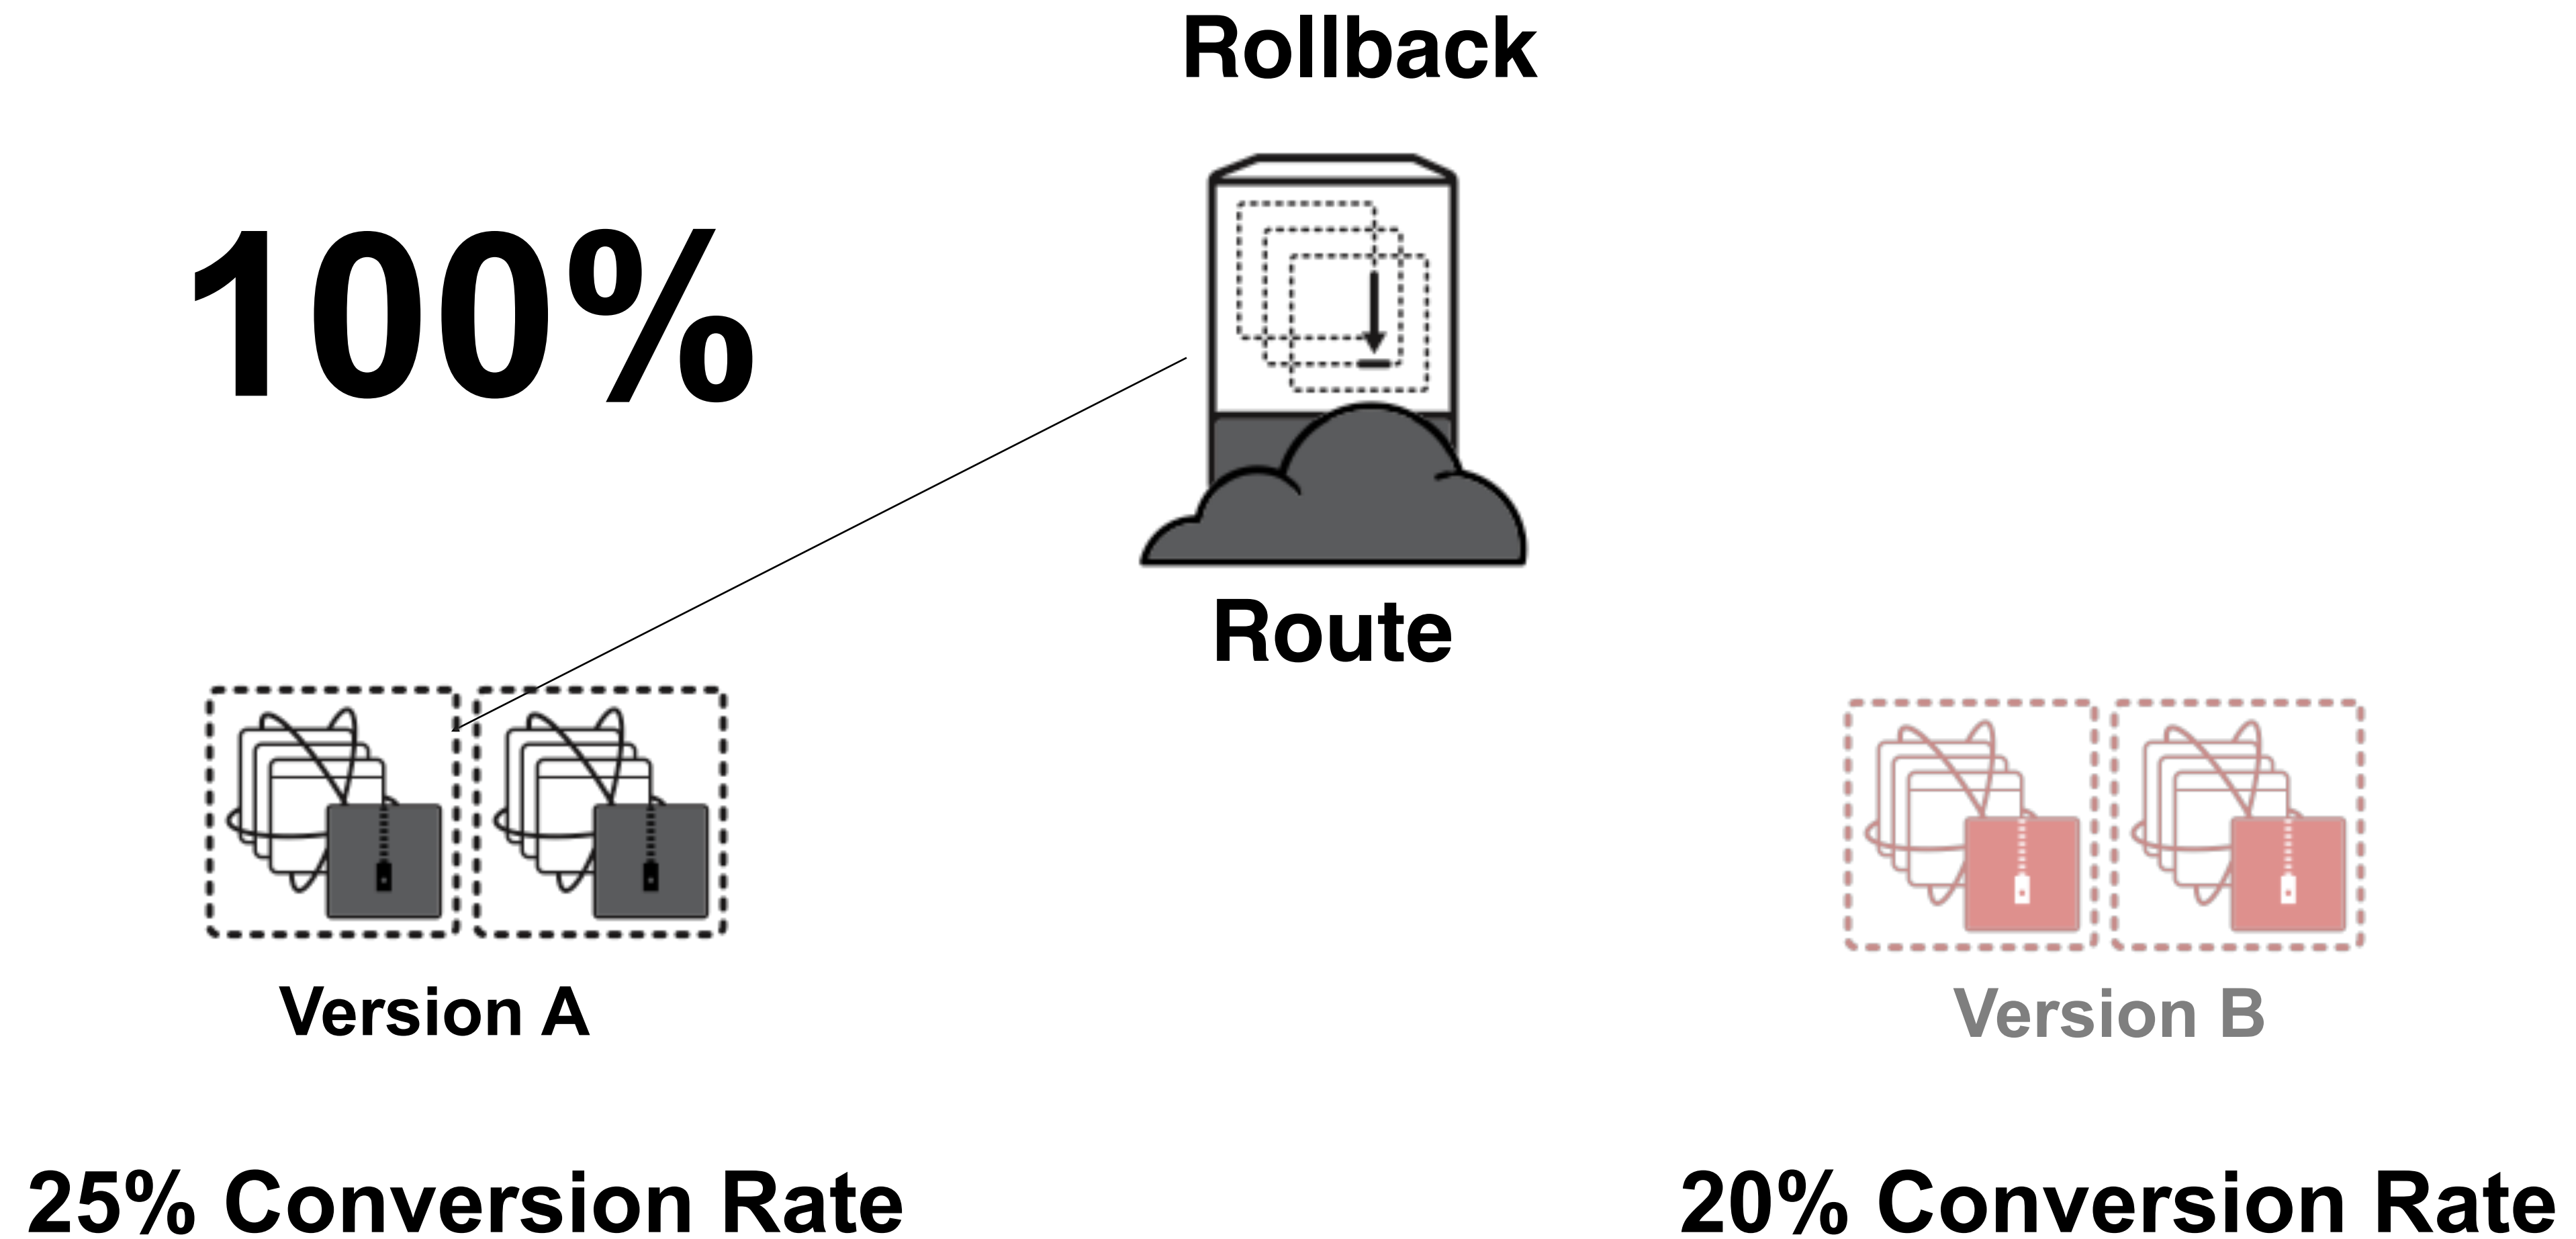
# CANARY DEPLOYMENTS



# CANARY DEPLOYMENTS




# CANARY DEPLOYMENTS







# SECURING YOUR CONTAINER ENVIRONMENT




**Builds**




**Images**




**Registry**



**Container  
host**



