

Kubernetes for Newbies



Greg Horie

Overview

- Linux Containers
 - Quick intro / recap
- Kubernetes - What is it?
- Kubernetes Architecture
- Live Demo
 - Kubernetes Pod
 - Kubernetes Deployment
 - Kubernetes Service
- Summary
- Future Topics




Linux Containers - Quick Intro / Recap

- A lightweight method for running multiple applications under one Linux host
- **Feature - Portability**
 - Each container encapsulates its dependencies
 - Ensures consistent behavior across environments
- **Feature - Efficiency**
 - Containers share the host's kernel
 - Lower overhead compared to a virtual machine
- **Feature - Scalability**
 - Containers can be spun up or shut down quickly
 - Enables rapid deployment and scaling



Containers - The Challenges

- Container runtime runs on a single host - great for dev, but not prod
- Making containers production-worthy requires more
- Challenges to consider:
 - Redundancy - i.e. multiple hosts in case of failure
 - Networking across redundant hosts
 - Shared file systems, configurations, secrets
 - Load balancing
 - Scheduling workloads
 - etc.



Kubernetes - What is it?

- K8S = Kubernetes
- Kubernetes is Greek for helmsman or pilot
 - Following the container / Docker shipping metaphor
- K8S is an open-source container orchestration platform
- Automates the deployment, scaling, and management of containerized applications across a set of hosts (nodes)



kubernetes

Kubernetes - History

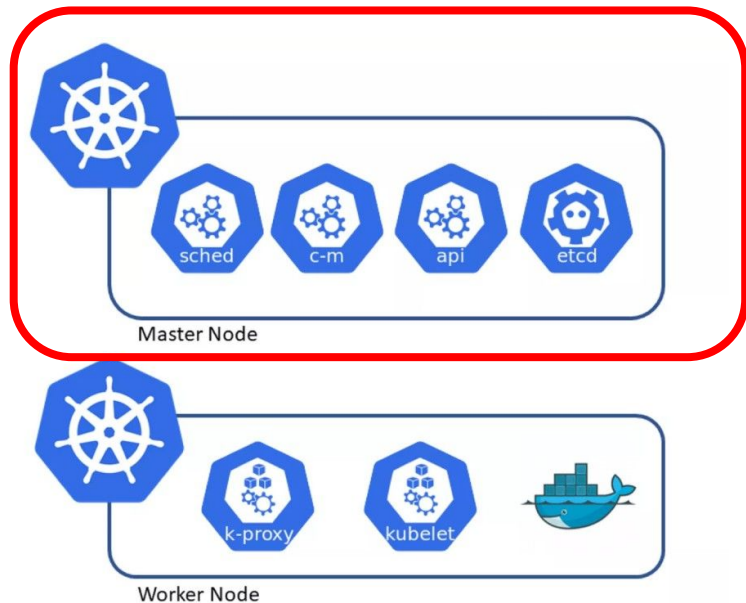
- **2003-2004 - Borg**
 - Early orchestration platform built at Google to manage container-based applications
- **2014 - Kubernetes open-sourced**
 - Google engineers open sourced K8S (based on Borg)
 - Written in the Go programming language
- **2015 - Kubernetes 1.0 released**
 - Google partners with the Linux Foundation
 - Forms the Cloud Native Computing Foundation (CNCF)
 - CNCF goes on to host many open source projects - containerd, prometheus, etcd, etc.
- **2017 - Winner of the container wars**
 - Industry rallies around K8S - Docker, Microsoft AKS, Amazon EKS, etc.
- **Today**
 - Continues to evolve with a strong community contributing to its future



