



**KubeCon**



**CloudNativeCon**

North America 2018

# Clusters all the way down Crazy Multi-cluster Topologies



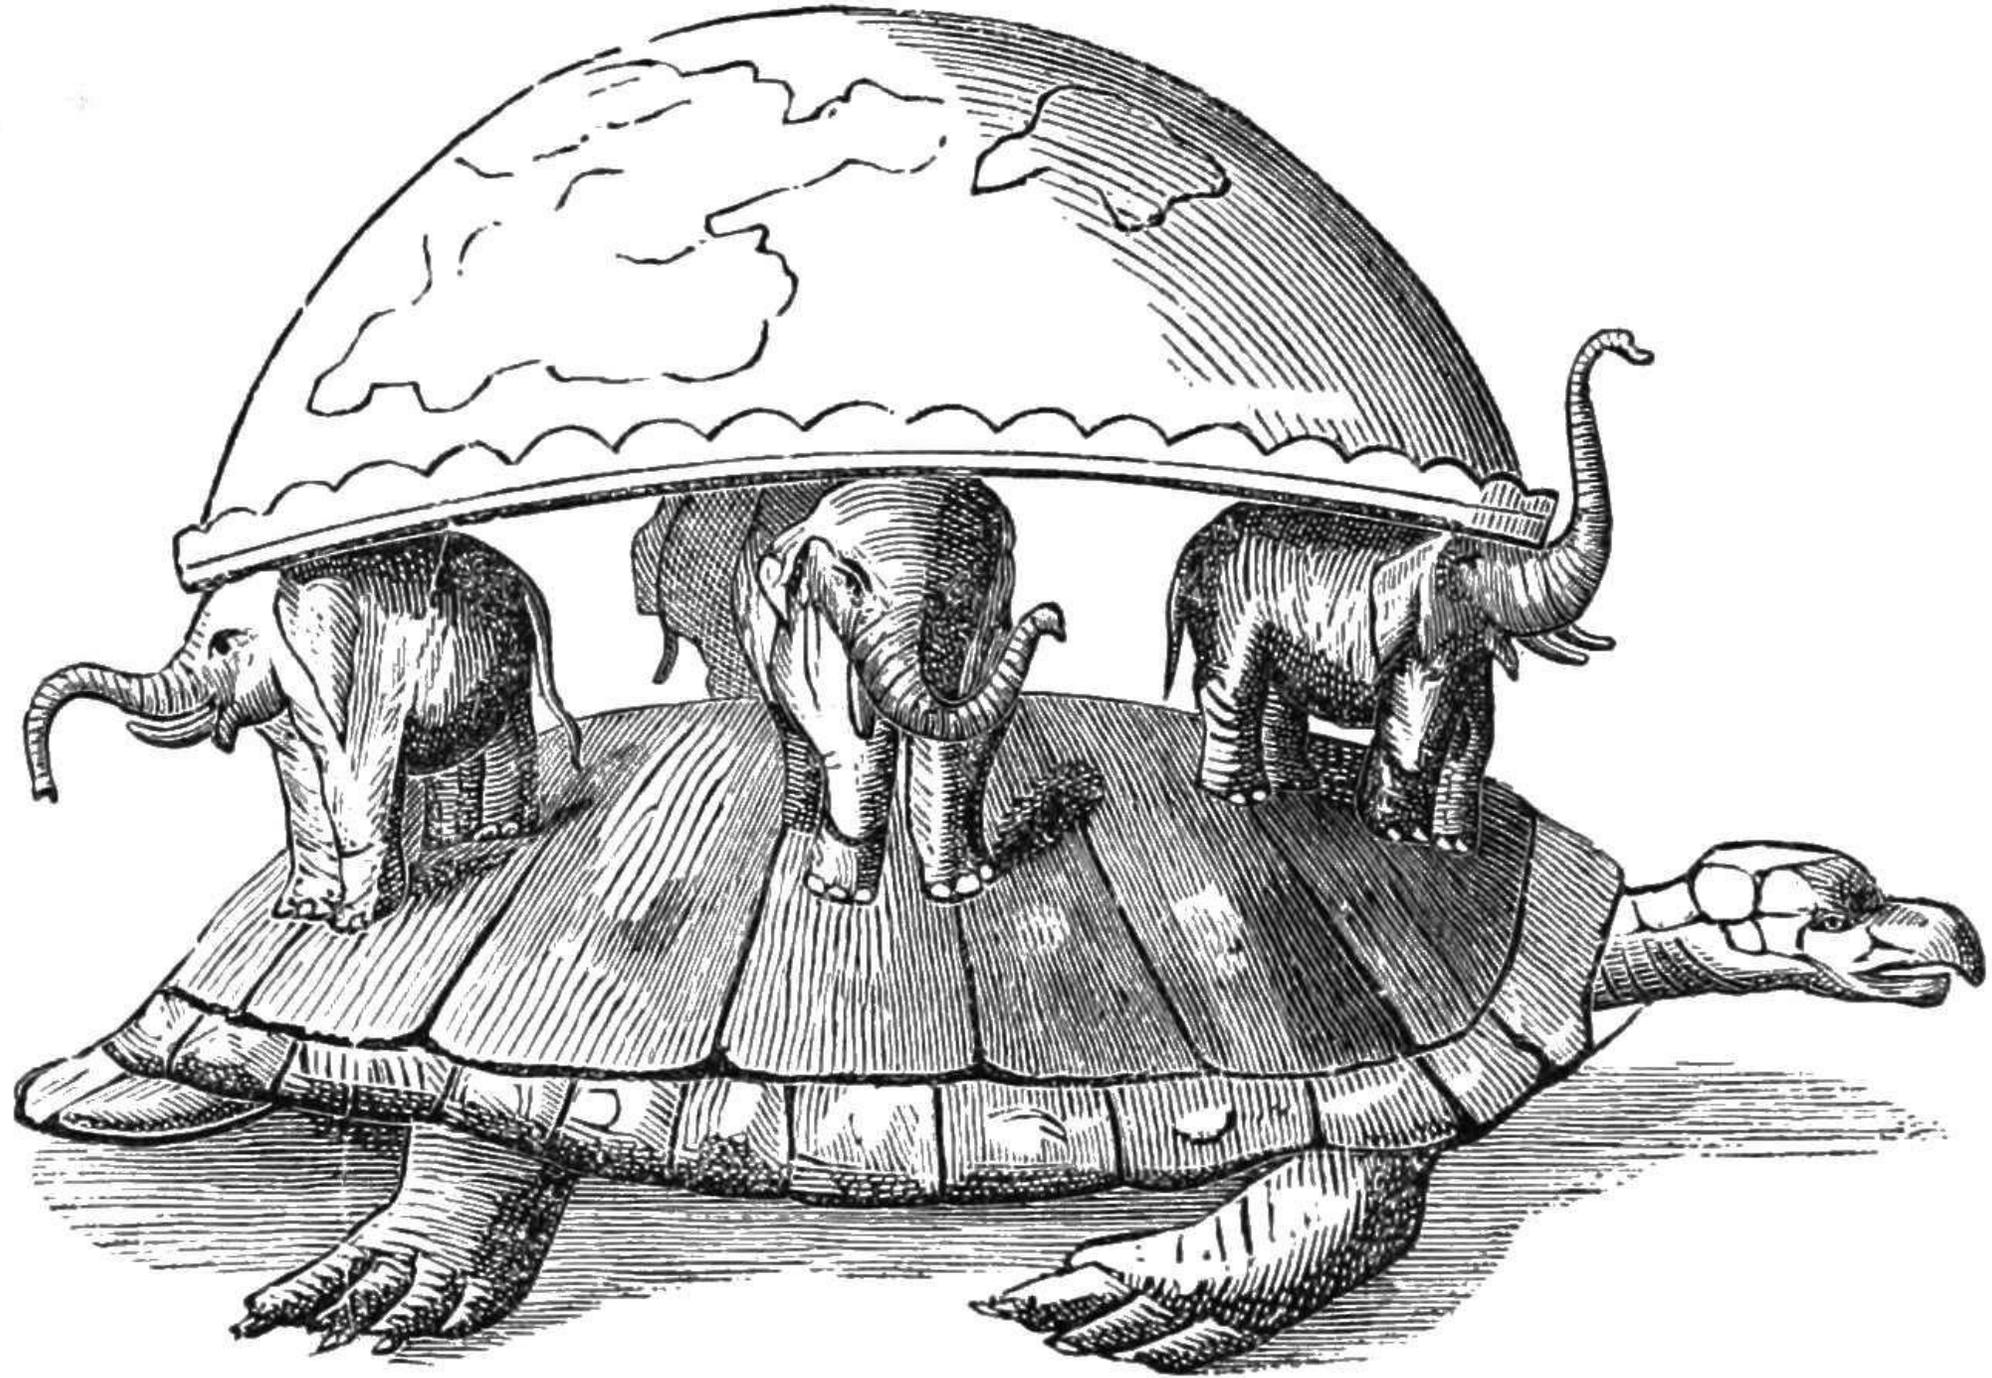


**Matt Caulfield**  
**Co-founder & CEO**

@

**oORT**

**Cloud Native Edge Computing**







**kubernetes**

**Cluster  
Management**



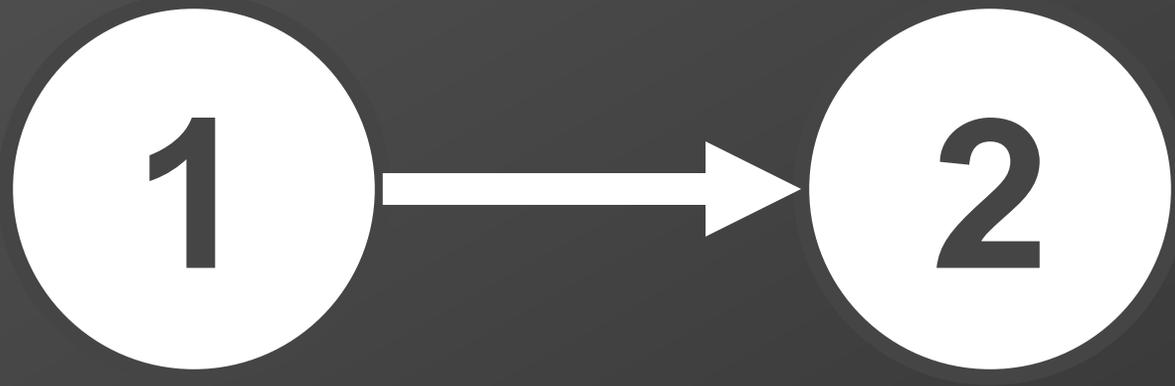


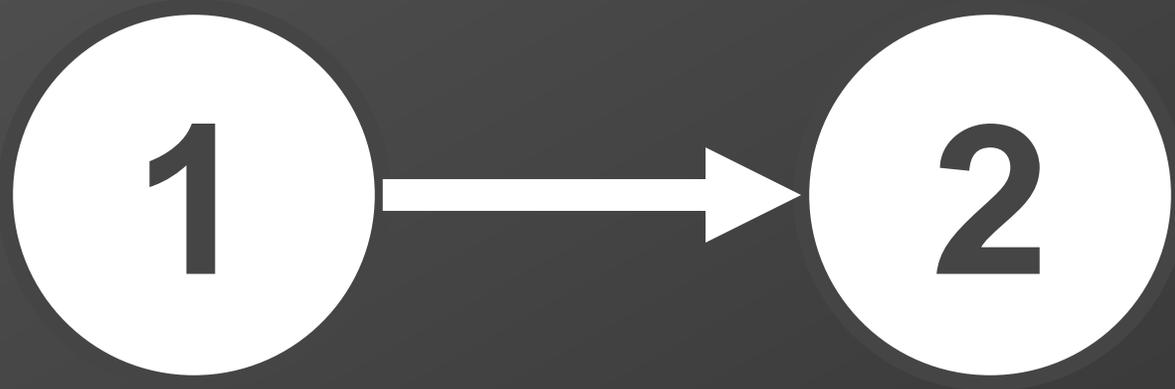
**kubernetes**

**Multi-Cluster  
Management**

