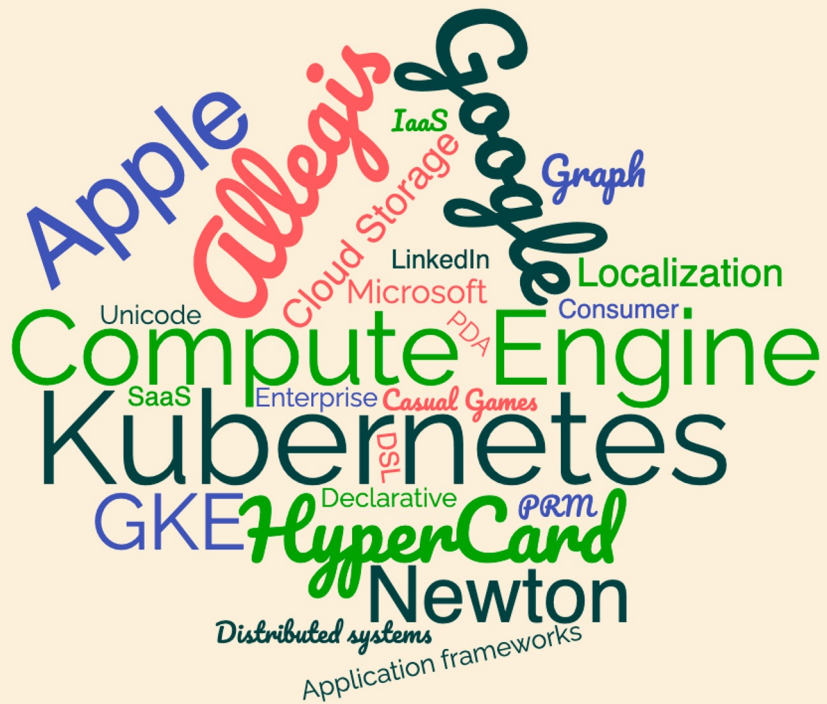


# Building Cloud infrastructure at Google

from VMs to Kubernetes

Martin Gannholm





Google Cloud

# Google's Proto-Cloud

- Container orchestration (Borg)
- Service discovery
- Distributed file system (GFS->Colossus)
- Blob storage (Blobstore)
- NoSQL (BigTable, Megastore->Spanner)
- Log processing
- MapReduce
- Unstructured query (BigQuery)
- Message queue
- Identity
- SDN - Ingress/Load balancing

# GCS - Google Cloud Storage

- HA control plane
- Customer account and billing
- Simplified model over Blobstore, auto provisioning
- Admin console
- Command line tool

# GCE - Google Compute Engine

- HA control plane
  - Customer account and billing
  - Admin console
  - Command line tool
- 
- High scale provisioning
  - Extra layers of security (dedicated machines, encryption)
  - Persistent disks
  - Encapsulated networking
  - Suppression of preemption
  - Live migration

# Google/Cloud Native Principles

# Expect Failures

Unplanned failures *will* happen, so focus on:

- Detection
- Failing gracefully
- Resilience
- Playbooks
- MTTR
- DiRT



# Avoid State

Cattle, not pets. Components should not retain state between requests

Benefits:

- No special recovery flow
- No codependency between component instances
- Deployment and scaling are trivial

# Separate Concerns

Factor system into components

- Loosely coupled
- Narrow contracts

Benefits:

- Independently deployable
- Independently scalable
- Swappable implementations

# Layer Security

- Namespaces
- Inter-component communication
- System monitoring
- Admission control

# Build with confidence

- Build testable code
- Automate verification
- Monitored rollout

Infrastructure OS for the rest of us

# Kubernetes

Docker opened the door

Open Source - drive ecosystem

Agnostic

Innovations:

- State and Spec - Controller pattern
- Extension mechanisms: CRDs, aggregated APIs
- Pods - co-scheduled containers, addressable IP and shared network and storage

# Serverless compute abstraction: Knative

Make common patterns of stateless services easy

Google's Cloud Run conforms to the same API, but runs on Google's infra

# Summary

Kubernetes is the portable Infrastructure OS for the Cloud.

By adopting Kubernetes you get many of the Cloud Native principles implemented out of the box. You can also build your own abstractions and tooling, or leverage the rich ecosystem that exists. See the CNCF landscape