

# Stories from the Playbook

KubeCon Europe  
3 May 2018

Fred van den Driessche [fvdd@google.com](mailto:fvdd@google.com)  
Tina Zhang [ttz@google.com](mailto:ttz@google.com)

# Hi!

- GKE Site Reliability Engineers
- Currently part of Google SRE in London
  - 200+ people
  - More than 20 different teams: Ads, Cloud, Mobile, ...
- Sister team in Seattle

# GKE in Numbers

- 1 SRE team, 2 shards
- 15 Regions, 45 Zones
- 99.5% Uptime SLO
- Support last 3 minor versions
- Weekly release cycle

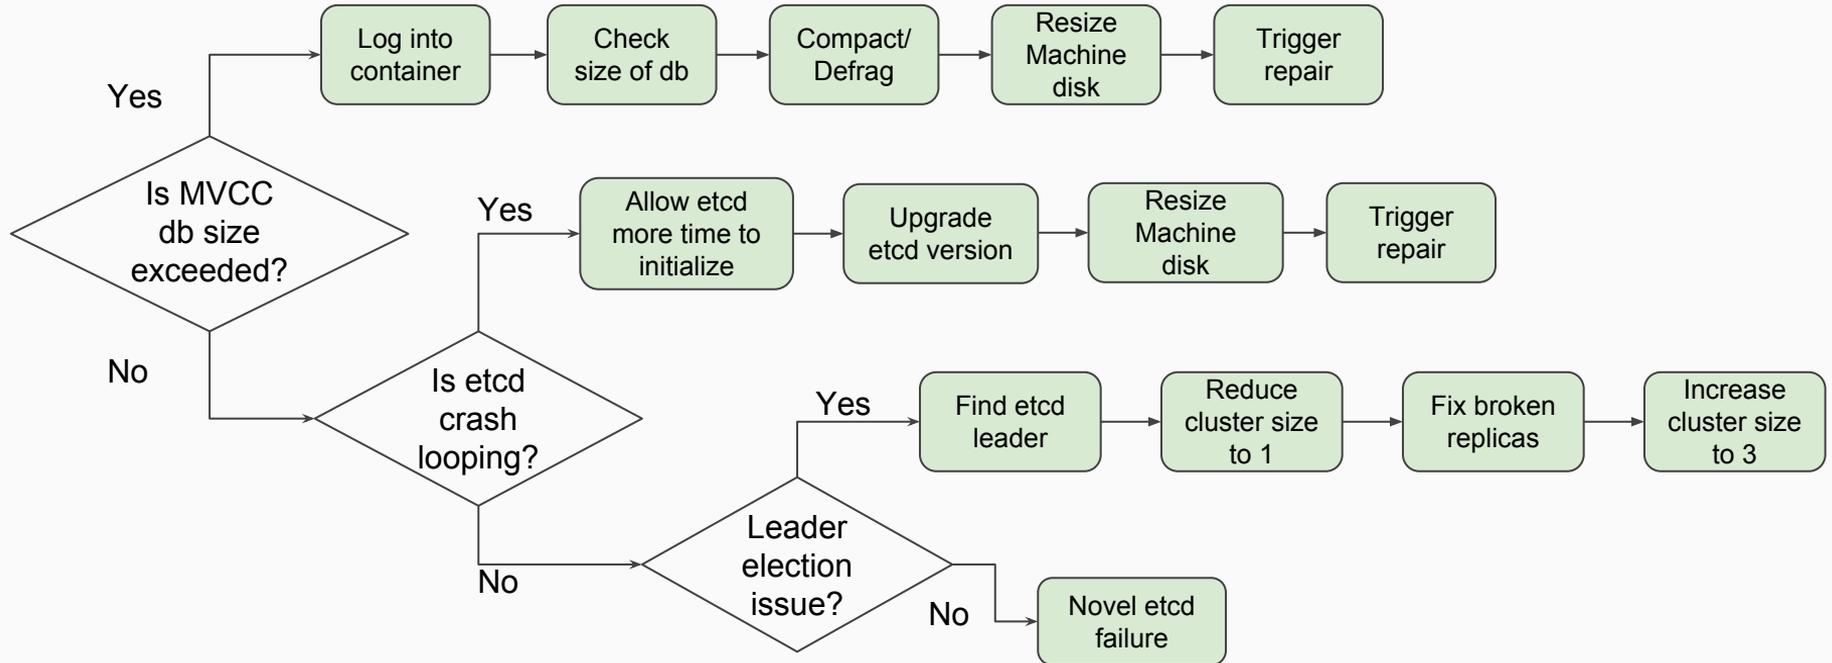


# Playbooks?

*'thinking through and recording the best practices ahead of time in a "playbook" produces roughly a 3x improvement in MTTR as compared to the strategy of "winging it."*