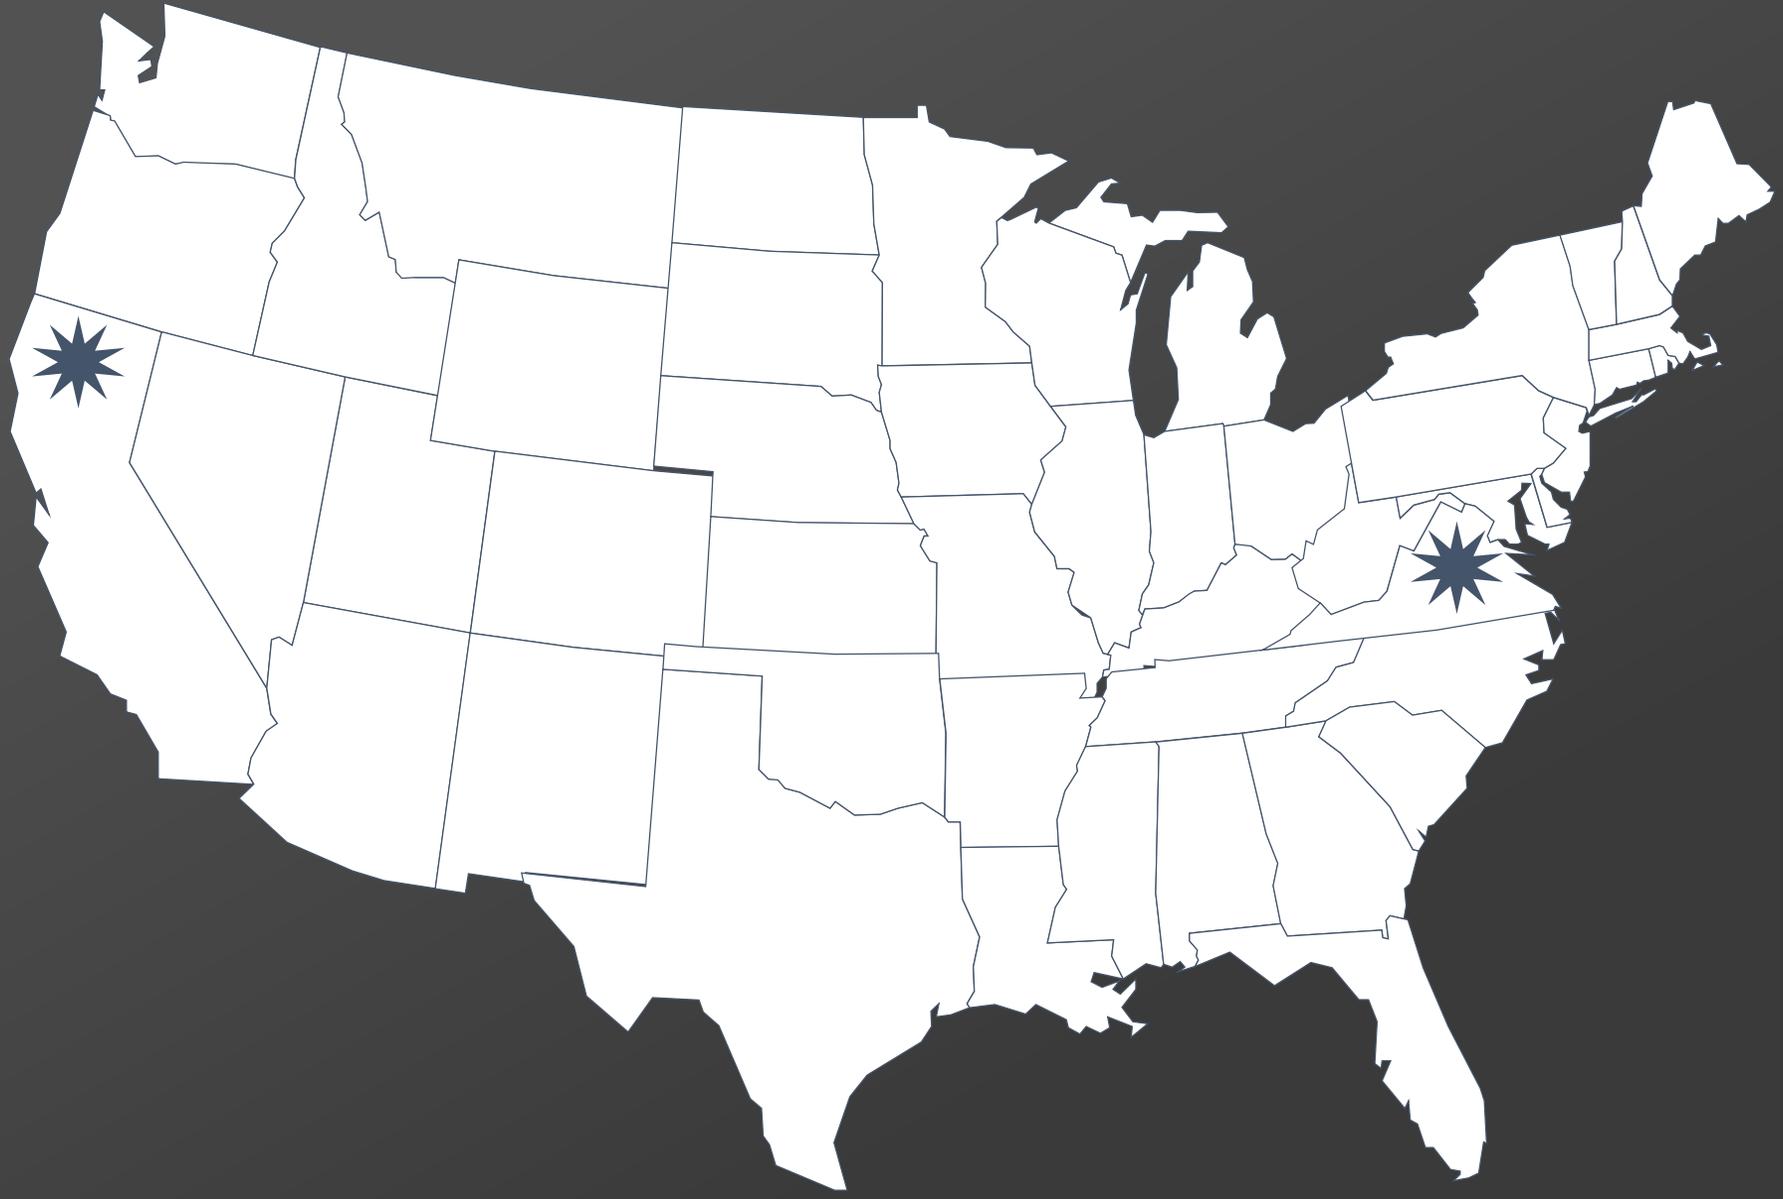


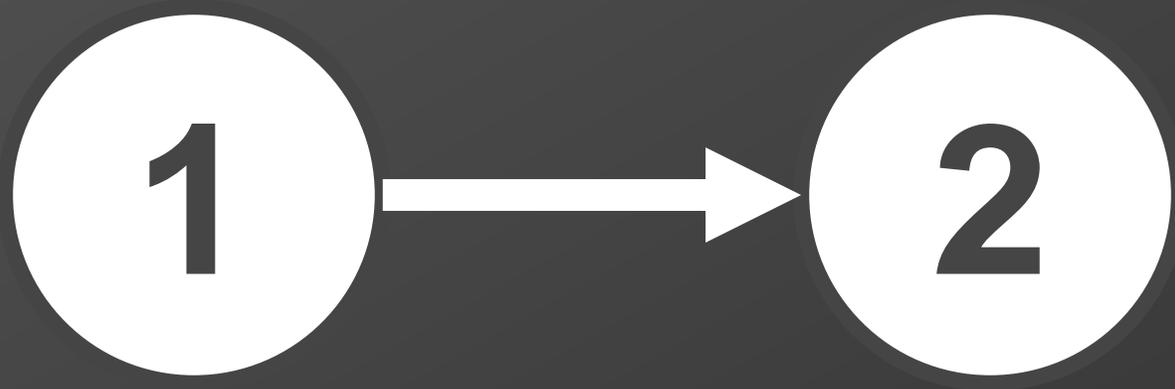
# MULTI-CLUSTER CONFIDENCE



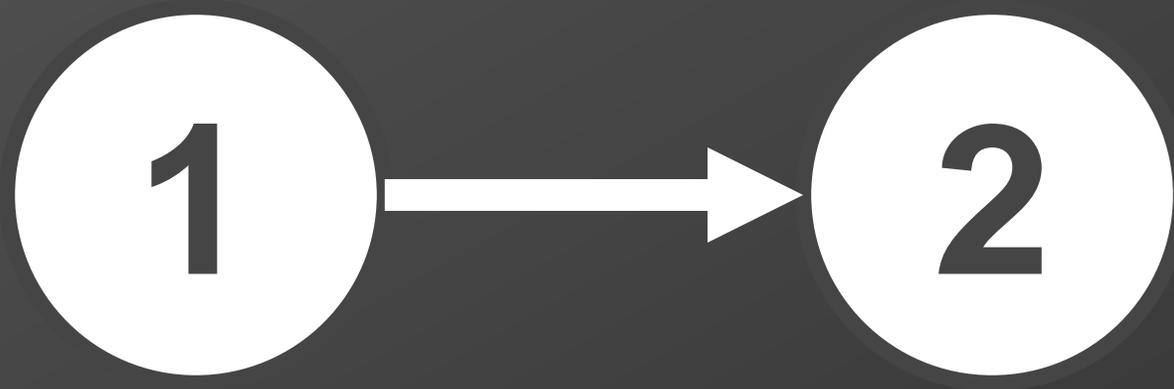


**Why?**





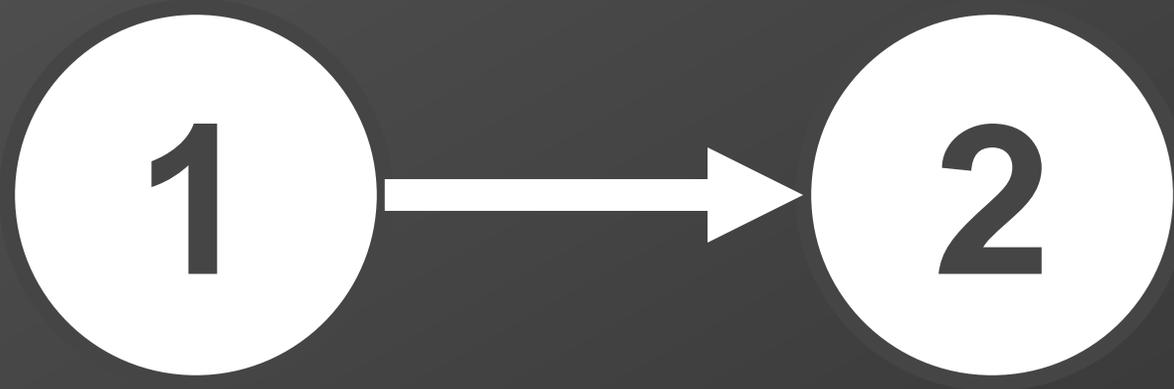
**Scalability?**



**High Availability**

**1 Cluster : 1 Region**

**(we'll revisit this later)**



**High Availability**

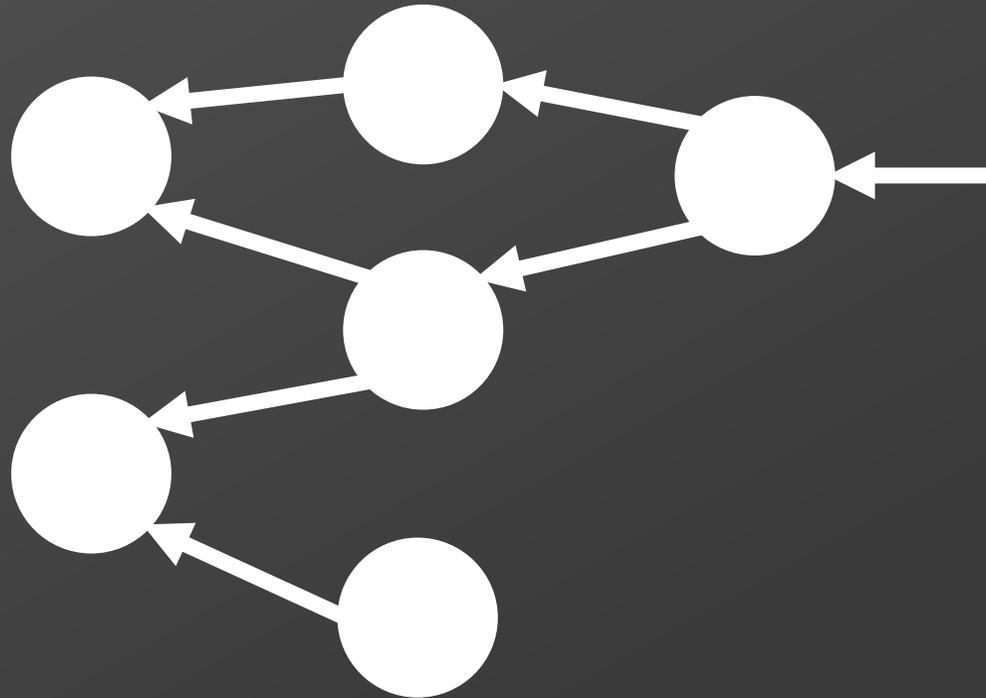
**What's so hard about  
geo-redundancy?**

