

# ONE

One Network Engineer

James Paussa

network infrastructure engineer

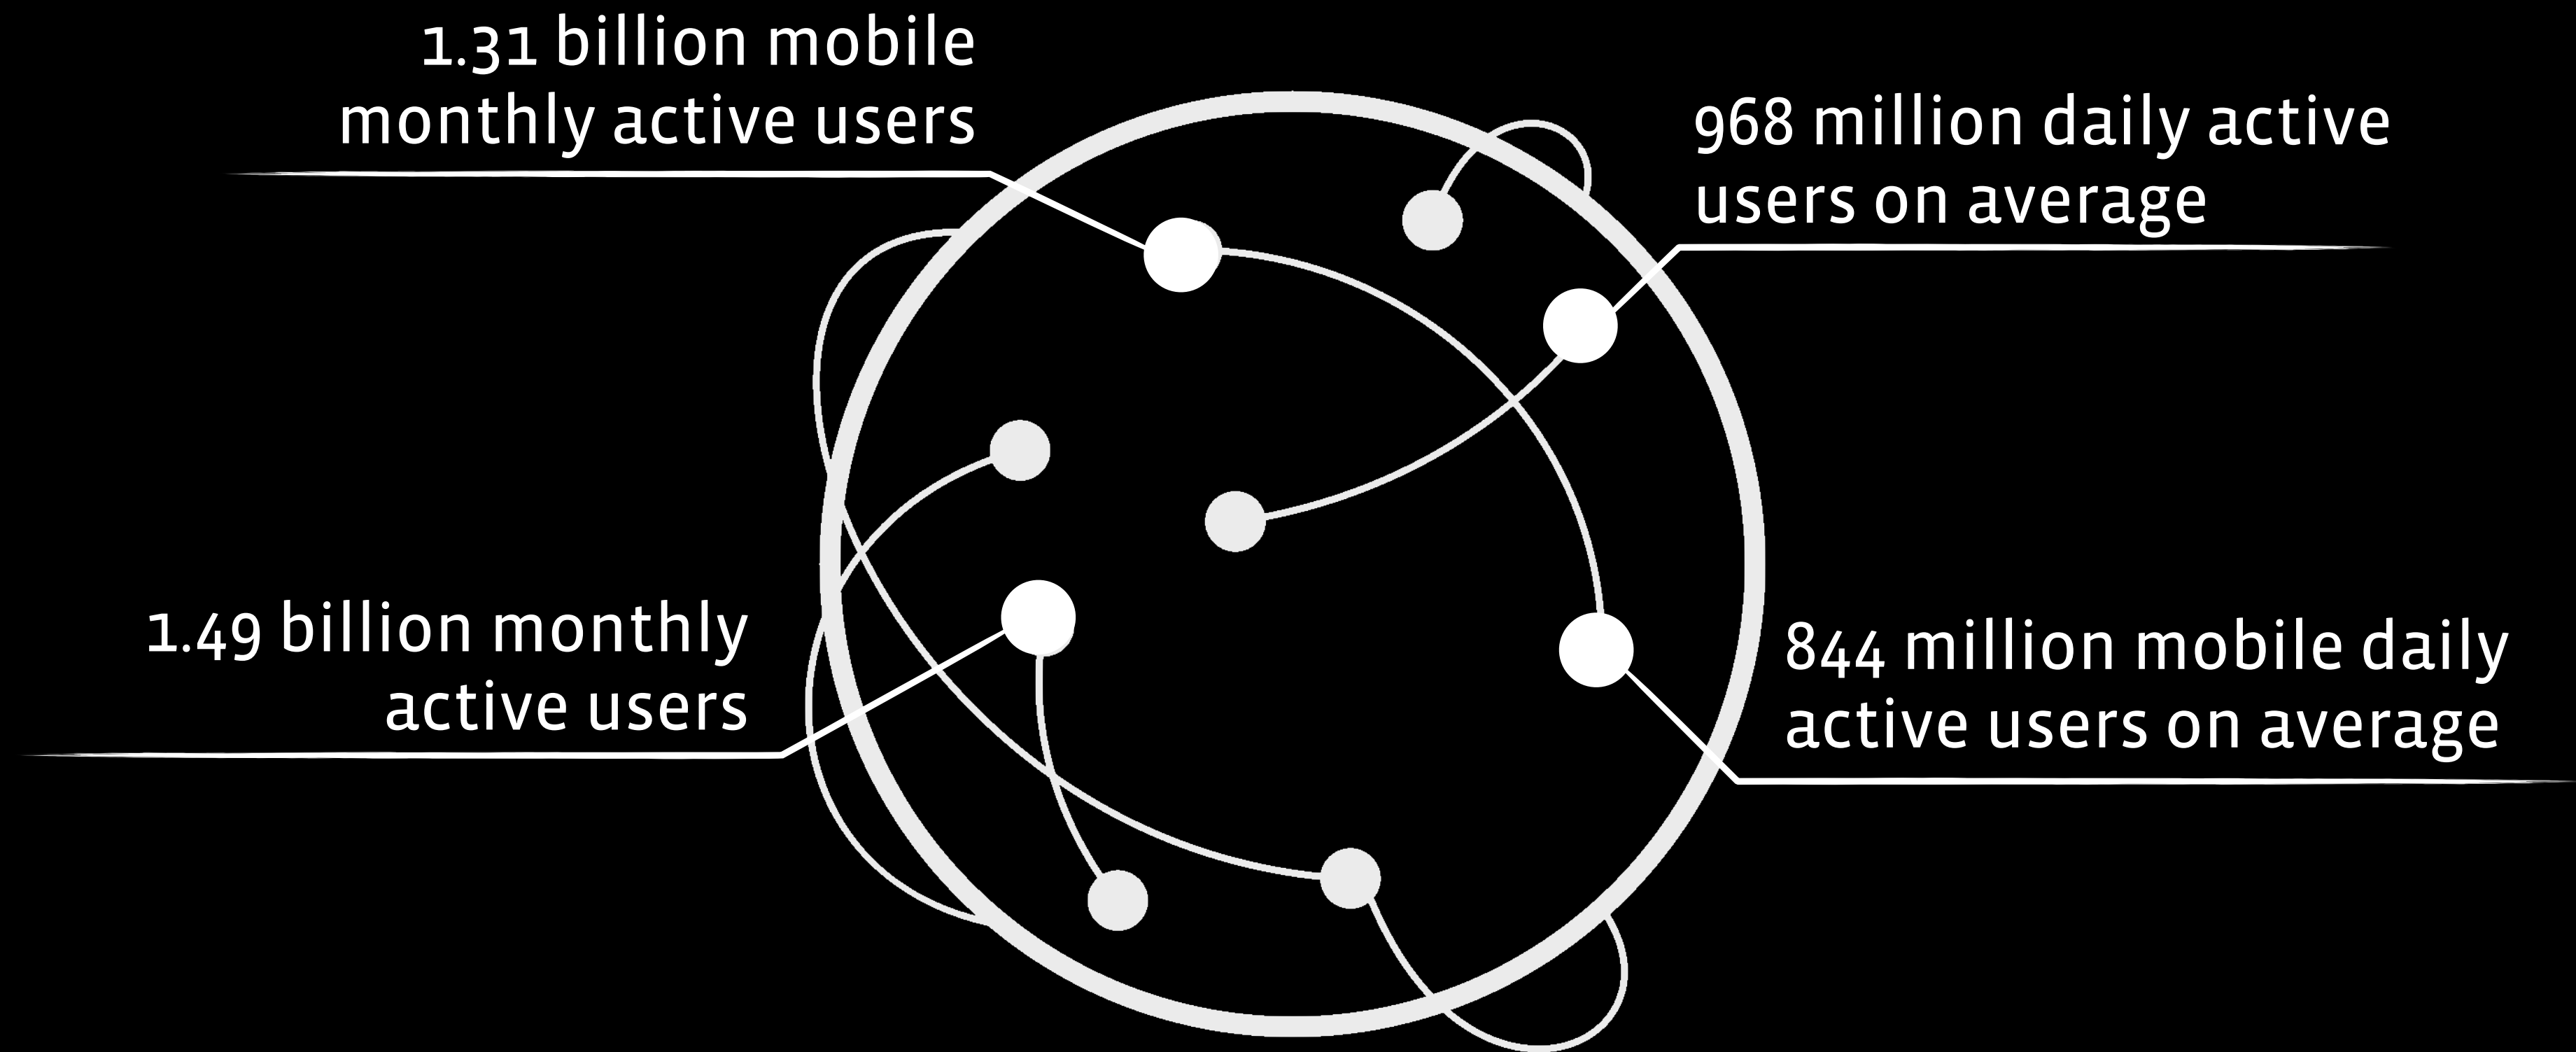


