



Designing for Failure:

How to Manage Thousands of Hosts Through Automation

Brandon Bercovich

October 29, 2018

Uber

We will be digging into three problem areas:

1. Utilization
2. Fleet management
3. Host management

We will be limiting the scope of this talk to stateless services.

Utilization

Our first try at sharing hosts:

No containerization and manual placement.



Why fix two problem when you can fix one?

Containerization with manual placement.




Automatic placement


Failed hosts no longer impact services.




Scheduling a service



Host failure



Service is rescheduled









kubernetes



MESOS




openstack®




Fleet Management

Puppet


T0



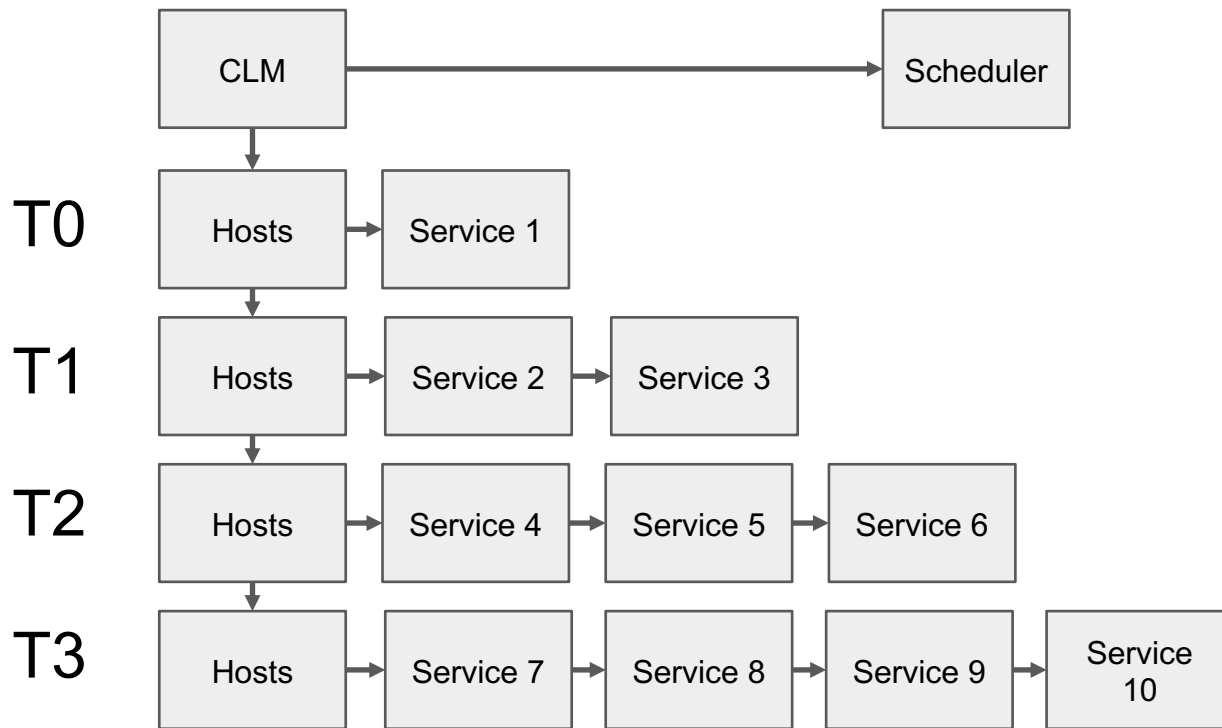
