



+



Puppet user since early 2009  
~1000 hardware clients  
~250 virtual clients

Nearly 100% production puppetized  
~31,000 lines of puppet + erb  
67 contributors

# Early Tech Community: No staff







**Current Ops Community:  
No new volunteers**

Ops





500,000,000 readers

100,000 editors

650 developers

10-15 operations






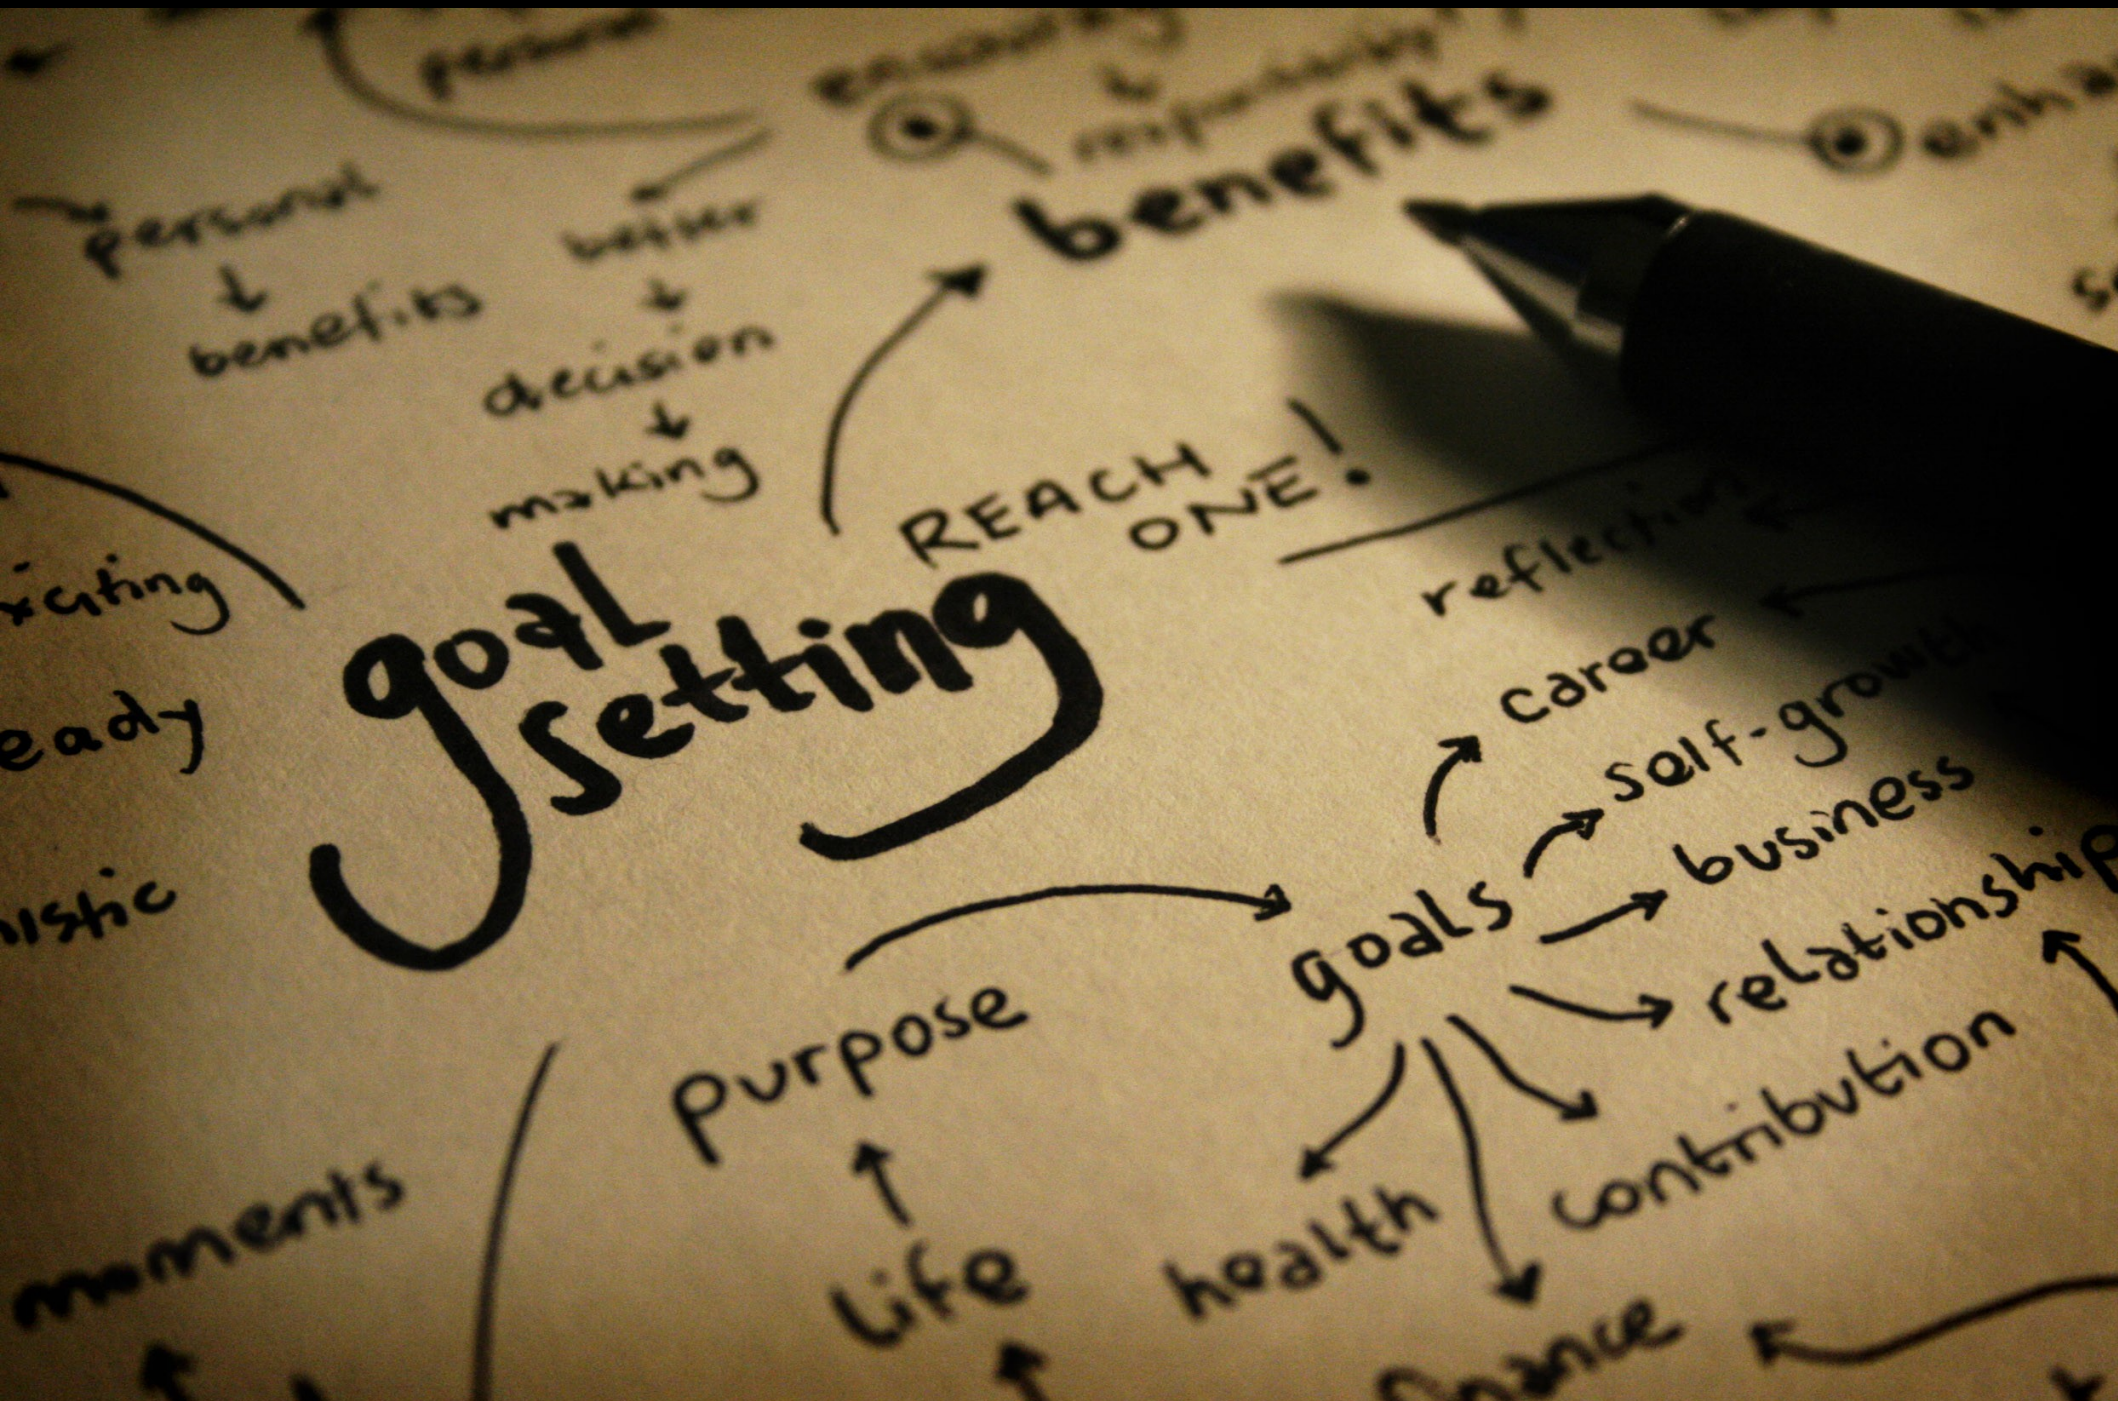
