

ATC Apps Journey to Private Cloud

Matt Skipton World Wide Technology
Advanced Technology Center

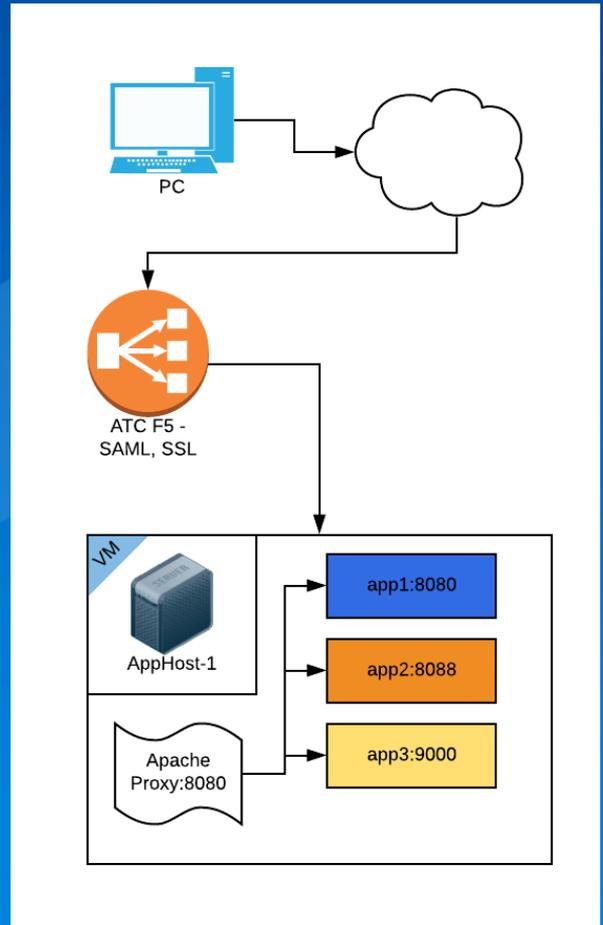
Background

- **The ATC has specific, custom, apps that need to run in the ATC**
 - POC management
 - ATC utilities
 - Platform proxies
 - Lab management
- **Different languages and stacks**
- **Microservices framework creates lots of small apps – eases development and reusability but makes traditional hosting problems worse**