Kubernetes - Architecture

Control plane (master) nodes

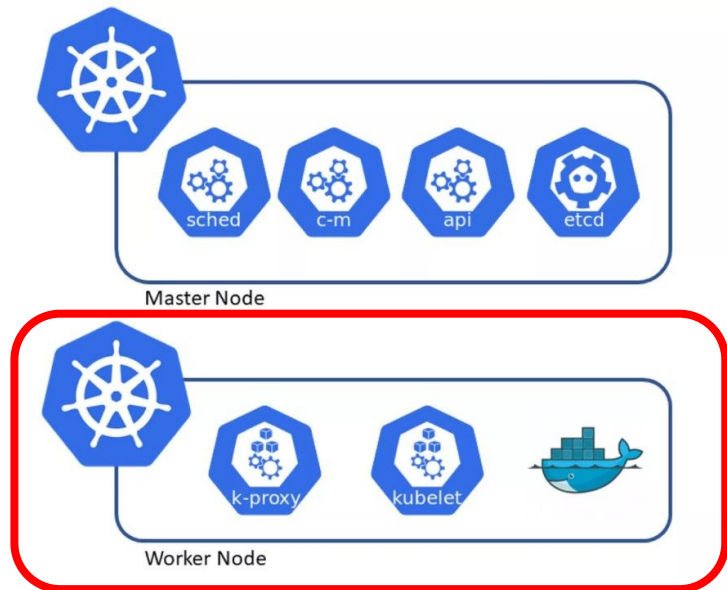
- **kube-apiserver** - Exposes Kubernetes API for cluster management.
- **kube-controller-manager** - Manages desired state via controllers.
- **kube-scheduler** - Assigns Pods to nodes based on resource requirements.
- **etcd** - Distributed key-value store for cluster data.
- ... and many more depending on your environment.



Kubernetes - Architecture

Worker nodes

- **kubelet** - Agent running on workers, executing instructions from the control plane.
- **container runtime** - Software responsible for running containers (e.g. Docker).
- **kube-proxy** - Maintains network rules for load balancing inside the cluster.
- **CNI (Container Network Interface)** - Overlay network enabling Pod communications across nodes.



Minikube & Project Prep



Note - Assumes container runtime installed - e.g. Docker

Install minikube

```
$ curl -LO \  
https://storage.googleapis.com/minikube/releases/latest/minikube-linux-amd64  
  
$ sudo install minikube-linux-amd64 /usr/local/bin/minikube
```

Install kubectl

```
$ curl -LO "https://dl.k8s.io/release/$(curl -L -s  
https://dl.k8s.io/release/stable.txt)/bin/linux/amd64/kubectl"  
  
$ sudo install kubectl /usr/local/bin/kubectl
```

GitHub Repo

```
$ git clone \  
https://github.com/netserf/netsig-presentation-kubernetes-for-newbies.git
```

Minikube Init



Start Cluster

```
$ minikube start --nodes 2 -p newbie-demo
```

```
$ minikube status -p newbie-demo
```

```
$ minikube ssh -p newbie-demo -n newbie-demo # control plane node  
$ ps -e | grep -E 'kube|etcd'
```

```
$ minikube ssh -p newbie-demo -n newbie-demo-m02 # worker node  
$ ps -e | grep -E 'kube|docker'
```

Nodes Provisioned

```
$ kubectl get nodes
```

```
$ kubectl describe nodes
```

Namespaces

- **Resource Partitioning** - Divides cluster resources into logical groups.
- **Isolation** - Securely share cluster with multiple teams.

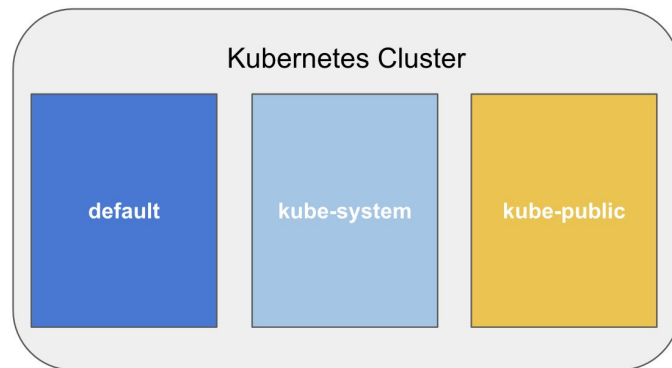
Create a namespace:

```
$ kubectl get namespace
```

```
$ kubectl create namespace newbie-ns
```

```
$ kubectl describe namespace
```

```
$ kubectl config set-context \
  newbie-demo --namespace=newbie-ns
```



Namespace: kube-system

- **System Components** - Dedicated namespace for management components.
- **Critical Operations** - Hosts management components (schedulers, controllers, network plugins, etc.) essential for cluster operations.
- **Isolated** - Separates critical system components from user workloads.


```
$ kubectl get namespace kube-system
```

```
$ kubectl get pods --namespace=kube-system
```



K8S Pod

- **Basic Unit** -The smallest unit in the Kubernetes object model
 - Each Pod contains 1-to-many containers
- **Isolation** - Pod processes and resources are isolated from other Pods
 - Like a mini virtual machine
- **Shared Resources** - Containers within a Pod share the same IP addresses, ports, volumes, configs, etc.



