

Google SRE: Chasing Uptime

What do Google clusters look like?

How do we manage them?

Site Reliability Engineering


- Manage Google's serving infrastructure
- Plan and execute new capacity deployment (e.g. new datacenters)
- Performance tuning
- Handle serious system problems and outages

Commodity Hardware


- IDE Drives, Midrange CPUs, non-redundant power supplies
- Cheap and readily available parts
- Outstanding bang for the buck



Google, circa 1996



Much improved, Google rack, but...



...tangled wires, cardboard, cork, bent motherboards!

Commodity Hardware

- IDE Drives, Midrange CPUs, non-redundant power supplies
- Cheap and readily available parts
- Outstanding bang for the buck
- Unreliable, temperamental, flaky

Site Reliability Engineering

- Automate common failure cases: Disk, memory, CPU errors, misconfiguration
- ...and the dreaded “unexplained server down”
- Deploy and maintain monitoring and automation infrastructure

Strength in numbers: shard

- Dataset is huge, but divided up among many machines, giving us:
- Subset of data fits on one machine
- Splitting processing gives lower latency
- More CPUs gives higher throughput

Site Reliability Engineering

- How is the decision to divide up data among machines made?
- Correct for uneven workload and changes in query mix
- Design, deploy, and maintain automation

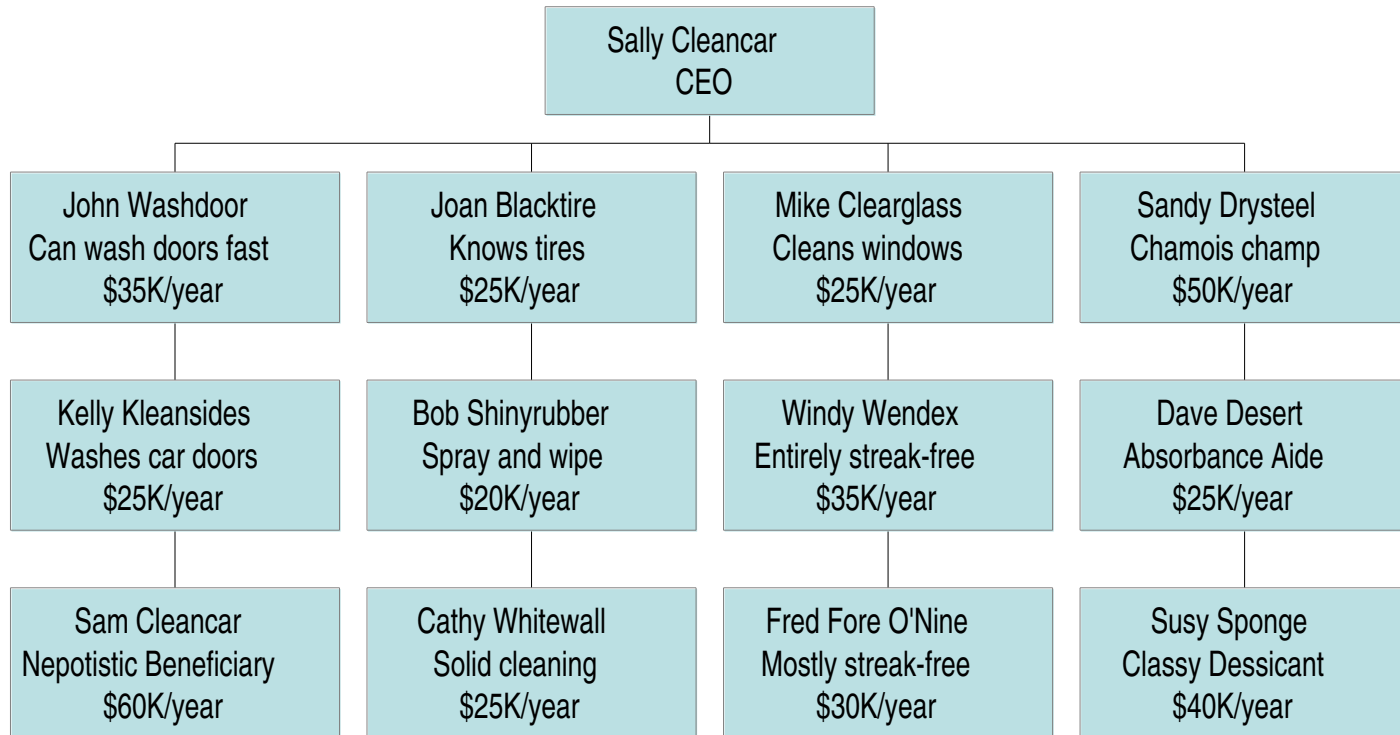
Strength in numbers: clone

- Each server has a number of clones that serve the same subset of data
- Many queries to be processed in parallel, giving scalability
- If any one clone goes down, others pick up, giving reliability

Sally's carwash



Sally's carwash



Employee cost: \$395K/year

Site Reliability Engineering

- Diagnose and fix performance problems on live serving systems
- Plan and execute deployment of new capacity
- Work with new projects deployments to ensure they meet production criteria
- ...and much more!

Questions?

A few starters:

- How can I learn more?
- Is Google hiring?
- Contacts:

Angus Lees alees@google.com

Eric Pollmann pollmann@google.com

Michelle Lo mlo@google.com

Ravi Pindiprolu rpindiprolu@google.com