**facebook**

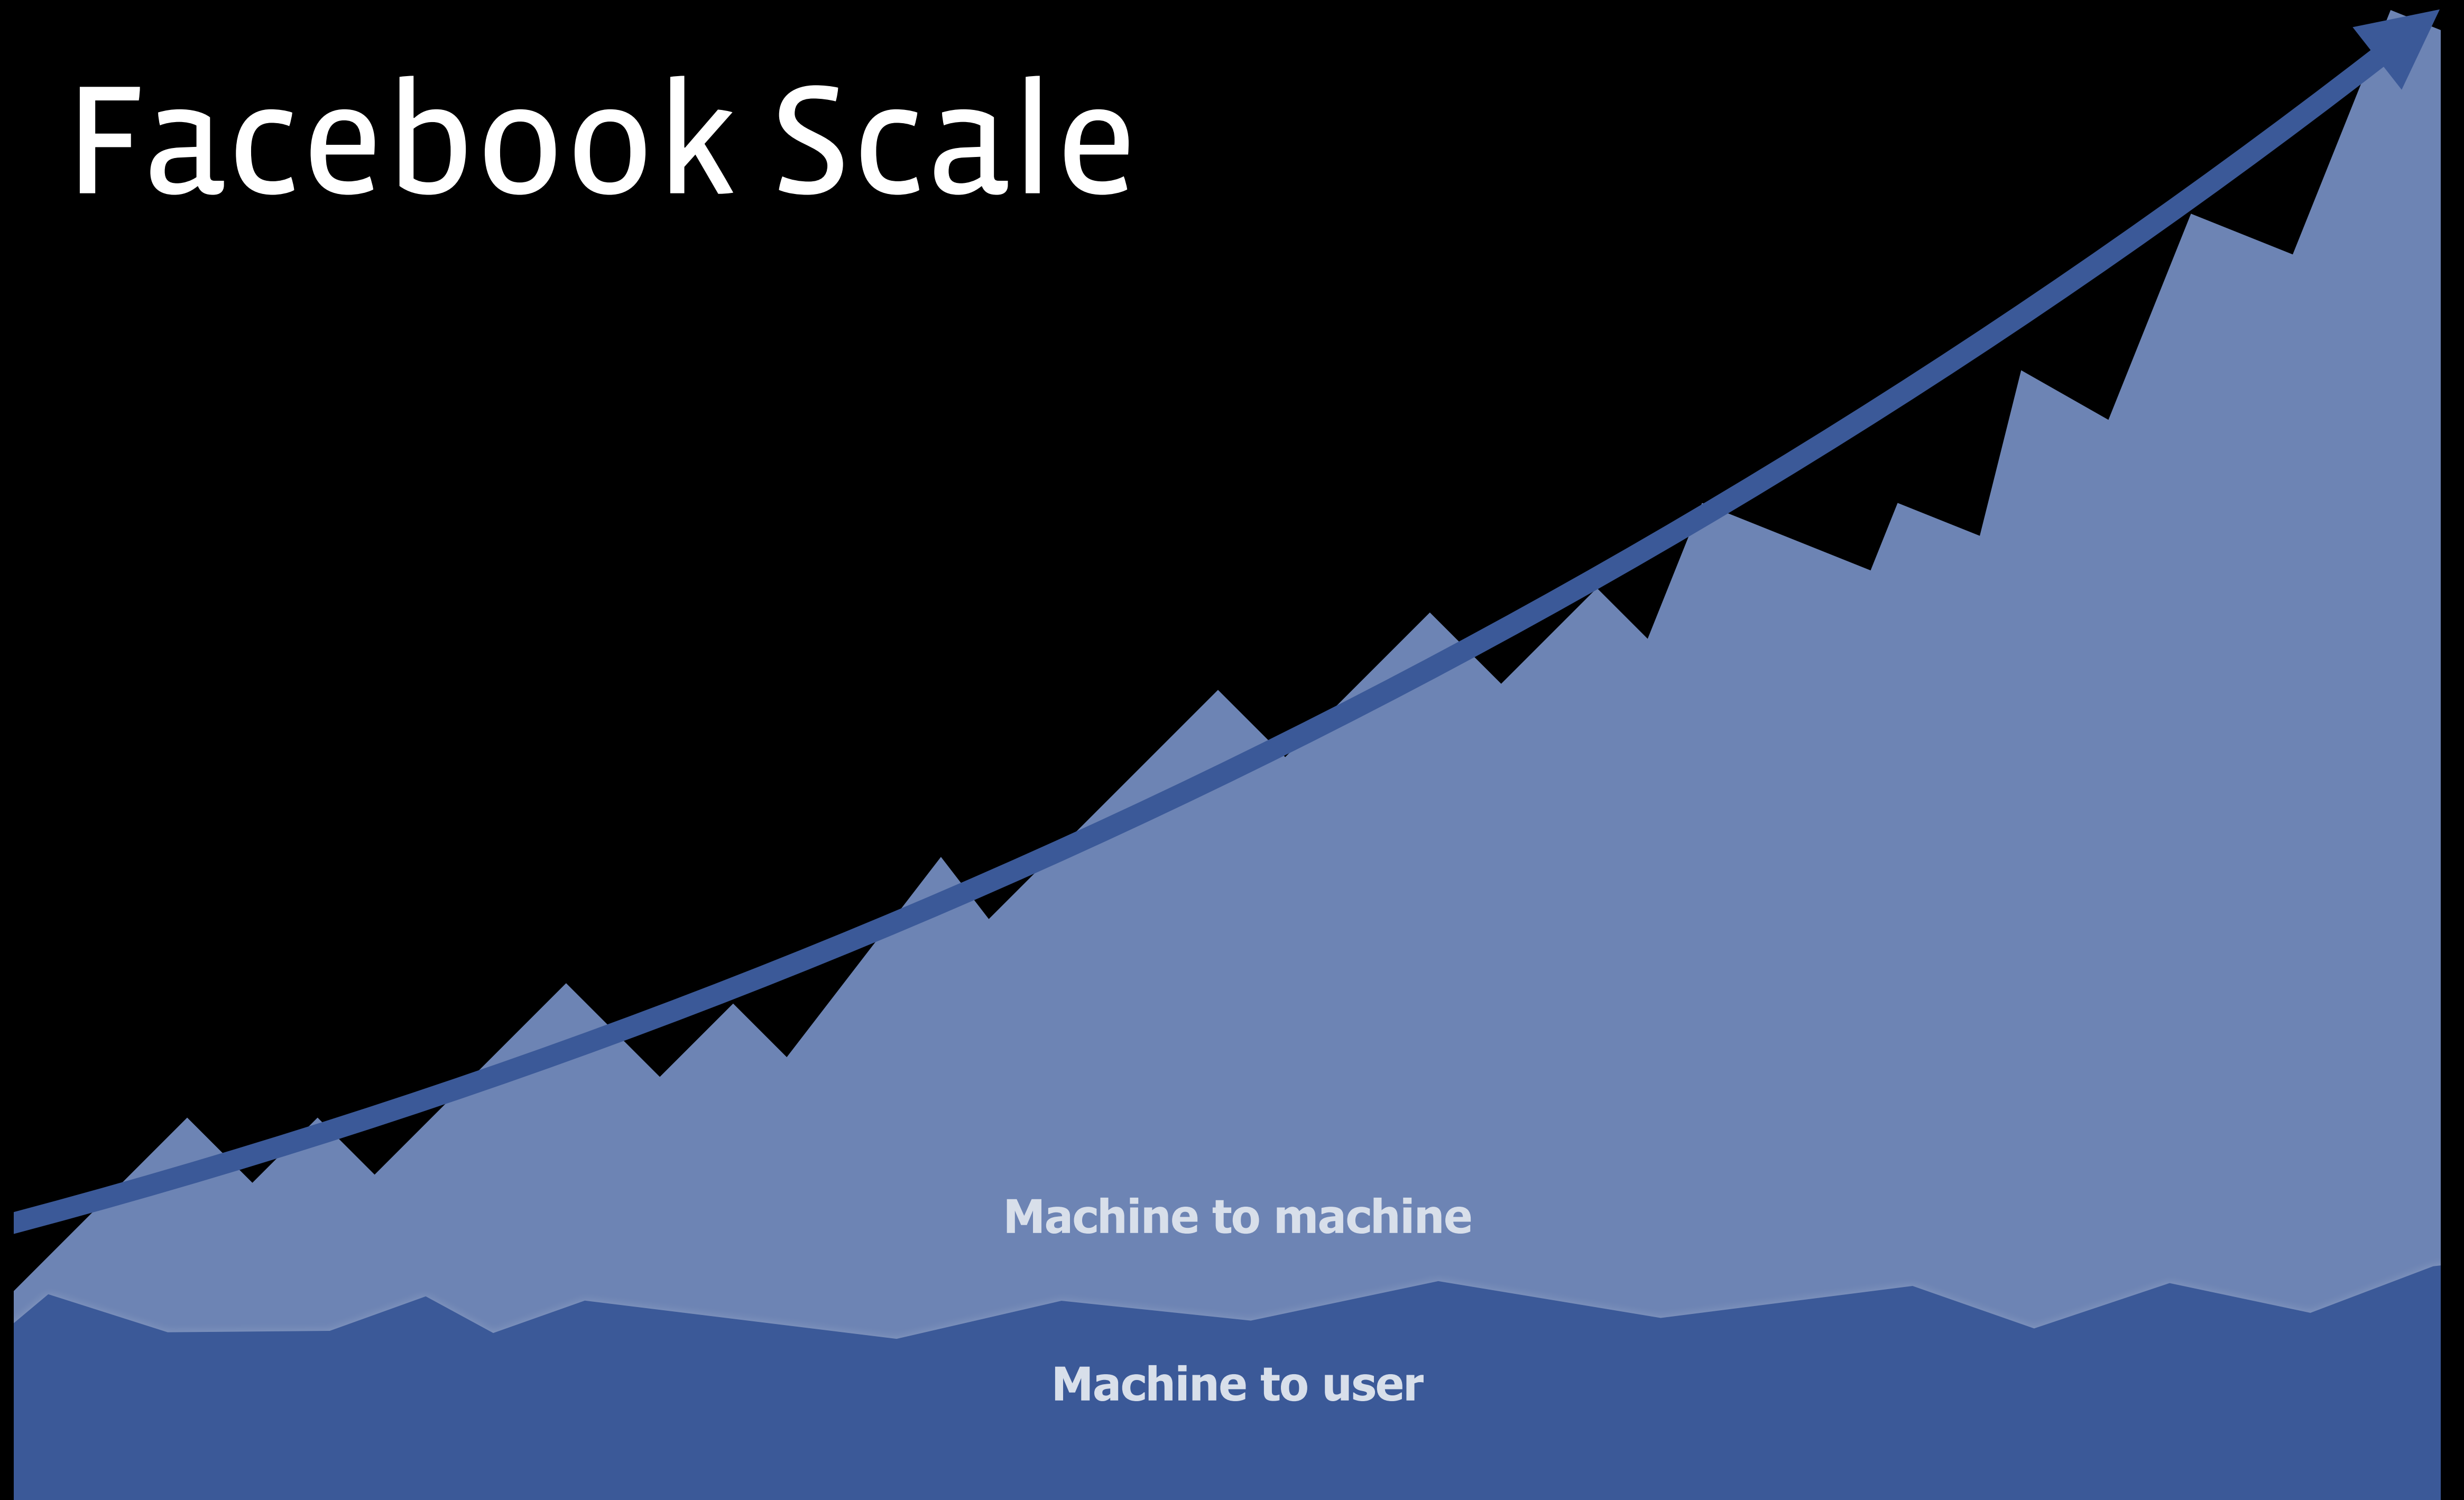
# What Do We Do?