# Open Infrastructure to Volunteers



# WIKIMEDIA LABS

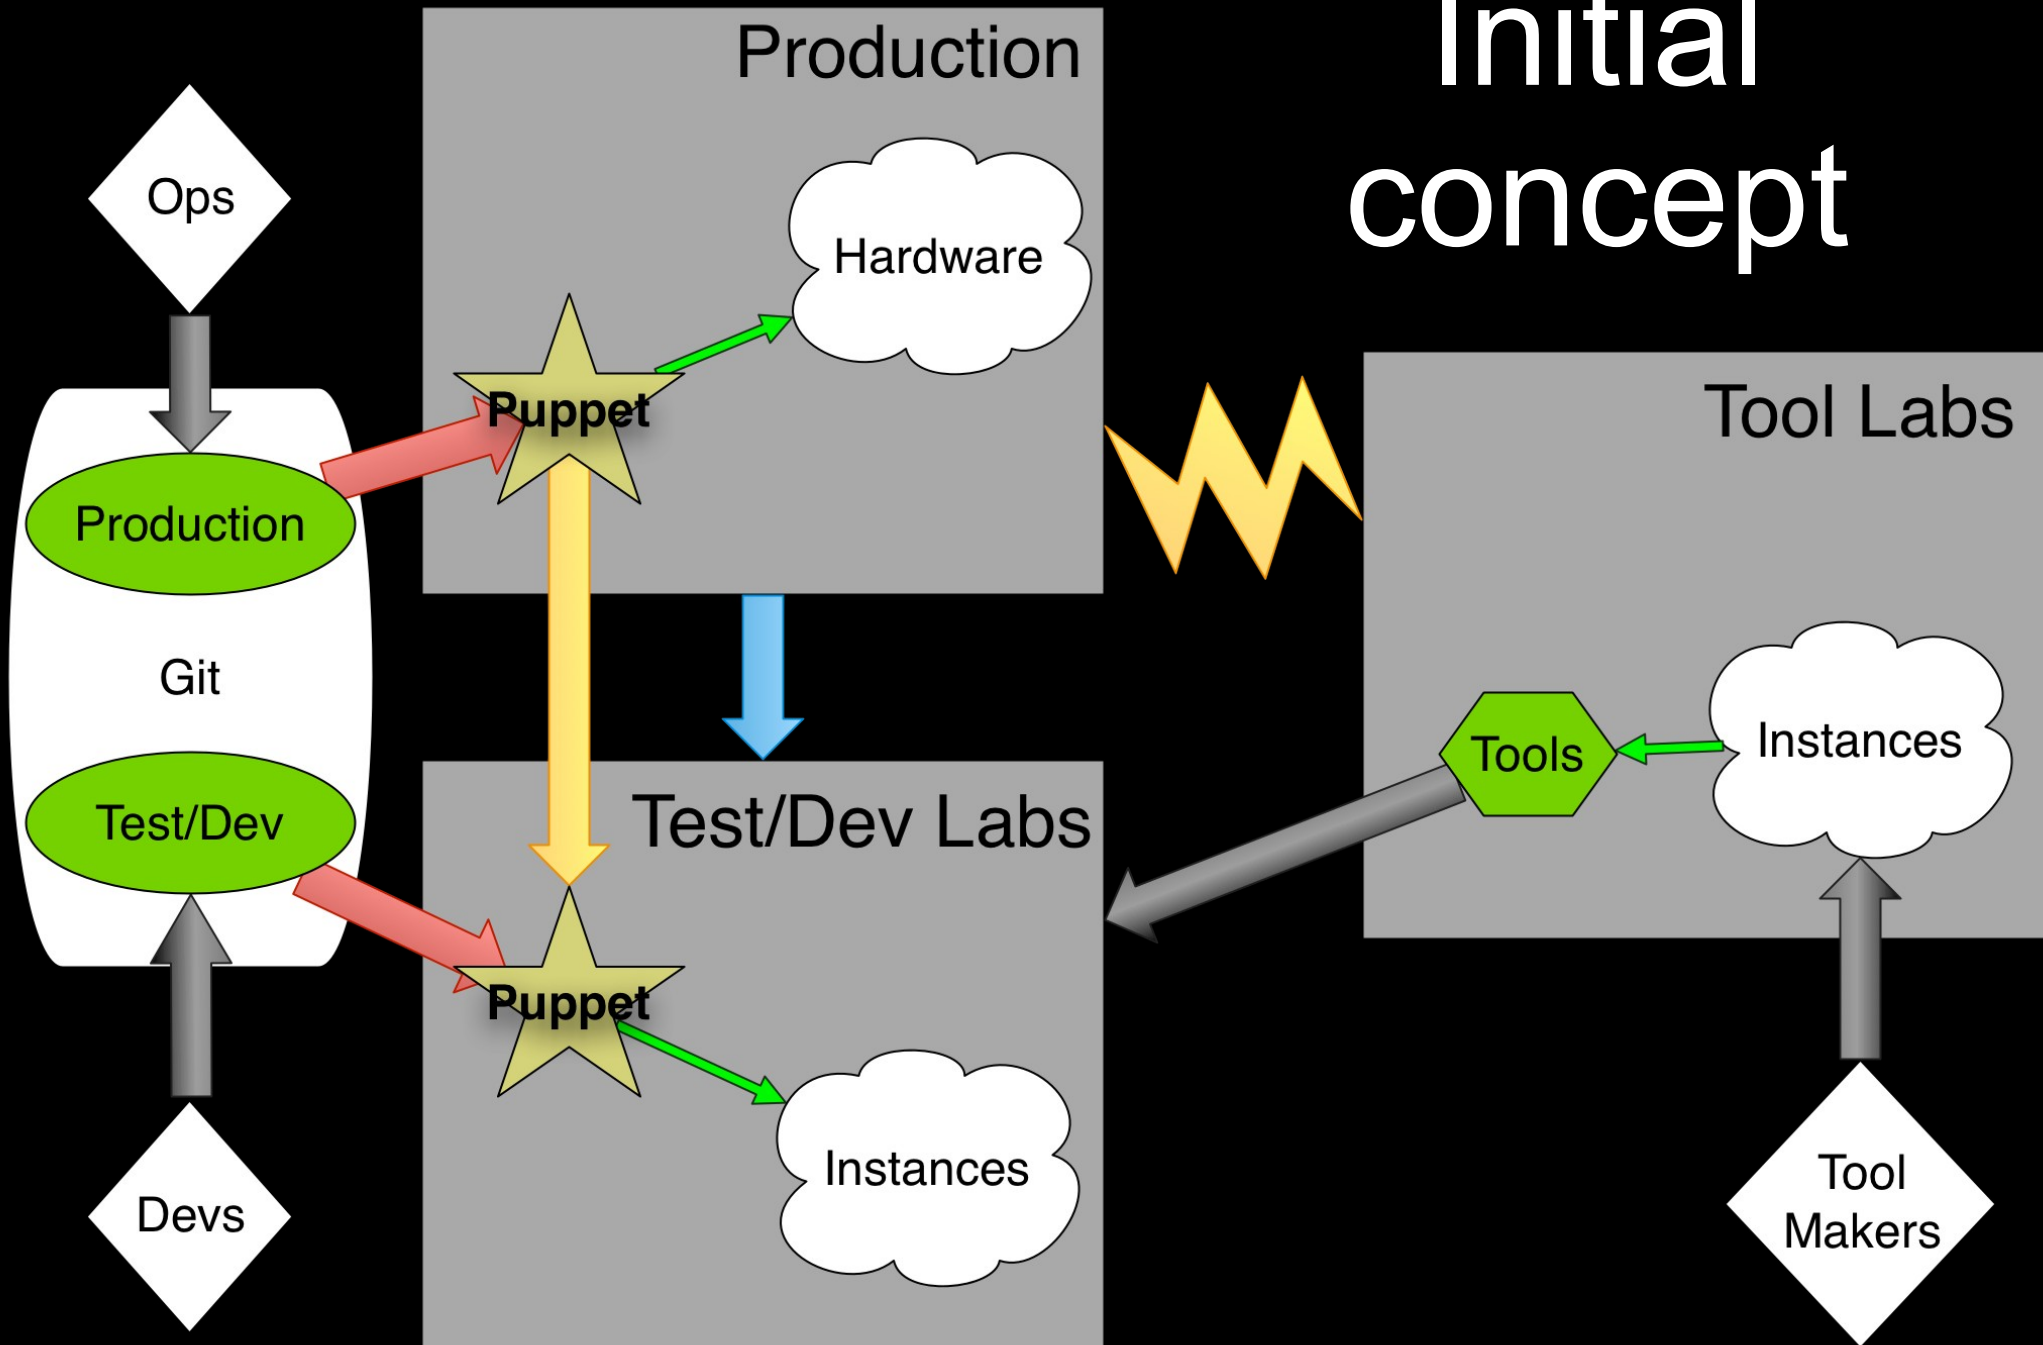


# Goals





# Initial concept





# Release the puppet repo





Release  
the puppet repo





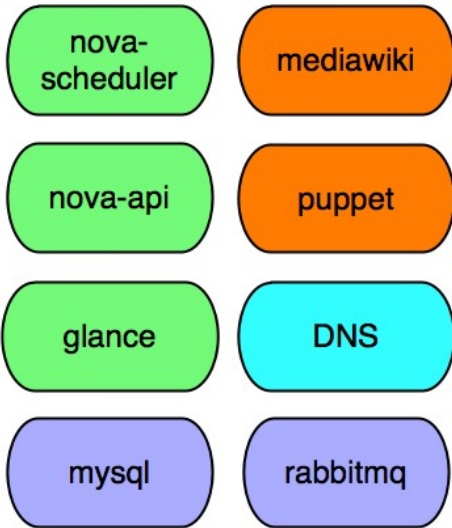
openstack™

CLOUD SOFTWARE

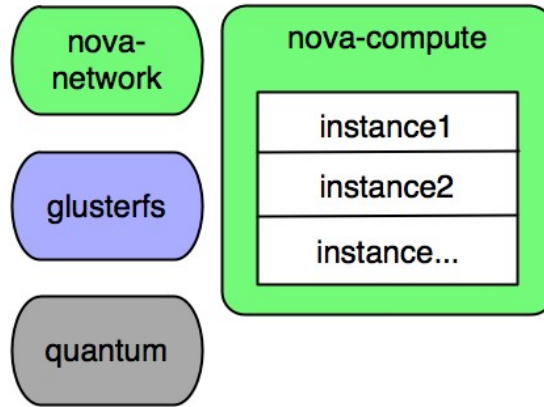
Projects.

Not customer accounts.

### Controller



### Compute nodes



### Analytics clusters



### Replicated and user database cluster