Platform Lifecycle Manager




T3



Cluster Lifecycle Manager

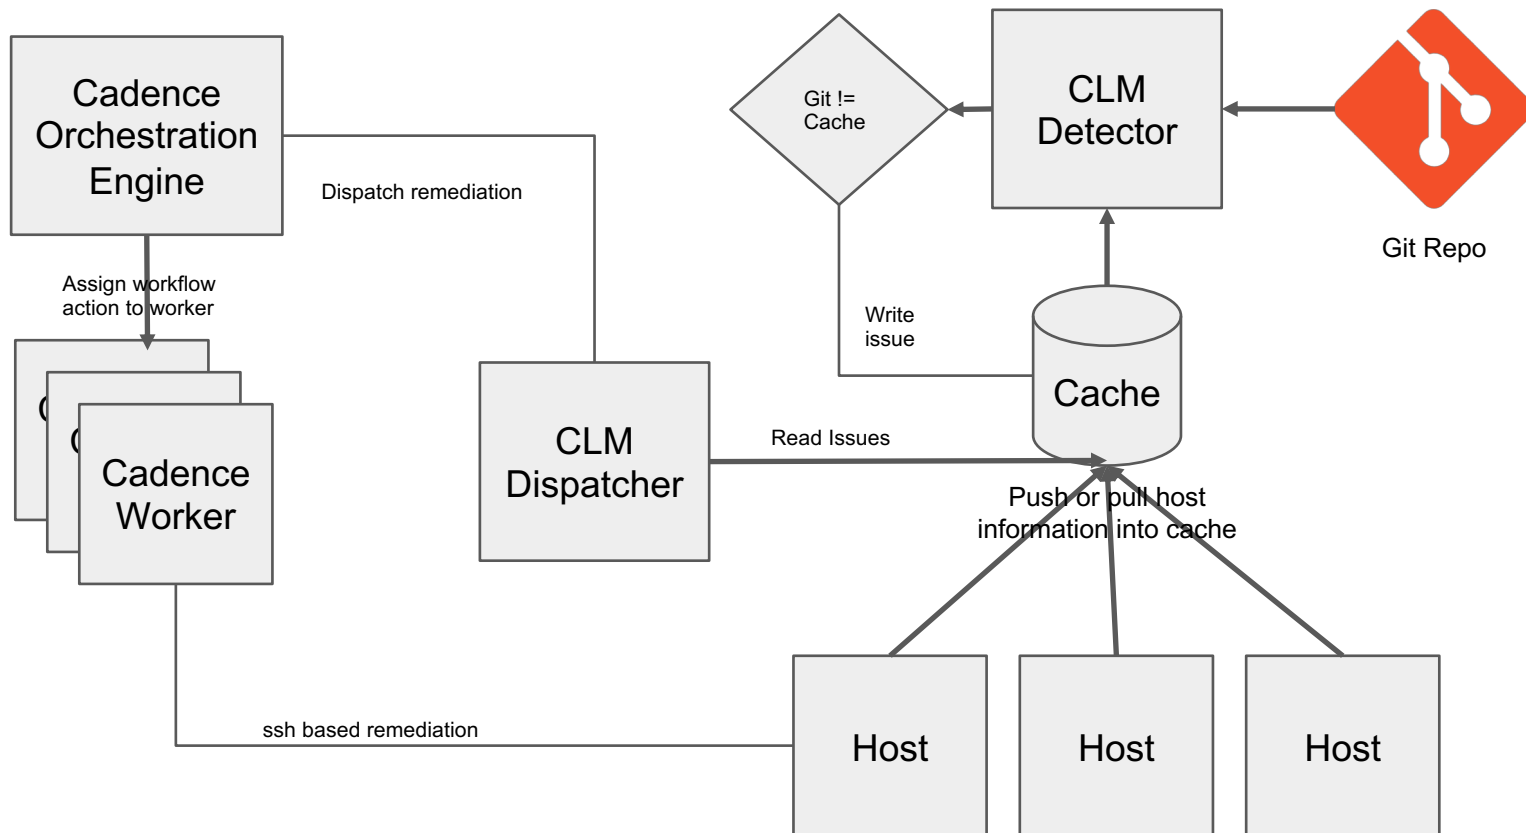


Cluster Lifecycle Manager

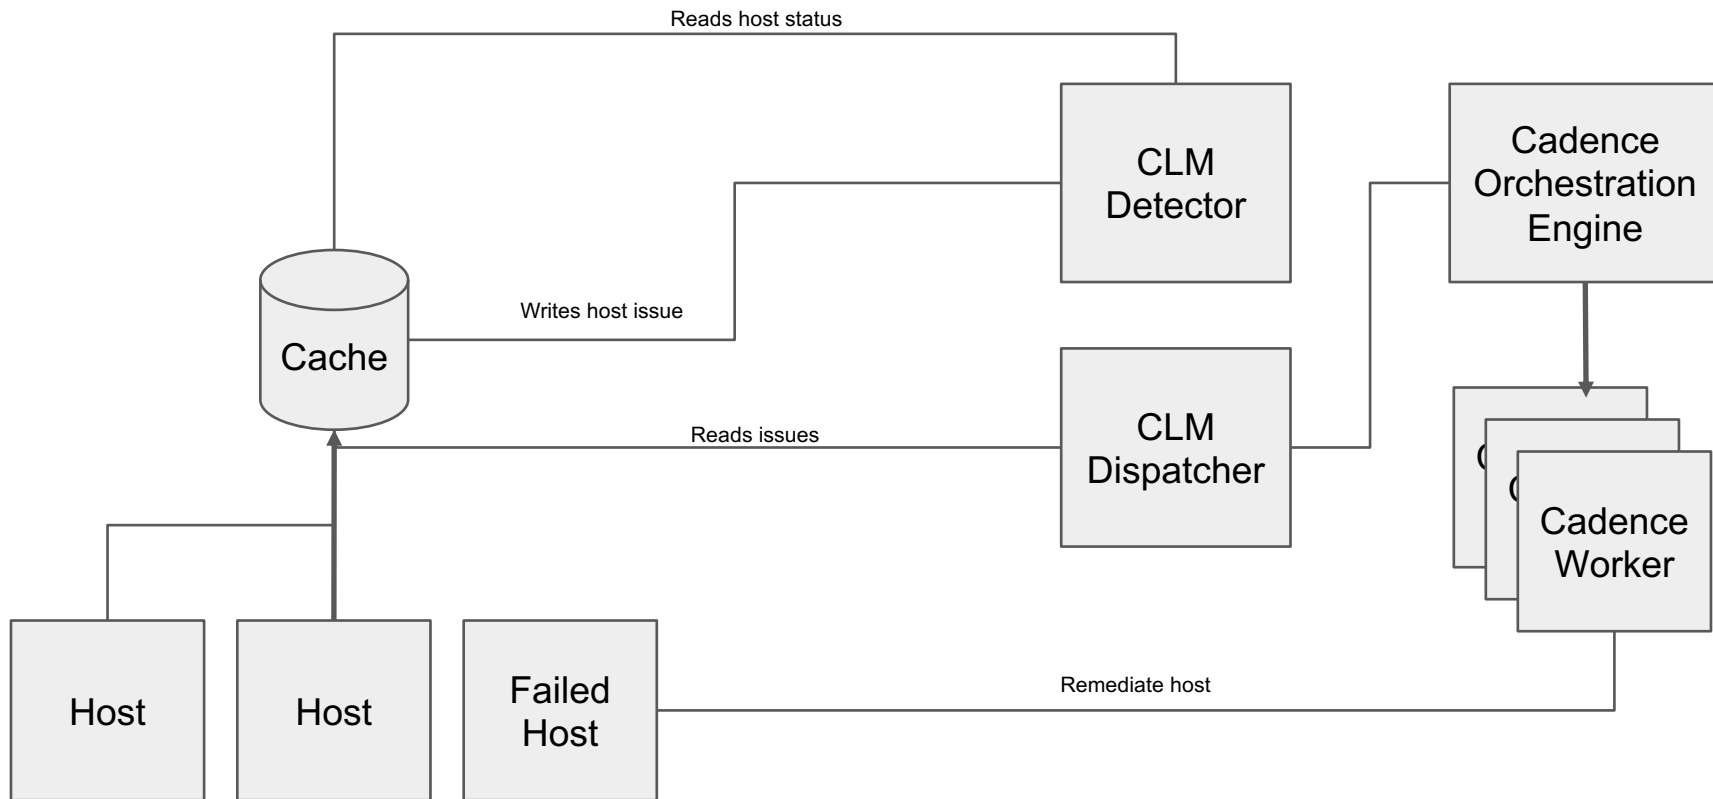


Host Management