**CI/CD**




**Network  
isolation**




**Monitoring &  
Logging**



**Storage**



**API & Platform  
access**

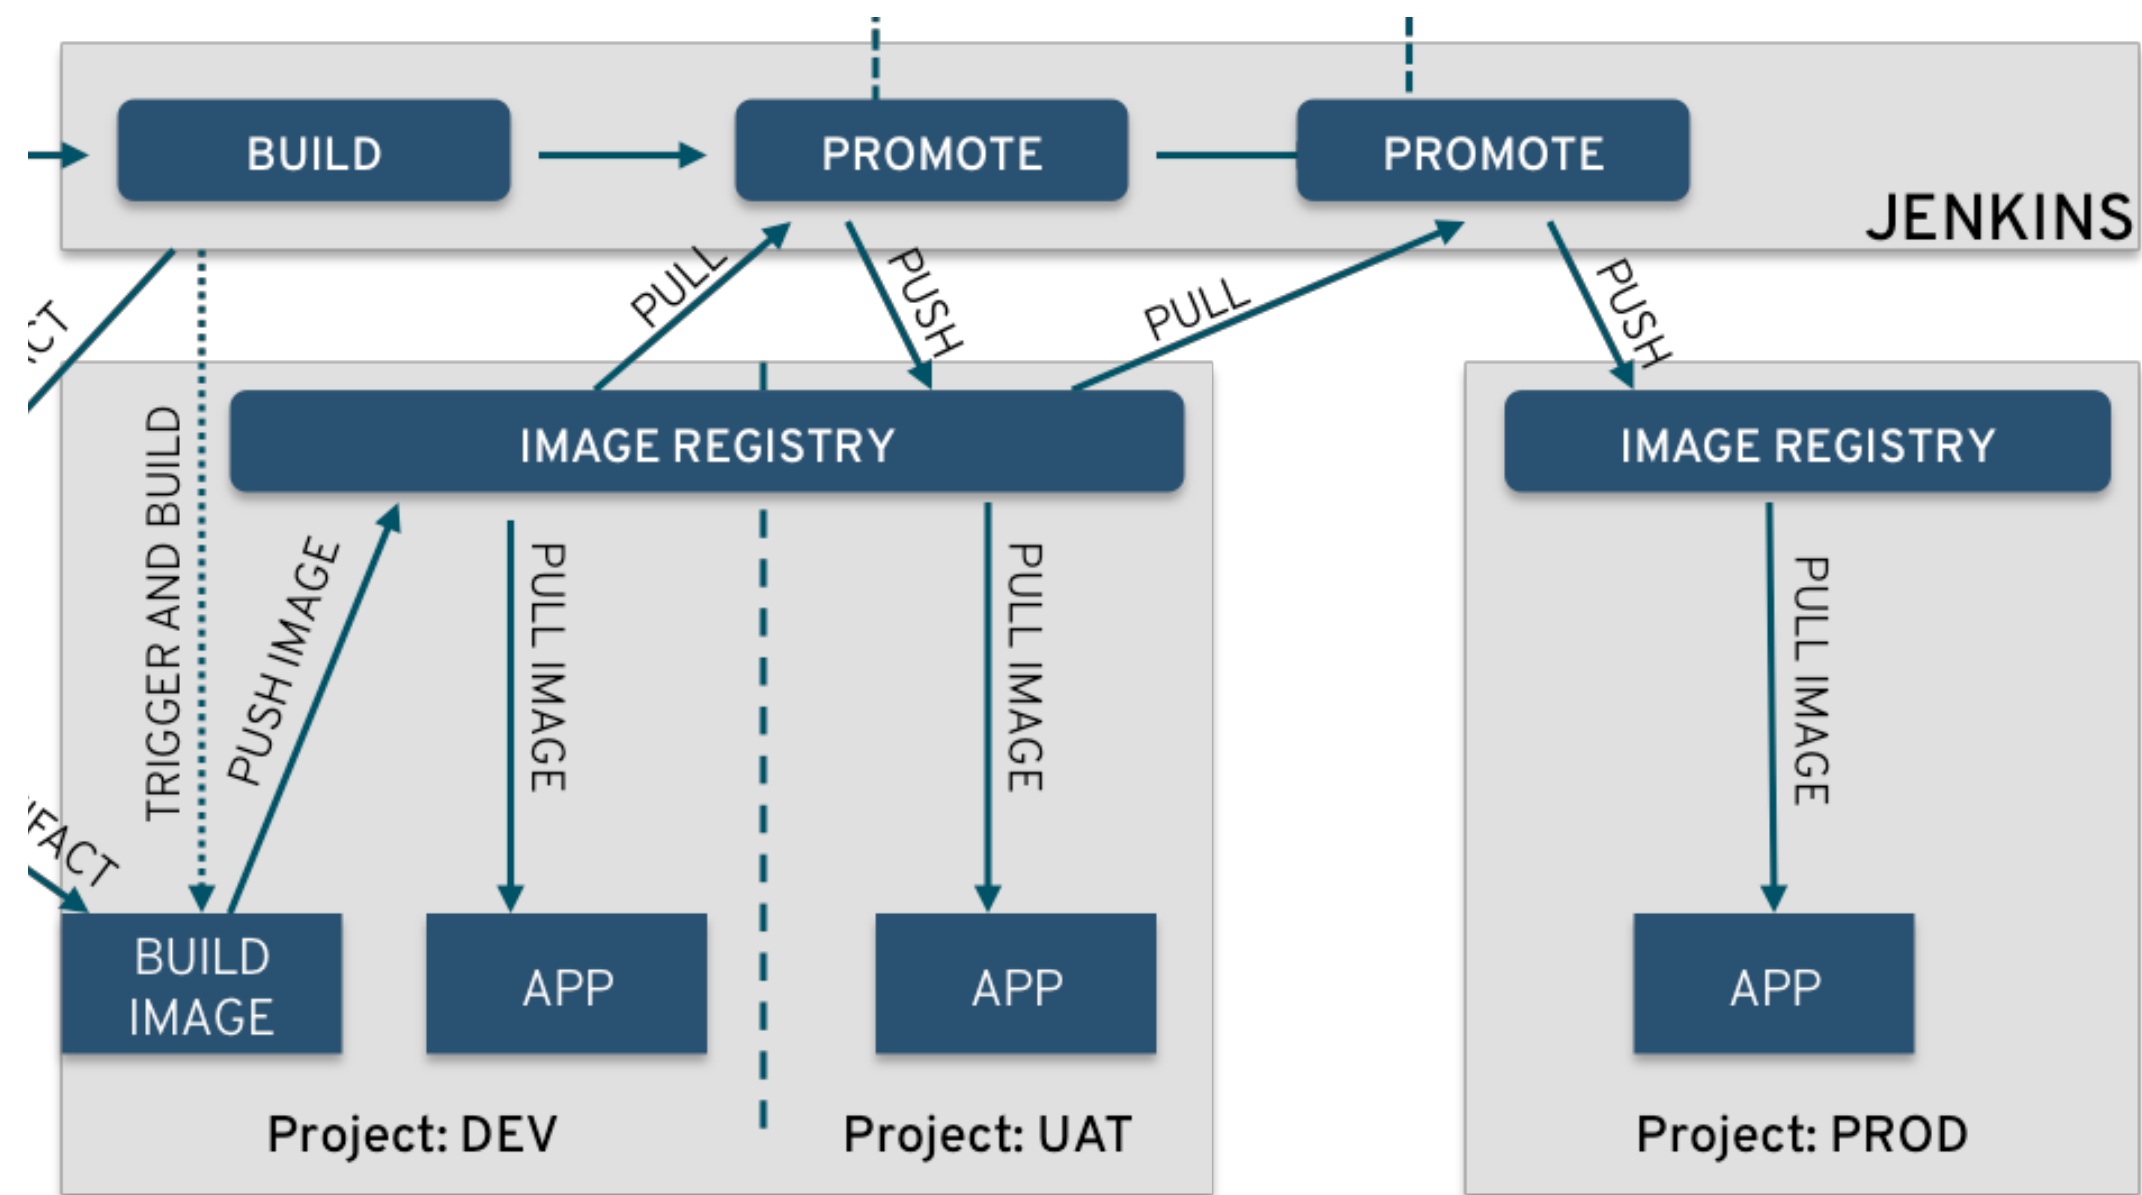


**Federated  
clusters**

# NETWORK SECURITY

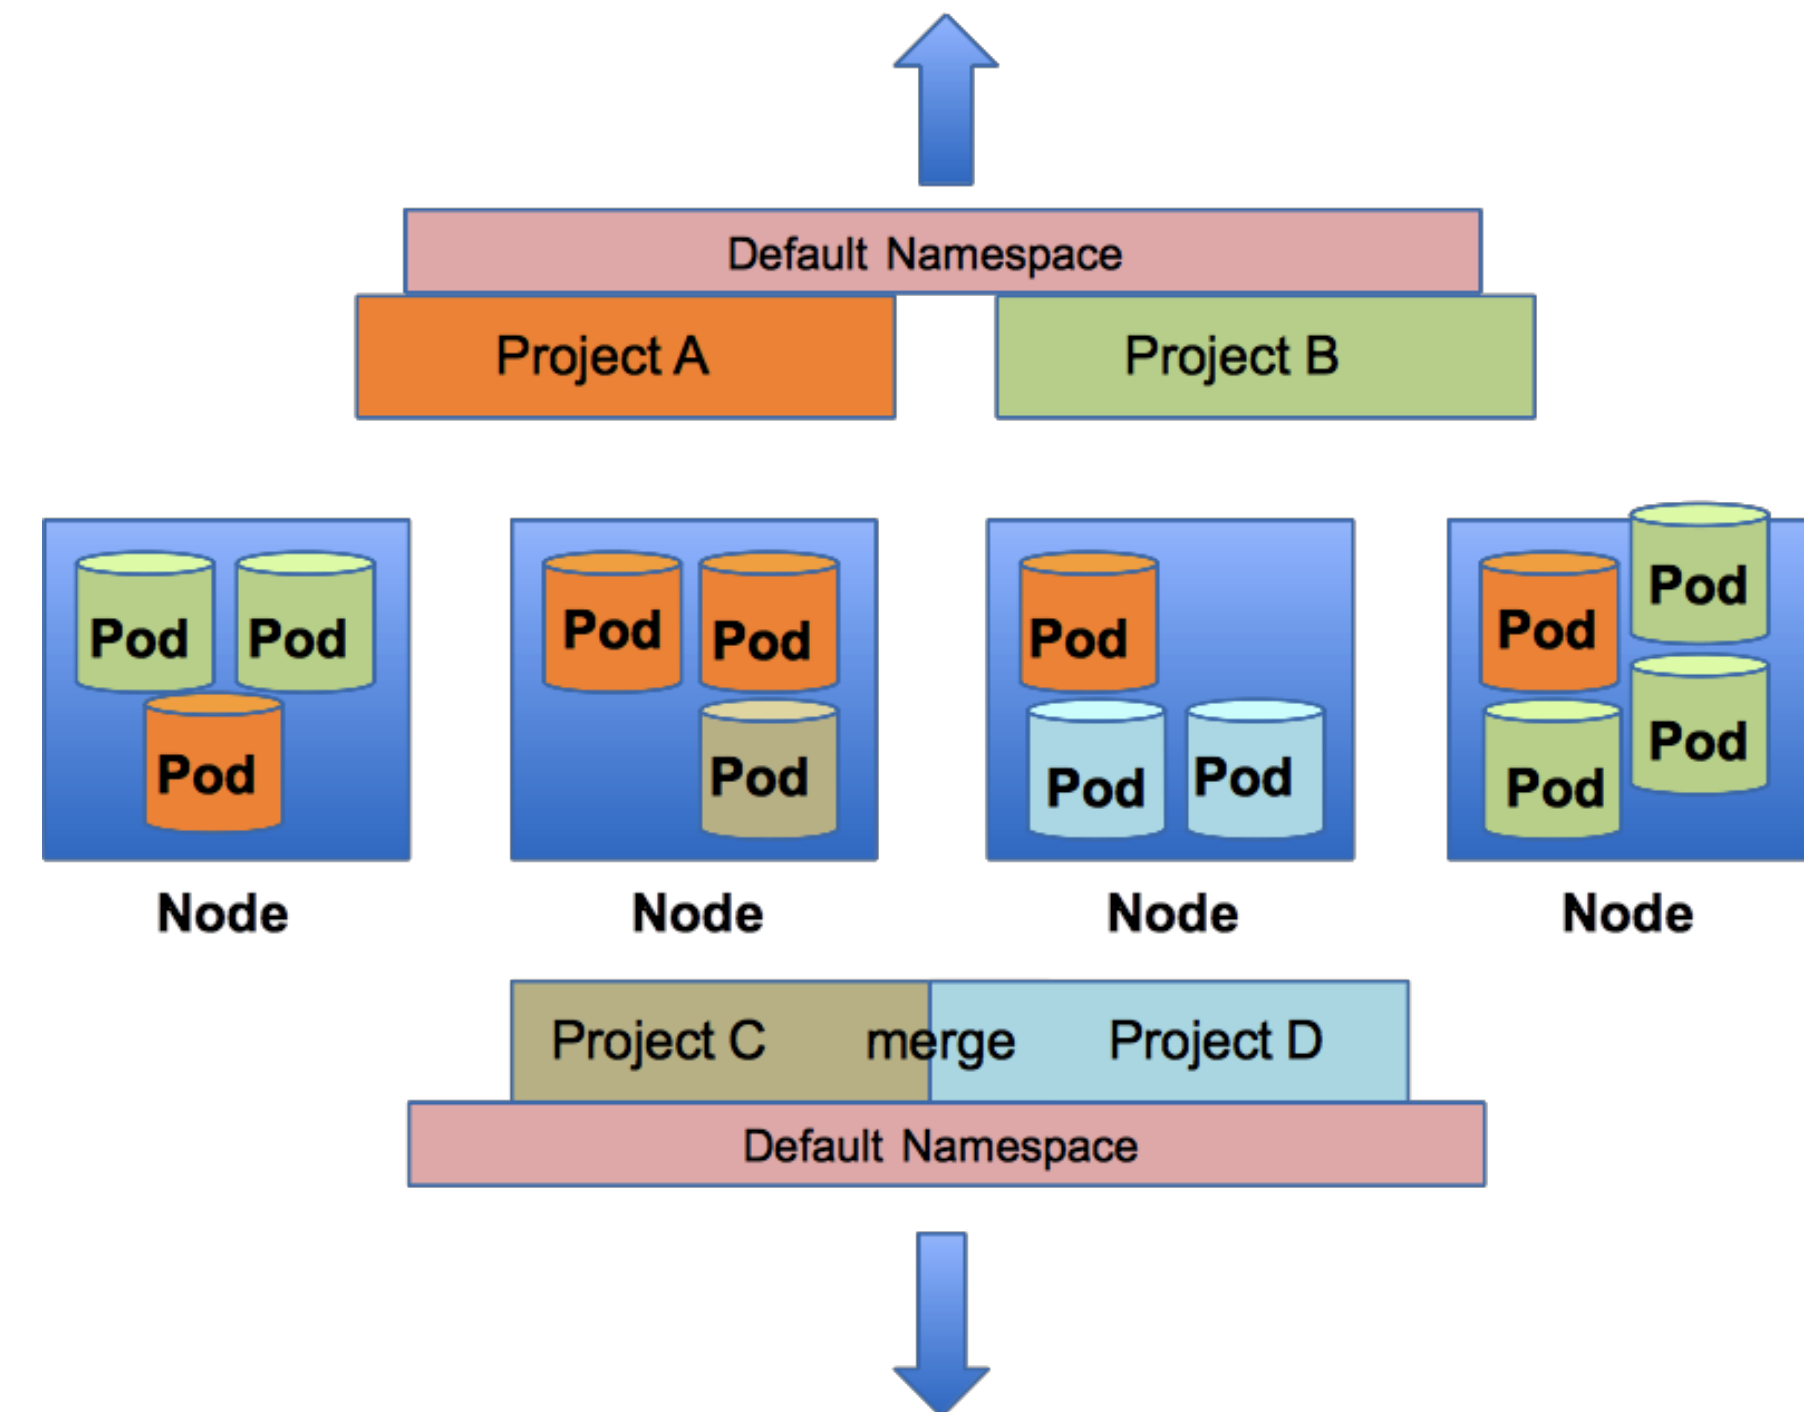