### Test automation



### Repositories/ code review



### Authentication/ authorization



### Per-project shared storage nodes



# Implementation



User creates instance



LDAP entry created



Instance created



User data installs puppet



Puppet master signs request



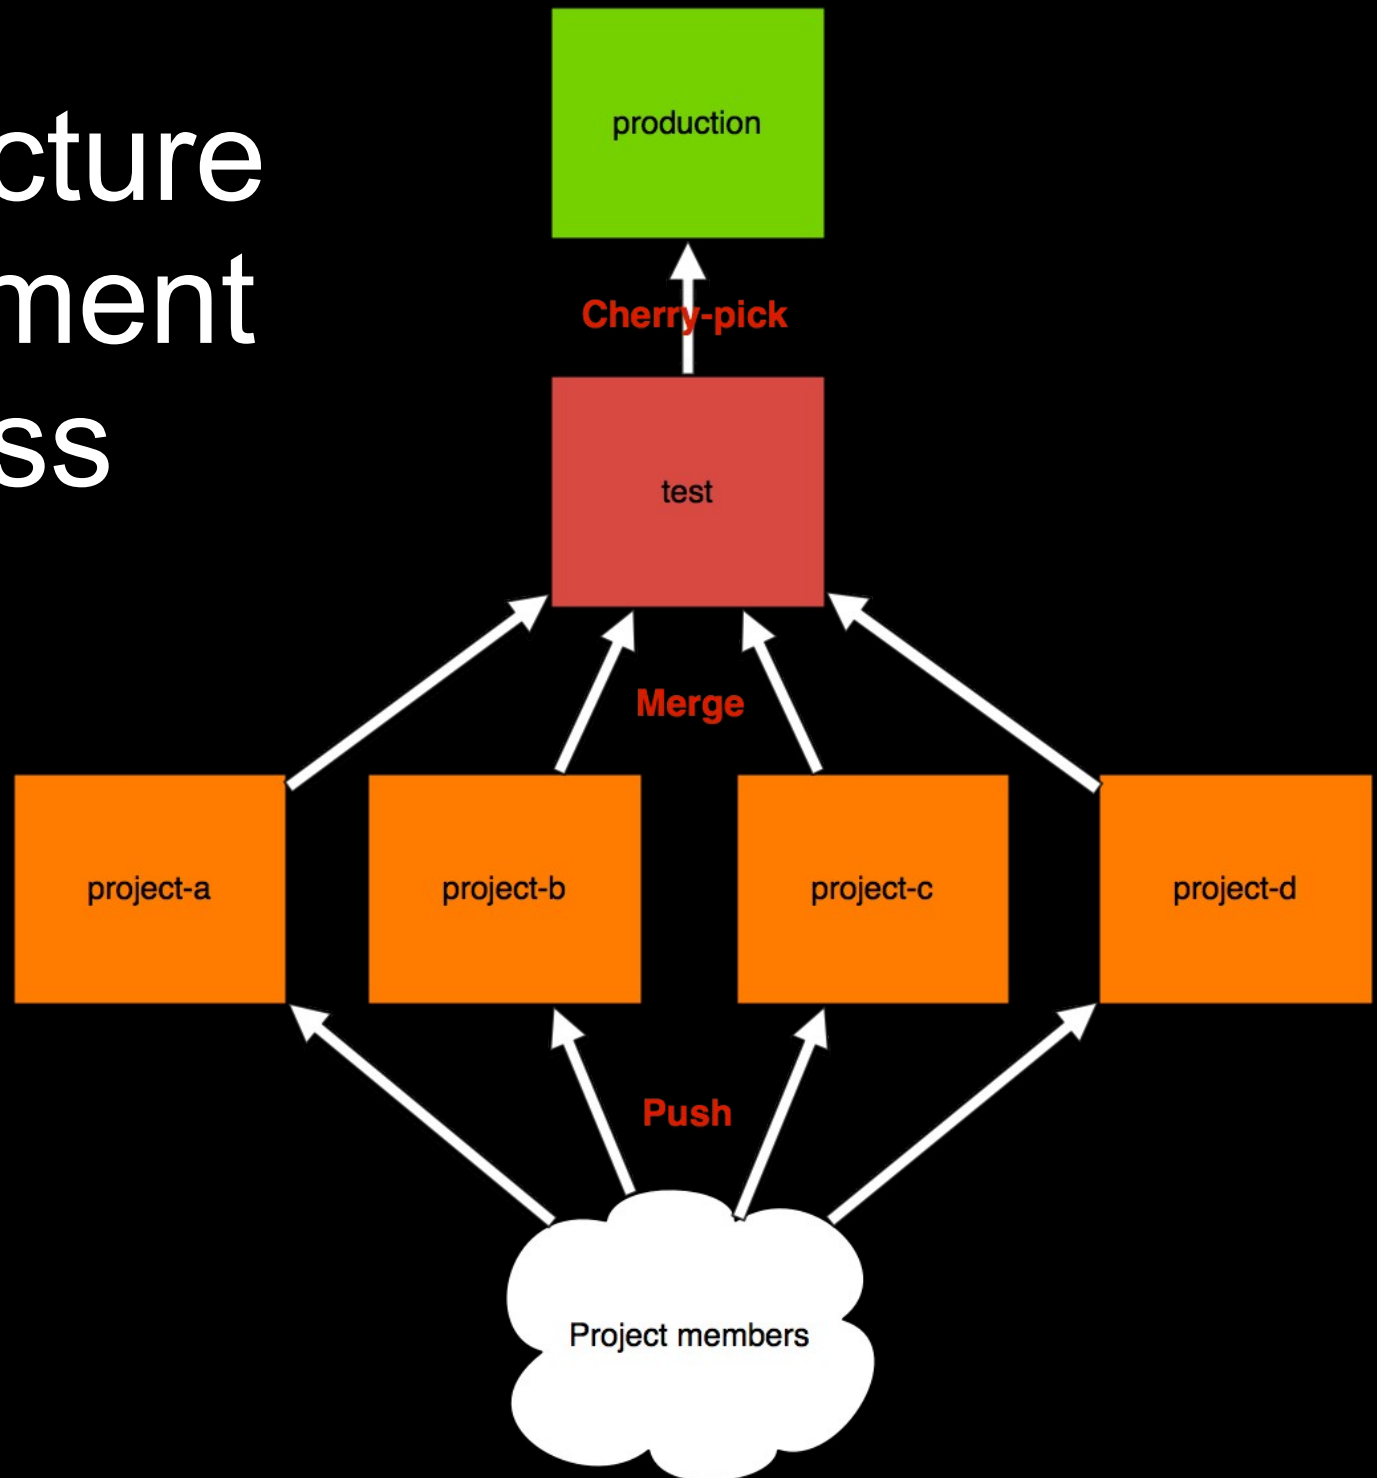
Instance builds

puppet apply



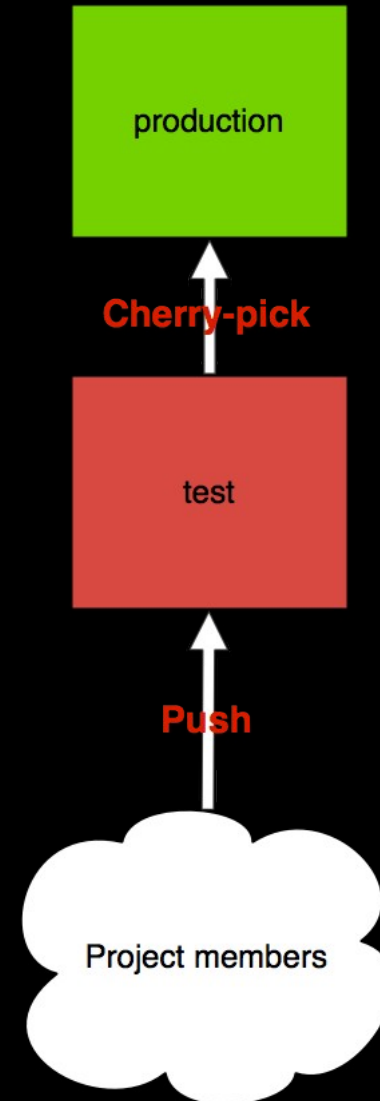
We didn't use modules!

