



**KubeCon**



**CloudNativeCon**

North America 2018

# Demystifying Data-Intensive Systems on Kubernetes



Lena Hall

@lenadroid

# Data-Intensive Systems



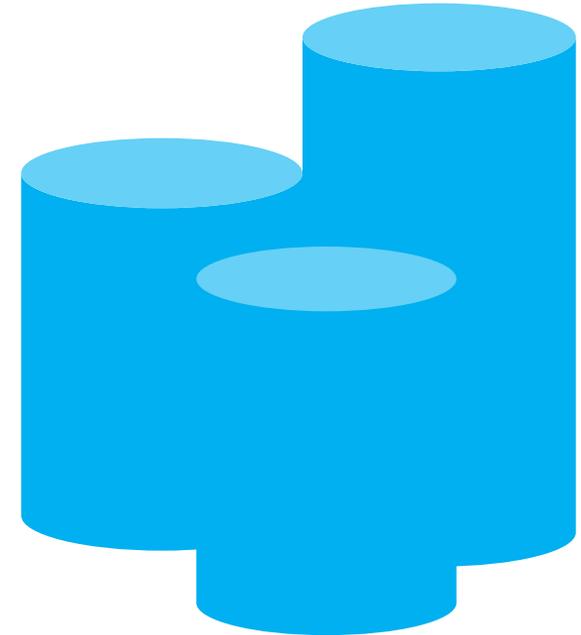
KubeCon



CloudNativeCon

North America 2018

- Databases
- Caches
- Stream-processing systems
- Any system that works with data



# Data-Intensive Systems



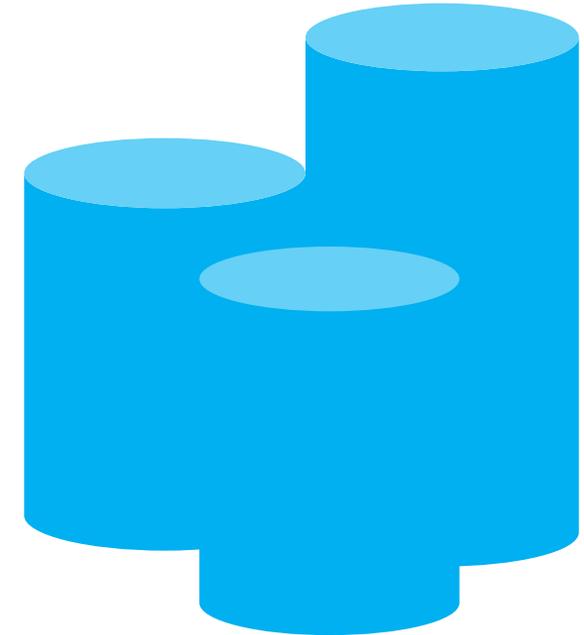
KubeCon



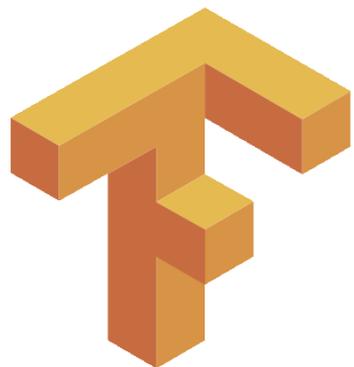
CloudNativeCon

North America 2018

- Databases
- Caches
- Stream-processing systems
- Any system that works with data



... or other non-trivial, not necessarily stateful systems



# Big Difference



KubeCon



CloudNativeCon

North America 2018

Stateless microservice

**VS**

Distributed stateful system, or other non-trivial system

on Kubernetes

# Why should you care

- Software Engineers and Solution Architects
- Engineering Managers and CTOs
- Maintainers and Contributors of Distributed Systems

# Talk Spoilers



KubeCon



CloudNativeCon

North America 2018

- **DECISIONS**

Making Decisions– is Kubernetes a good choice

- **SOLUTIONS**

Current state, challenges, and solutions

- **FUTURE**

The Future of systems on Kubernetes

# Decision Making



KubeCon



CloudNativeCon

North America 2018

**Decision** (p1, p2, p3, ... ) = **Yes** | **No**

p1 = required guarantees

p2 = existing skills and resources

p3 = acceptable risks

pN = ...

# Decision Making Best Practices



KubeCon



CloudNativeCon

North America 2018

- ? What are downsides and challenges of your current environment for running your system
- ? What problems will switch to Kubernetes solve
- ? What new problems will it create
- ? How big will increase/decrease in costs be
- ? What team or process changes will need to happen

# Decision Making Variables

For example:

- Ability to afford **resources/time** to troubleshoot issues
- Requirement to be **independent** from a cloud provider or environment
- Consistency/performance/availability/other **guarantees**
- Readiness to accept **possible risks**

... many more



KubeCon



CloudNativeCon

North America 2018

# Examples of motivation and challenges

# Common motivation

## Examples:

- Workload portability
- Convenience of deployment, operations, automation
- Independence from a cloud provider or environment
- Open and rapidly developing, rich ecosystem
- Flexibility and cost savings
- Faster start-up times
- Extensible and open API
- To benefit from existing Kubernetes infrastructure

# Common challenges

## Examples:

- Limited options for running certain workloads
- Limited examples of production-ready architectures
- Limited time, or resources to support the system
- Limited functionality/integrations for storage/networking
- Stability, reliability, etc. of existing solutions
- Need to build the solution almost from scratch
- Need to gain new skills to troubleshoot new environment

# Decision Making Best Practices



KubeCon



CloudNativeCon

North America 2018

How to determine possible risks and challenges?