**Issue #1**

**Application  
Topology**

# Issue #1 – Application Topology

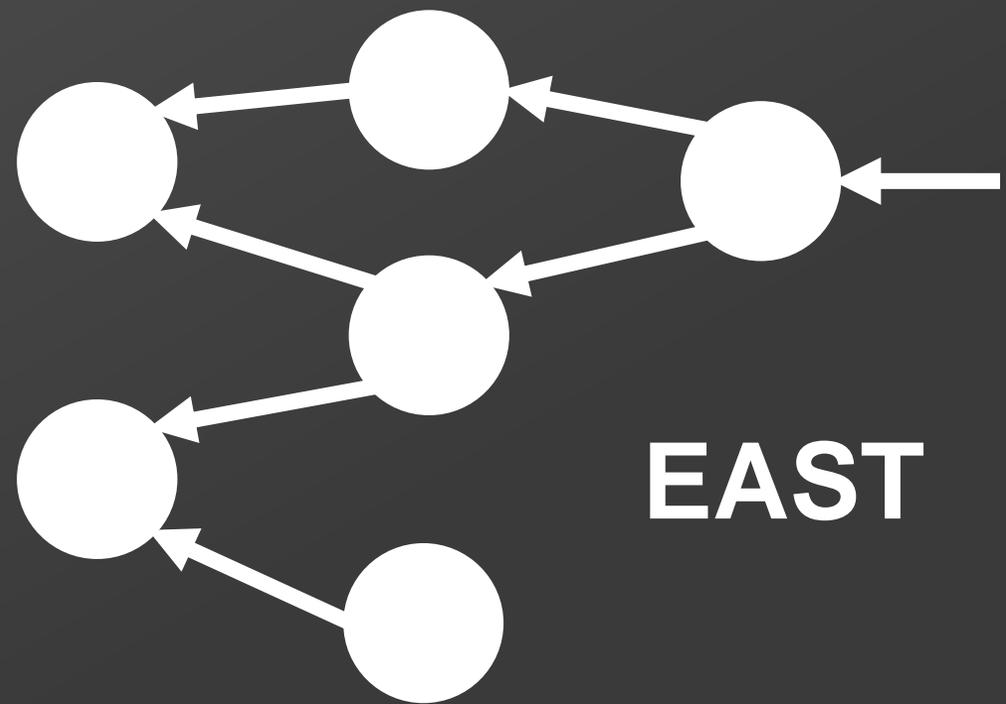
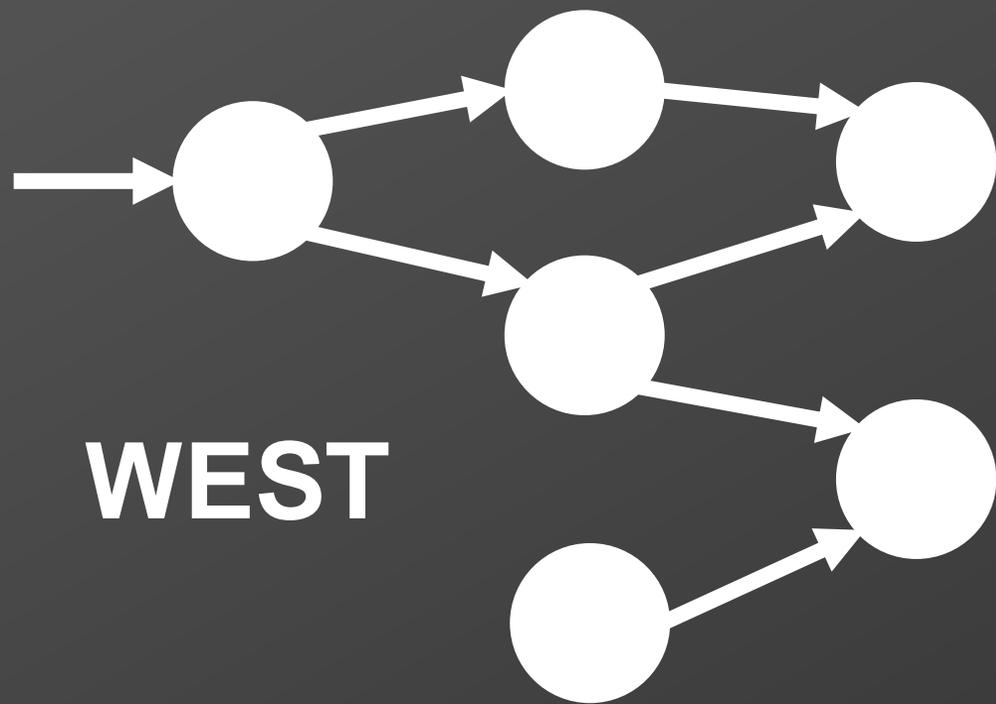
i.e. which parts of your application run where?





# Issue #1 – Application Topology

Clone It?

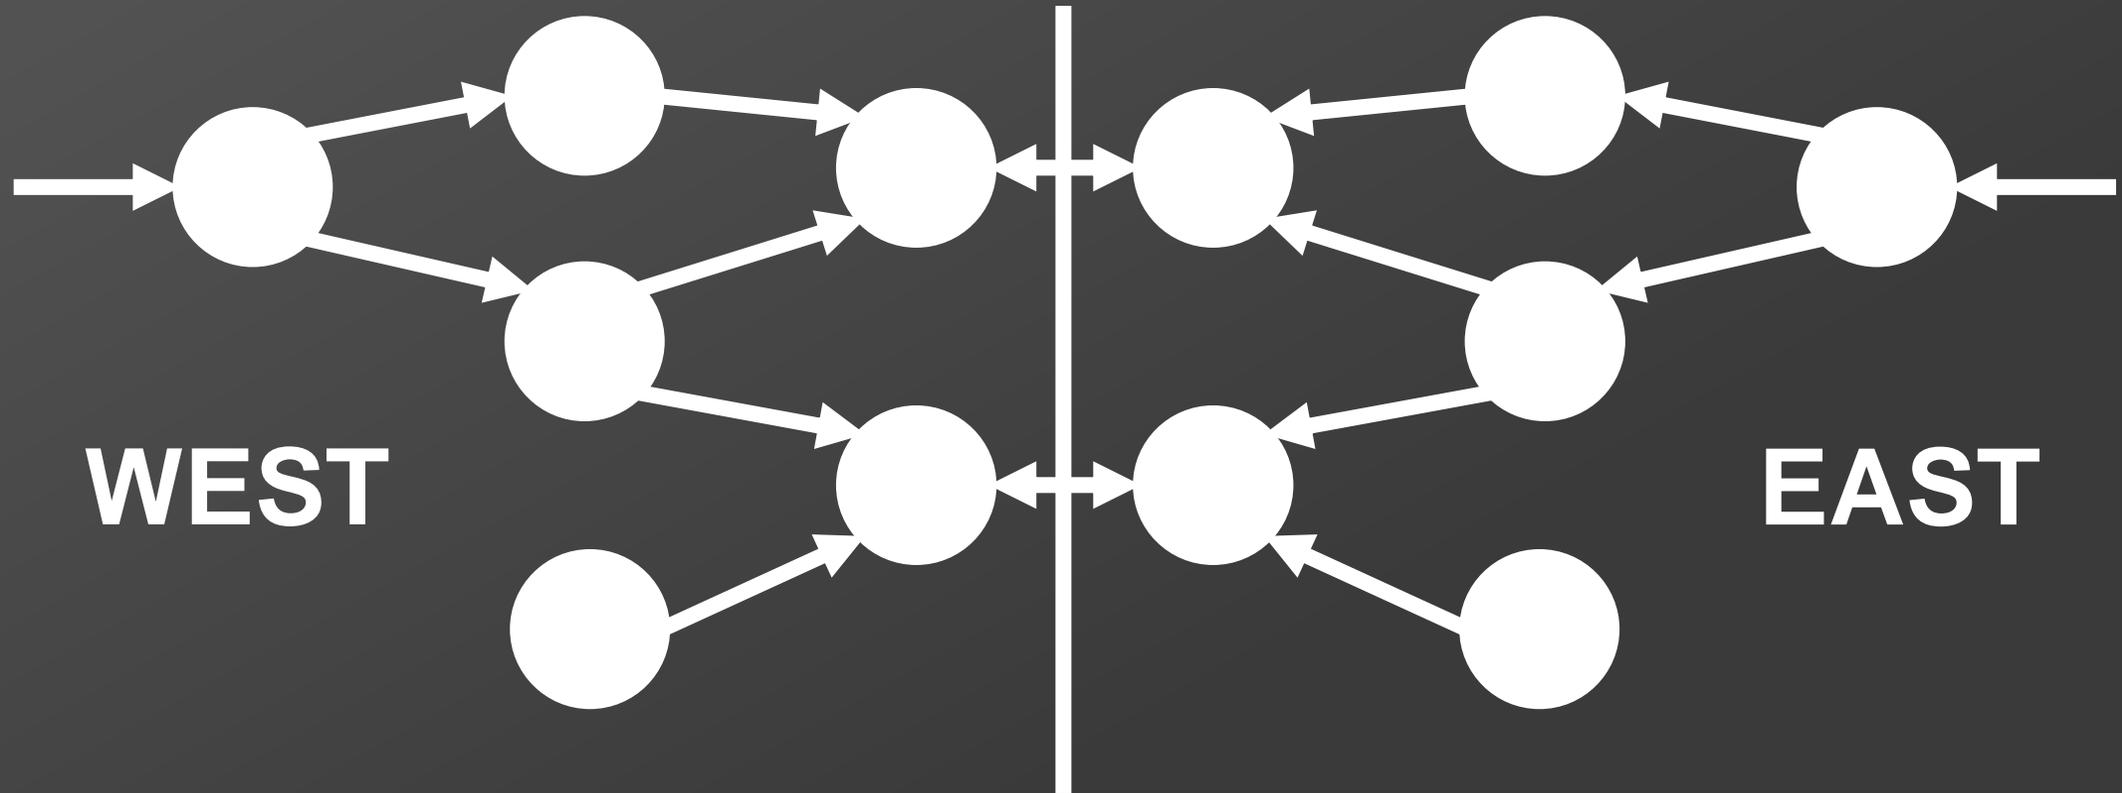


**Issue #2**

# **Data Replication**

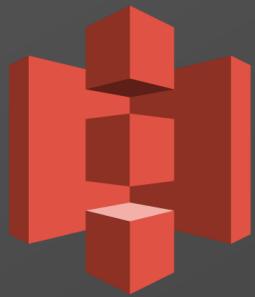
# Issue #2 – Data Replication

i.e. how is your data shared across clusters?