# NETWORK ISOLATION

# Multi-Environment



## Network Namespace provides resource isolation

# Multi-Tenant

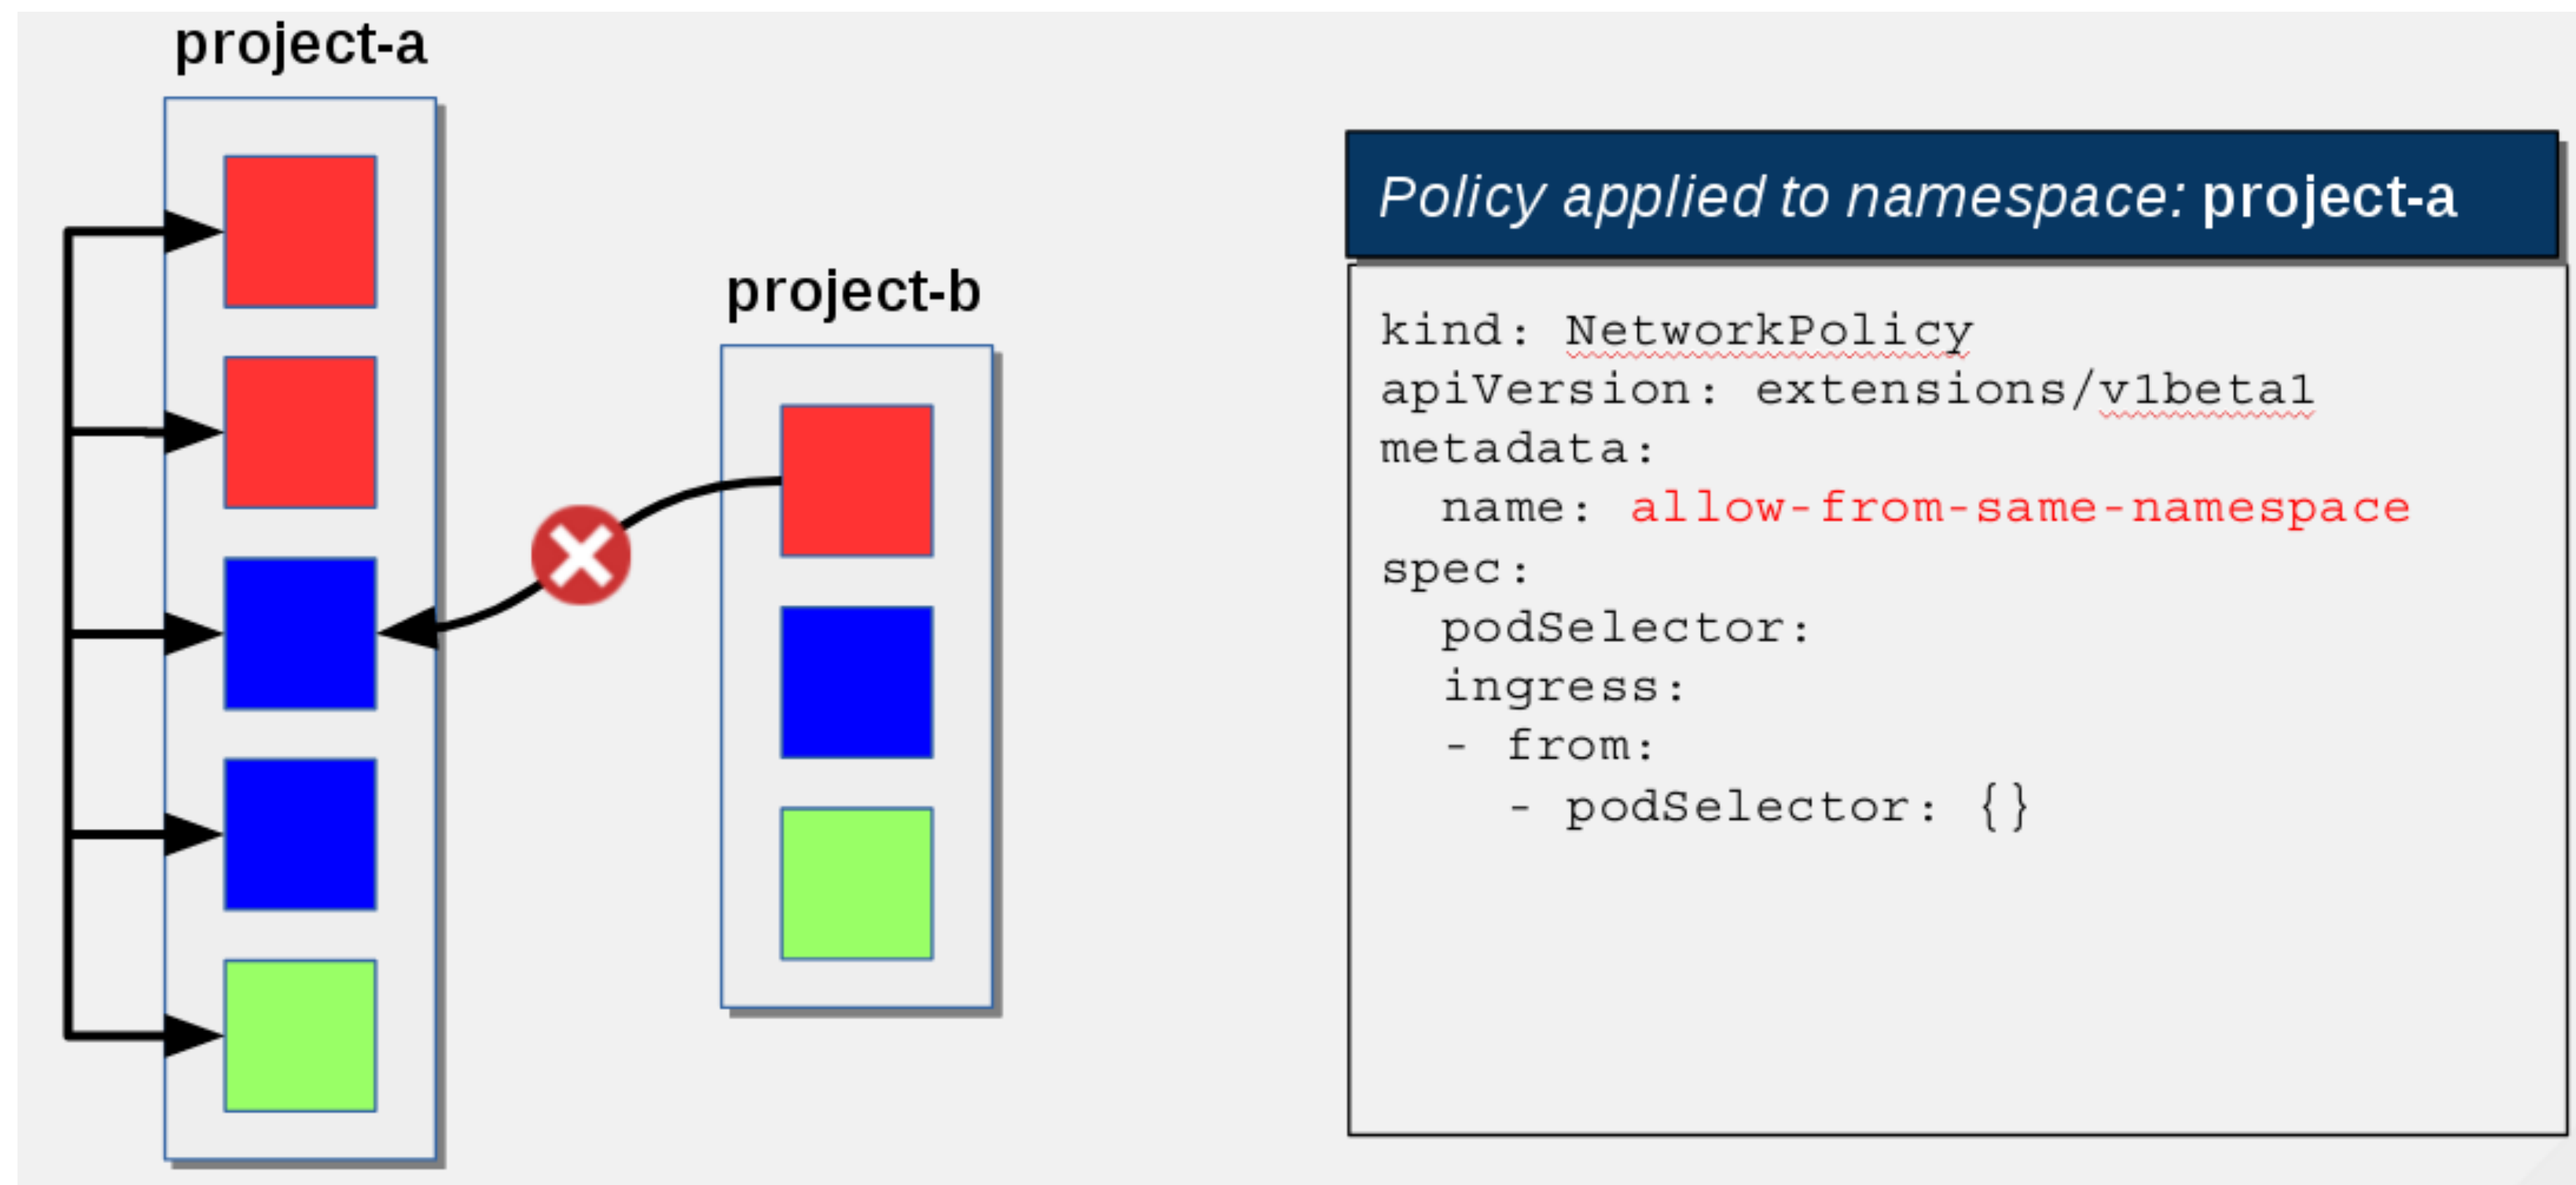




# NETWORK POLICY


example:

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




# NETWORK SECURITY

## Traditional Physical Network Model



- Each layer represents a Zone with increased trust - DMZ > App > DB, interzone flow generally one direction
- Intrazone traffic generally unrestricted