- A single engineer runs the network at a time
- Network monitoring and alarming
- Responsible for the entire production network

# Facebook Scale



# Facebook Scale





# Automation





# What We Don't Do

- On site work (cleaning fibres, LCs, etc)
- Working with remote hands
- Device deployment


# Myths

- Automation fixes everything
- We've fixed everything
- Doesn't apply


# Link Imbalance



# Link Imbalance




# Link Imbalance






# Link Imbalance




# Link Imbalance





# Link Imbalance



# Link Imbalance

srsly wtf.



# Link Imbalance

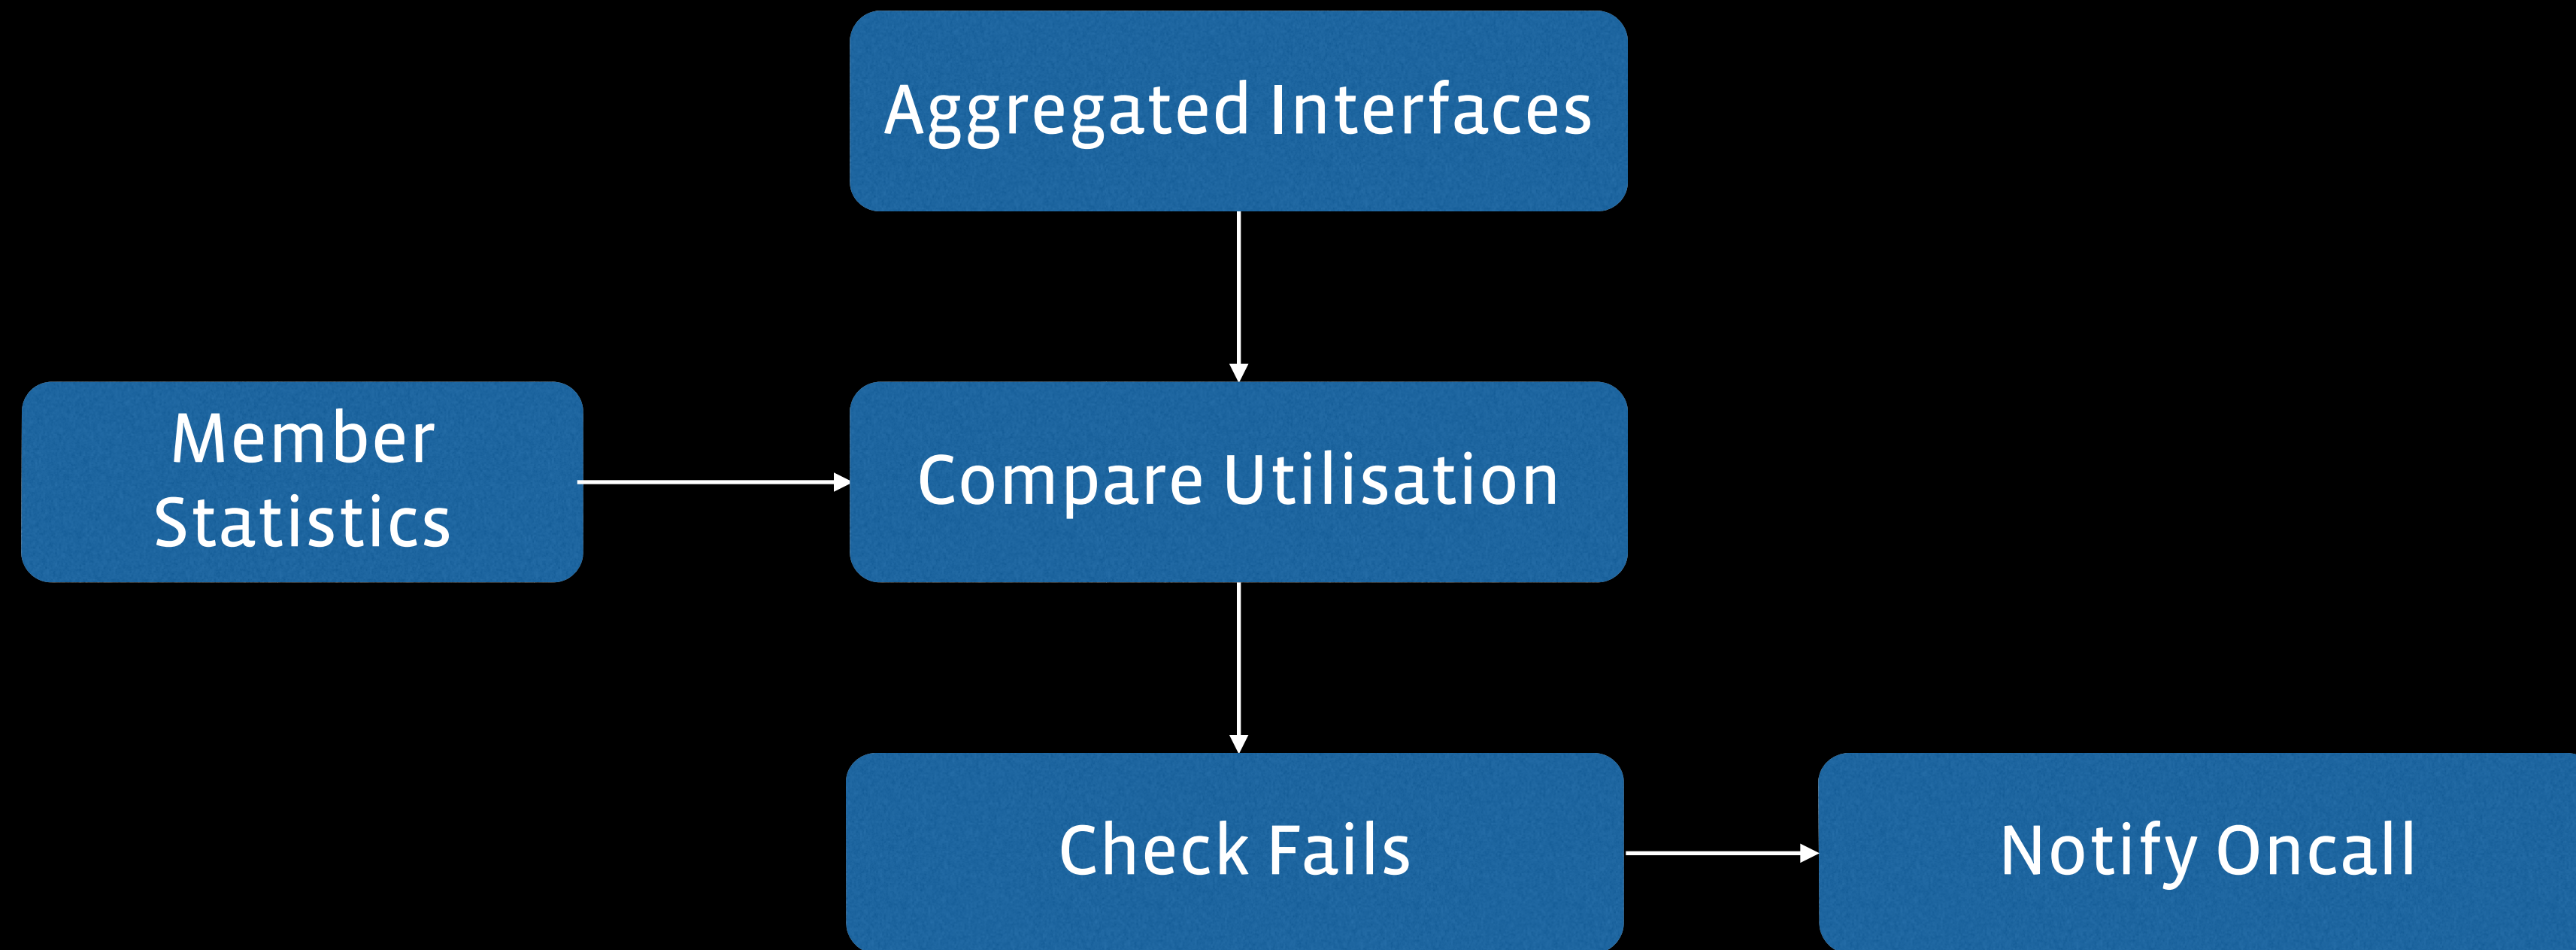
- Intercluster links only
- >80%+ Util
- Migratory