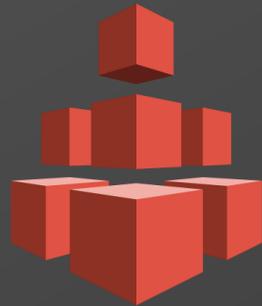


# Issue #2 – Data Replication

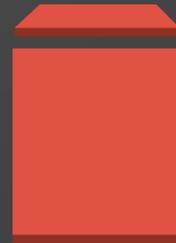
## AWS Toolbox



Amazon  
S3



Amazon  
EFS



Amazon  
EBS



Amazon  
RDS



Amazon  
DynamoDB

# Issue #2 – Data Replication

## GCP Toolbox

### Storage and Databases



**Cloud  
Storage**



**Cloud  
Bigtable**



**Cloud  
Datastore**



**Cloud SQL**



**Persistent  
Disk**

**Issue #3**

# **Load Balancing**

**Issue #3**

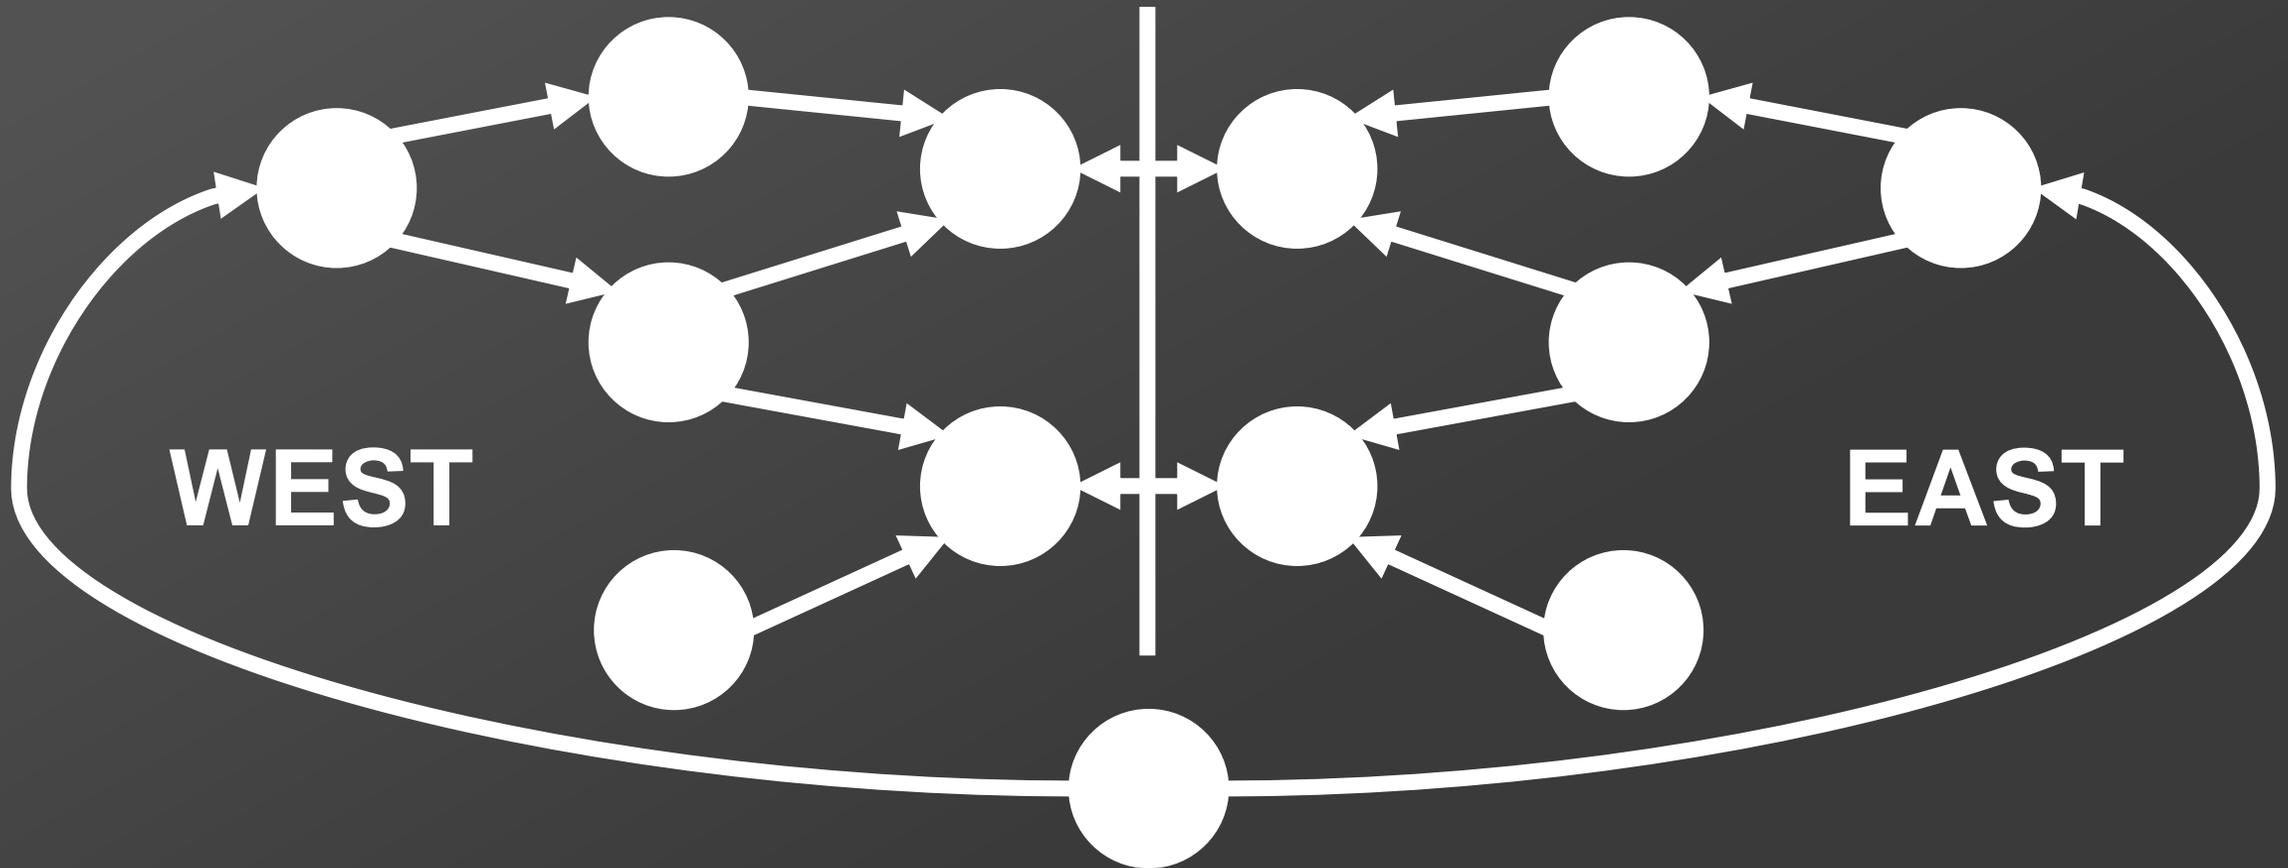
**~~Load Balancing~~**

**Issue #3**

~~Load Balancing~~  
**Traffic Routing**

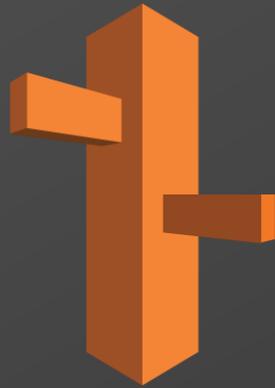
# Issue #3 – Traffic Routing

i.e. how is traffic shared across clusters?



# Issue #3 – Traffic Routing

## AWS Toolbox



**Amazon  
Route 53**



**Elastic Load  
Balancing**

# Issue #3 – Traffic Routing

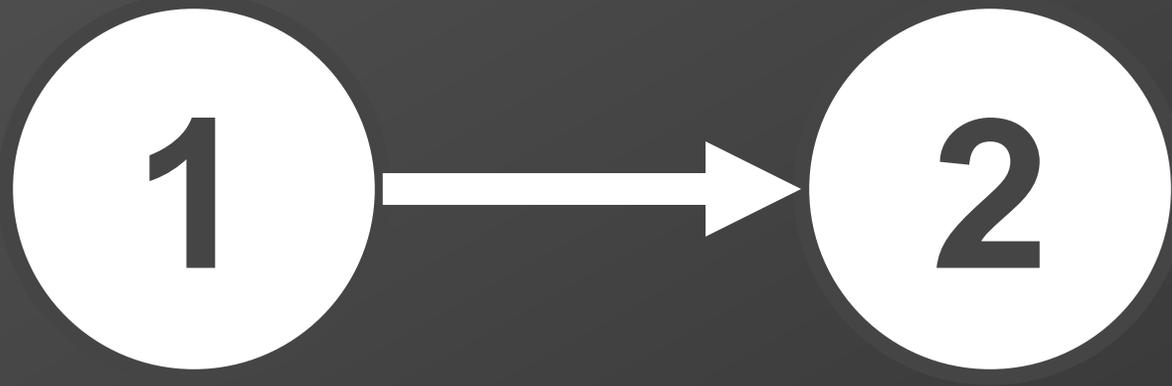
## GCP Toolbox



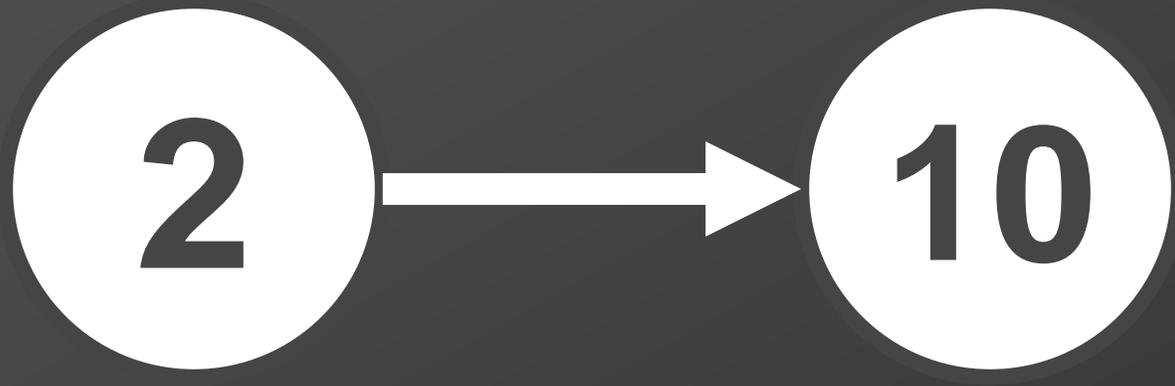
**Cloud Load  
Balancing**

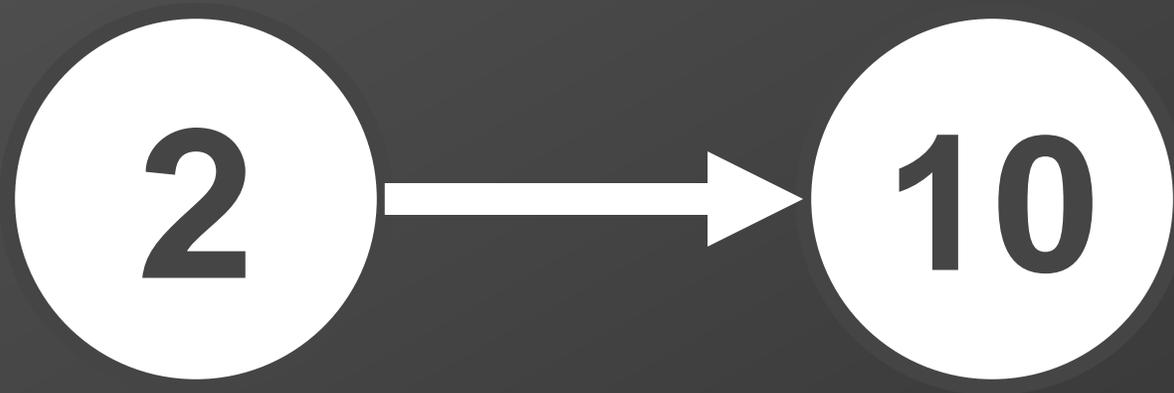


**Cloud  
DNS**



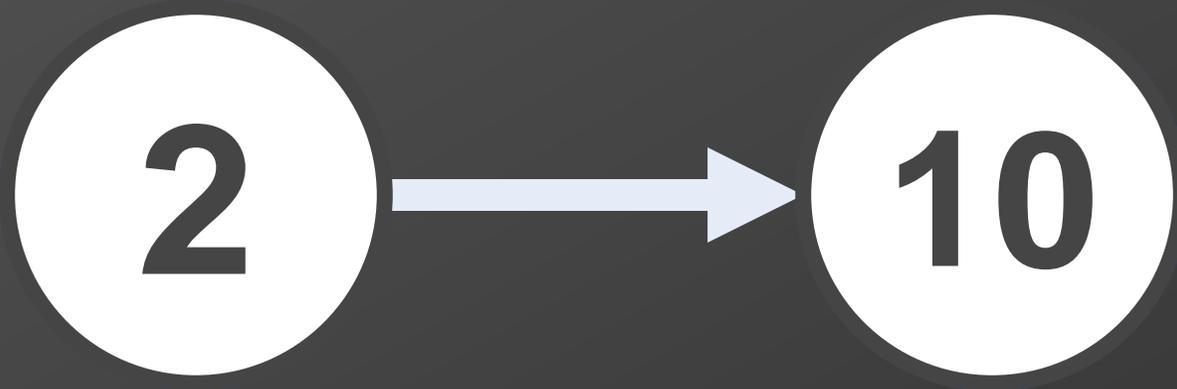
**COMPLETE**





**Why?**

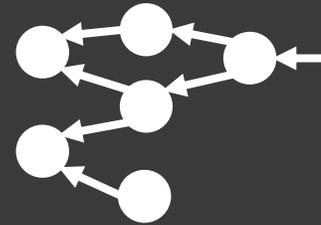
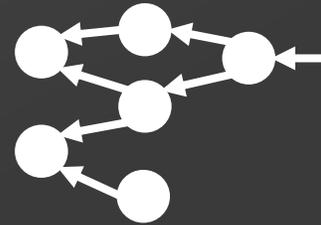
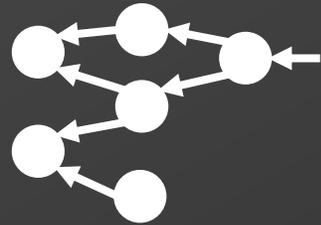
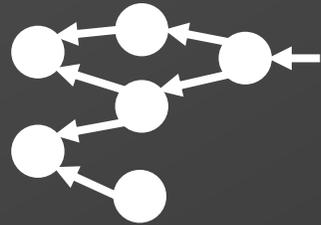
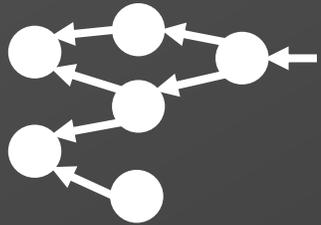
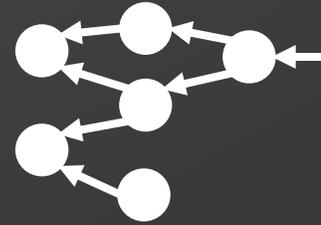
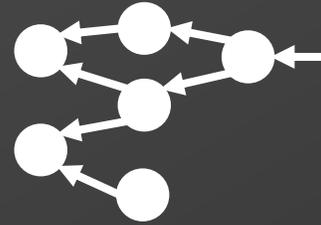
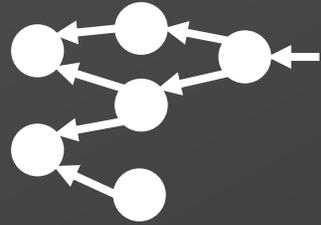
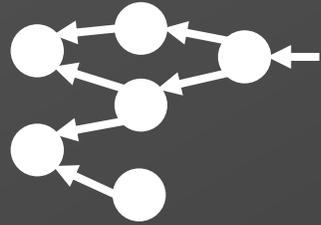
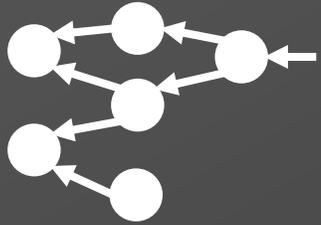




**User Experience**  
**Security & Data Sovereignty**

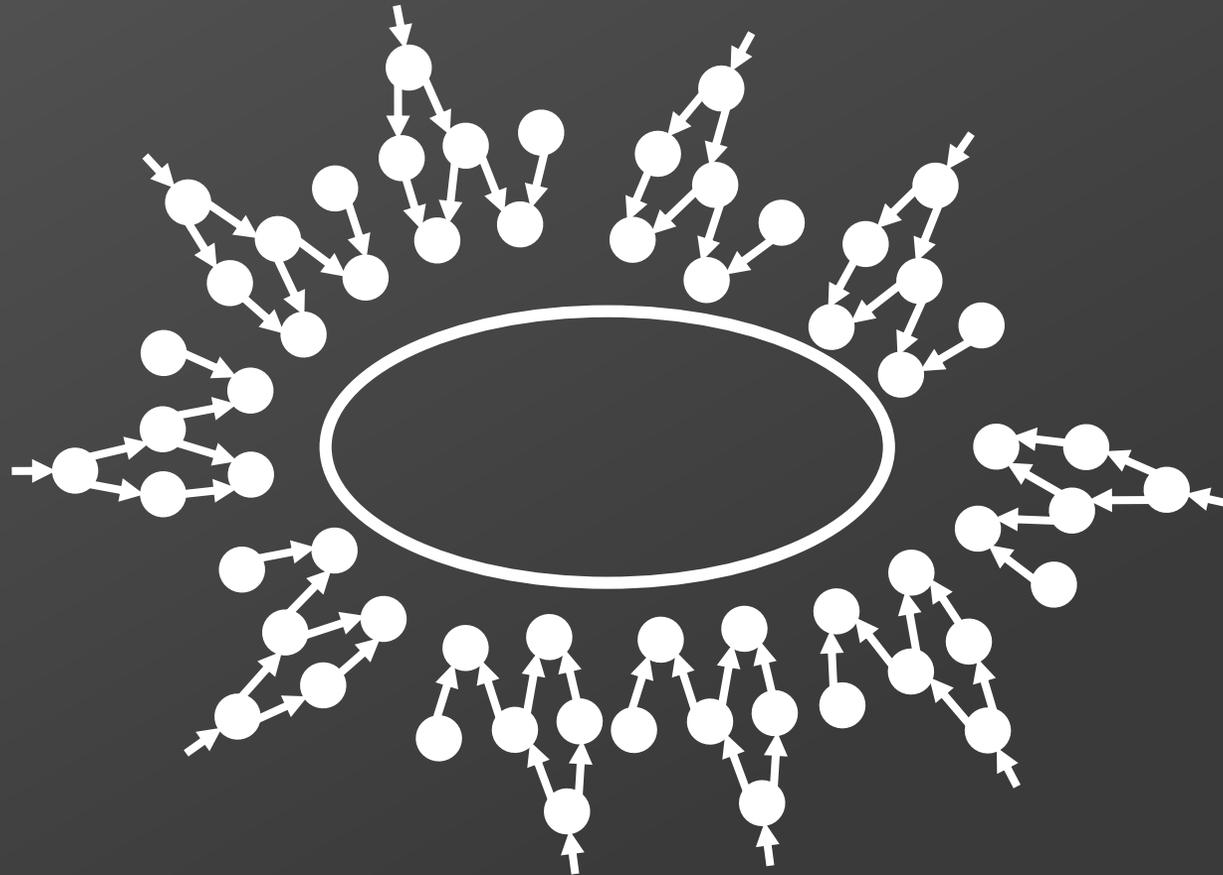
# Issue #1 – Application Topology

Clone It?