# Infrastructure development process

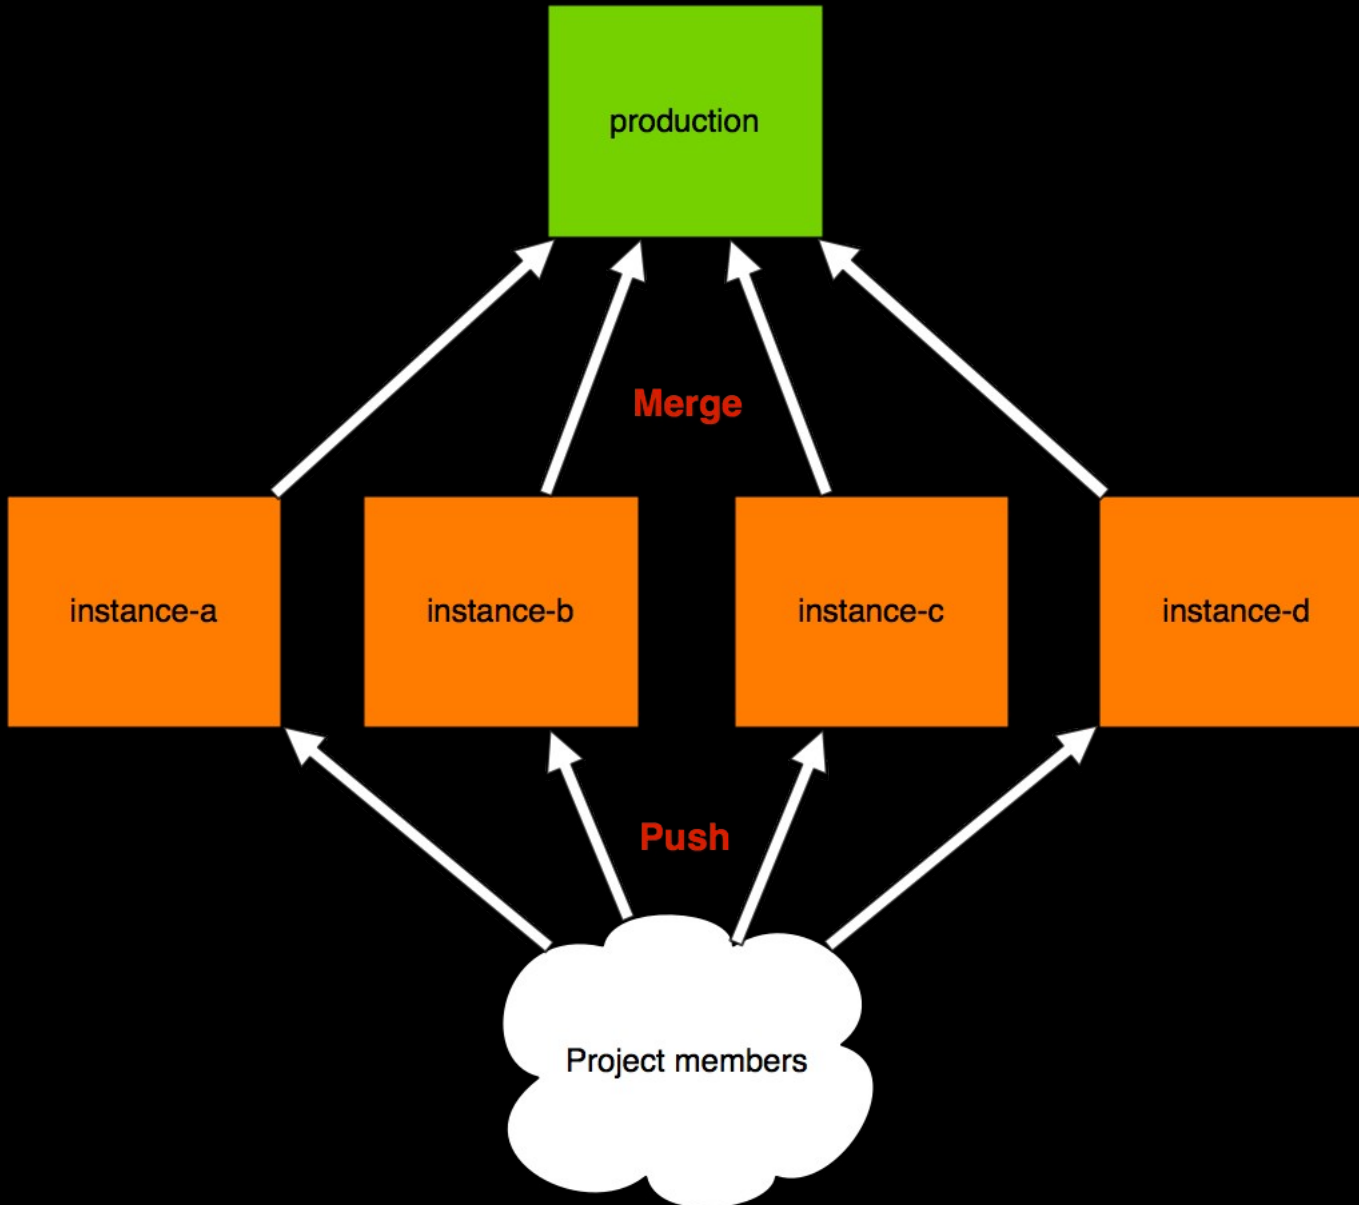


We didn't use modules!

