

# **Kubernetes fundamentals**

(for sysadmins)

# About me

- Sysadmin
- [Nimium](#)
- Working with Kubernetes for the past year
- [@0x6976](#)

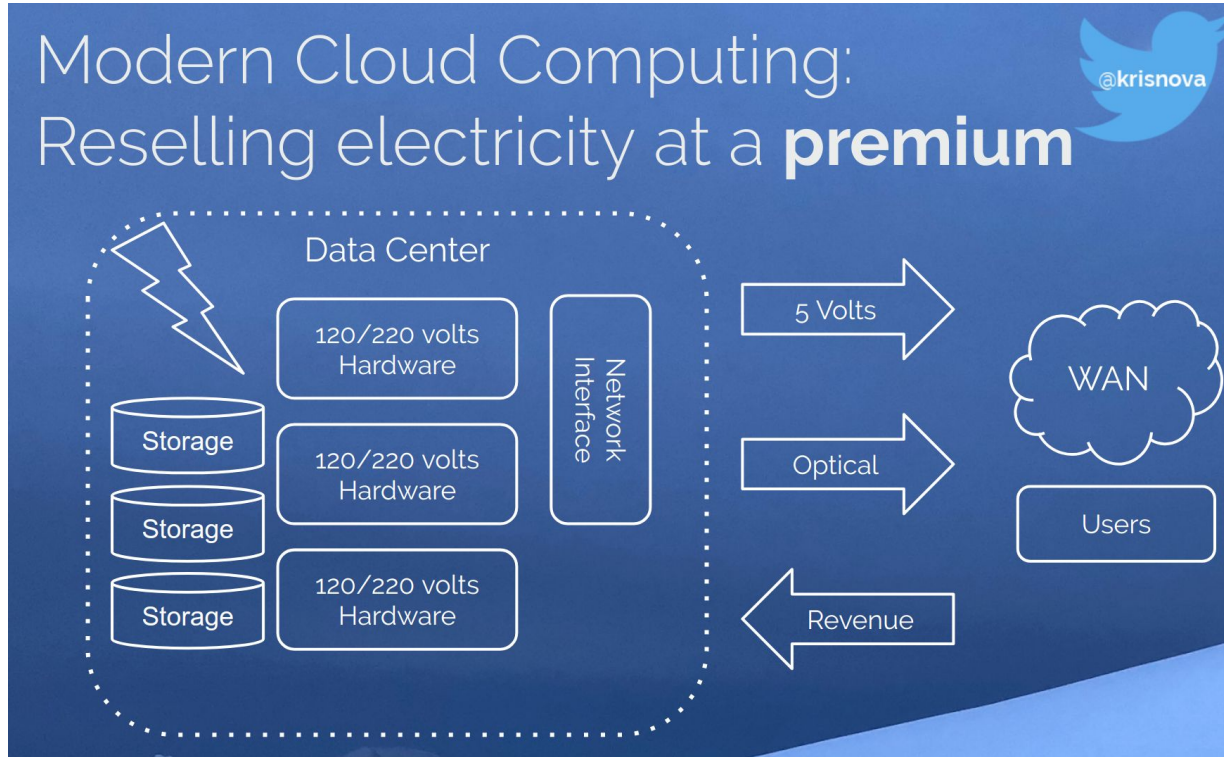
# Overview

- What is Kubernetes
- The building blocks of Kubernetes
- Some hiccups you might run into on your way to Kubernetes
- Q&A

# What is Kubernetes?

- “Kubernetes is named after the Greek god of spending money on cloud services”
  - Corey Quinn ([@QuinnyPig](#))

# What is the cloud?



- Kris Nova ([@krisnova](https://twitter.com/krisnova))

# What is Kubernetes?

- Three main questions
- Where?
  - Public cloud
  - **On premises**
  - Some kind of a stretched configuration
- How?
  - **Self-supported**
  - Outsourcing
  - The cloud(TM)
- What actually is Kubernetes?

# What is Kubernetes?

- A container orchestrator
- A set of software components that manages application lifecycle
- A scheduler
- A container runtime
- A set of pods (container groups)
- A set of replication and scaling rules for pods

# “Not my monkey, not my circus” line

- Separation of concerns (KUAR book)
- Application developer
  - Uses the API
- API reliability engineer
  - Maintains the container orchestration API
- OS reliability engineer
  - Takes care of the operating system
- HW reliability engineer
  - Takes care of the hardware



- **What are the building blocks in Kubernetes?**
- etcd
- API server
- Storage
- DNS
- Ingress
- Container runtime
- Networking

- **What are the building blocks in Kubernetes?**
- Monitoring
- Deployment
- Service mesh
- Container registries
- Resource limits
- Upgrade story
- Software ecosystem

# etcd

- Key-value store
- Make sure you have enough hardware
  - <https://github.com/etcd-io/etcd/blob/master/Documentation/op-guide/hardware.md#hardware-recommendations>
- **SSDs**
- Run a **multi-node cluster**
  - Five-member cluster recommended

# etcd

- **Backup**
- **Security**
  - Firewall
  - PKI
- **Restrict access to etcd**
  - Having access to etcd == having root access to the cluster
- **Scaling an etcd cluster is done for reliability, not performance**

# API server

- The front-end for the Kubernetes control plane
- Talks to etcd
- Can be scaled horizontally

# Storage

- Well abstracted in Kubernetes
  - Volumes
  - Storage claims
- Numerous options
  - **Ceph**
  - NFS
  - Hardware-specific integrations