# What Now?

- Detection
- Mitigation




# Imbalance Detection



# Imbalance Mitigation


SIP + DIP + SPORT + DPORT + Hash Key

# Imbalance Mitigation







# Imbalance Mitigation




# Link Imbalance



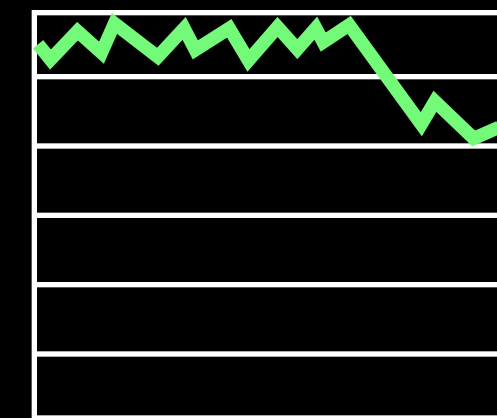
Detect



FBAR




Roll Hash



Resolved





SUIT GAS DIVERTER  
PUSH-CAM  
PULL-ARM

VENTILATION CONTROL

VENTILATION CONTROL

VENTILATION CONTROL

VENTILATION CONTROL

VENTILATION CONTROL

VENTILATION CONTROL

VENTILATION CONTROL

VENTILATION CONTROL

VENTILATION CONTROL

VENTILATION CONTROL

VENTILATION CONTROL



Link Imbalance

ZOMG!! WTF!! NO, NO,  
NO! &(#\$&\*()