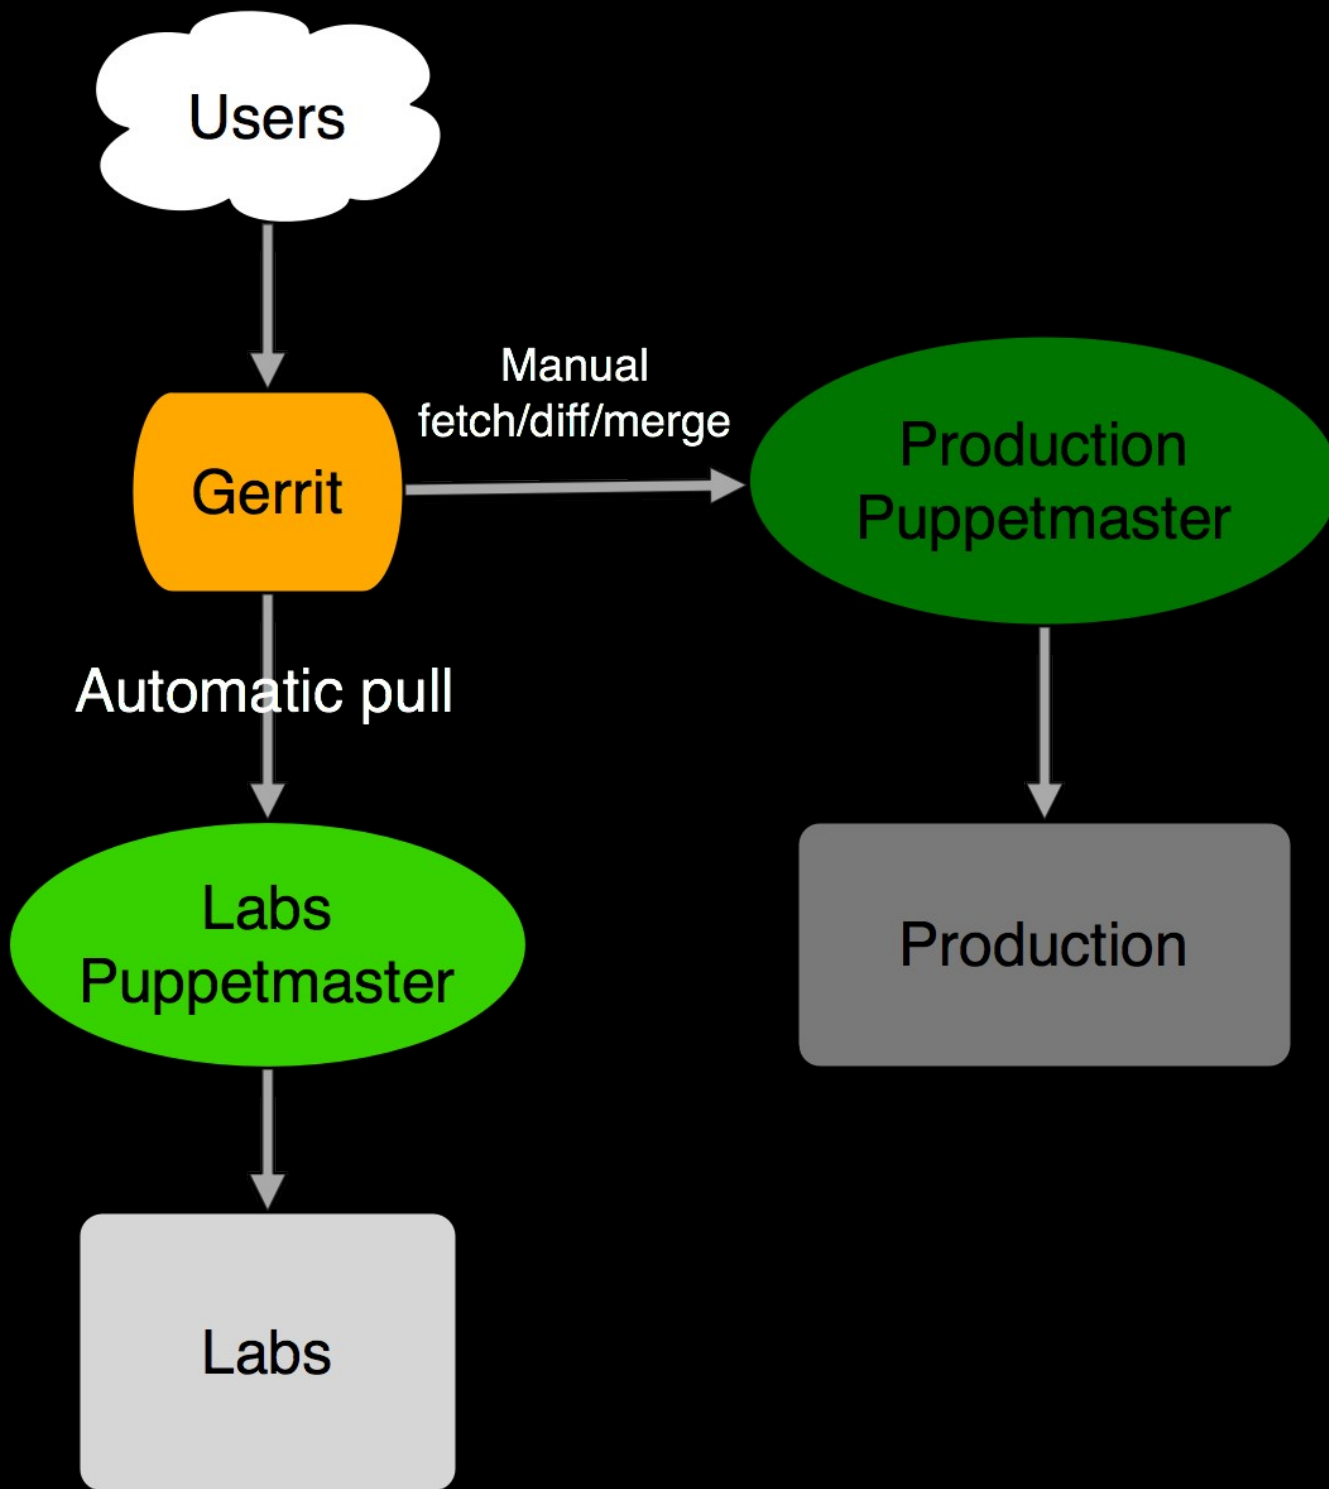
# Run Labs from test branch



# puppetmaster::self



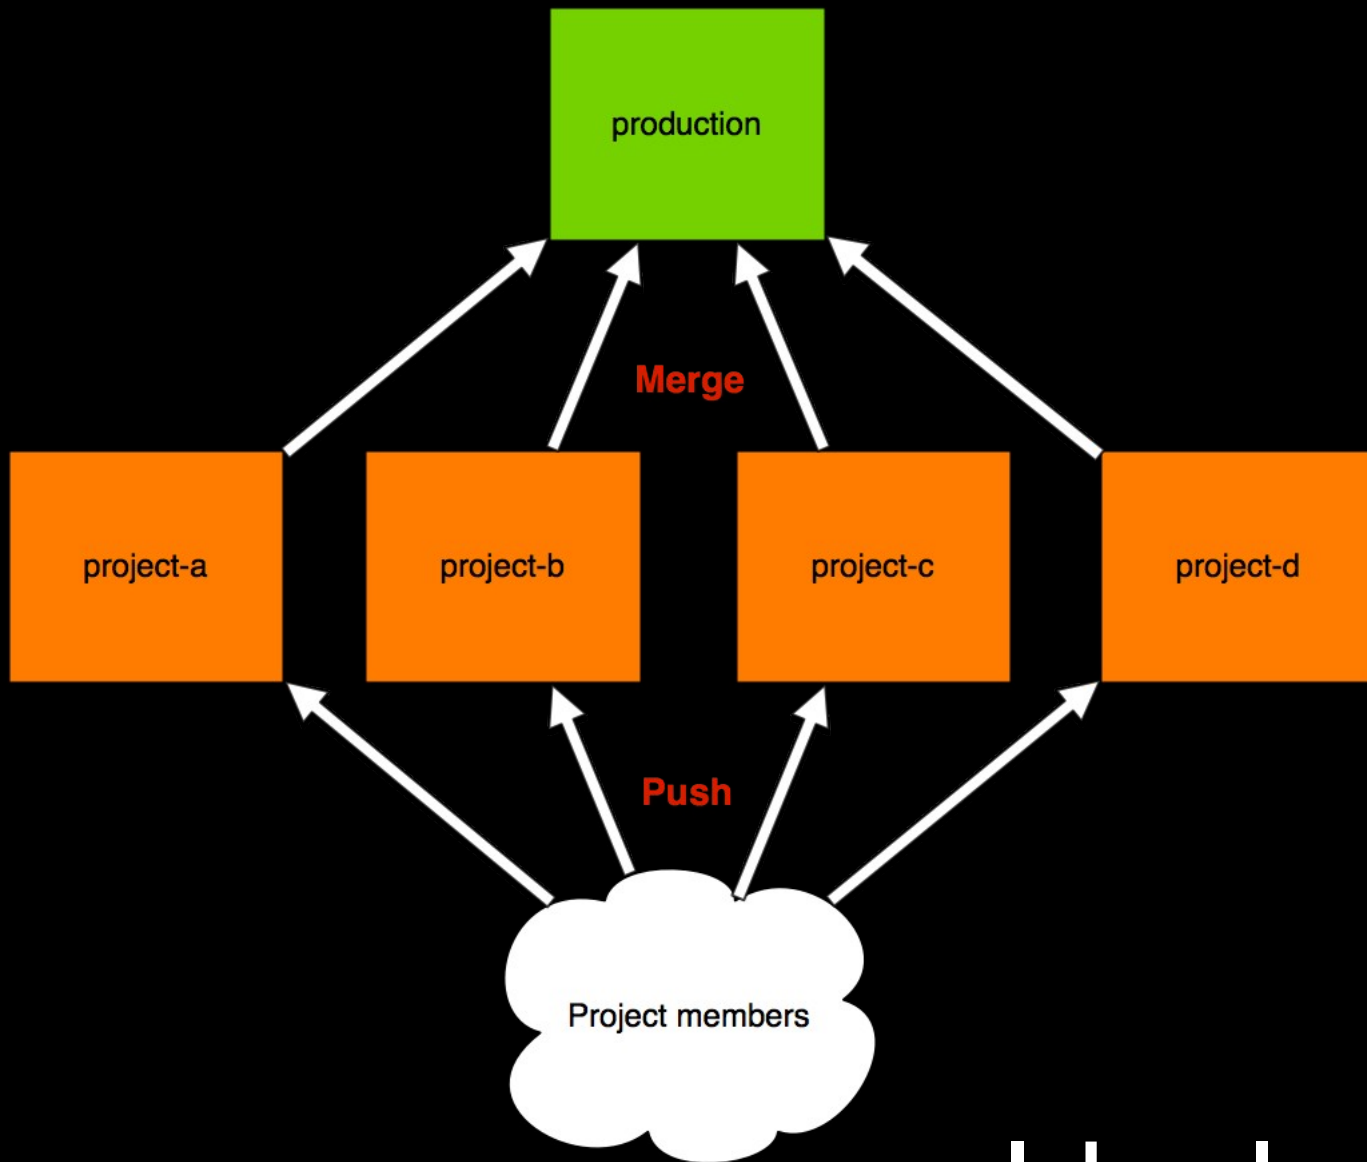
Modifying node definitions



Updating  
masters

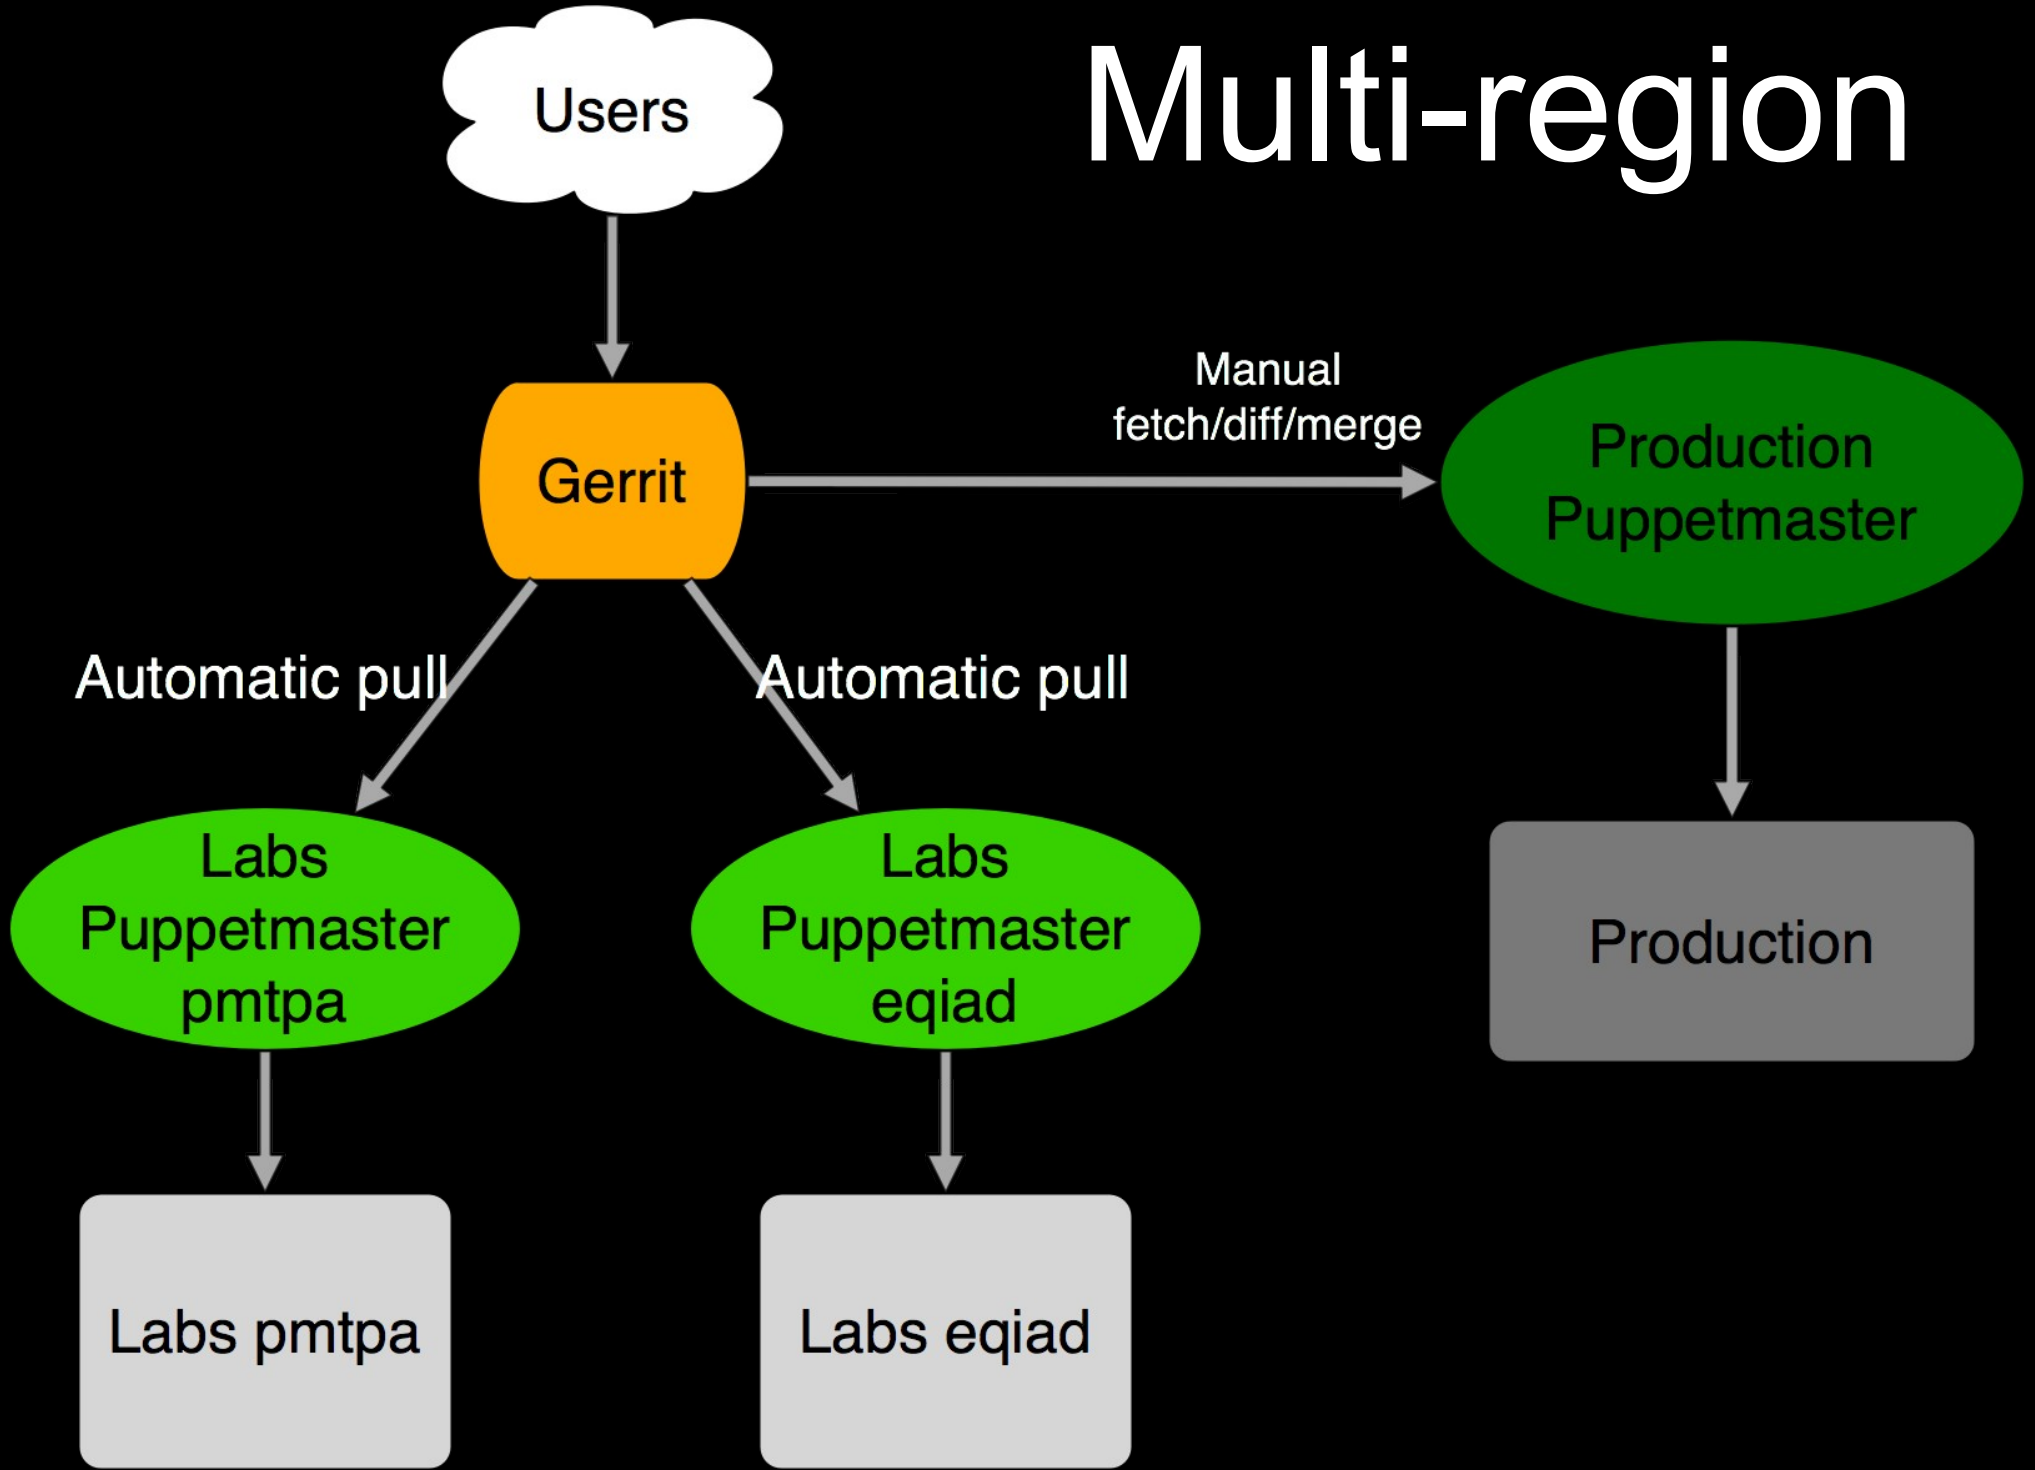


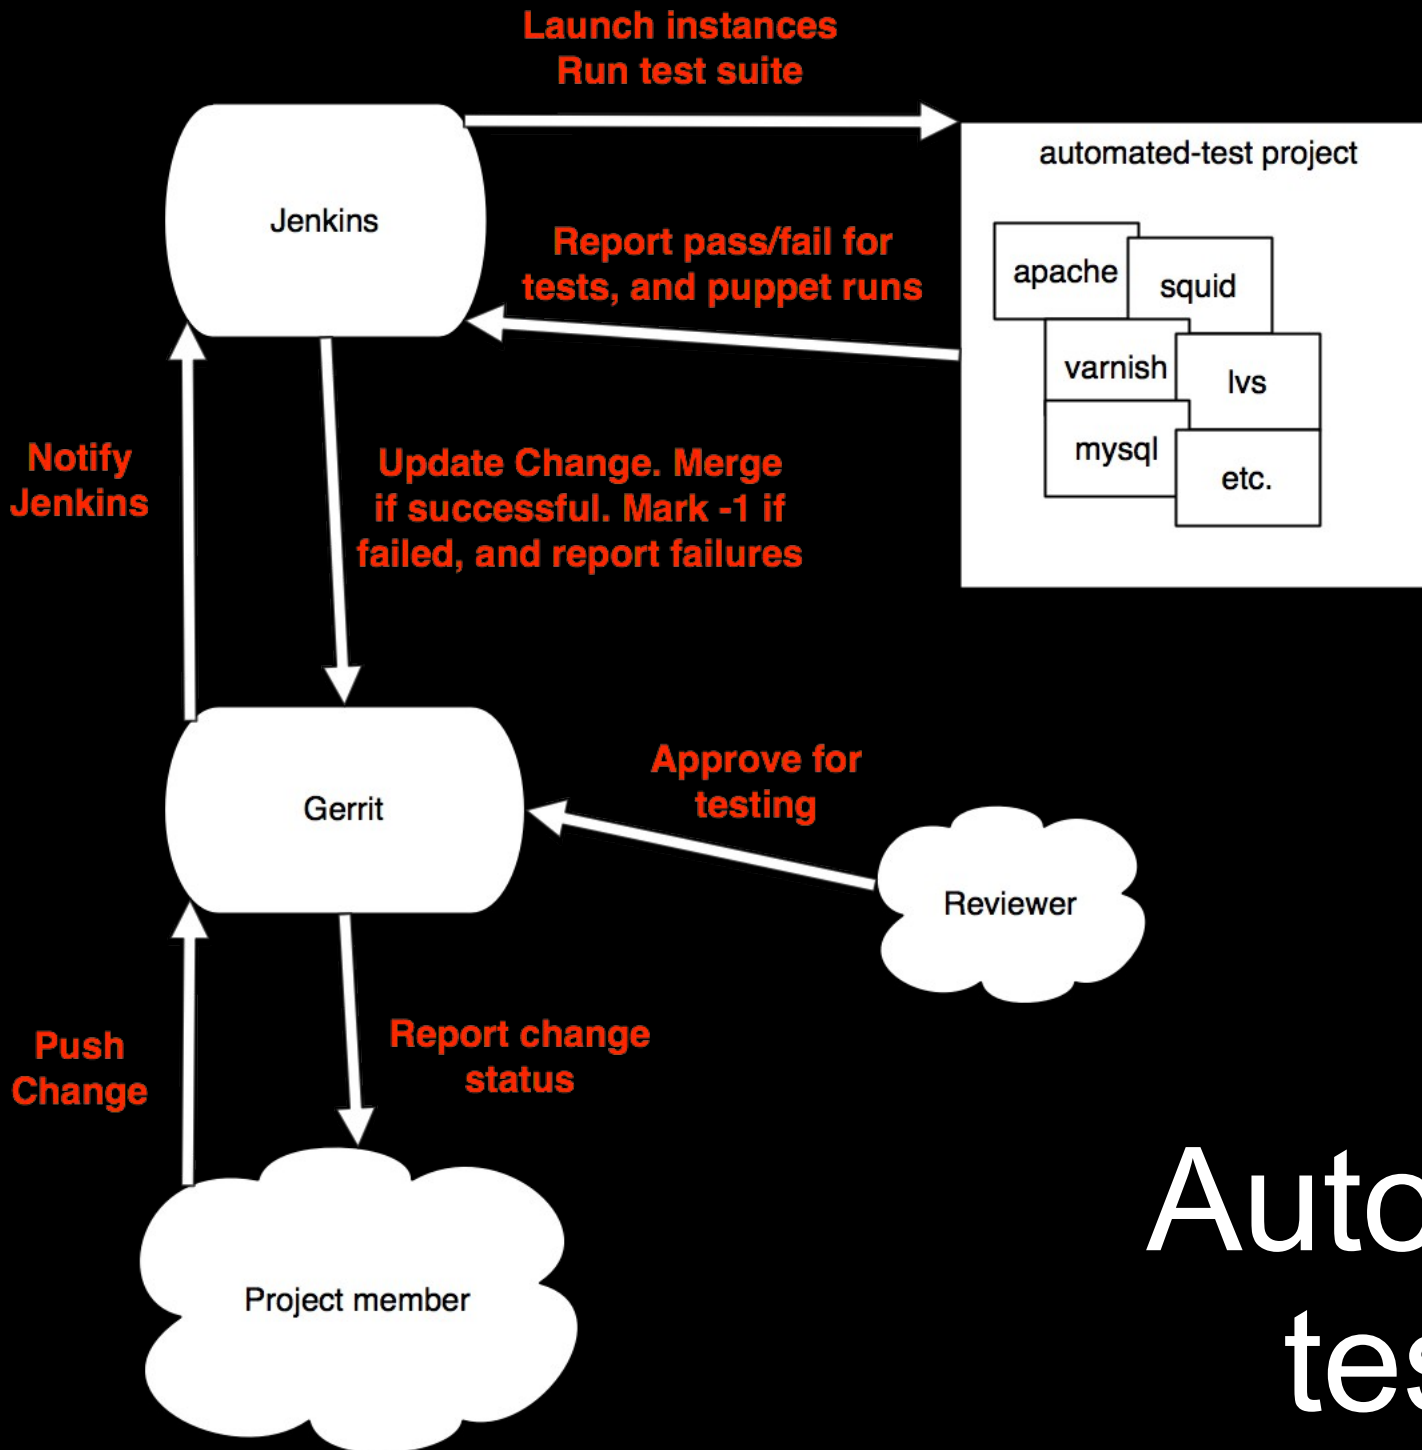
Future



Ideal: short-lived  
project branches

# Multi-region





# Automated testing



# Get Involved!

- Help with development
- Help with operations
- Help with community



Ryan Lane

Wikimedia Foundation

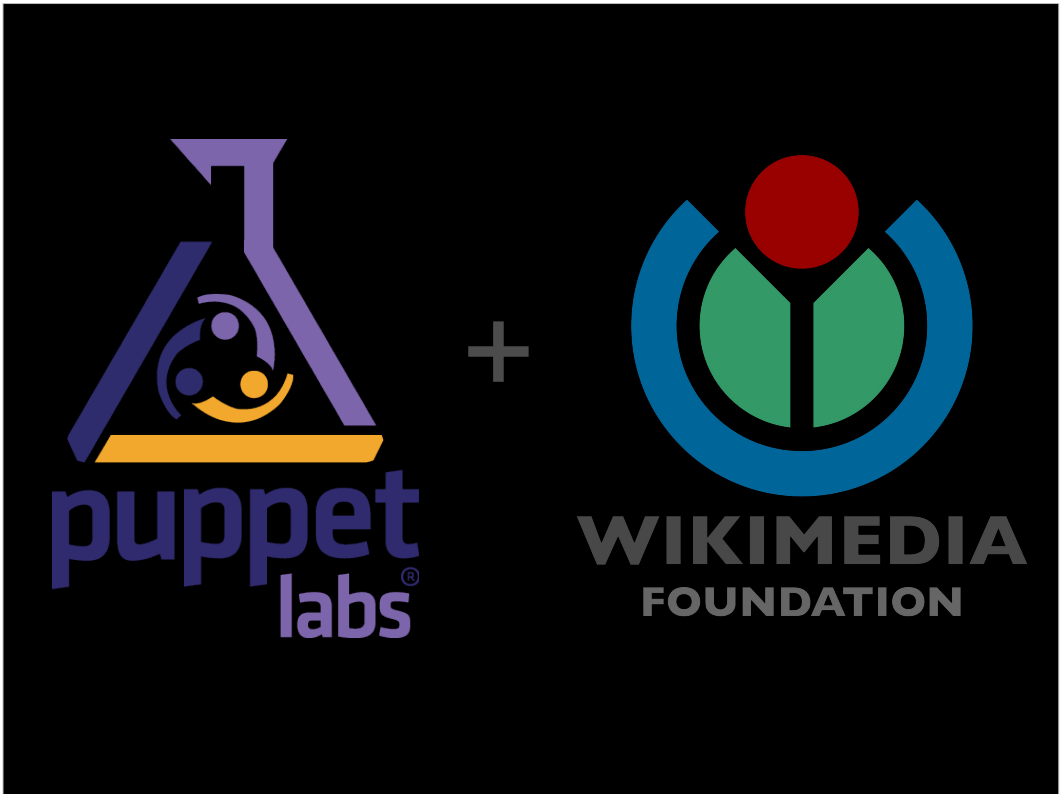
[ryan@wikimedia.org](mailto:ryan@wikimedia.org)

IRC: Ryan\_Lane on

Freenode; channels:

- [#mediawiki](#)
- [#wikimedia-labs](#)
- [#wikimedia-tech](#)
- [#openstack](#)
- [#openstack-dev](#)