# Decision Making Best Practices



KubeCon



CloudNativeCon

North America 2018

How to determine possible risks and challenges?

Understand what Kubernetes *can* or *can't* do.  
What it *is* or *isn't* responsible for.

# Decision Making Best Practices



KubeCon



CloudNativeCon

North America 2018

How to determine if Kubernetes can satisfy guarantees required by my system?

# Decision Making Best Practices



KubeCon



CloudNativeCon

North America 2018

How to determine if Kubernetes can satisfy guarantees required by my system?

Learn what abstractions and instruments Kubernetes and its API have to guarantee or implement your system requirements.



KubeCon



CloudNativeCon

North America 2018

To make the right decision is to  
understand how things work

# What Kubernetes can and can't do with built-in objects



KubeCon



CloudNativeCon

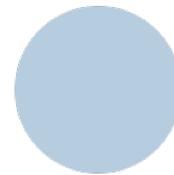
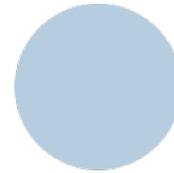
North America 2018



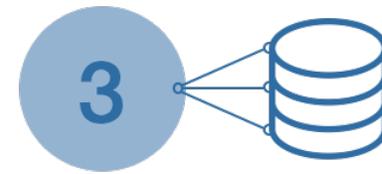
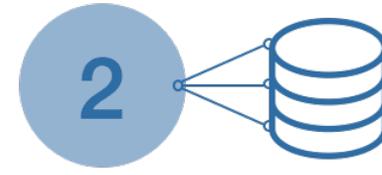
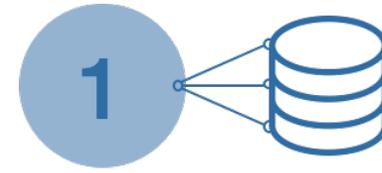
Pod



Job



Deployment



Stateful Set

and more...



KubeCon



CloudNativeCon

North America 2018

# Stateful Systems == Stateful Sets?



KubeCon



CloudNativeCon

North America 2018

**What if none of primitive Kubernetes types fully work for our systems?**



KubeCon



CloudNativeCon

North America 2018

# Things that need special care



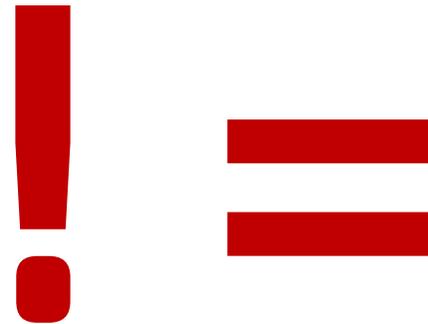
KubeCon



CloudNativeCon

North America 2018

# UP AND RUNNING



# OPERATING CORRECTLY

# Example: Scaling Clusters



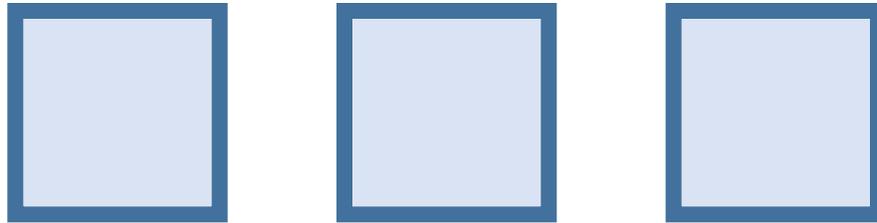
KubeCon



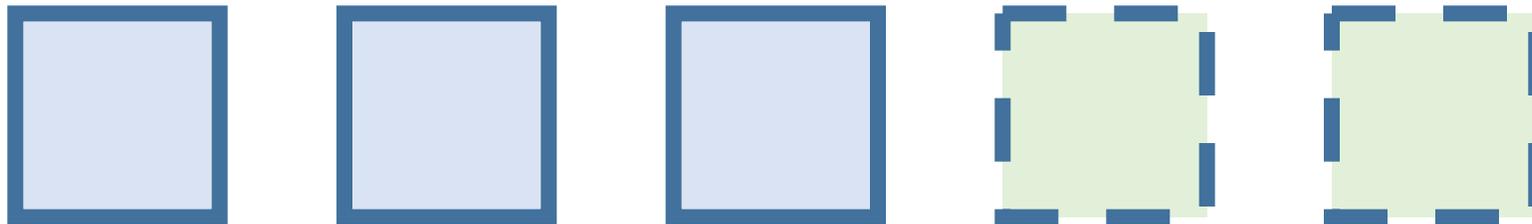
CloudNativeCon

North America 2018

From



To



- ✓ Manually assign partitions to new nodes
- ✓ Rebalance data to maintain even load
- ✓ Apply cluster configuration to new nodes