Run a K8S Pod

```
# Create a Pod imperatively
```

```
$ kubectl run nginx-pod --image=nginx:latest --restart=Never
```

```
$ kubectl delete pod nginx-pod
```

```
# Create a Pod declaratively
```

```
$ kubectl apply -f k8s/nginx-pod.yaml
```



Take a Look at the Pod

```
$ kubectl get pod nginx-pod [-o wide] [-o yaml]
```

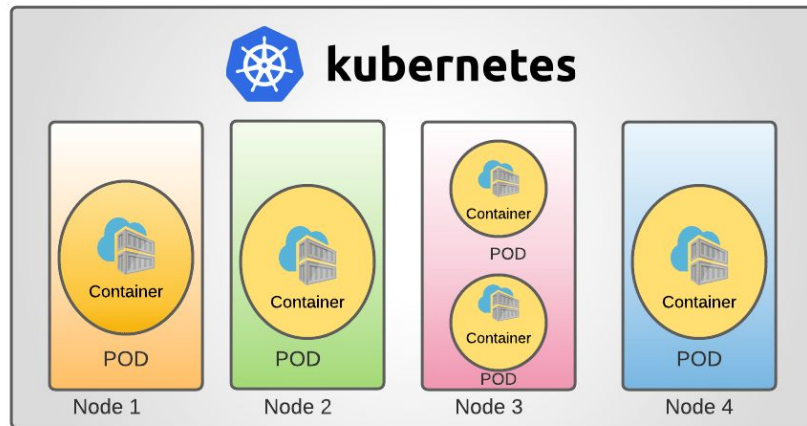
```
$ kubectl describe pod nginx-pod
```

```
$ kubectl logs nginx-pod
```

```
$ kubectl exec -it nginx-pod -- /bin/bash  
$ curl <container-ipaddr>:80
```

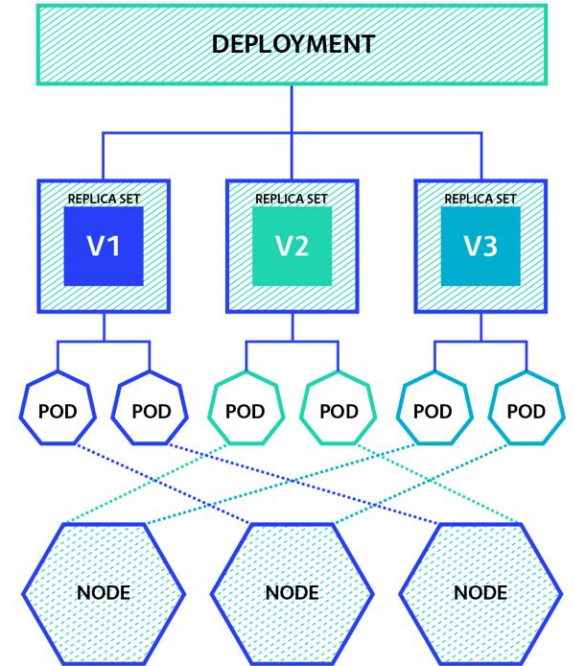
```
$ kubectl port-forward nginx-pod 8080:80
```

```
$ kubectl delete pod nginx-pod
```



K8S Deployment

- **Scalability** - Scale applications up or down by adjusting replica counts.
- **Self-Healing** - Health checks and automatic replacement of unhealthy pods.
- **Rolling Updates** - Updates without downtime, with quick rollback options.
- ... this is where K8S value starts to show.