## Kubernetes Logical Network Model

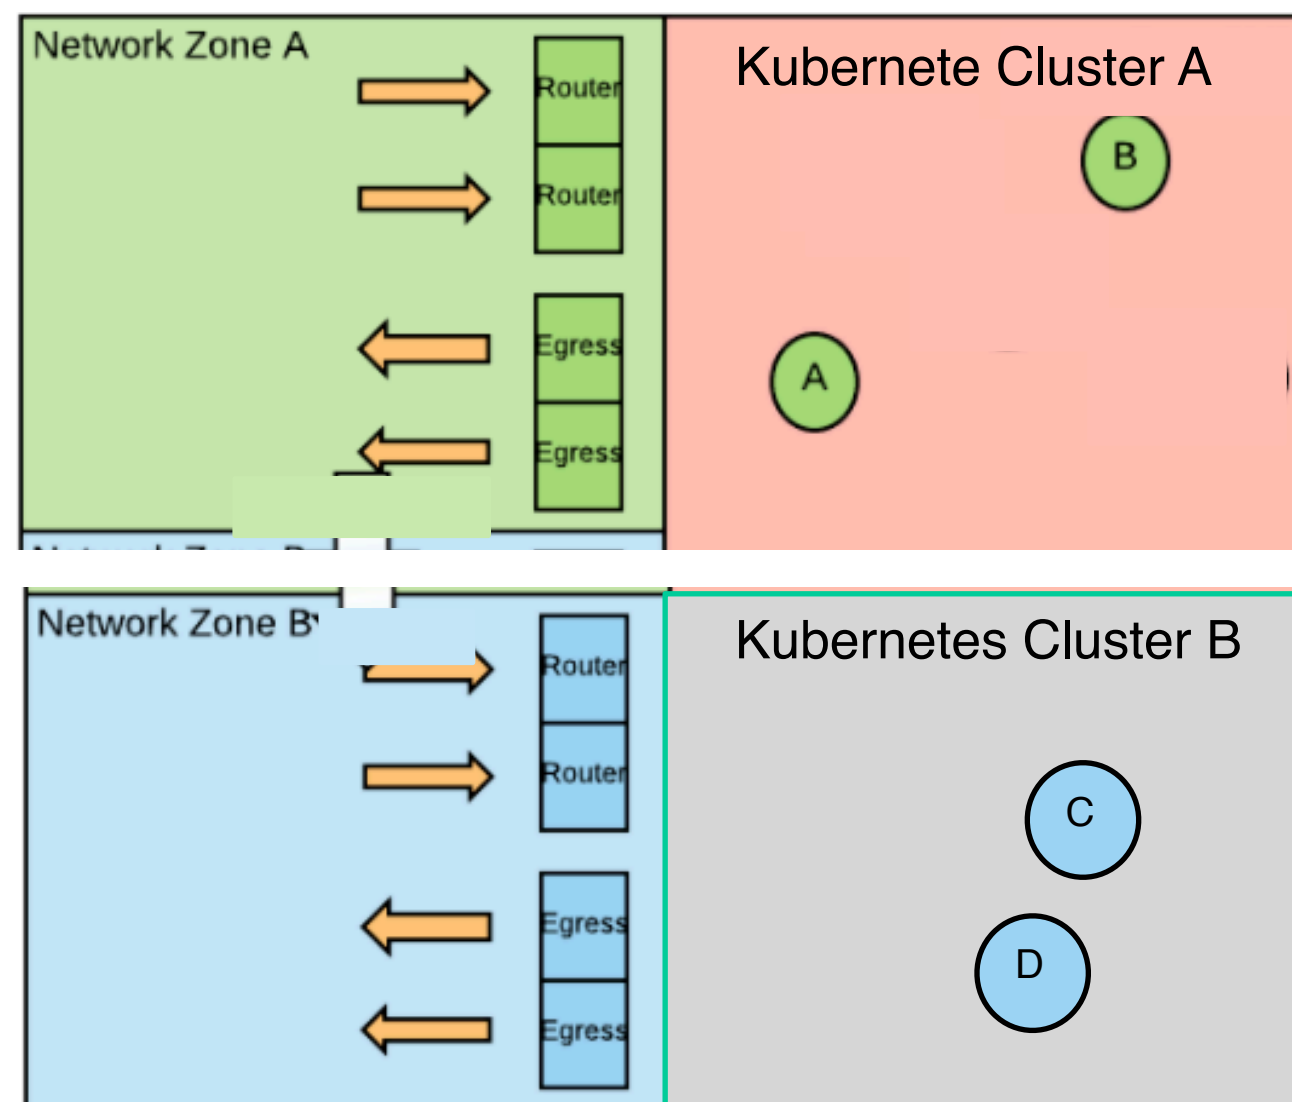


- Kubernetes uses a flat SDN model
  - All pods get IP from same CIDR
  - And live on same logical network
  - Assumes all nodes communicate