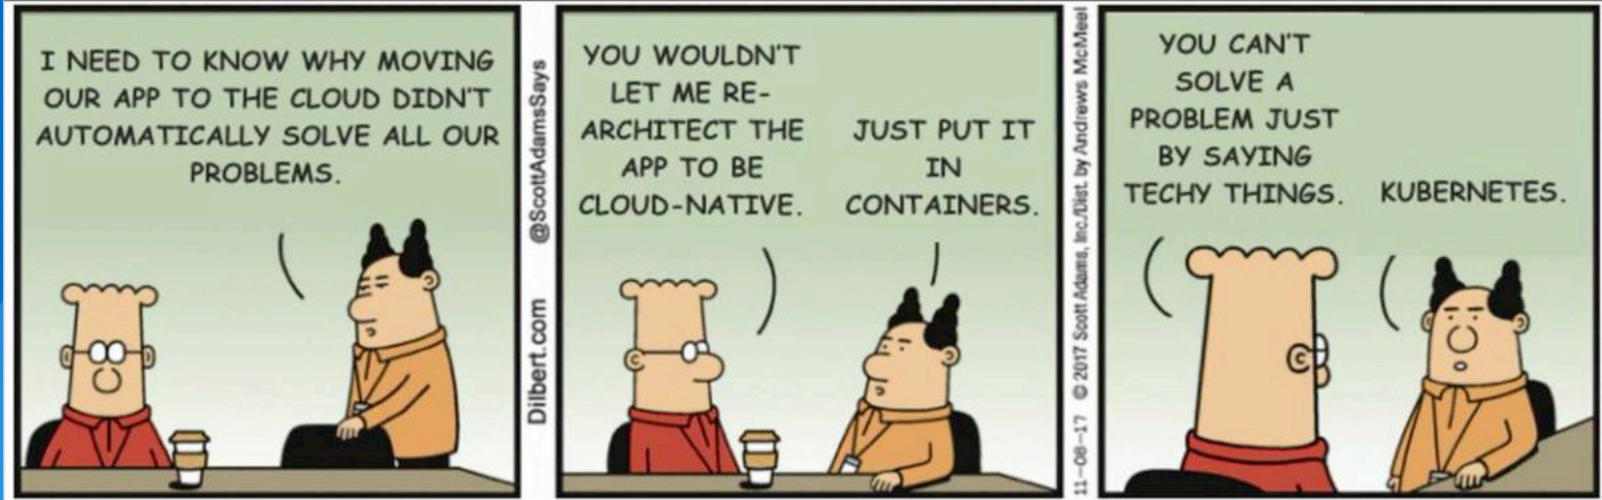


In the beginning....

- **Traditional app hosting**
- **VM's for dev/test/prod**
- **Problems with this method**
 - No app HA/scalability
 - Different versions of frameworks could conflict
 - Logging & monitoring inconsistent
 - Host(vm) maintenance required all apps to go down
 - VM's with full server OS
 - Mapping ports to apps
 - Complex/Disjointed CD pipeline
 - Complex stack more brittle
 - Slow to provision
 - Microservice routing complexity



What should we do?



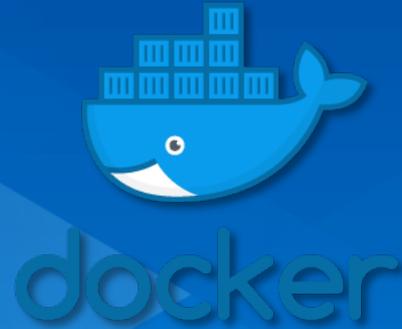
Lets Go to Containers!

- **Problems containers solve**

- Framework conflicts
- Provisioning speed
- Simplified CD Pipeline
- App build config lives with code
- Consistent deployment on dev laptop and prod

- **Unsolved Problems**

- No app HA/scalability
- Logging & monitoring inconsistent
- Host(vm) maintenance required all apps to go down
- Mapping ports to apps



Lets add a container orchestration layer!

- **There are many container orchestration solutions**
- **MesoSphere**
- **Cloud Foundry**
- **Docker Swarm**
- **Kubernetes (the industry leader)**
 - Problems Kubernetes solves
 - App HA / Scalability
 - Host(vm) maintenance without downtime
 - Automatic app communication routing
 - Unsolved problems
 - Logging becomes consistent but still needs to be sent somewhere for aggregation
 - No tracing of packets for troubleshooting
 - Possible security concerns
 - No host monitoring
 - Complex to deploy
 - RBAC is hard!



kubernetes



What is Kubernetes (k8s)?

- Started at Google
- Open Source
- Maintained by the supported by the Cloud Native Computing Foundation <https://www.cncf.io>
- Applications are defined through a declarative model in YAML
- kubectl – CLI management app
- Many GUI management options
- <https://www.cncf.io/the-childrens-illustrated-guide-to-kubernetes/>

Kubernetes (K8s) is an open-source system for automating deployment, scaling, and management of containerized applications.