- SRE Book

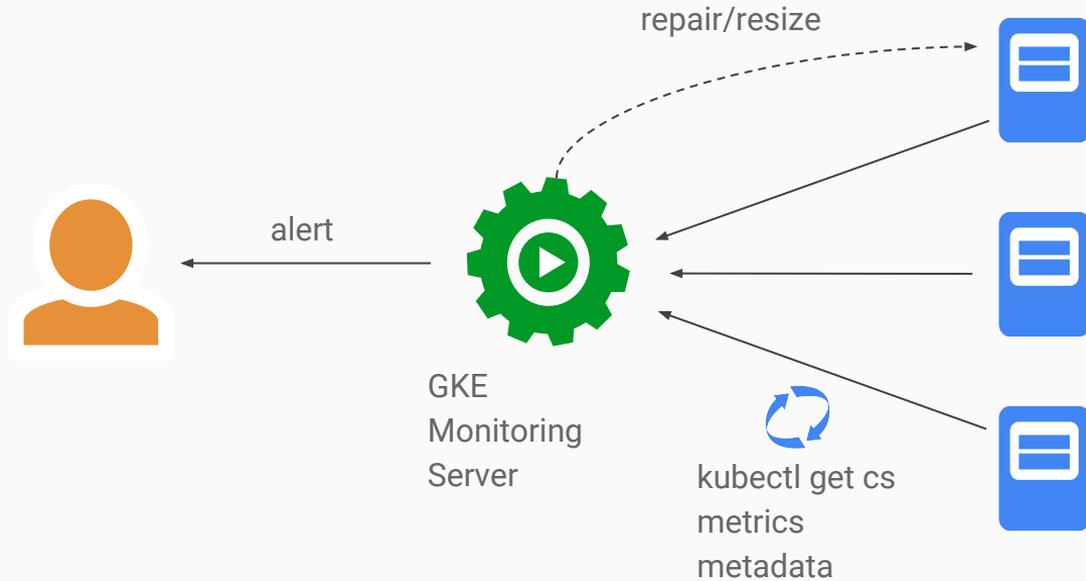
# Playbook Example: etcd debugging



# Observability

- Metrics
  - GKE Control Plane
  - On-cluster
- Dashboards
  - Control Plane service
  - Per-cluster
  - Fleet-wide
- Probers

# How do we monitor clusters?

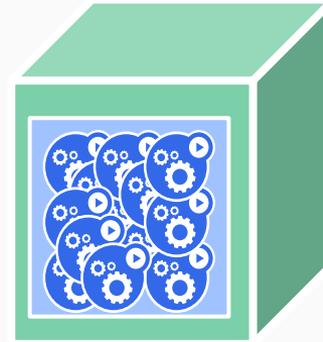
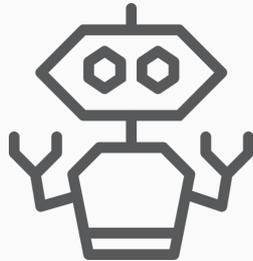