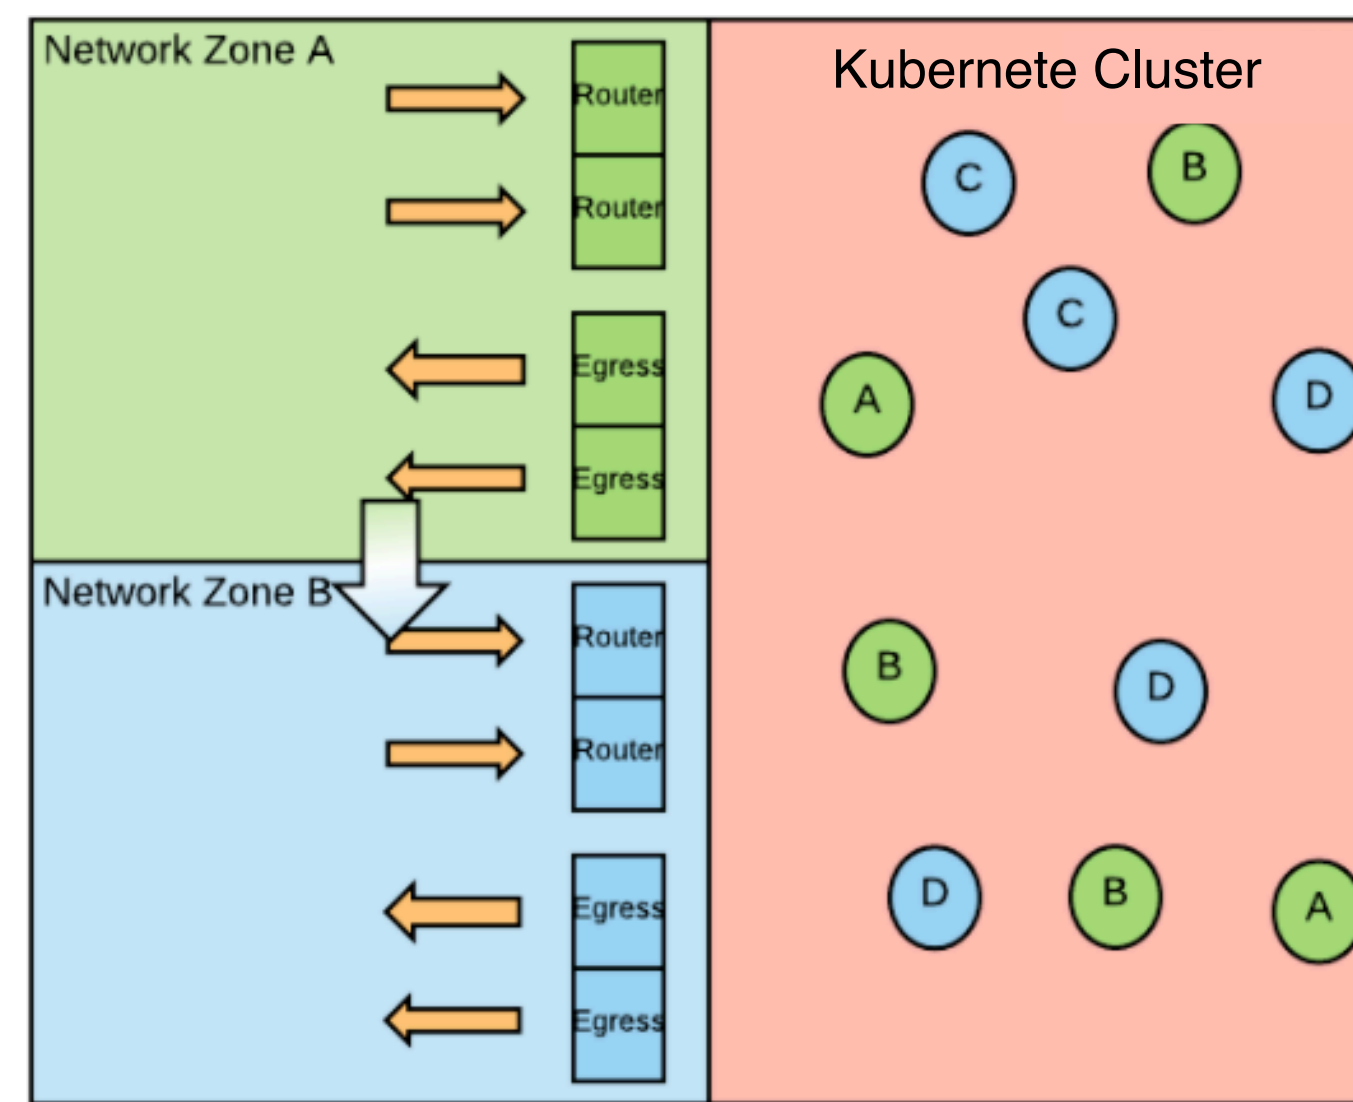
# NETWORK SECURITY MODELS

## Co-Existence Approaches

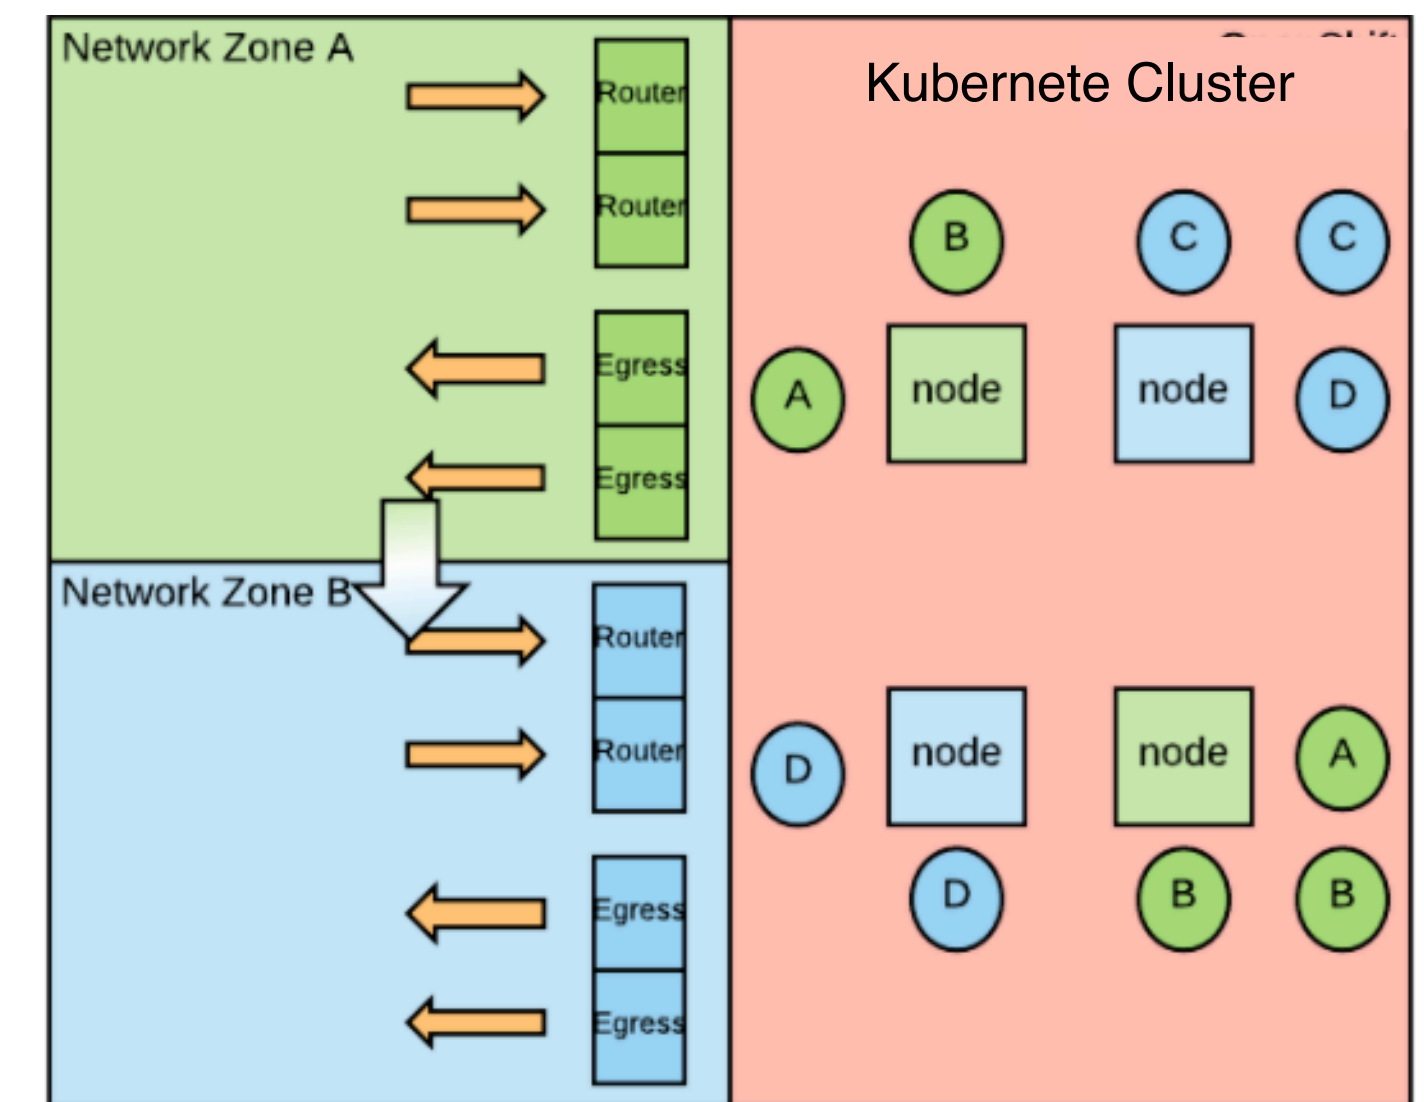
### One Cluster Per Zone



### One Cluster Multiple Zones



### Physical Compute isolation based on Network Zones



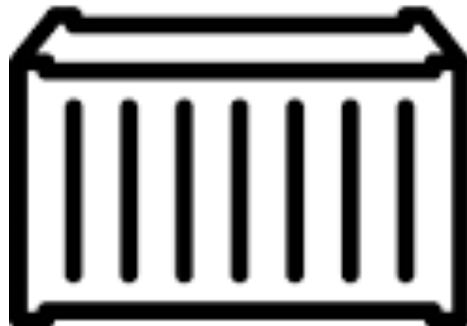
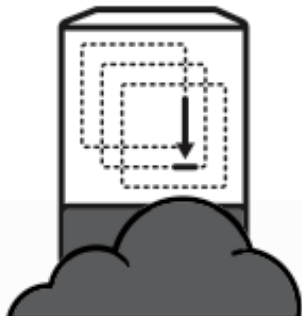


<https://blog.openshift.com/openshift-and-network-security-zones-coexistence-approaches/>



# MONITORING & LOGGING

# KUBERNETES MONITORING CONSIDERATIONS

	<u>Stack</u>	<u>Metrics</u>	<u>Tool</u>
<b>Application</b>		Distributed applications <ul style="list-style-type: none"><li>- traditional app metrics</li><li>- service discovery</li><li>- distributed tracing</li></ul>	prometheus + grafana jaeger tracing istio
<b>Kubernetes*</b>		Cluster services, services, pods, deployments metrics	prometheus + grafana kubernetes-state-metrics probes
<b>Container*</b>		Container native metrics	kubelet:cAdvisor
<b>Host</b>		Traditional resource metrics <ul style="list-style-type: none"><li>- cpu, memory, network, storage</li></ul>	node-exporter

# LOGGING

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






# STORAGE SECURITY

# STORAGE SECURITY

## Local Storage Quota



Node's /var/lib/origin

Sometimes we can also have storage isolation requirements: pods in a network zone must use different storage endpoints than pods in other network zones.

We can create one storage class per storage endpoint and then control which storage class(es) a project can use