Run a K8S Deployment

```
# Create a deployment imperatively
```

```
$ kubectl create deployment nginx-deployment --image=nginx:latest
```

```
$ kubectl scale deployment nginx-deployment --replicas=3
```

```
$ kubectl delete deployment nginx-deployment
```

```
# Create a deployment declaratively
```

```
$ kubectl apply -f k8s/nginx-deployment.yaml
```



Take a Look at the Deployment



```
$ kubectl get deployment nginx-deployment [-o wide] [-o yaml]
```


```
$ kubectl describe deployment nginx-deployment
```

```
$ kubectl get pods # look at pods supporting this deployment
```

```
$ kubectl logs <pod-name> # use pod names discovered
```

K8S Service

- **Load Balancing** - Distributes incoming traffic among pods.
- **Stable Endpoint** - Provides a single, stable endpoint for communications.
- **Flexible** - Facilitates both internal cluster and external requests.



Run a K8S Service



```
# Expose the deployment imperatively
```

```
$ kubectl expose deployment nginx-deployment \  
  --name=nginx-service --port=80 --type=NodePort
```

```
$ kubectl delete service nginx-service
```

```
# Expose the deployment declaratively
```

```
$ kubectl apply -f k8s/nginx-service.yaml
```

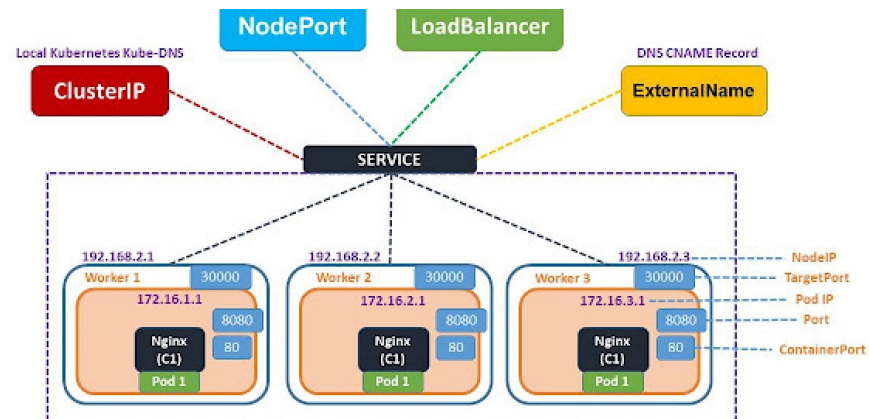
Take a Look at the Service

```
$ kubectl get service nginx-service
```

```
$ kubectl describe service nginx-service
```

```
$ minikube service nginx-service --url \  
-p newbie-demo -n newbie-ns  
$ curl <url>
```

```
$ kubectl logs <pod-name>
```



Clean-up

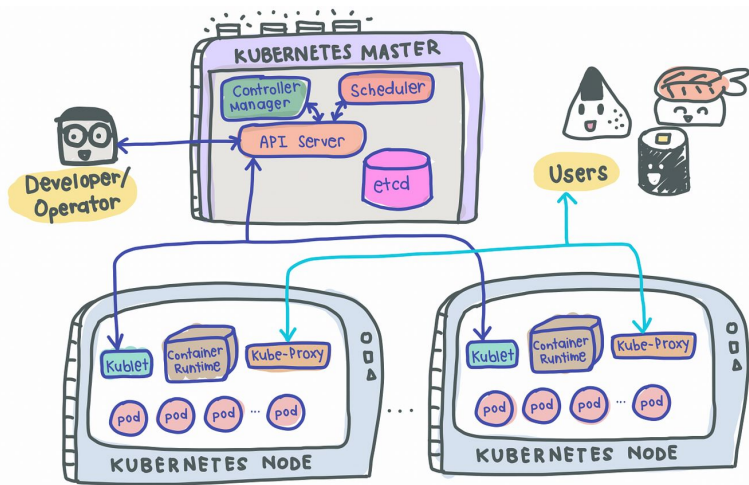
```
$ minikube delete -p newbie-demo
```

```
$ minikube status -p newbie-demo
```



Summary

- Kubernetes is a container orchestration platform
- Organizes machines (nodes) into clusters
- Streamlines the lifecycle management of containerized applications



Possible Future Discussions

- Orchestration
 - K8S ConfigMaps, Secrets, Persistent Volumes
 - K8S Ingress, Gateway
 - Istio Service Mesh
 - Hashicorp Nomad
- Observability
 - Prometheus / Grafana
 - ELK / EFK
 - Loki
- Messaging
 - RabbitMQ / ActiveMQ
- Data Pipelines
 - Airflow / Dagster
- Other ideas welcome!