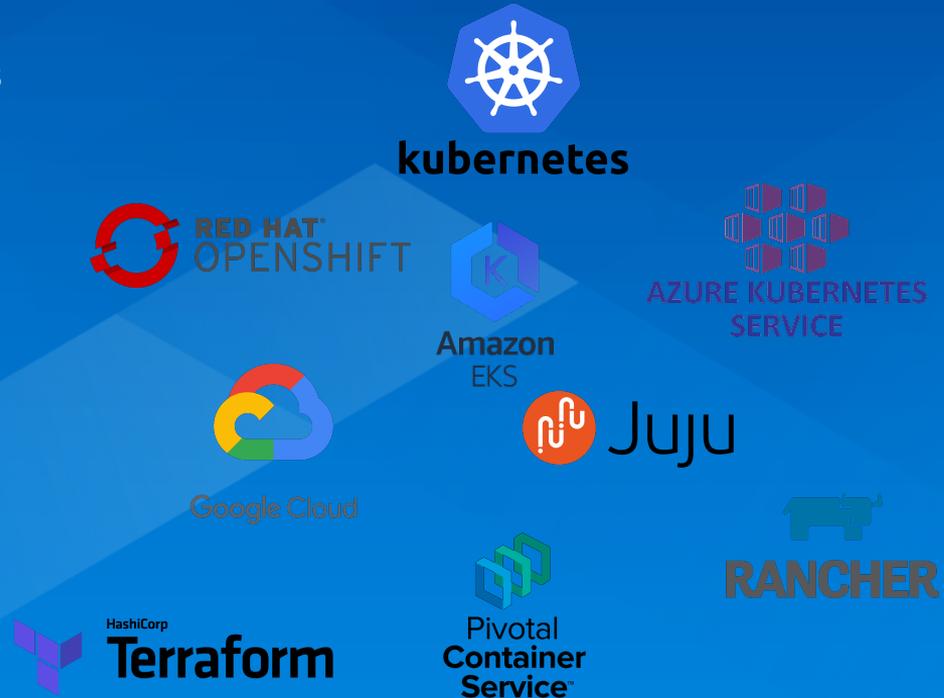
It groups containers that make up an application into logical units for easy management and discovery. Kubernetes builds upon 15 years of experience of running production workloads at Google, combined with best-of-breed ideas and practices from the community.

- <https://kubernetes.io>



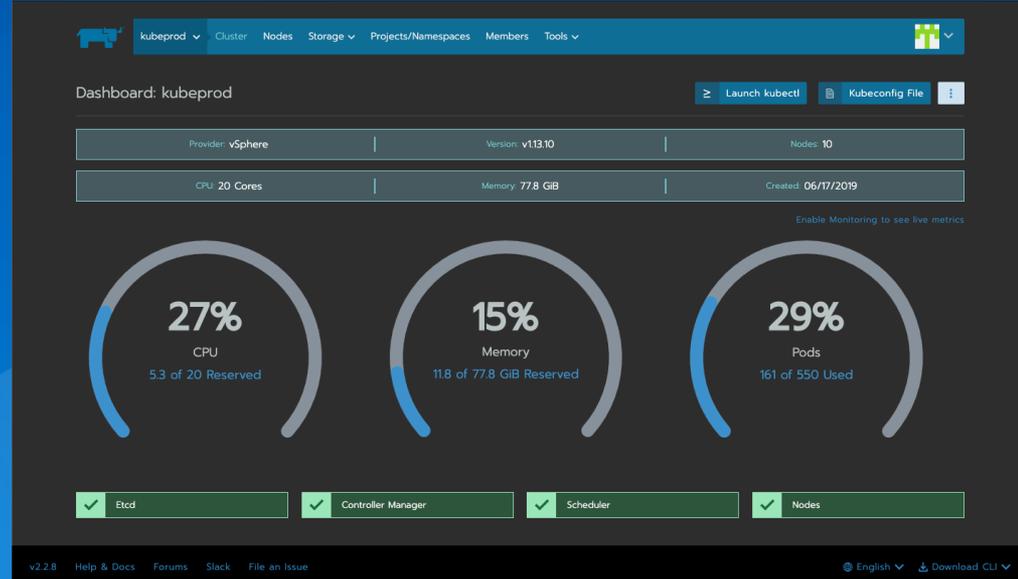
Kubernetes it is, but how to deploy & maintain?