## Security Context Constraints

# API & PLATFORM ACCESS



# API & PLATFORM ACCESS

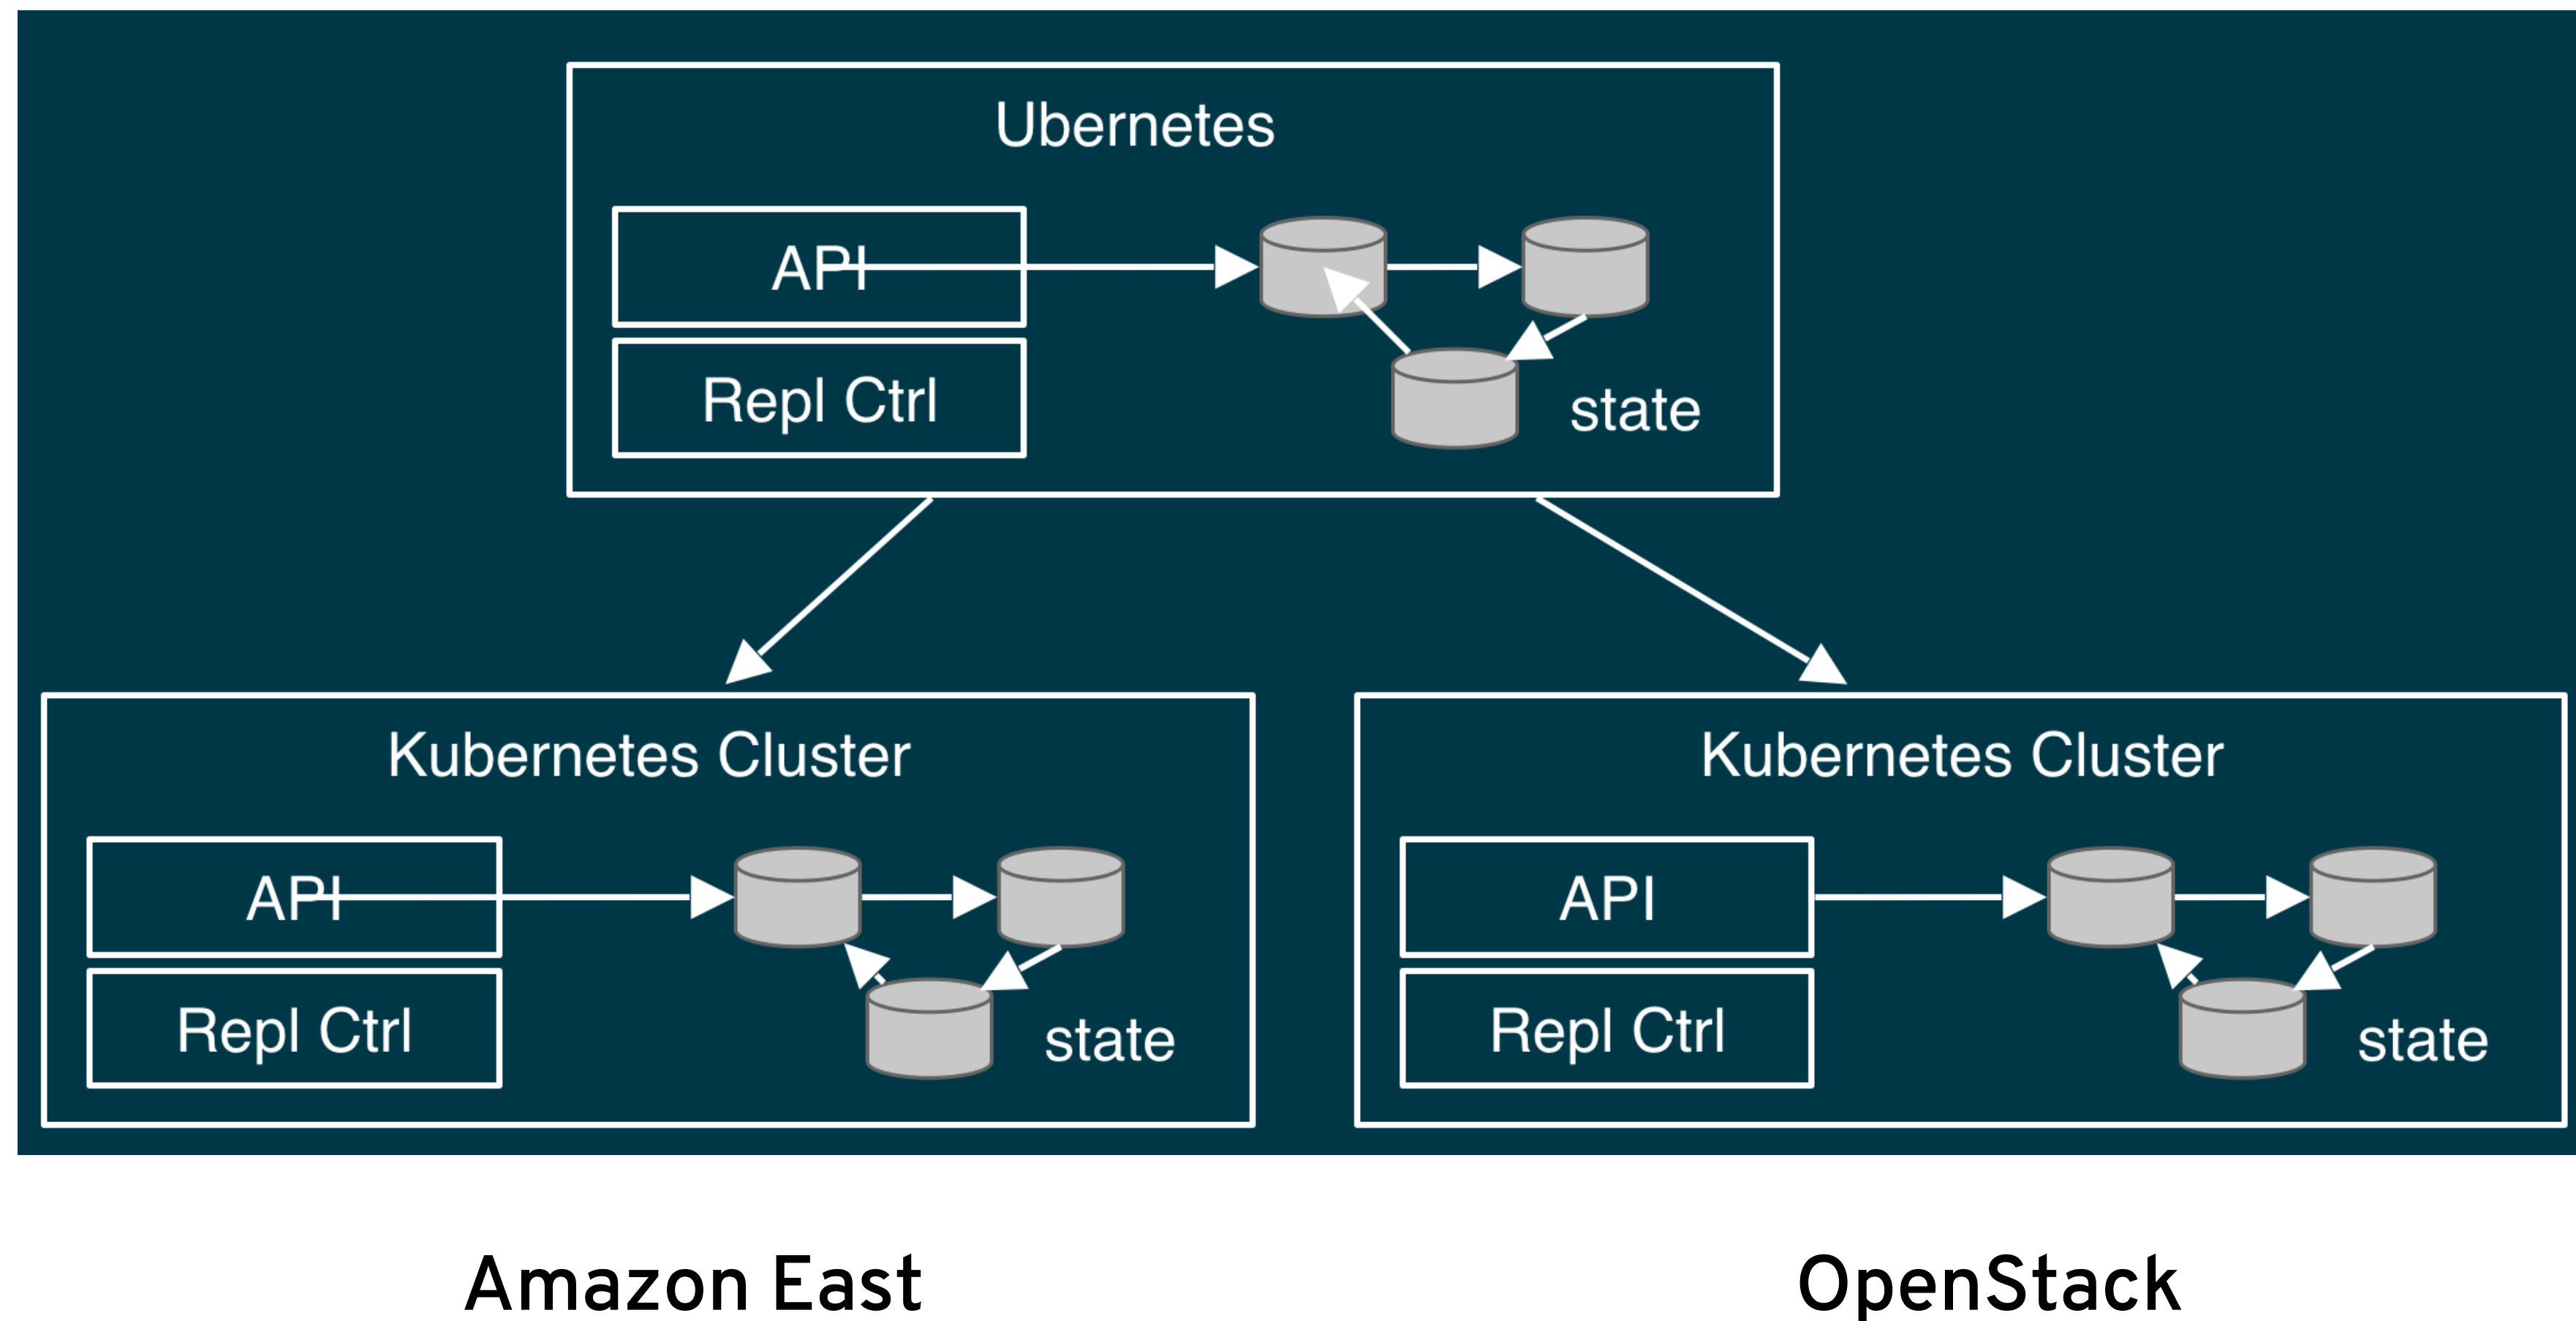
Authentication  
via  
OAuth tokens and  
SSL certificate

Authorization  
via  
Policy Engine  
checks  
User/Group  
Defined Roles

# FEDERATION

# FEDERATED CLUSTERS

Roles & access management (in-dev)