Backup Slides



System Containers vs. Application Containers

System Containers:

- Running system-level processes and services
 - Like a mini virtual machine
- A lightweight environment for system-level tasks
- Designed to encapsulate and deploy components of the operating system
 - System daemons - systemd, cron, rsyslog, etc.
- No goal to decouple services any more than a traditional VM or bare-metal host
- Examples
 - LXC / LXN

```
CPU [|||||] 3.4% Tasks: 107, 390 thr; 1 running
Mem [|||||] 1.08G/3.84G Load average: 0.35 0.12 0.07
Swap [|||||] 0K/923M Uptime: 00:51:58

PID  Name  PPID  NI  VIRT  RSS  SHM  S  CPU%  MEM%  TTW  Command
1895  phoenixna  20  0  3627M  344M  125M  S  1.4  8.6  0:52.82 /usr/bin/gnome-shell
829  phoenixna  20  0  547M  88088  4948  S  1.4  2.2  0:17.11 /usr/lib/xorg/xorg vt2 -
2911  phoenixna  20  0  10844  4116  3308  R  1.4  0.1  0:00.17 htop
1975  phoenixna  20  0  795M  51016  38188  S  0.0  1.3  0:03.45 /usr/libexec/gnome-tern
824  phoenixna  20  0  11568  8688  3876  S  0.0  0.2  0:00.97 /usr/bin/dbus-daemon --s
777  kernoops  20  0  11264  448  0  S  0.0  0.0  0:00.07 /usr/sbin/kerneloops
1  root  20  0  163M  11468  8380  S  0.0  0.3  0:08.63 /sbin/init splash
222  root  19  1  1724  17832  6292  S  0.0  0.4  0:00.67 /lib/systemd/systemd-jou
266  root  20  0  24676  7188  5932  S  0.0  0.2  0:00.27 /lib/systemd/systemd-ude
537  systemd-r  20  0  24036  13328  9244  S  0.0  0.3  0:00.30 /lib/systemd/systemd-res
547  systemd-t  20  0  90260  6264  5492  S  0.0  0.2  0:00.00 /lib/systemd/systemd-tn
538  systemd-t  20  0  90260  6264  5492  S  0.0  0.2  0:00.14 /lib/systemd/systemd-tn
574  root  20  0  232M  7388  6536  S  0.0  0.2  0:00.05 /usr/lib/accountservitc
639  root  20  0  232M  7388  6536  S  0.0  0.2  0:00.00 /usr/lib/accountservitc
570  root  20  0  232M  7388  6536  S  0.0  0.2  0:00.09 /usr/lib/accountservitc
640  root  20  0  235M  18240  7680  S  0.0  0.3  0:00.06 /usr/lib/policykit-1/pol
598  root  20  0  235M  18240  7680  S  0.0  0.3  0:00.38 /usr/lib/policykit-1/pol
623  syslog  20  0  219M  4488  3716  S  0.0  0.1  0:00.02 /usr/sbin/rsyslogd n -
624  syslog  20  0  219M  4488  3716  S  0.0  0.1  0:00.08 /usr/sbin/rsyslogd n -
625  syslog  20  0  219M  4488  3716  S  0.0  0.1  0:00.02 /usr/sbin/rsyslogd n -
F1  halt  F2  setup  F3  search  F4  filter  F5  tree  F6  home  F7  vice  F8  vice  F9  kill  F10  quit
```

System Containers vs. Application Containers

Application Containers:

- Designed to encapsulate individual applications and their dependencies
 - Fosters portability across different environments
 - Enables a microservices architecture
- Optimizes performance and scalability for the application itself
 - Avoids oversubscribing resources for more efficient use of host resources
- Examples
 - Docker
 - Containerd
 - Podman
 - CRI-O
 - Kubernetes