# Issue #2 – Data Replication

Full mesh replication? Ring topology? Clos?

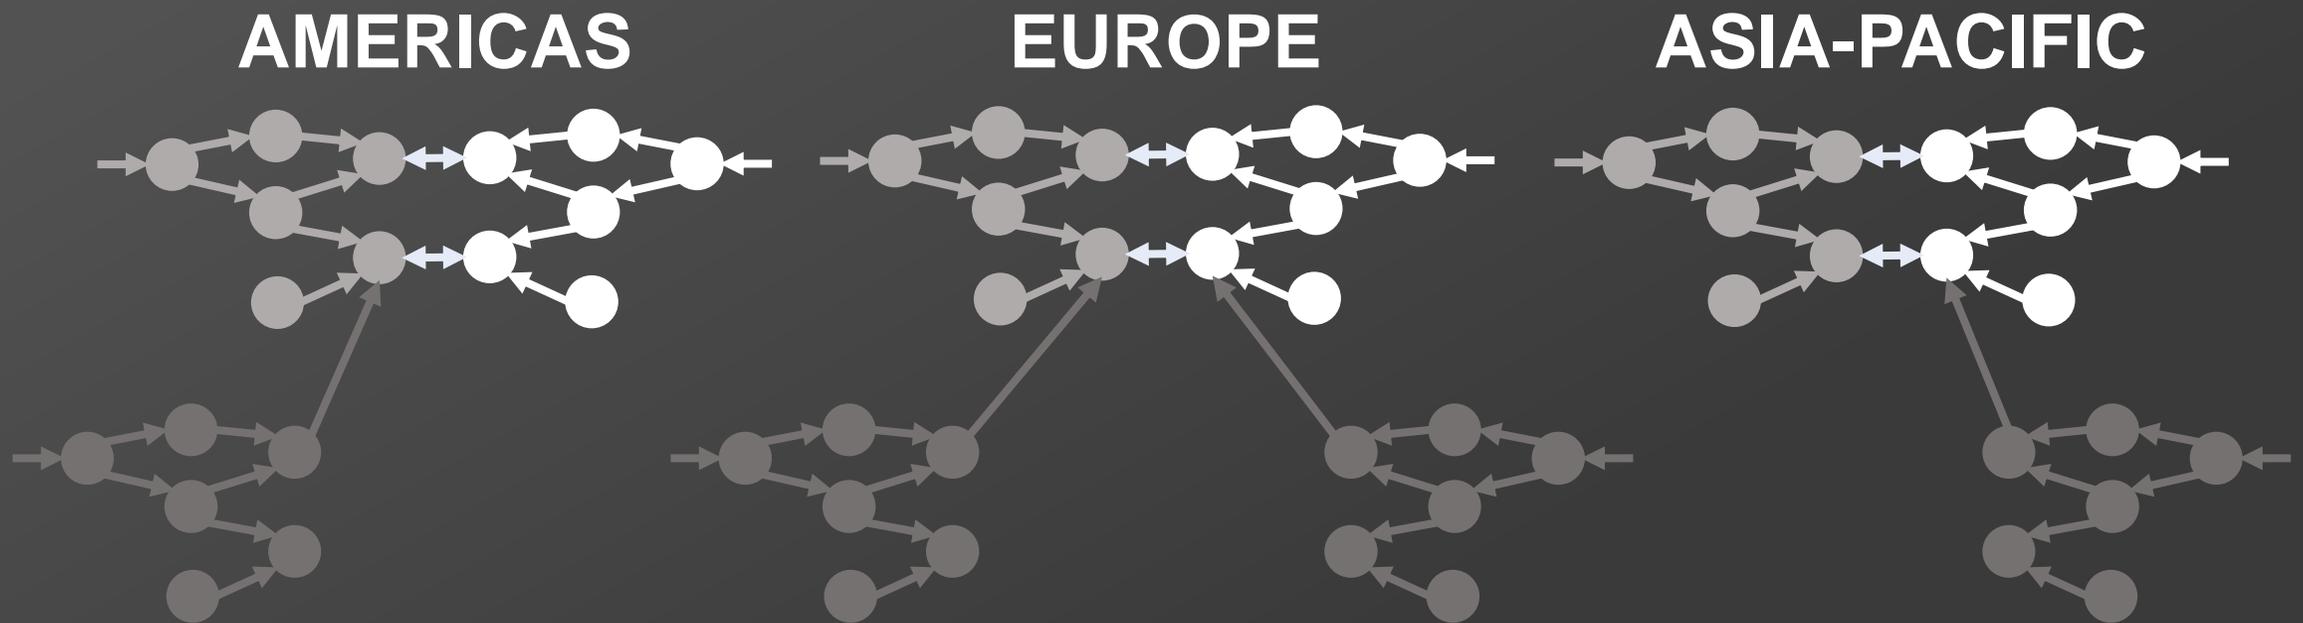


**Concept #1**

# **Hierarchical Topology**

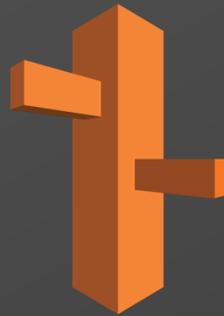
# Hierarchical Topology

App and Data Topology are Inextricably Linked



# Issue #3 – Traffic Routing

AWS, GCP, OSS / CNCF



**Amazon  
Route 53**



**GCP Load  
Balancing**



**CoreDNS**

**Issue #4**

**Service  
Discovery**

# Issue #4 – Service Discovery

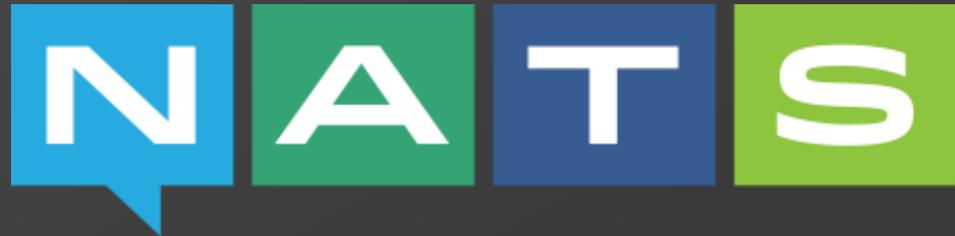
AWS, GCP, OSS / CNCF – Message Queues



Amazon  
SQS



GCP  
Pub/Sub



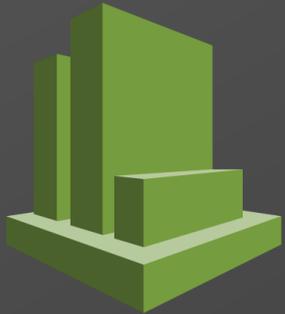
NATS

**Issue #5**

**Logging &  
Monitoring**

# Issue #5 – Logging & Monitoring

AWS, GCP, OSS / CNCF



**Amazon  
CloudWatch**



**GCP Monitoring  
& Logging**



**Fluentd &  
Prometheus**

**Issue #6**

**Networking**

# Issue #6 – Networking

AWS, GCP, OSS / CNCF

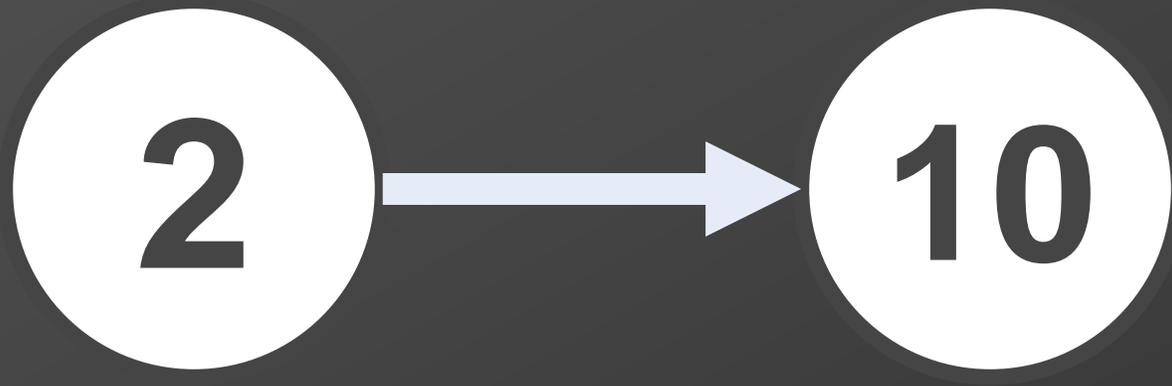


Amazon  
VPC

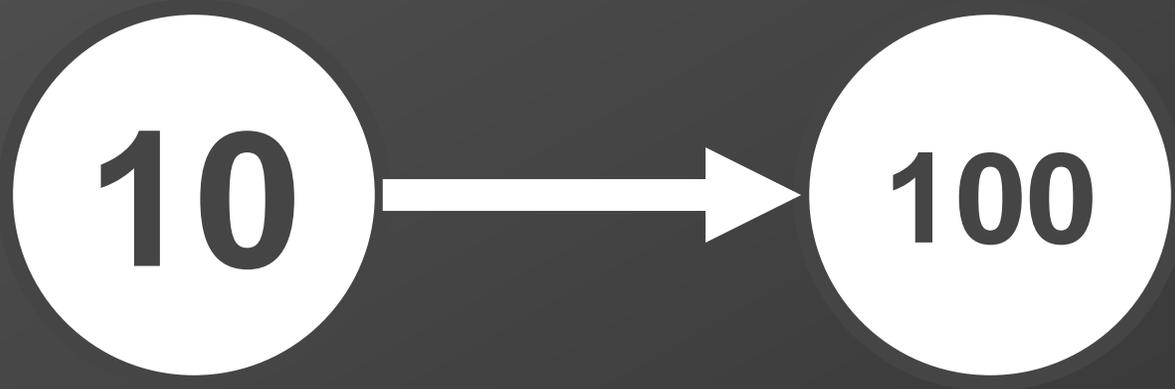


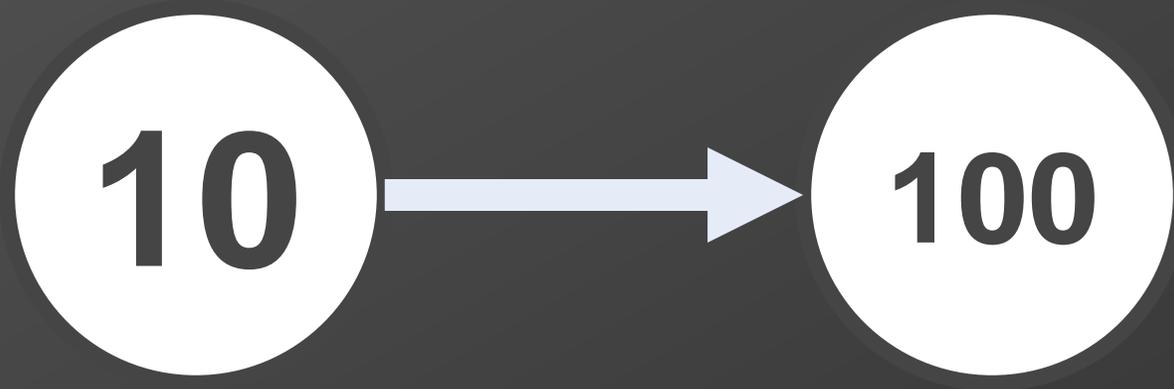
GCP Virtual  
Network





**COMPLETE**





**Why?**



ROAD



WRONG  
WAY

ROAD  
CLOSED

DETOUR

**1 Cluster : 1 Region**

**(let's revisit this)**

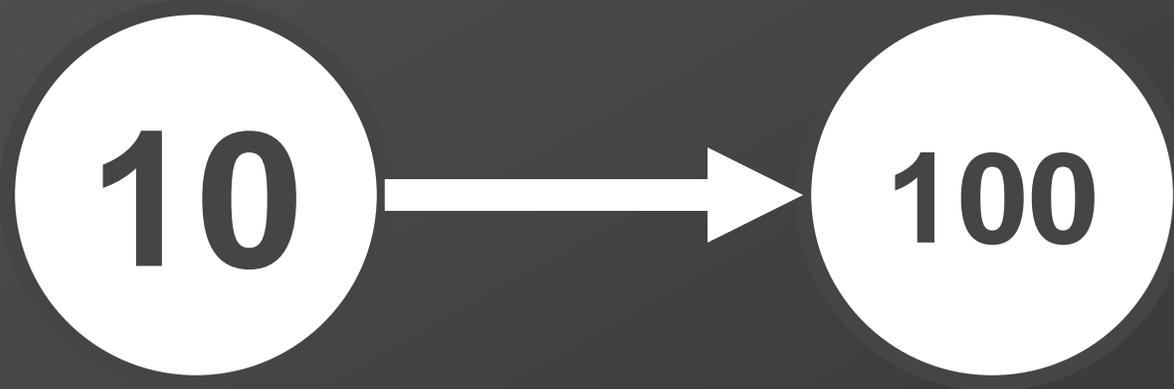
**> 1 Cluster : 1 Region ?**

# Reasons for >1C:1R

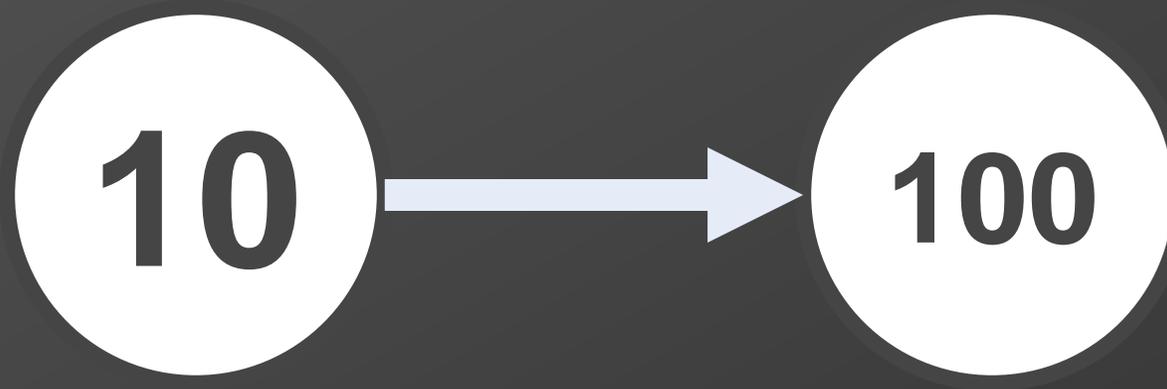
**#1 – Isolation of Teams or Special Resources**