# What could possibly go wrong?

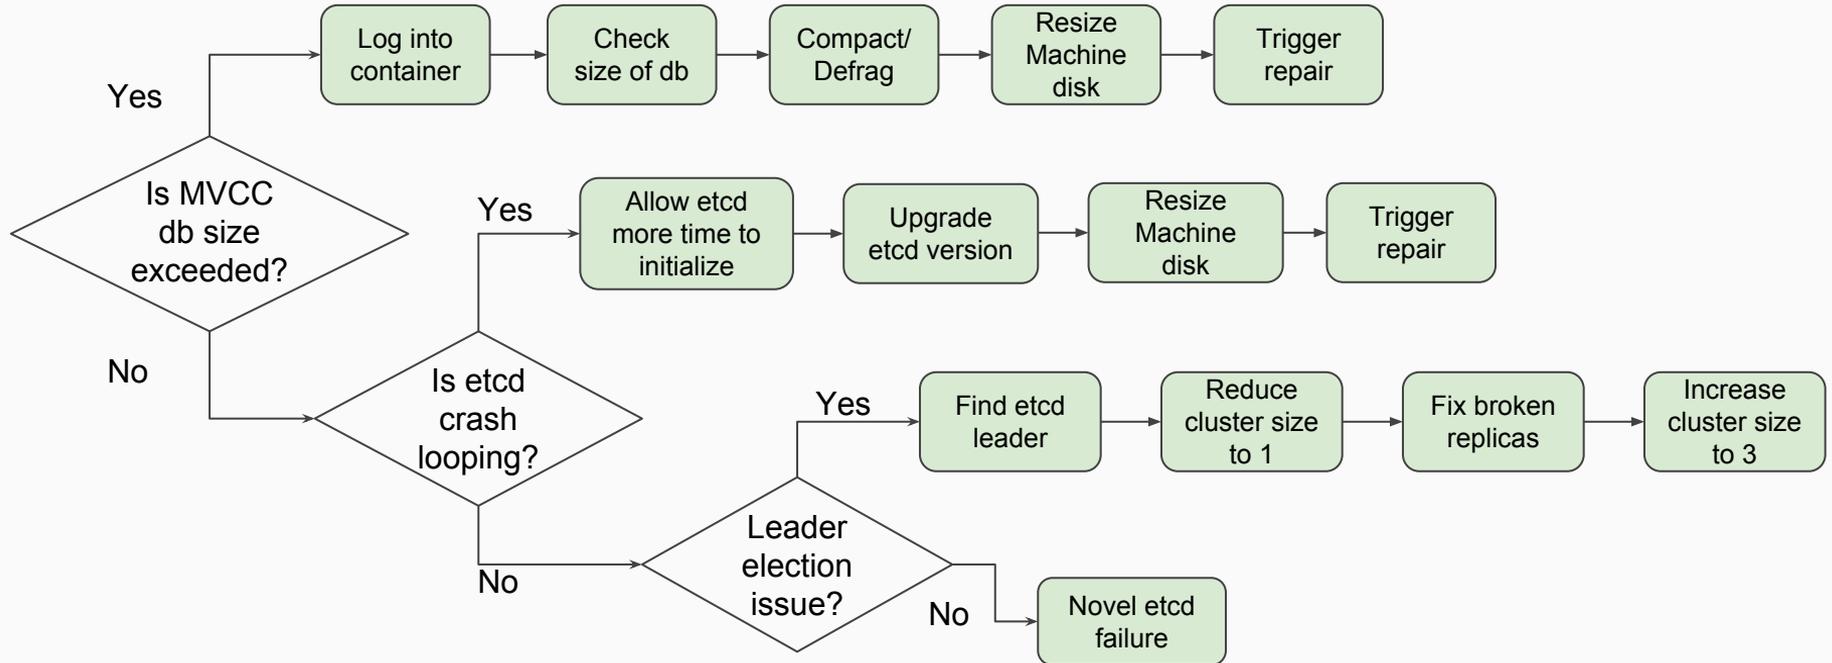
**PANIC!**



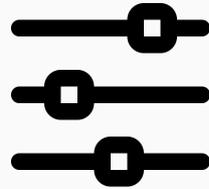
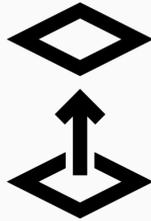
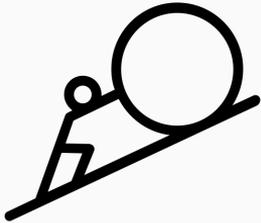
fluentd



# Playbook Example: etcd debugging



# Common Procedures



# Top Tips for a Happy ~~Onealer~~ Cluster

- Provision nodes to reflect workload
- Disk I/O
- Use resource limits on Pods
- Frequent etcd backups
- Upgrades
- Keep your service accounts enabled!

# Incident Management

- Coordinate, Communication, Control
- Established escalation paths
- Blameless Postmortem
  - What went well?
  - What went badly?
  - Where we got lucky
- Action items

# Docker Live Restore

1	<b>Symptoms:</b>	Pods not scheduled when docker restarts
2	<b>Mitigation:</b>	Rollback to previous OS image Downgrade existing nodes
3	<b>Root Cause:</b>	docker <code>live-restore</code> flag
4	<b>Fix:</b>	Disable <code>live-restore</code> flag in new image

# Image Corruption

1	<b>Symptoms:</b>	"Creation" prober failures in Europe geo
2	<b>Mitigation:</b>	???
3	<b>Root Cause:</b>	Internal GCR mirror serving corrupt image
4	<b>Fix:</b>	Purged internal mirror

# dnsmasq vulnerabilities

1	<b>Symptoms:</b>	???
2	<b>Mitigation:</b>	???
3	<b>Root Cause:</b>	Several vulnerabilities discovered in dnsmasq
4	<b>Fix:</b>	Create/Rollout patch releases for each K8s version

# What are we currently working on?

- Scaling GKE 100x
- Problem Fingerprinting
- Smarter Repairs and Resizing
- Private Clusters
- Cluster API