## Early Tech Community: No staff



- When we started
  - Volunteers had root
  - Built the site and the software
  - Lenient access policy – lead to flexibility
  - Downtime wasn't considered as major an issue then



## Current Ops Community: No new volunteers

- Growth has made us wary of giving root or shell to volunteers
- Size makes us slower as changes are harder
- Lack of volunteers help makes us less flexible and slower



- This makes operations a bottleneck
  - Growth and speed of dev community is held back by operations
- Access policy makes it difficult for volunteers to understand how infrastructure works
- Also makes it hard to make changes that make it through code review

<http://www.flickr.com/photos/tswestendorp/4631910301/sizes/o/in/photostream/>

330 000 000 readers

100 000 editors

500 news agents

70 000 operators





- Solution: open the infrastructure back up to volunteers
  - Opening infrastructure development allows developers to bypass operations for common tasks
  - Thanks to orchestration, configuration management, and cloud computing operations can work like a software project
  - If we can treat operations like a software project, we can also treat it like an open source project
  - Still have the same issues with access – must take a new approach

<http://www.flickr.com/photos/nanoprobe67/5761031999/sizes/o/in/photostream/>



Wikimedia Labs is a virtualized test, development, and semi-production environment for Wikimedia Foundation staff and for Wikimedia volunteers.