**#2 – Grid Computing, Clusters as Cattle**

**#3 – Product Packaging, Shipping a Pod**



**Why?**



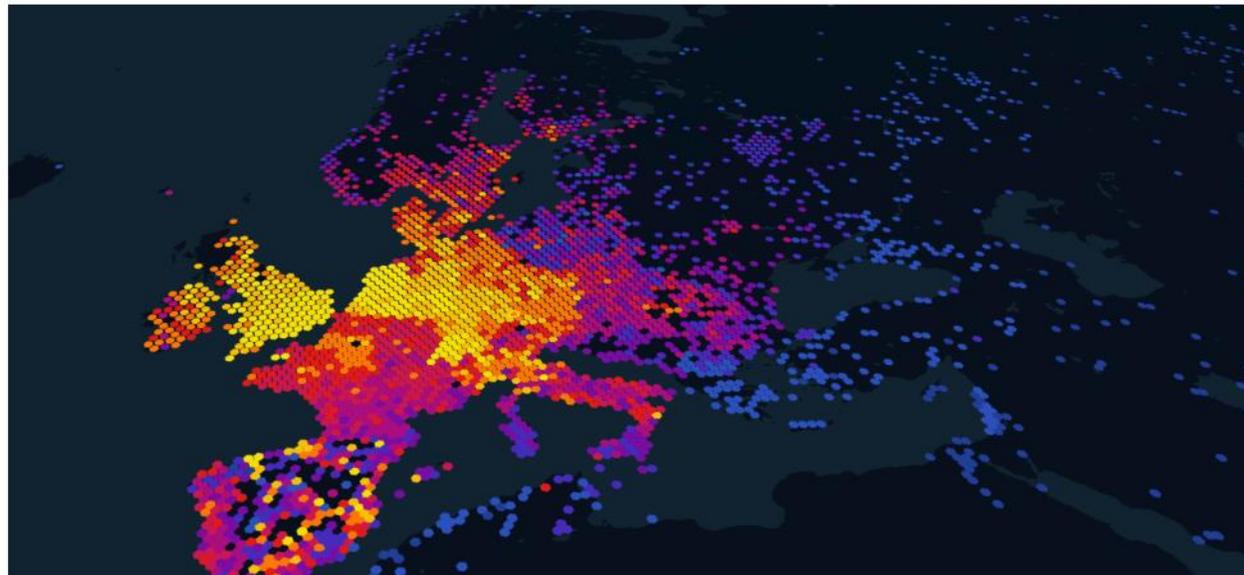
**Performance  
& Scalability**

# Dropbox Example (20 PoPs)

## Dropbox traffic infrastructure: Edge network

Oleg Guba and Alexey Ivanov | October 10, 2018

  436  0 



**#1 - App Topology**

**#2 - Data Replication**

**#3 - Traffic Routing**

**#4 - Service Discovery**

**#5 - Logging & Monitoring**

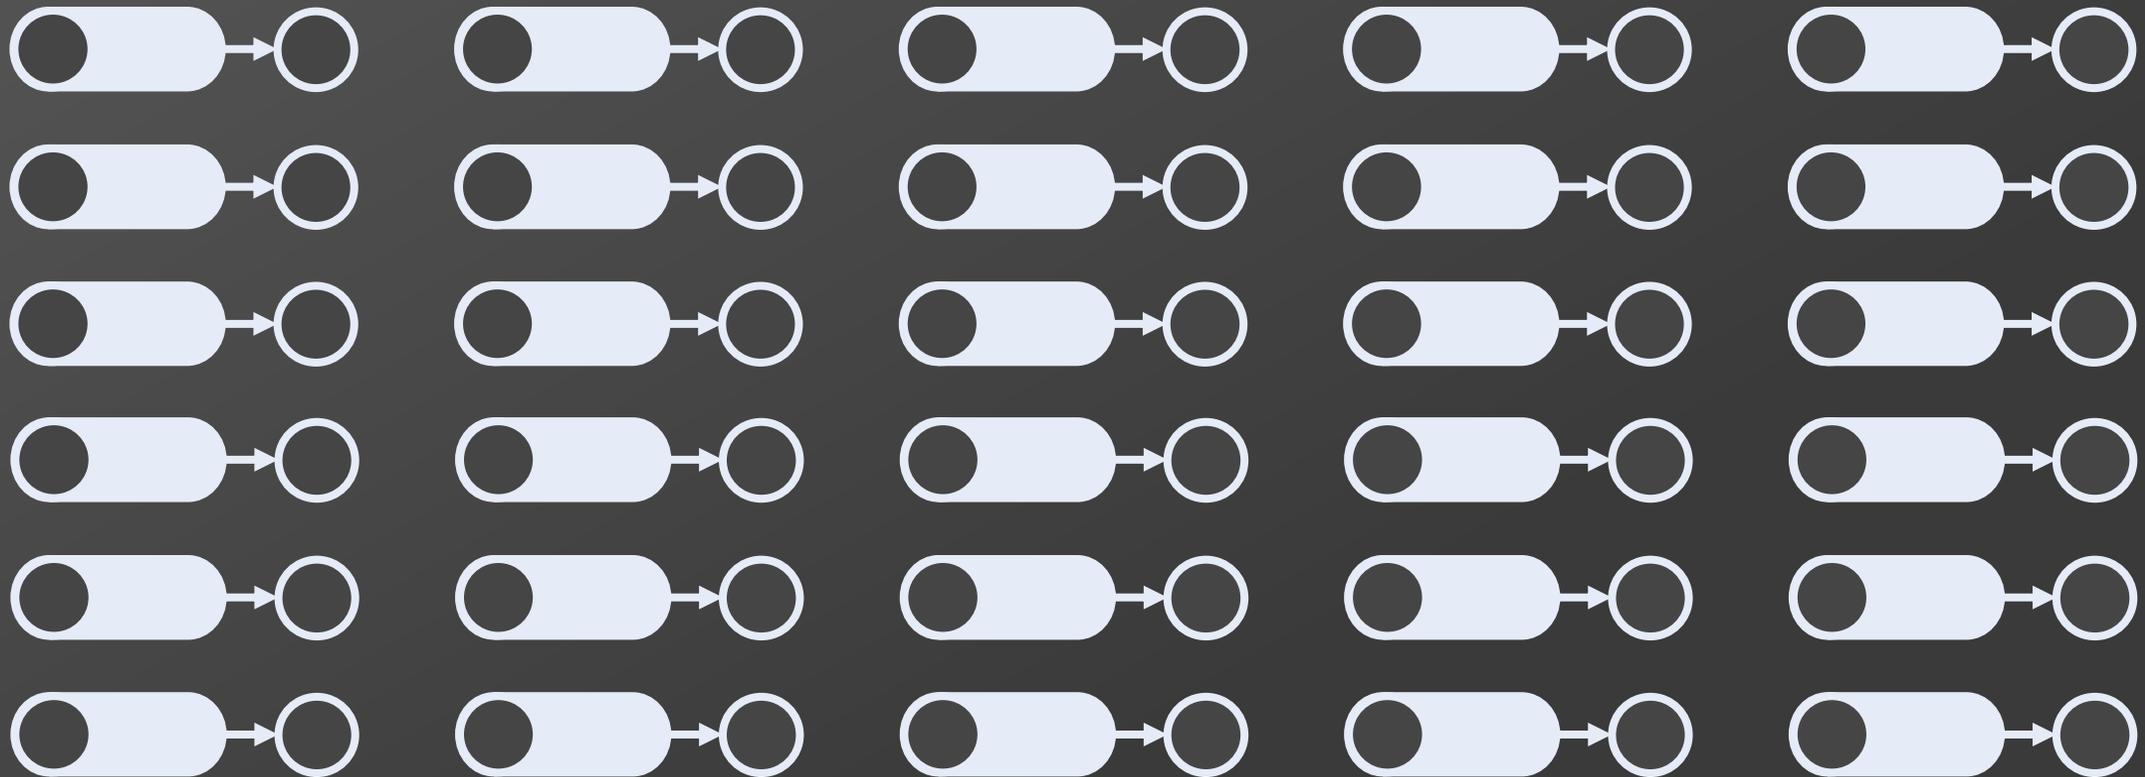
**#6 - Networking**

**Issue #7**

**CI / CD**

# Issue #7 – CI / CD

Maintain a unique pipeline per cluster?

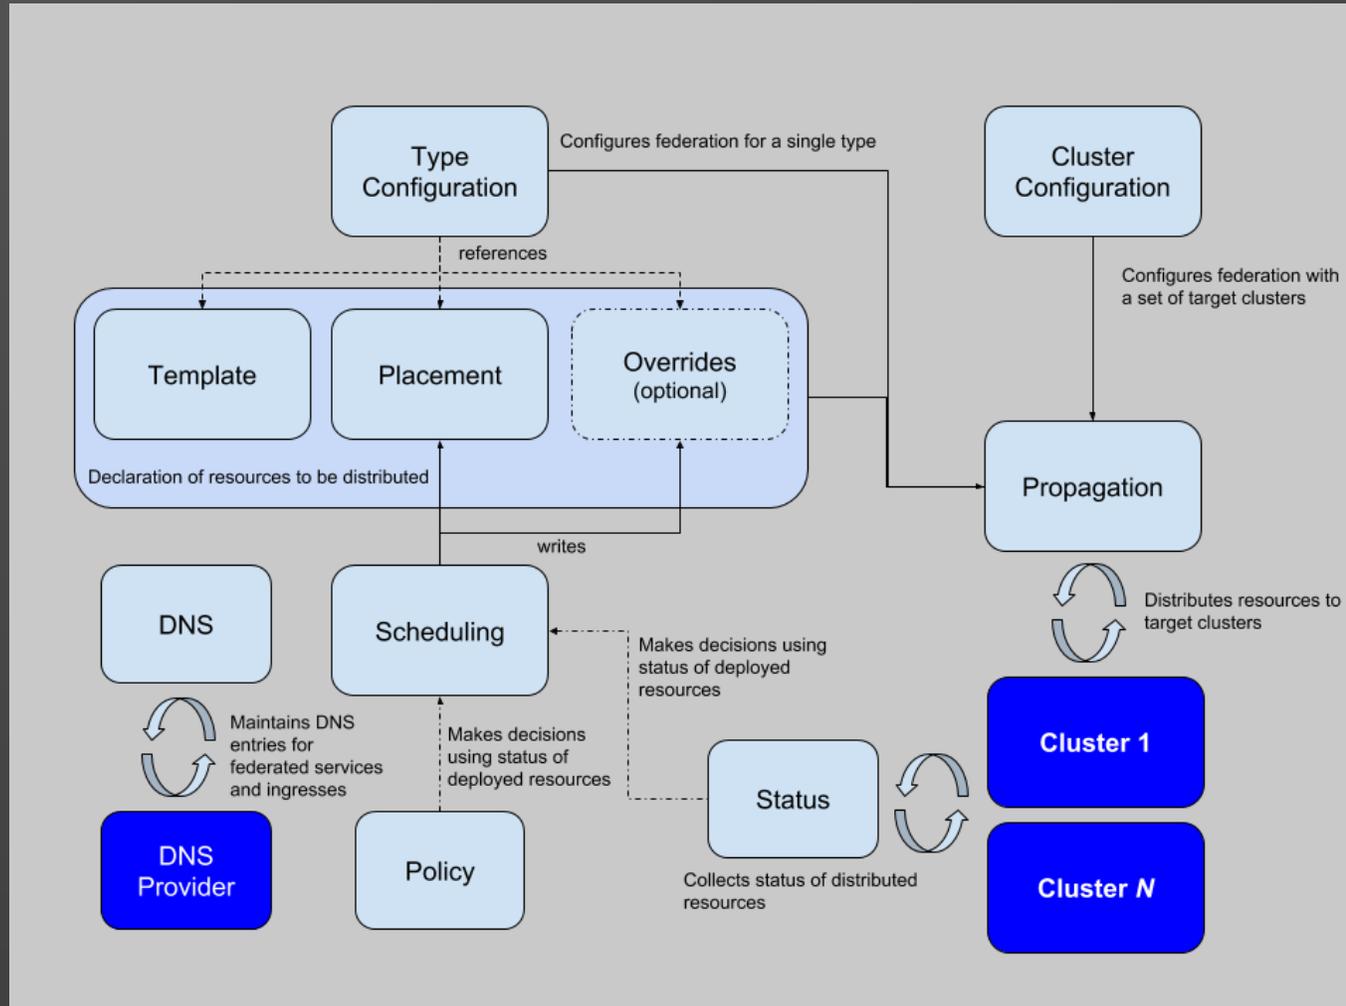


**Concept #2**

**Resource  
Propagation**



# Kubernetes Federation v2



**Issue #8**

**Security**

# Issue #8 – Security

Trust the integrity of 100 clusters?