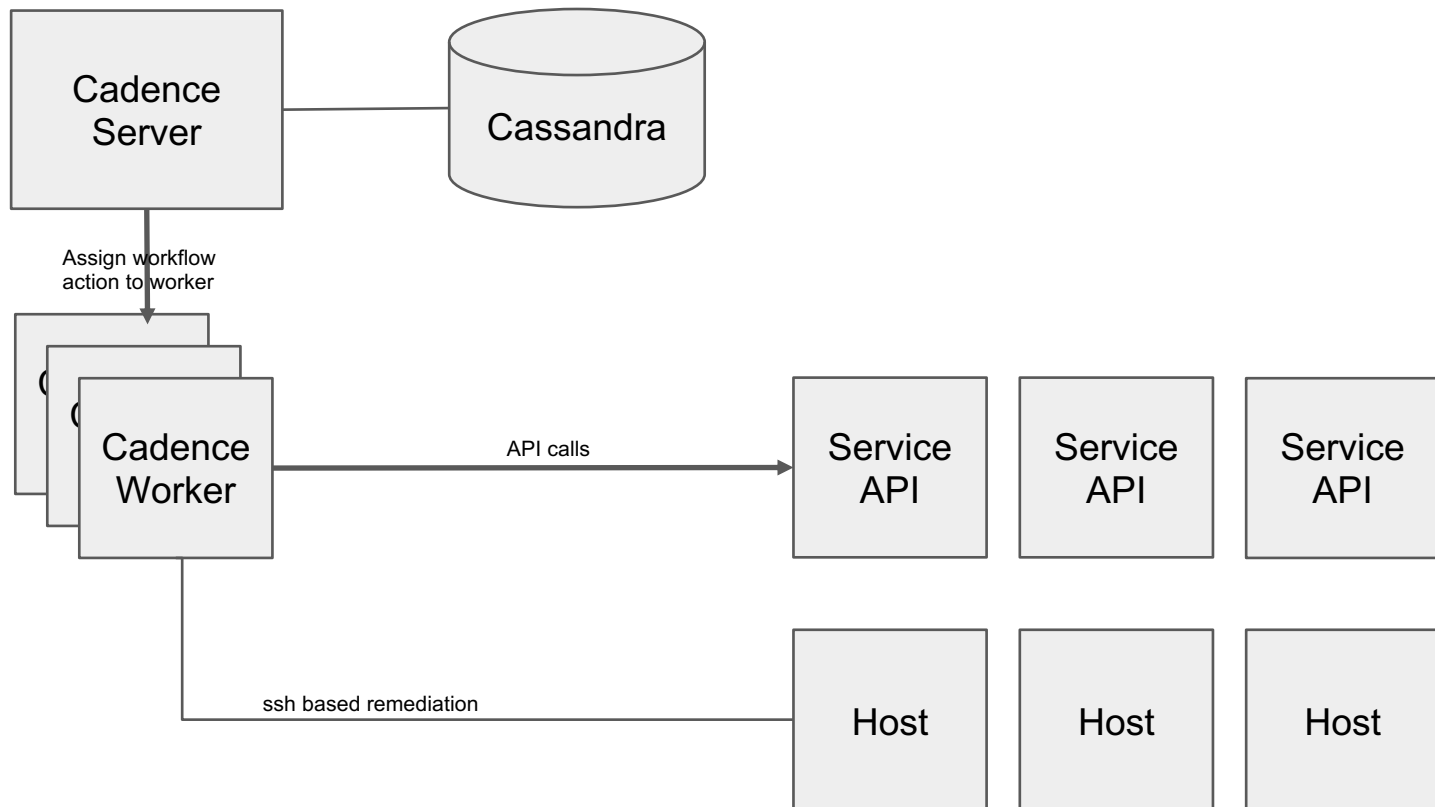
Architecture



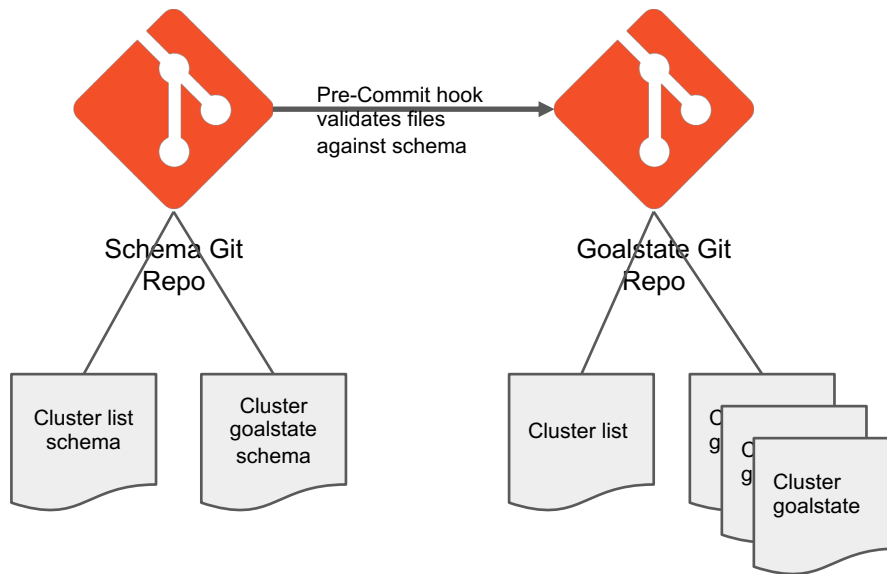
Architecture



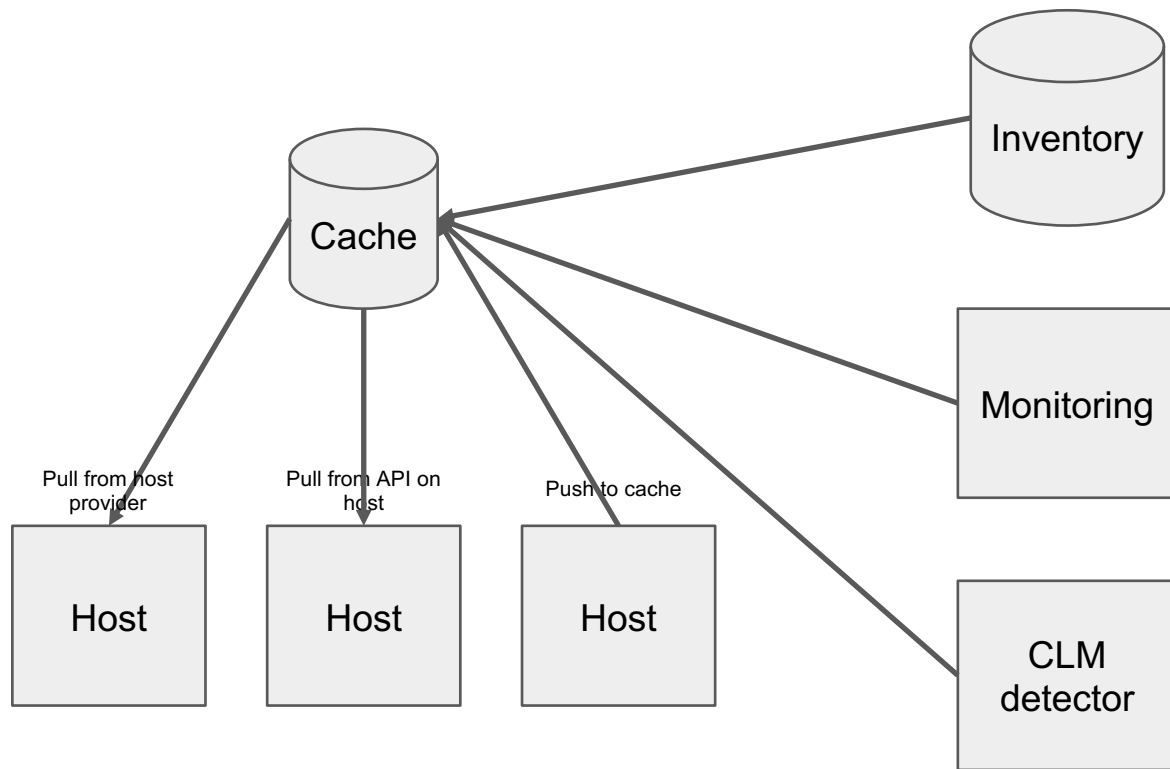
Cadence Orchestration Engine




Store Goalsate



Cache



Tying it all together with CLM



Q & A