# DNS

- Runs inside the cluster
- **Make sure your networking is set up properly**
  - Your worker nodes need to be able to talk to the DNS
  - Your master nodes need to be able to talk to the DNS
- **Make sure your pods and nodes are configured properly**
  - Queries which do not match the configured cluster domain suffix will be forwarded to the upstream DNS defined on the node

# Ingress

- A lot of available options
  - **NGINX**
  - F5
  - Contour
  - HAProxy
  - Traefik
  - Istio




# Container runtime

- Several options
  - Docker
  - rkt
  - CRI-O
  - frakti
- **Docker Engine is hard to avoid**

# Networking

- Several options available
  - **Calico**
  - Flannel
- 2018 Public Cloud Performance Benchmark Report
  - <https://www.thousandeyes.com/press-releases/2018-public-cloud-performance-benchmark-report>

# Networking



- Dave Anderson ([@dave\\_universetf](https://twitter.com/dave_universetf))

# Monitoring

- Some great options here
  - Prometheus
  - TICK stack
- Absolutely crucial if you want to run a cluster by yourself
- **Prometheus + Alertmanager + Grafana + various exporters**

# Monitoring

- Prometheus collects the data and stores it
  - Also provides a basic UI for your PromQL queries
- Alertmanager can route alerts
  - Also does deduplication and silencing of alerts
  - Supports various ways of sending out alerts
    - Slack, Opsgenie, PagerDuty, email
- **Try to figure out your monitoring high availability story early on**

# Deployment

- **Helm**
- Roll out YAML by hand
- Terraform
  - Has a Kubernetes provider
- Pulumi

# Container registries

- Nexus
- Artifactory
- Something that can handle both Helm charts and Docker images?
  - VMware Harbor
- **Handling images/charts is important**

# Service mesh

- Service mesh brings control, security and observability of services, API calls and traffic for your Kubernetes clusters
- **Istio**



# Resource limits

- Use limits and requests
  - CPU
  - Memory
- **Use reasonable values**
- Take care of your storage
- **Plan for failure**
  - OOM killer

# Upgrade story

- From the OS perspective
  - New OS versions
  - Security patching
  - Switching to another OS
- From the Kubernetes perspective
  - Kubernetes version upgrades

# Software ecosystem

- **There's a lot of stuff out there**
- Heptio (now VMware) things (Ark/Velero, Sonobuoy, Gimbal)
- kube-hunter
- kube-bench (CIS benchmark)
- Possible OpenStack ecosystem parallels in the future?

# Software ecosystem

CNCF serves as the vendor-neutral home for many of the fastest-growing projects on GitHub, including Kubernetes, Prometheus and Envoy, fostering collaboration between the industry's top developers, end users, and vendors.



50,399

# of contributors to  
CNCF projects



60,610

Registered for free  
Kubernetes EdX  
course



79

Certified  
Kubernetes  
Distributions and  
Platforms



95,912

CNCF Meetup  
members

- CNCF (<https://www.cncf.io>)

# RBAC

- If you don't implement this you're going to have a bad time
- Really hard to get right for your organization
- **Plan time to do this properly**
  - It can impact the processes

# RBAC



The image shows a screenshot of a Twitter profile and a tweet thread. On the left is the profile card for Kelsey Hightower, a verified user with handle @kelseyhightower, located in Portland, OR, and joined in June 2010. The main part of the image shows a tweet from Kelsey Hightower (@kelseyhightower) dated 6:28 AM - 8 Dec 2017. The tweet text is: "The less people using kubectl directly. The less people you need to create RBAC policies for." This tweet has 32 retweets and 146 likes. Below the tweet are two replies, both from Kelsey Hightower (@kelseyhightower) on 8 Dec 2017. The first reply says: "Limit the scope of access to a Kubernetes cluster to automation tools and cluster administrators who may have to debug it or keep it running." The second reply says: "This has nothing to do with job titles. If you are responsible for running or debugging a Kubernetes then you'll need access to kubectl at some point, but why force that to be every developer from day one?"

**Kelsey Hightower** ✓  
@kelseyhightower

Following

The less people using kubectl directly.  
The less people you need to create  
RBAC policies for.

6:28 AM - 8 Dec 2017

32 Retweets 146 Likes

10 32 146

Tweet your reply

**Kelsey Hightower** ✓ @kelseyhightower · 8 Dec 2017  
Limit the scope of access to a Kubernetes cluster to automation tools and cluster administrators who may have to debug it or keep it running.

1 19 55

**Kelsey Hightower** ✓ @kelseyhightower · 8 Dec 2017  
This has nothing to do with job titles. If you are responsible for running or debugging a Kubernetes then you'll need access to kubectl at some point, but why force that to be every developer from day one?

7 9 63

- Kelsey Hightower ([@kelseyhightower](https://twitter.com/kelseyhightower))

# Cluster validation

- **Both when deploying and while it is in production**
- Serverspec/InSpec tests
- Running e2e tests from Kubernetes
- Sonobuoy
- OpenSCAP
- You may also want to think about performance testing

# Security

- There are some vendors in the ecosystem
  - Aqua Security
  - NeuVector
  - Twistlock
  - Isovalent (Cilium)
- **Auditing**



# Security



- Jessie Frazelle ([@jessfraz](https://twitter.com/jessfraz))

# Conclusion

- Kubernetes is good
- The ecosystem is rich
  - Vendors
  - Open source projects that are available
- **You need to make sure you're not just deploying to Kubernetes**
  - Make sure you have a complete infrastructure that you can leverage in day-to-day operations

**Thank you!**  
Questions?