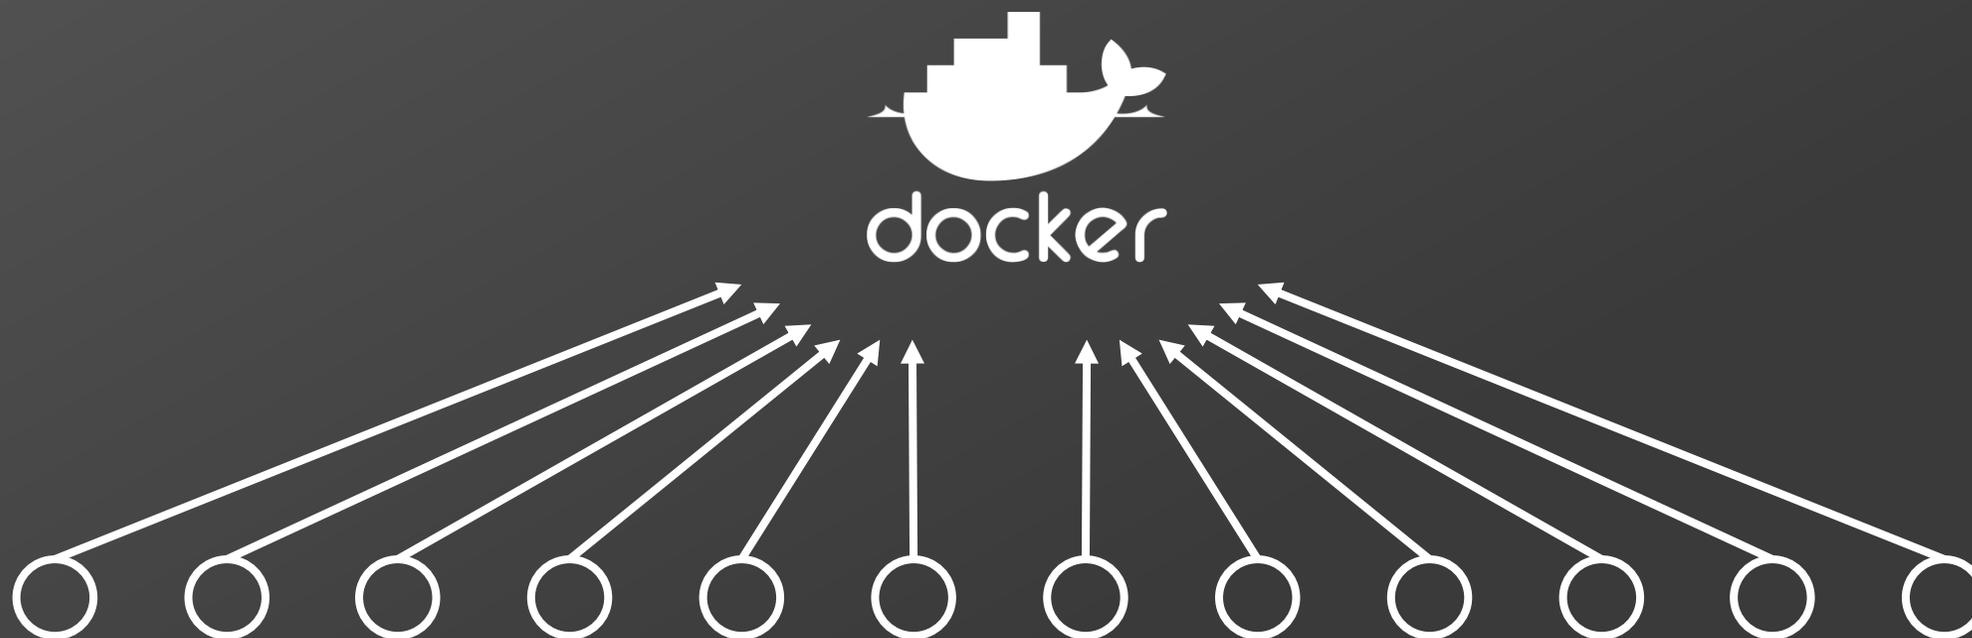


**Issue #9**

# **Image Registry**

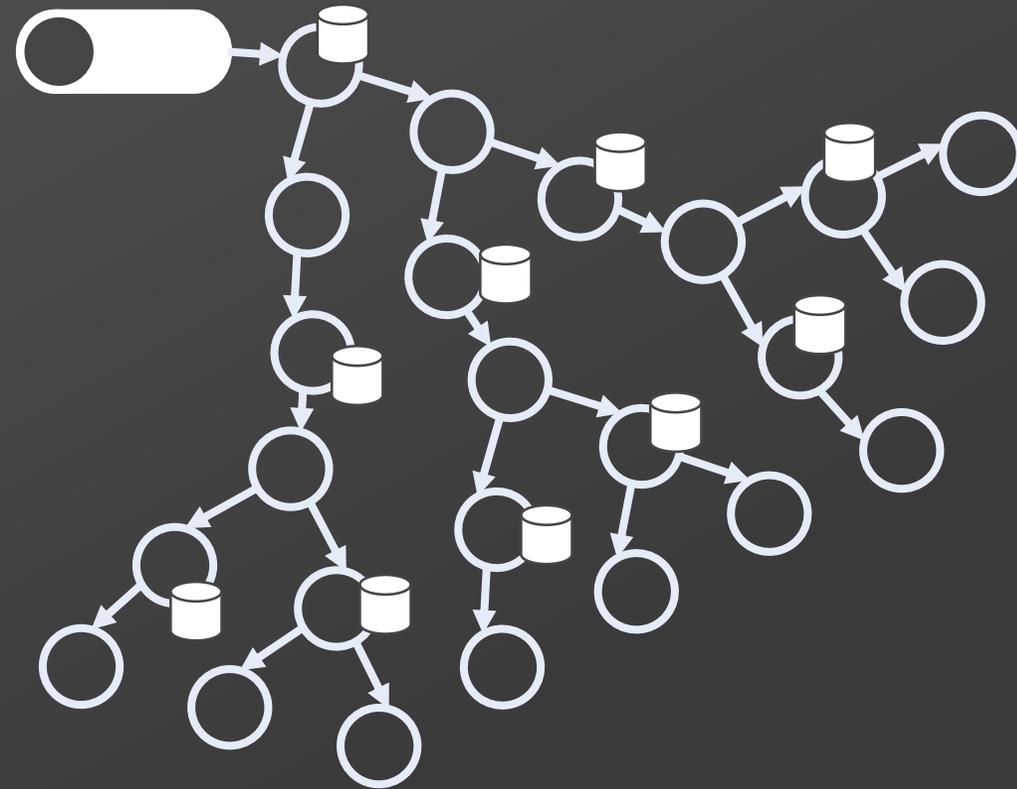
# Issue #9 – Image Registry

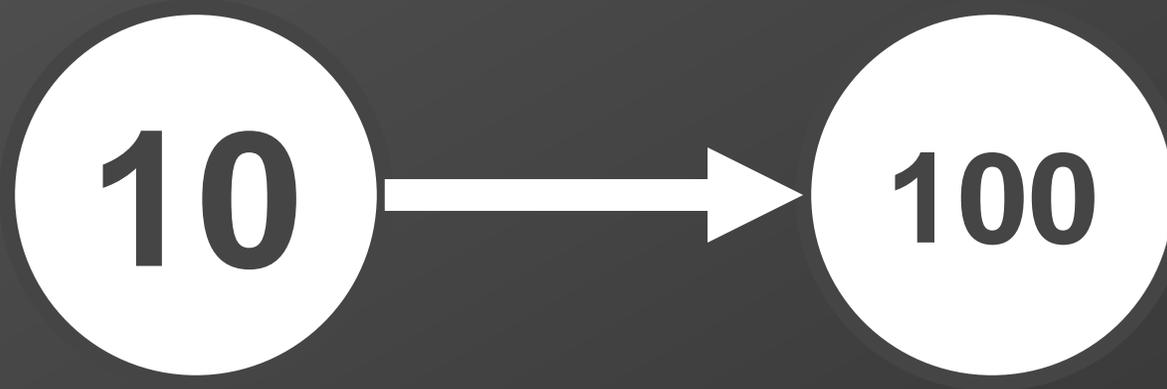
Rely on Docker Hub for image distribution?



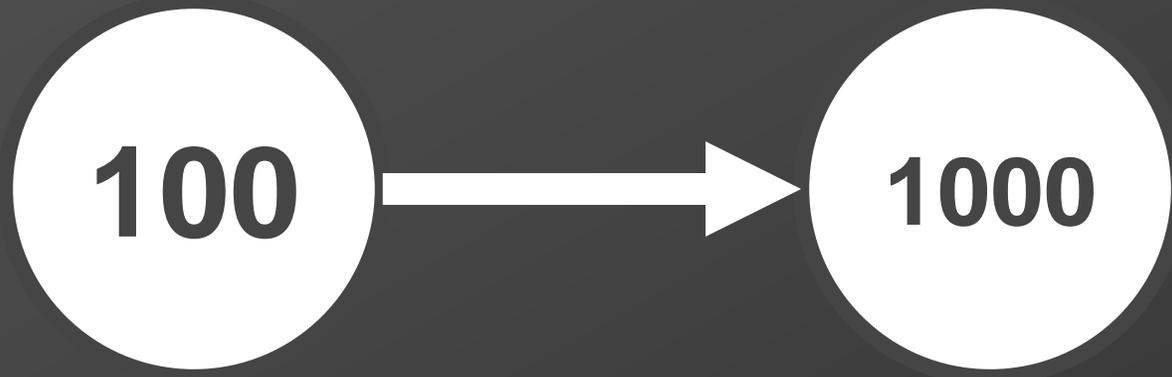
# Issue #9 – Image Distribution

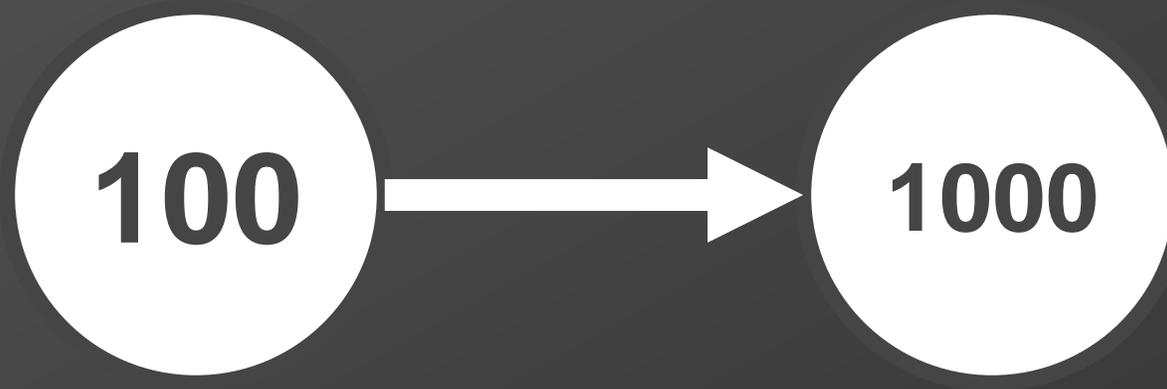
## Build Image / Artifact Caching into Topology





**COMPLETE**





**Why?**

# Chick-fil-A Example (2000 clusters)

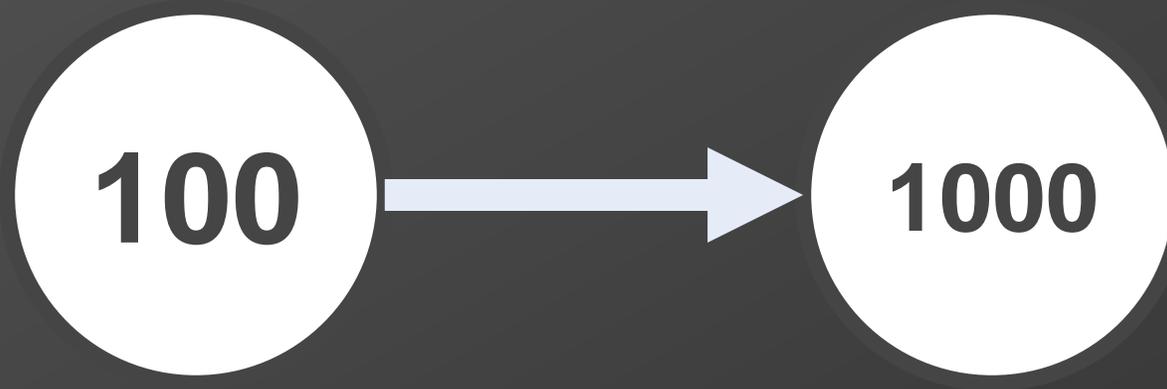


Chick-fil-A Tech Blog  
Jun 25 · 6 min read

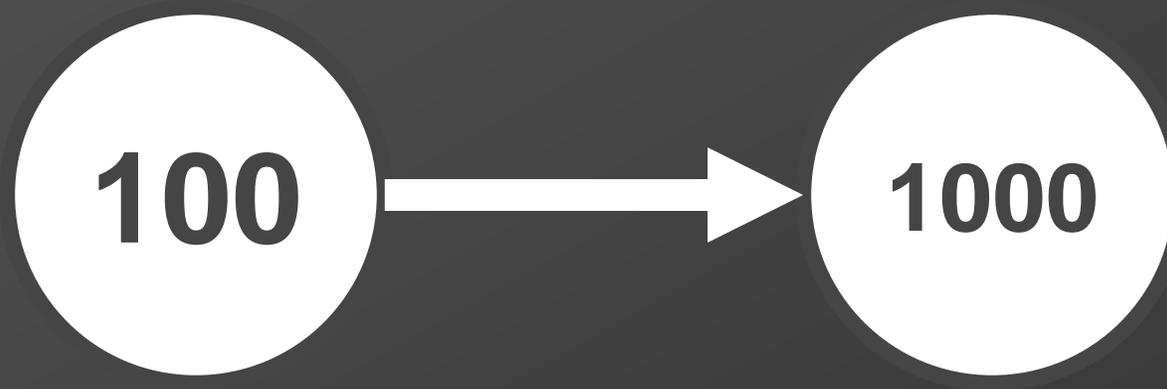
## Bare Metal K8s Clustering at Chick-fil-A Scale

by [Brian Chambers](#), [Caleb Hurd](#), and Alex Crane





**Latency, Bandwidth,  
Security, Privacy, Resiliency**



**Edge Computing**

**Issue #10**

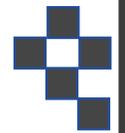
**Scarcity**

# Issue #10 – Scarcity

Run every service on every cluster all the time?