- Plain OSS “manual install” Kubernetes
- RedHat OpenShift
- Pivotal PKS
- Amazon EKS
- Azure AKS
- Google GKE
- Ubuntu/Cononical Juju/Conjure-Up
- Terraform
- Rancher



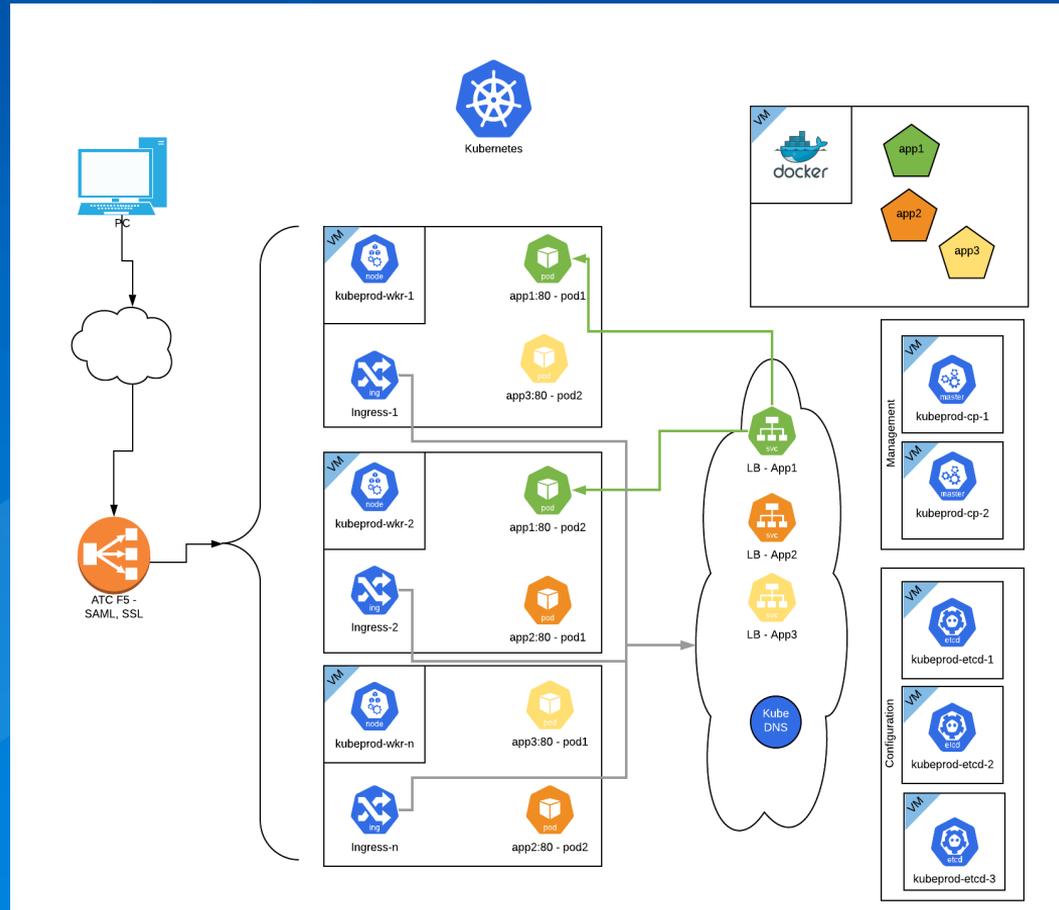
Why Rancher

- Works with on-prem vSphere and major cloud providers
- Easy to deploy / manage many Kubernetes clusters (we have 6 today)
- All VM's can be replaced at any time (cattle not pets)
- OSS Product is solid, can get support if we want
- Provides good management interface for users
- Makes Kubernetes RBAC much easier
- Provides host monitoring



How we are built

- Kube nodes in VM's for flexibility
- Worker nodes in dedicated subnet to ease load balancing
- All apps deploy with multiple replicas (with a couple exceptions)
- F5 for initial SSL termination, SAML Auth, and Load Balancing
- Custom CI/CD pipeline using Jenkins and custom code