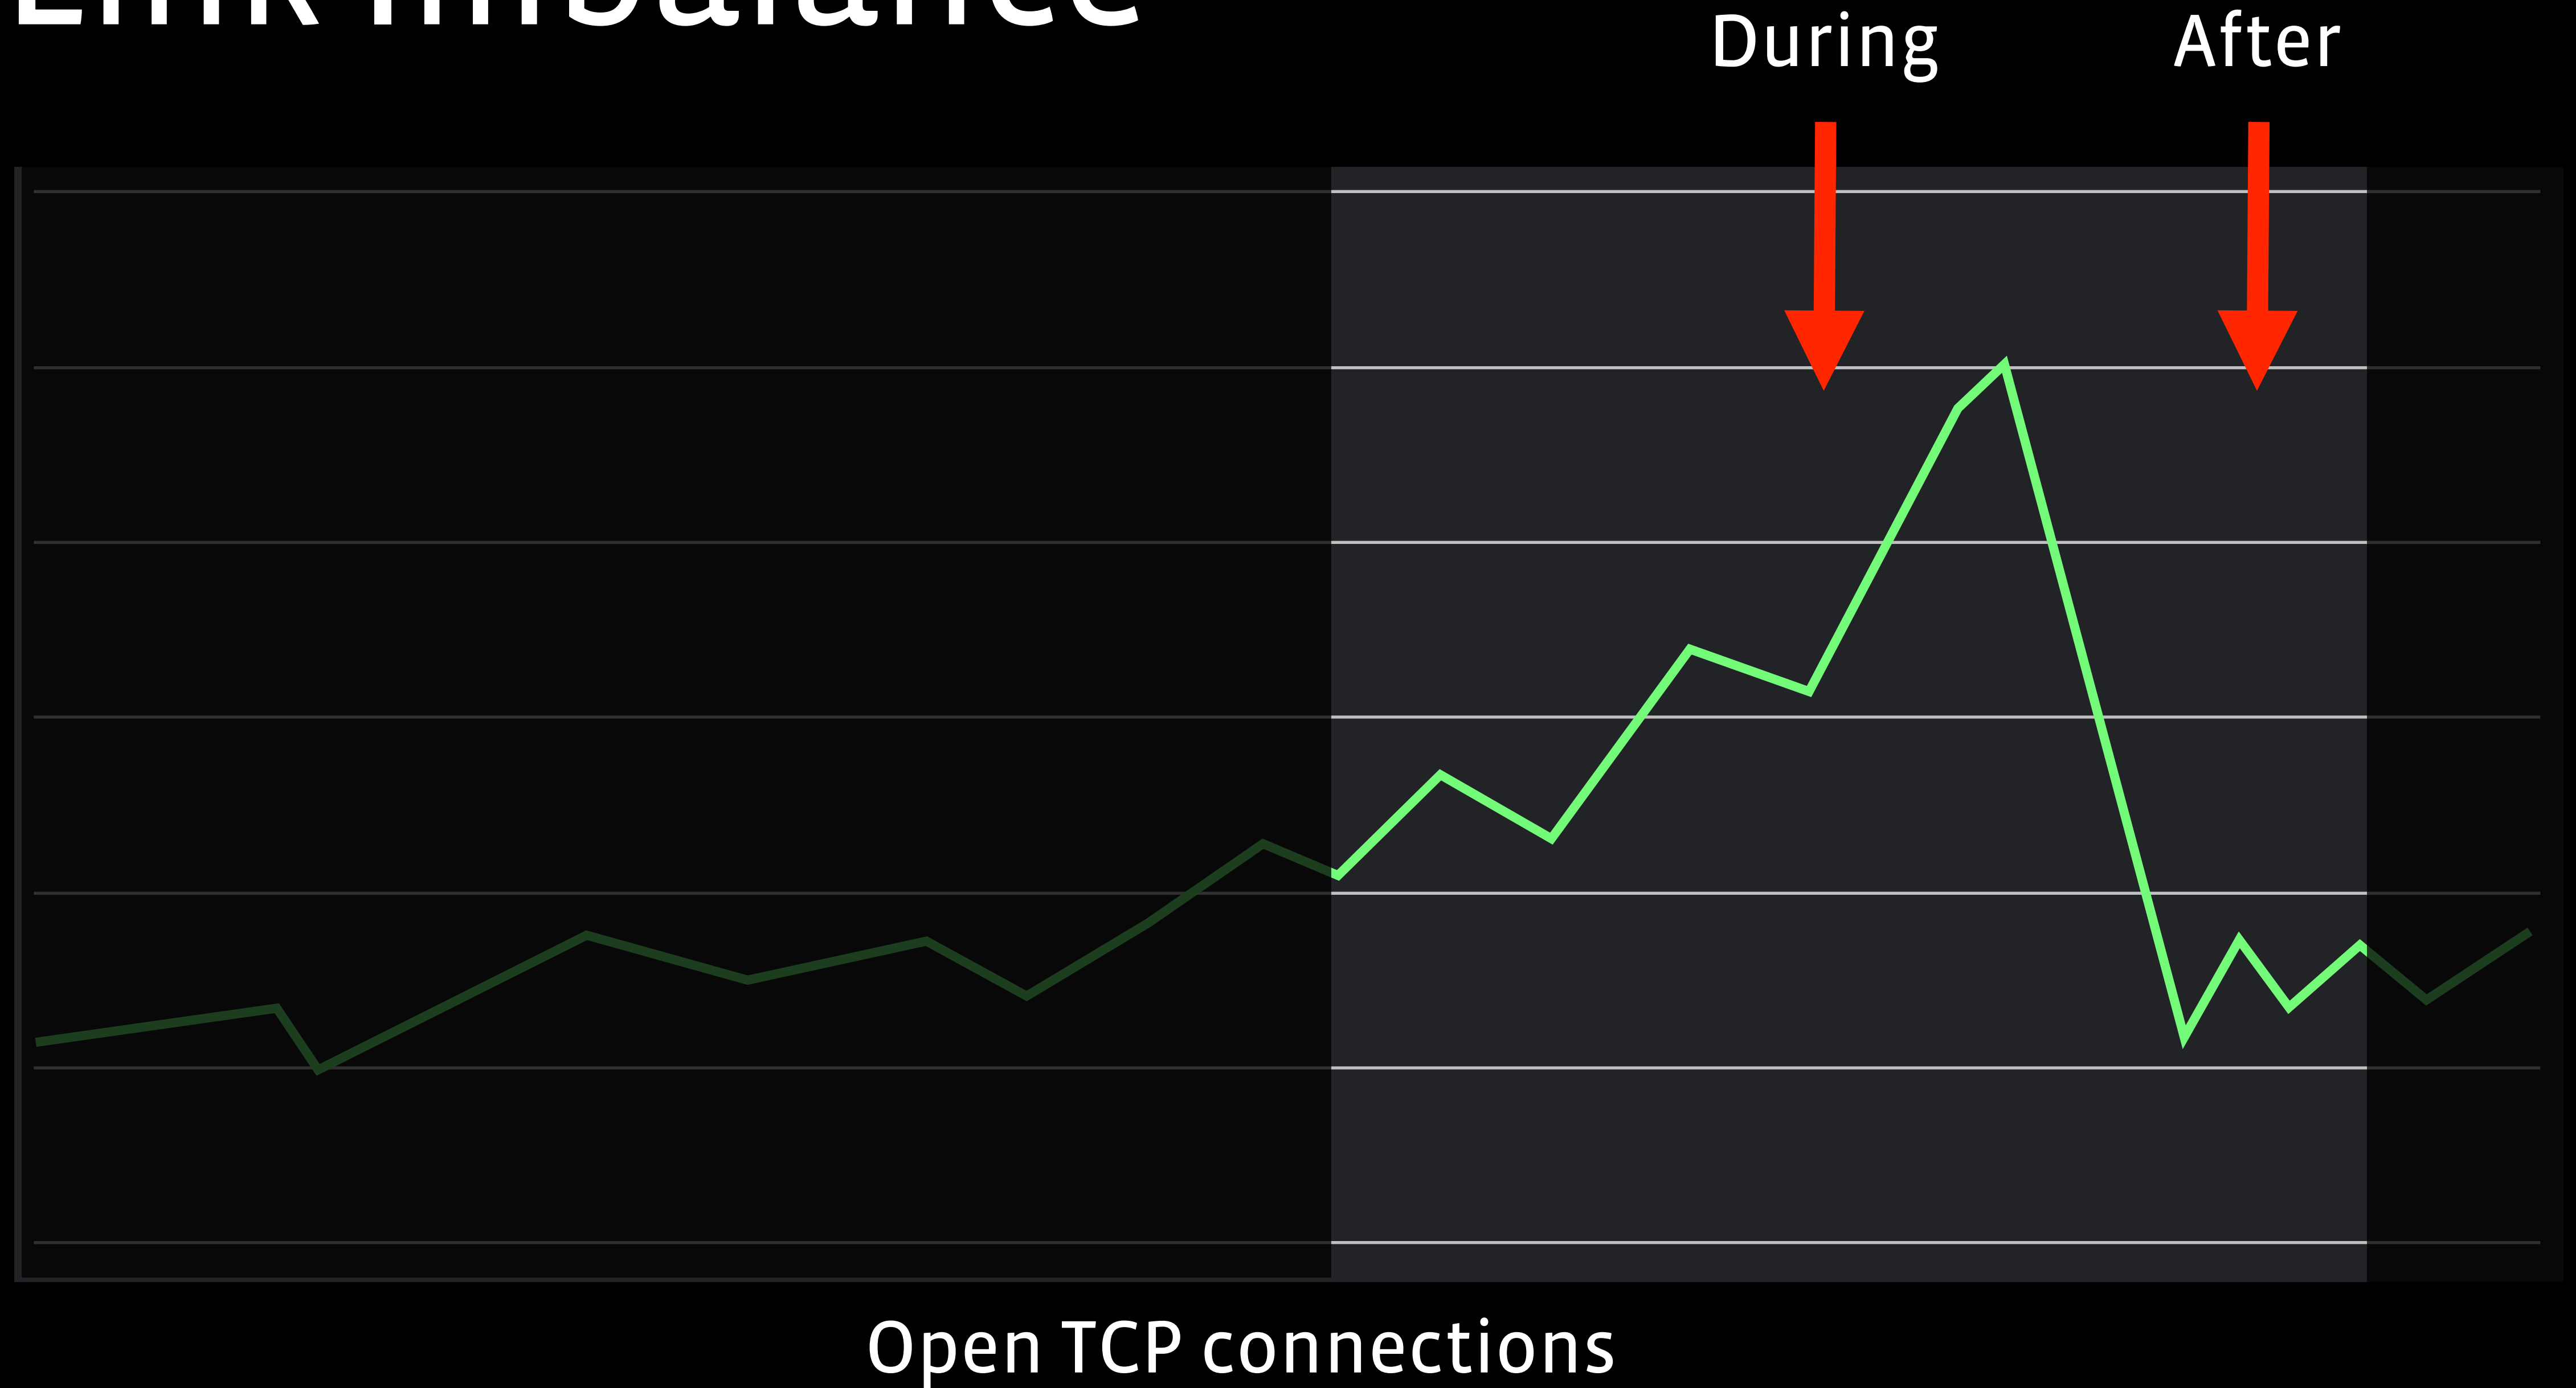
# Link Imbalance

During





RX Window

# Link Imbalance



# It Isn't (always) The Network






Cache


DB





Cache

DB



Cache


DB

# Link Imbalance - Lessons Learned


- Resolved issues
- Root Cause
- Software helps
- Service owner identified
- Resolution time
- Small loss, significant impact