# Example: Safe Cluster Restarts



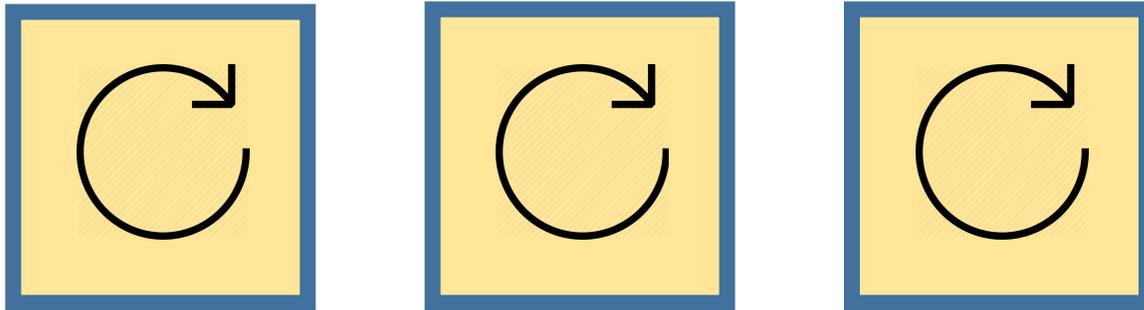
KubeCon



CloudNativeCon

North America 2018

- ? Are there any under-replicated partitions?
- ? Is the cluster in healthy state?



- ✓ Restart one node at a time
- ✓ Wait for each node to catch up



KubeCon



CloudNativeCon

North America 2018

# Custom Resource Definitions

## CRDs



KubeCon



CloudNativeCon

North America 2018

# Custom Controllers



KubeCon



CloudNativeCon

North America 2018

# OPERATOR

## CRD + Custom Controller

# Anatomy of an Operator



KubeCon



CloudNativeCon

North America 2018

Controller

Queue that the Controller subscribes to

Informer

Resource Structs/Types

CRD

# Example



KubeCon



CloudNativeCon

North America 2018

## Running a TensorFlow Job with KubeFlow and tf-job Operator on Azure Kubernetes Service



Behind the scenes

tf-operator



KubeCon



CloudNativeCon

North America 2018

# Ways to build an Operator

# Example Operators



KubeCon



CloudNativeCon

North America 2018

Apache Kafka/Confluent Operator

[confluent.io/confluent-operator](https://confluent.io/confluent-operator)

Apache Cassandra Operator

[github.com/instaclustr/cassandra-operator](https://github.com/instaclustr/cassandra-operator)

TensorFlow Operator

[github.com/kubeflow/tf-operator](https://github.com/kubeflow/tf-operator)

Apache Flink Operator

[osdir.com/apache-flink-development/msg09830.html](https://osdir.com/apache-flink-development/msg09830.html)

Apache Spark Operator

[github.com/googlecloudplatform/spark-on-k8s-operator](https://github.com/googlecloudplatform/spark-on-k8s-operator)

# Awesome Operators



KubeCon



CloudNativeCon

North America 2018

 /operator-framework/awesome-operators

# Takeaways



KubeCon



CloudNativeCon

North America 2018

Understand your goals, benefits, and challenges

Define “correct” operation for your system

Know when to use existing core Kubernetes types, and when to create custom resources



KubeCon



CloudNativeCon

North America 2018

# What's next?

# Future



KubeCon



CloudNativeCon

North America 2018

- More Operators
- Better Operators
- Easier to write Operators

# Future



KubeCon



CloudNativeCon

North America 2018

- Better workload portability
- More focus on writing your apps
- More automation

# Future



KubeCon



CloudNativeCon

North America 2018

- Multi-Cloud becomes reality
- More independence

# Share your data-intensive scenarios



KubeCon



CloudNativeCon

North America 2018

[bit.ly/data-k8s](https://bit.ly/data-k8s)

@lenadroid

# Resources



KubeCon



CloudNativeCon

North America 2018

Lena's talk from GOTO Chicago 2018 - running a distributed database on Kubernetes, specifics of Stateful Sets, with examples of using Cassandra and Spark:

[bit.ly/statefulsets-gotochgo](https://bit.ly/statefulsets-gotochgo)

Lena's blog and other talks:

[bit.ly/lena-blog](https://bit.ly/lena-blog) and [bit.ly/lena-talks](https://bit.ly/lena-talks)

# Lena Hall



# lenadroid



- ✓ Works on Azure at  Microsoft
- ✓ Lives in  Seattle
- ✓ F# Software Foundation Board of Trustees
- ✓ Organizes [@ML4ALL](https://twitter.com/ML4ALL) 
- ✓ Program Committee for Kafka Summit
- ✓ Has a channel: [YouTube /c/AlenaHall](https://www.youtube.com/c/AlenaHall)



**KubeCon**

**CloudNativeCon**

————— North America 2018 —————

**Thank you!**