CORE



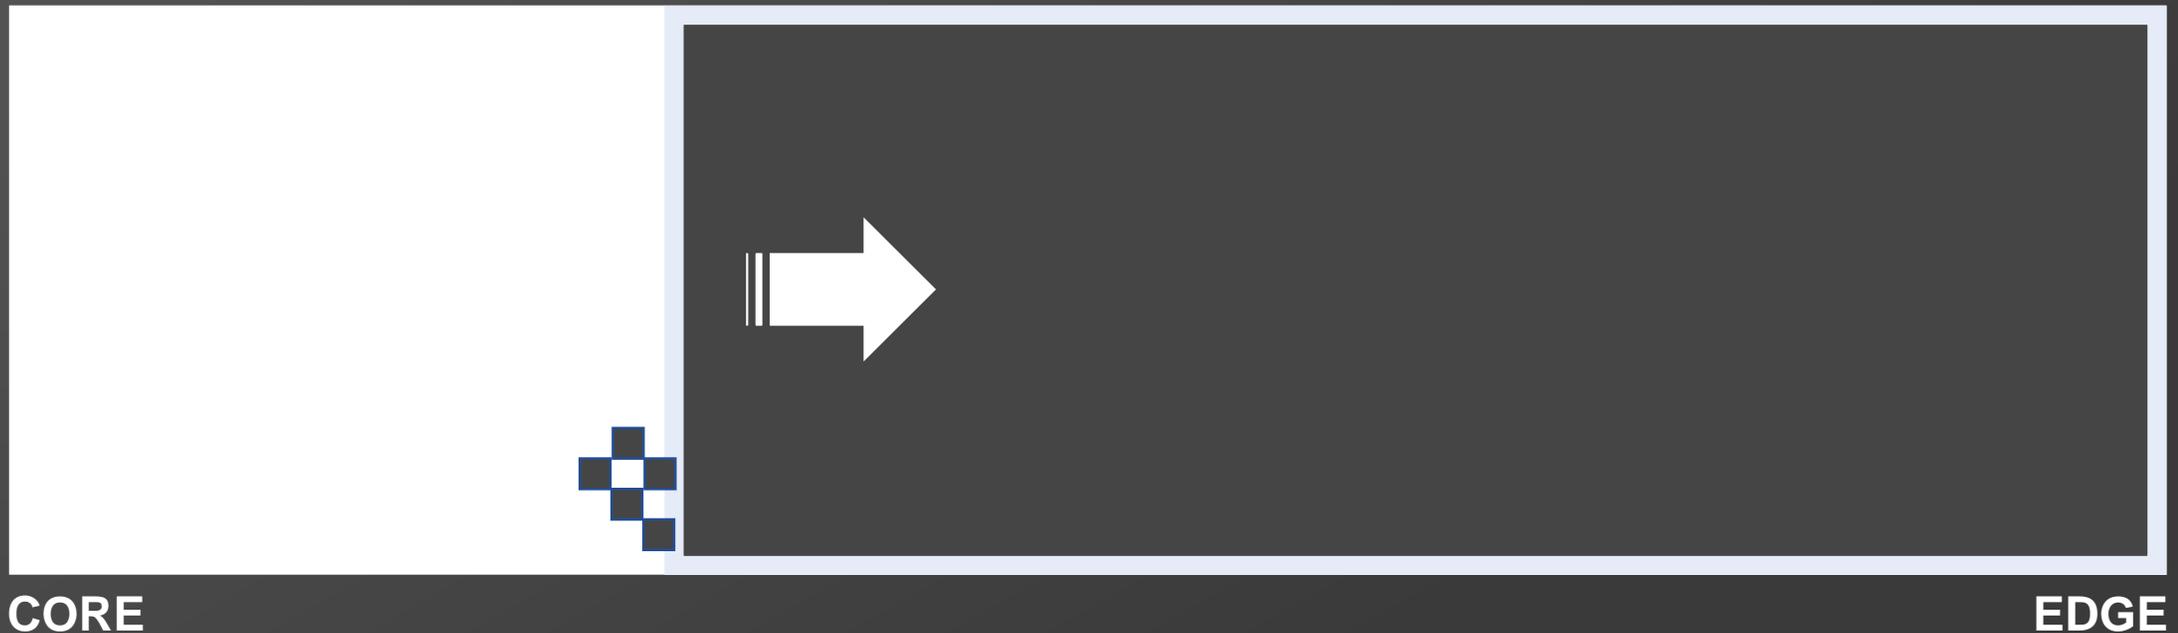
EDGE

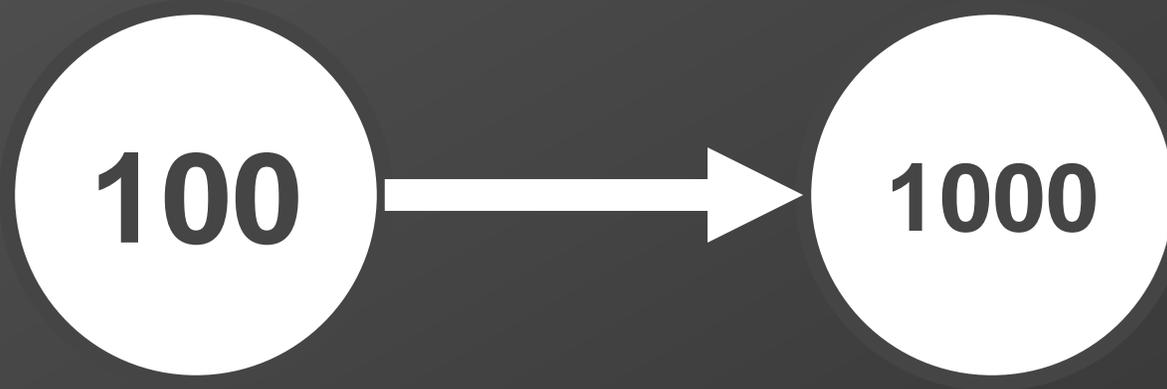
**Concept #3**

**Multi-dimensional  
Auto-scaling**

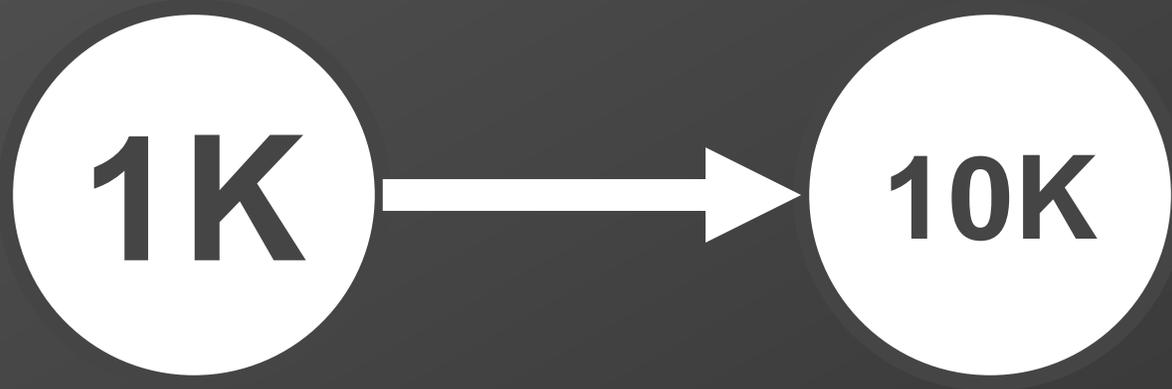
# Issue #10 – Scarcity

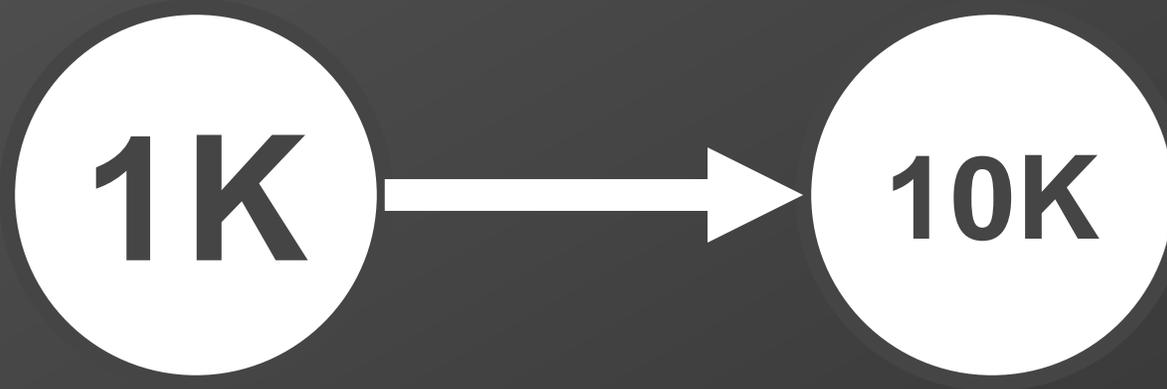
Auto-scale out to the Edge and back as needed



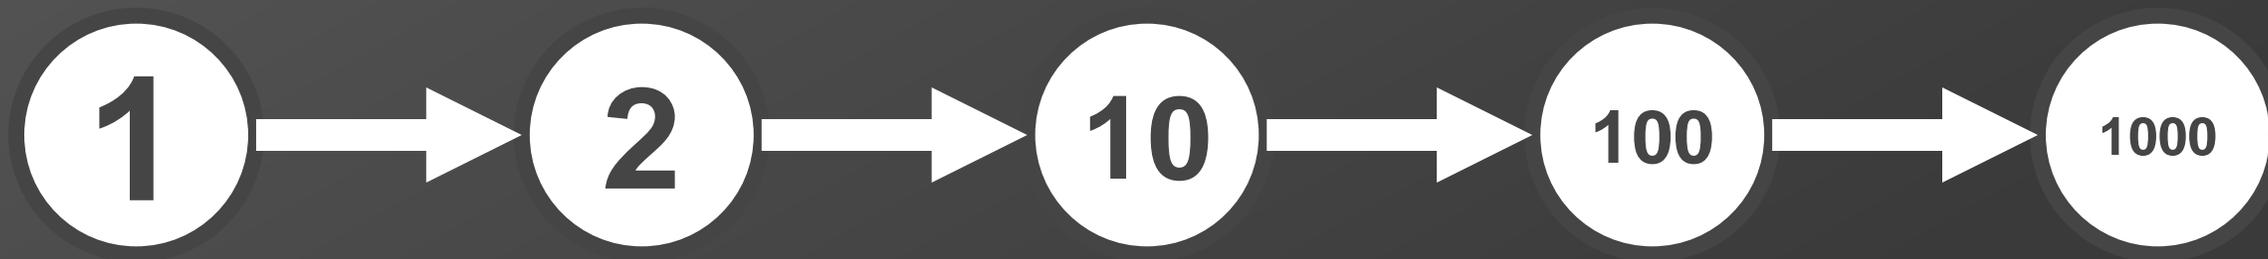


**COMPLETE**





**Future**



# 10 Multi-Cluster Issues

**#1 - App Topology**

**#6 - Networking**

**#2 - Data Replication**

**#7 - CI / CD**

**#3 - Traffic Routing**

**#8 - Security**

**#4 - Service Discovery**

**#9 - Image Distribution**

**#5 - Logging & Monitoring**

**#10 - Scarcity**

# **3 Multi-Cluster Concepts**

**#1 – Hierarchical Topology**

**#2 – Resource Propagation**

**#3 – Multi-dimensional Auto-scaling**

*Never use more  
clusters than  
you need*

# MULTI-CLUSTER CONFIDENCE

# OORT

Cloud Native Edge Computing

Email

Get updates

**Q & A**

<https://calendly.com/mcaulfield/kubecon>

# **Clusters all the way down Crazy Multi-cluster Topologies**

**presented by Matt Caulfield**



**KubeCon**

**CloudNativeCon**

————— **North America 2018** —————