# WHAT'S NEXT



**Istio**



**Traffic  
Control**



**Service  
Resiliency**



**Chaos  
Testing**

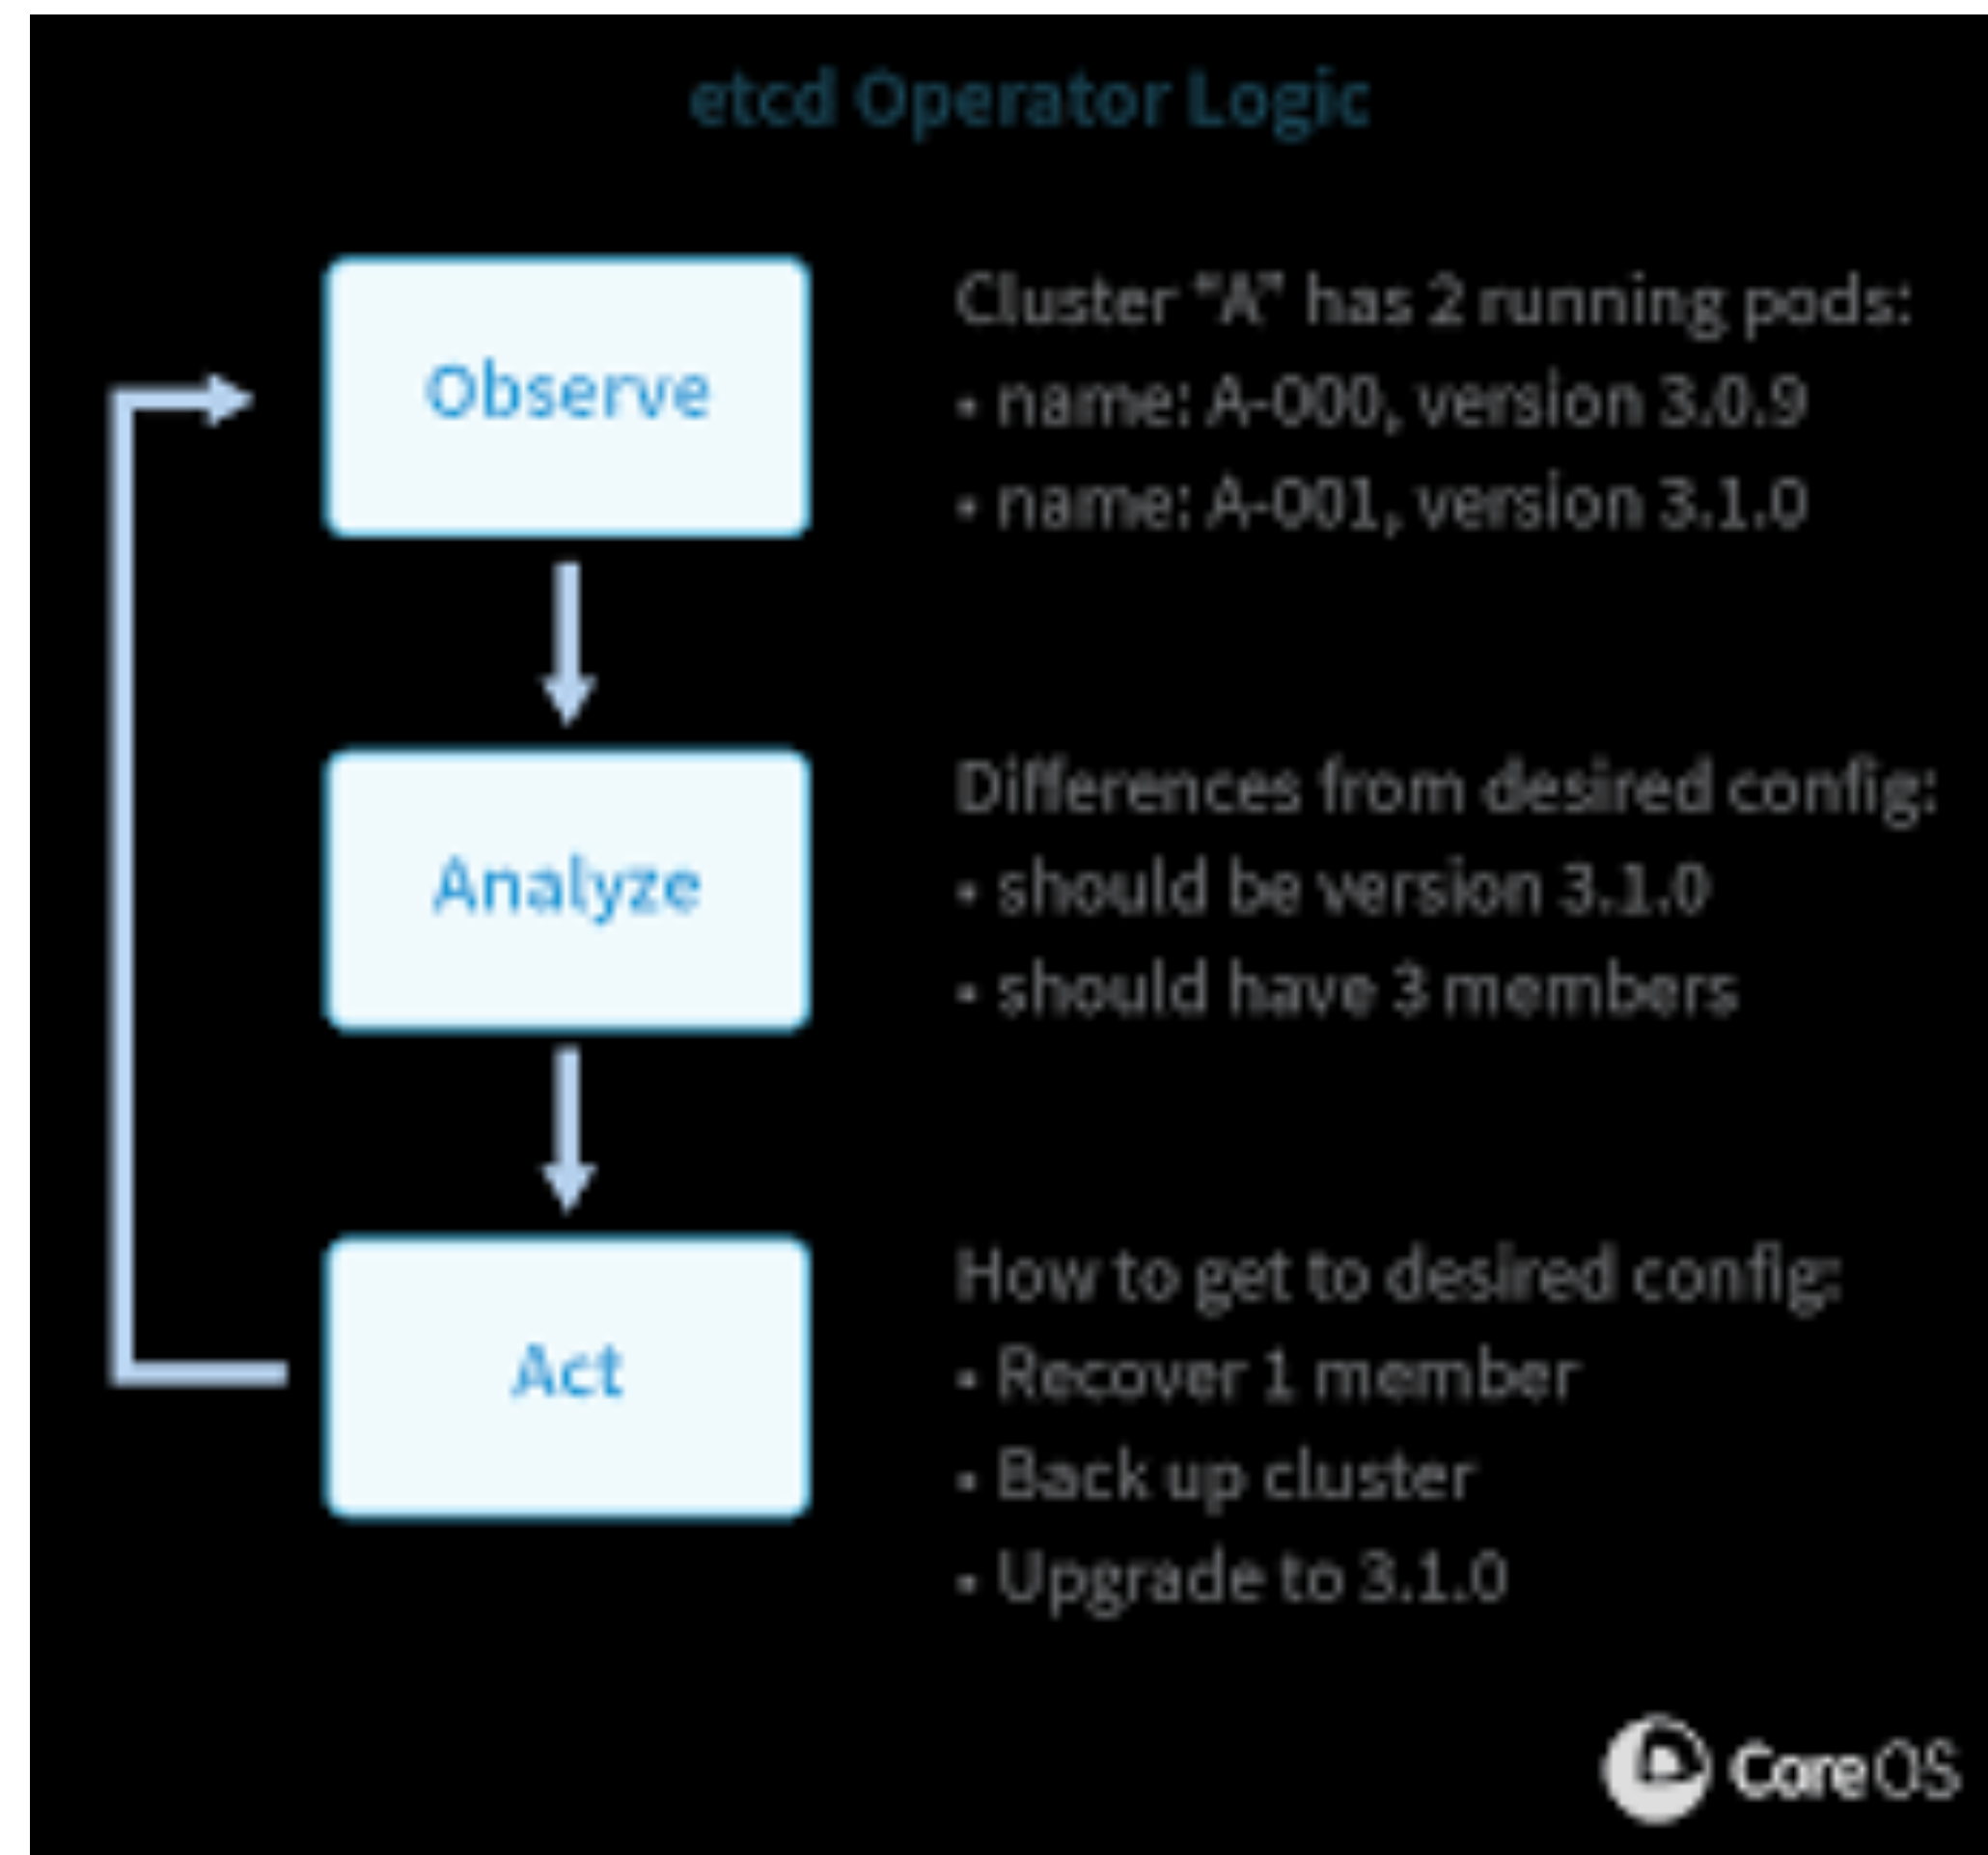


**Observ-  
ability**



**Security**

# OPERATORS



# 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)**