Let's build a cluster

1. Initial cluster build
2. Deploy a simple web app
3. Deploy from Rancher catalog / helm



RancherOS

A Lightweight Container Operating System

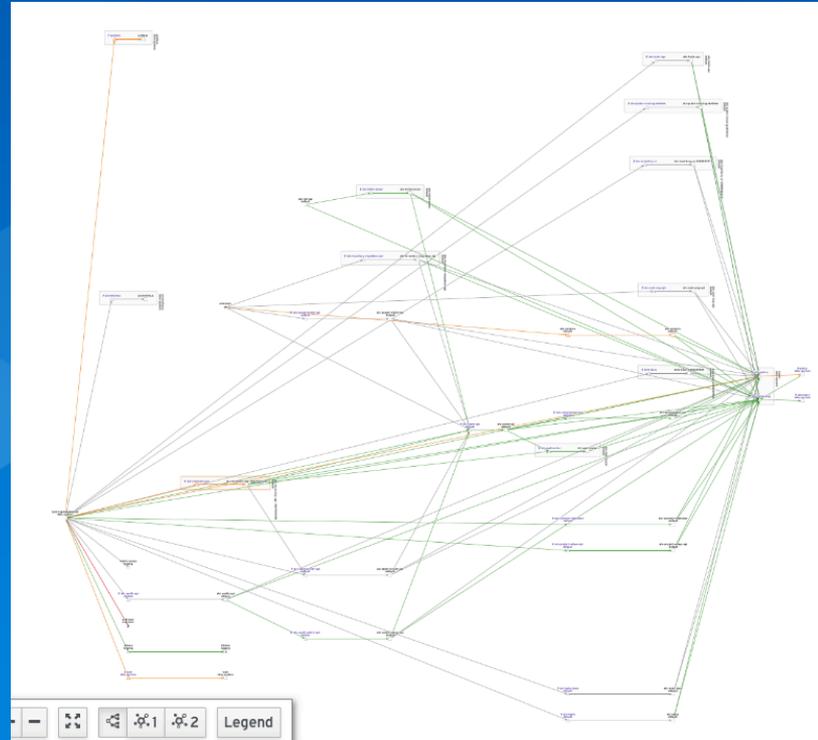
RancherOS includes the bare minimum amount of software needed to run Docker. Everything else can be pulled dynamically through Docker. RancherOS makes it simple to run containers at scale in development, test, and production. By containerizing system services and leveraging Docker for management, the operating system provides a very reliable and easy to manage container-ready environments.

- <https://rancher.com/rancher-os/>



Unsolved Problem – Understanding Microservices

- **Service Mesh – Istio**
 - Why we use it today
 - Understanding service dependencies
 - Network level tracing (up to layer 7)
 - Impose rules at layer 7 for specific apps
 - http status monitoring (are we erroring often)
 - Additional security
 - Other features for the future
 - Zero Trust environment support
 - Added code level features - circuit breakers



Benefits we have seen with this configuration

- We have re-deployed all nodes many times without service interruption
- Clean consistent CD pipeline
- Fast deployment of new apps in HA config
- Made us aware of application level issues we didn't have visibility to before
- Troubleshooting is easier



Conclusions

- **Containers and Kubernetes and Service mesh's can solve a lot of app deployment issues and they are advancing at a rapid pace**
- **They are made up of many different tools that require a lot of choices and automation to work well**
- **Kubernetes is very complex to deploy but, many of our partners have developed solutions to help get it running**
- **Service Mesh can add a lot of visibility/functionality/security to your current apps but requires a lot of configuration to implement**



Definitions

- **Container**

- A standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another. Container images are lightweight, standalone, executable packages. When a container is launched it is isolated as much as possible from the OS which provides security and consistent behavior.

- **Kubernetes**

- A container orchestration layer. It provides features on top of existing container benefits. Such as: network mapping, load balancing / HA, scalability, storage mapping, and management.

- **Pod**

- A grouping of containers that execute on the same host. They are managed as 1 unit.

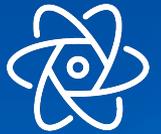
- **Sidecar**

- A container in a POD that is not the main application. It provides additional services to the main application. Ex: Istio proxy where all network traffic routes through it on the way to the main container.

- **Service Mesh**

- A grouping of proxies that live next to each application, have configurable rules, low-latency infrastructure, and high throughput. They provide rule driven routing, service discovery, load balancing, encryption, observability, traceability, authentication and authorization.





Questions?