# 3 Months In A Leaky Boat



# Memory Issues

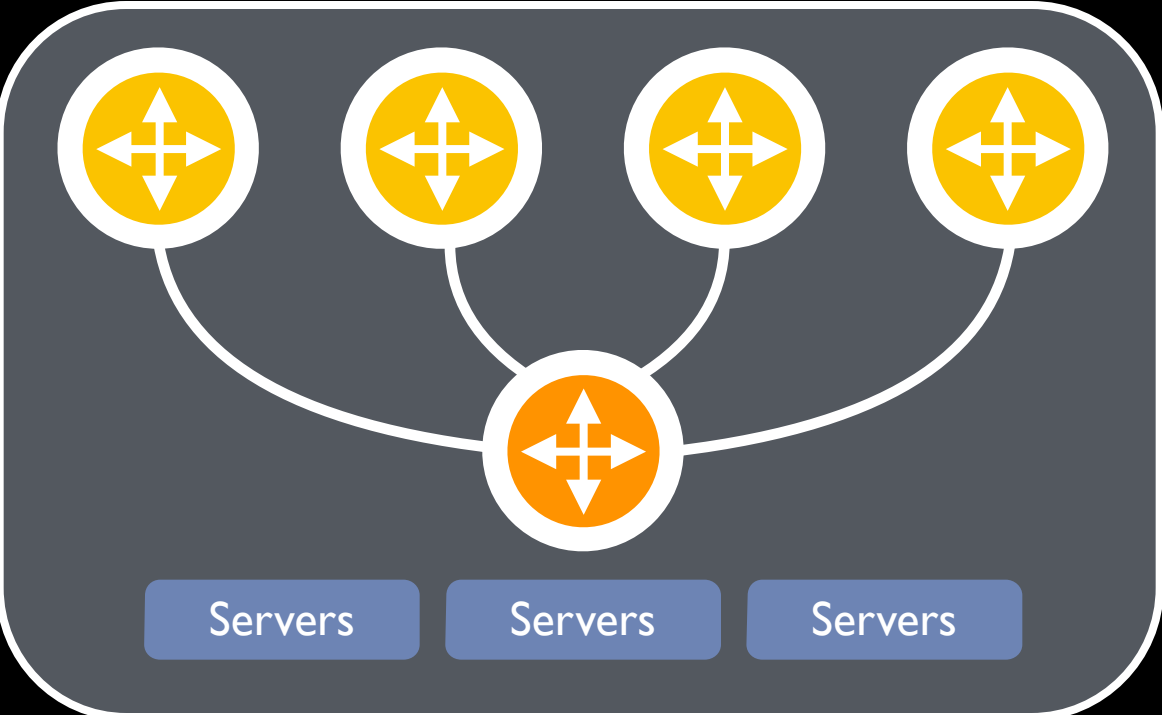




- Capacity
- Health
- Errors










# Detection




DIP	Loss	Latency
1.1.1.1	0.1	10
2.2.2.2	0	10




# Loss Effects on Throughput



# Different algos?



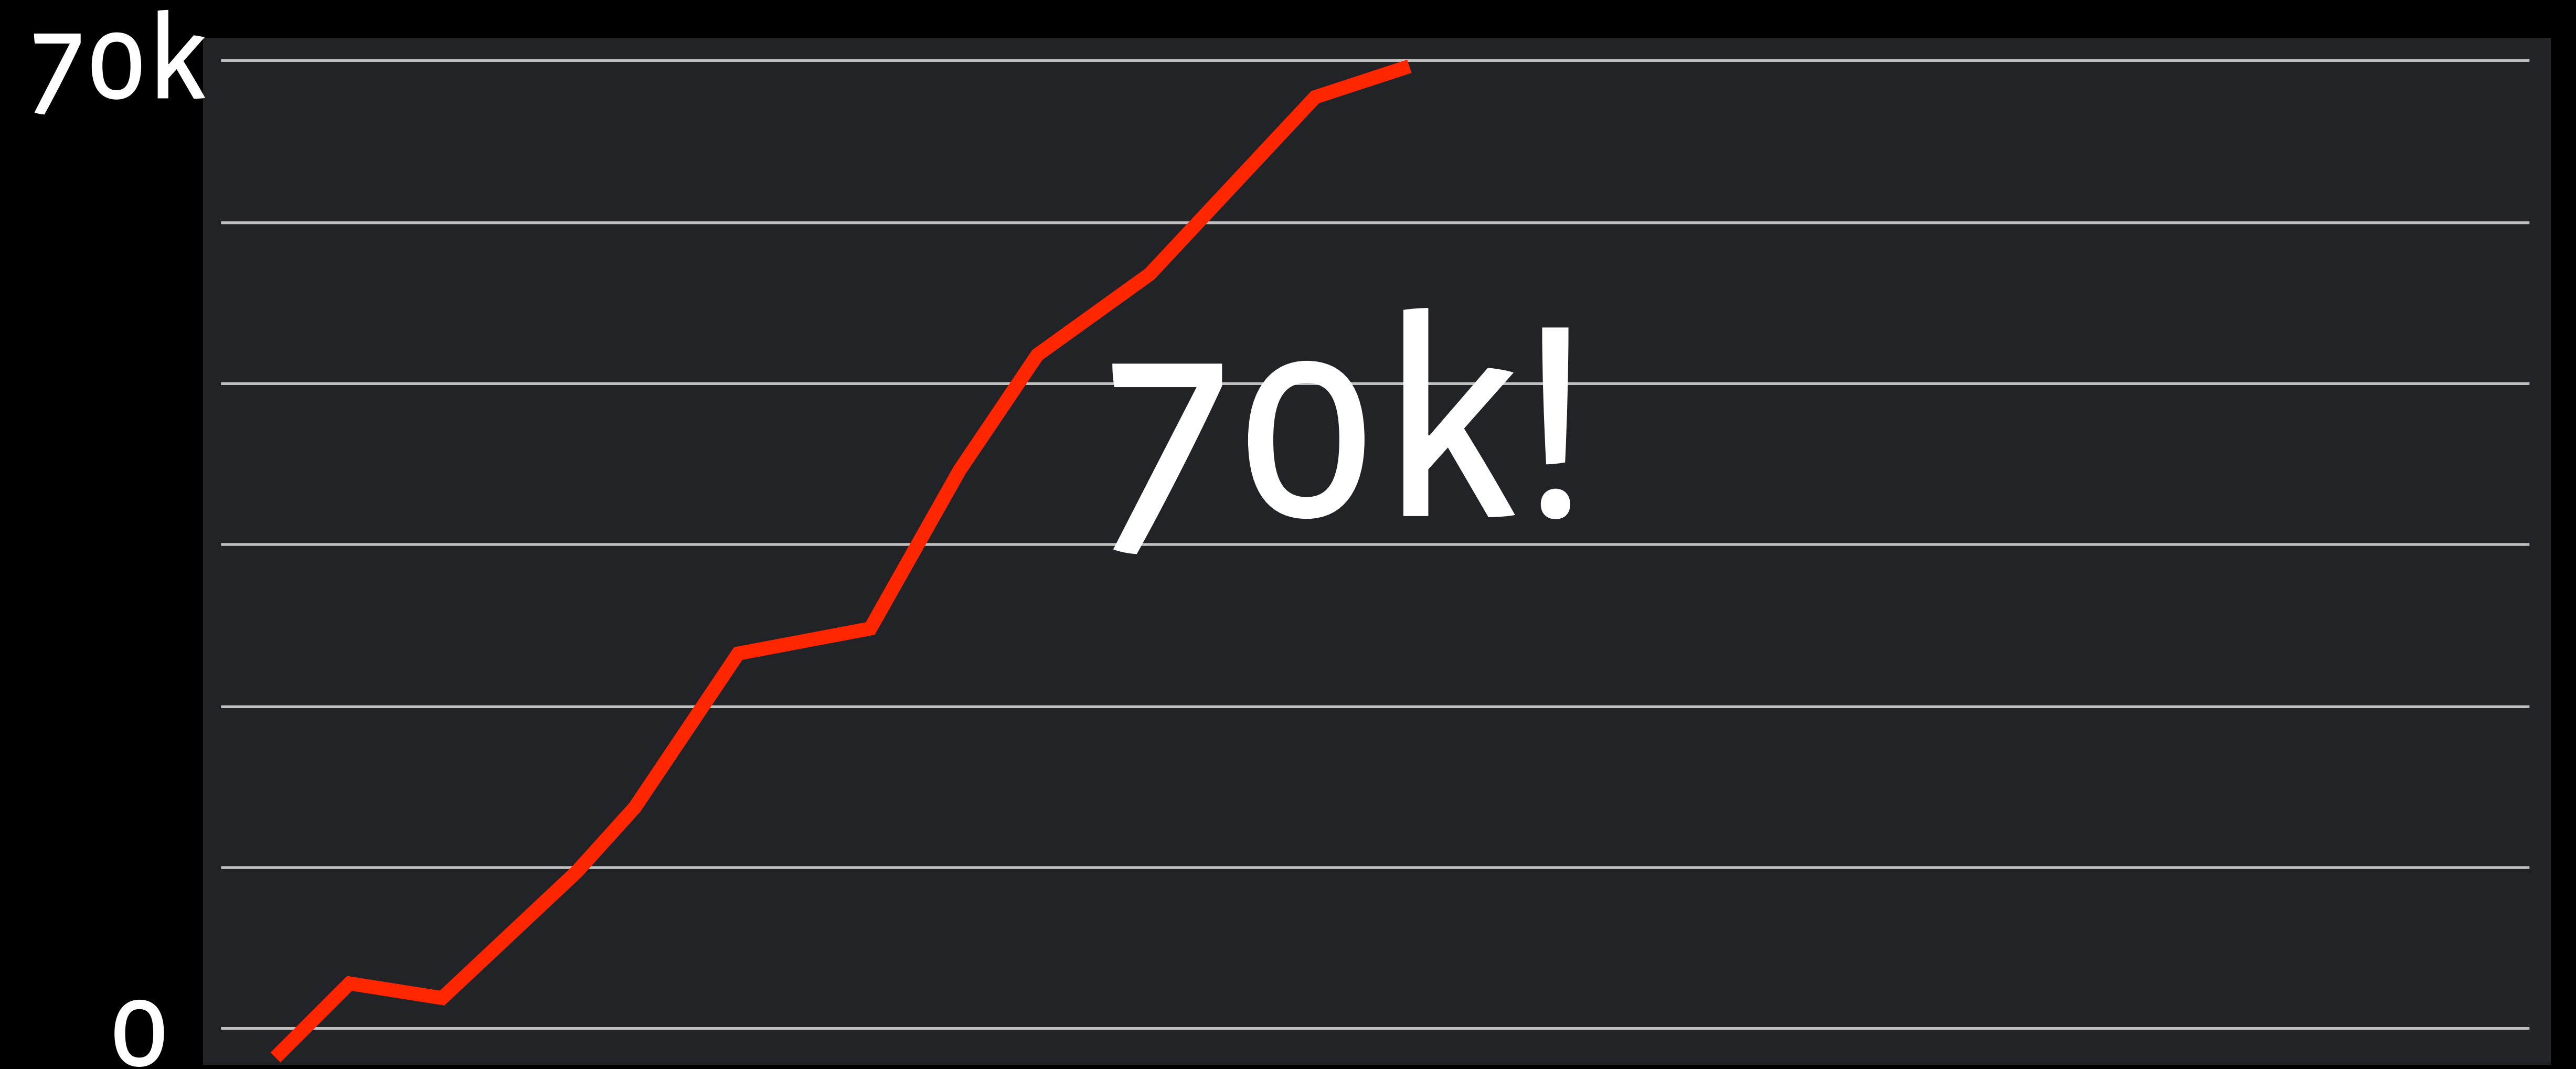
# Recovery time




So, wait, how does this apply to me?





# Alarms



# Interface Issues



- Capacity
- Health
- Errors



# Alarms Now



< 100

-99.999%

# Automation - One Month

3.37b

0.999%

750k

99.6%

Why?

Sleep

# Why?

$$750,000 * 2 = 1,500,000$$

$$1,500,000 / 60 = 25,000$$

# 150+

$$25,000 / 160$$



# Lessons Learned & Take Aways


Use what you've got

The sooner the robots  
take over the better


Prototype early,  
iterate often

Spend the time to  
root cause issues

Duct tape keeps  
things running







Why?



**DONE IS  
BETTER  
THAN  
PERFECT**

?