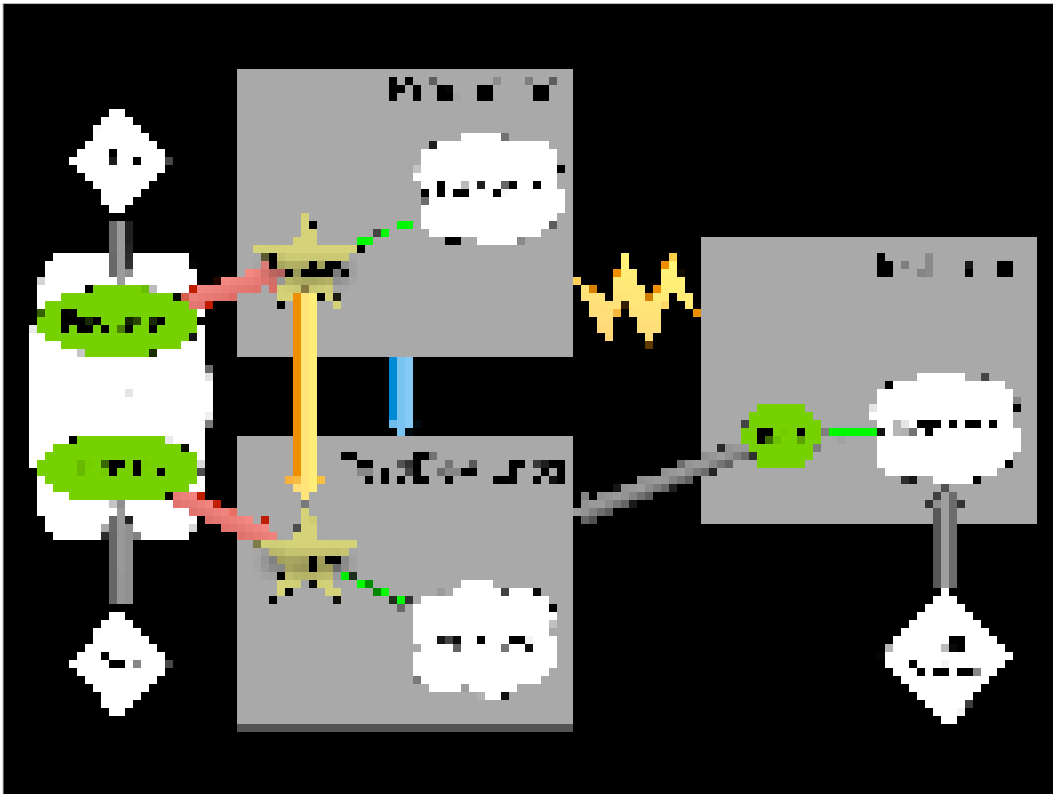


# Goals



I have three main goals:

1. Improve collaboration between staff and volunteers for software development
2. Have a process for providing higher levels of access for people who are not on the paid operation team staff. This includes staff developers, and all volunteers. I'd like to have an environment where anyone can eventually become root, even on our production cluster.
3. I'd like to have an environment where we can test major changes before we deploy them to the live site. We currently have no test environment.



- Puppet configuration kept in gerrit repository
  - Two branches used
    - Production for production
    - Test for test
- Two development environments to act as a ramp
  - Tool labs, intended for very new developers
    - Bots
    - Tools
    - Researchers
  - Test/Dev labs, intended for developers and ops volunteers
- Changes pushed and tested in test
- Code reviewed, then cherry-picked to production, then deployed



<http://www.flickr.com/photos/canihazit/6624625449/>



<http://www.flickr.com/photos/gotosira/4699302559/>

- First step: release our puppet configuration
  - Never planned on this
  - Spent weeks auditing configuration for sensitive things
  - Almost released \*.wikimedia.org private key




openstack™

cloud just made it



- We use nova projects as security separations between community maintained projects
  - Separation is very loose
  - Community controlled
  - Resources are shared per-project, not per instance
- Projects are global or specific
- Projects extended as posix security groups
  - Can manage permissions in global projects using specific projects





User calls instance



LDAP entry fetched



Instance created



User calls instance again




Supplier master signs request



Instance loaded



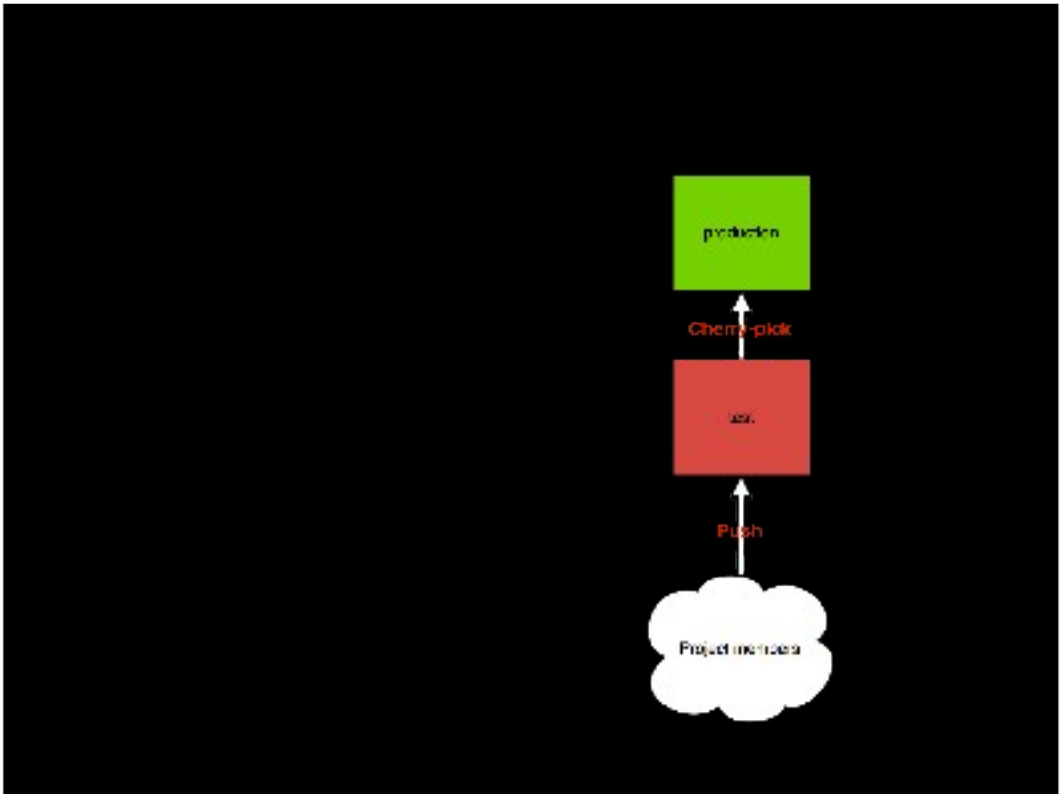





- Initial concept for moving changes from test to production was flawed. Too hard to test changes.
- Changed concept to make this easier
  - Every project gets a branch
  - Project members can switch between test and project branch
  - Can implement and test using project branch, and then merge changes to test branch
  - Test branch can be used to test changes in production clone before merging to production and deploying






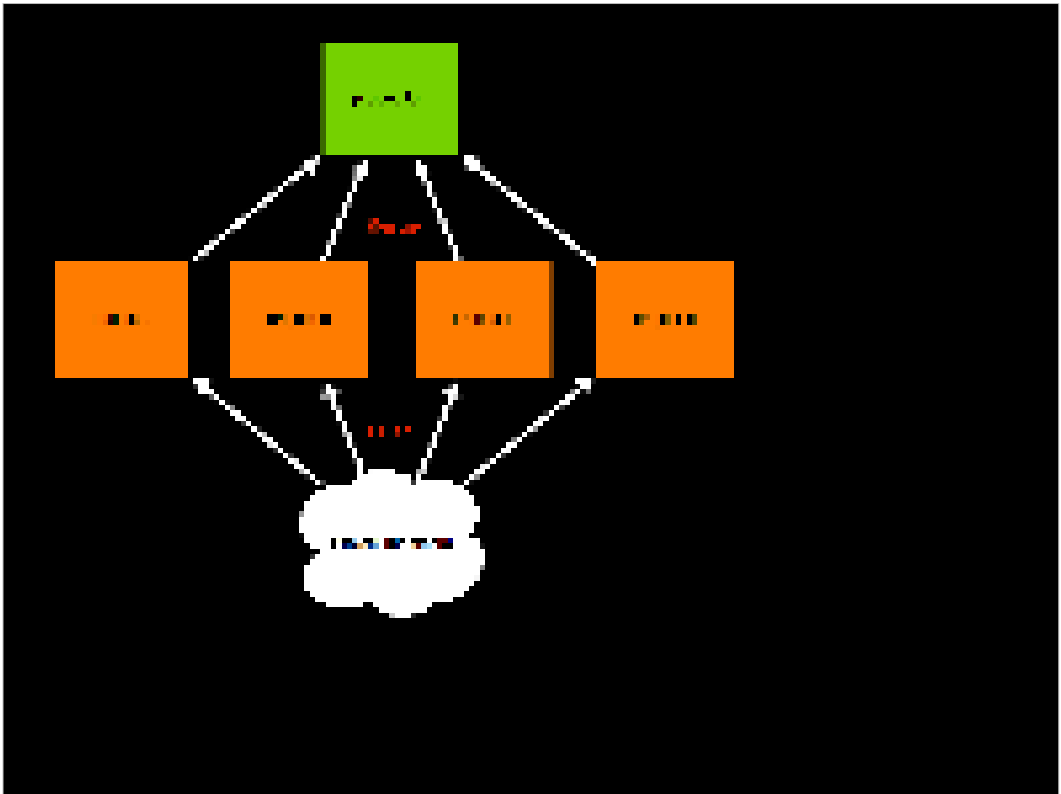




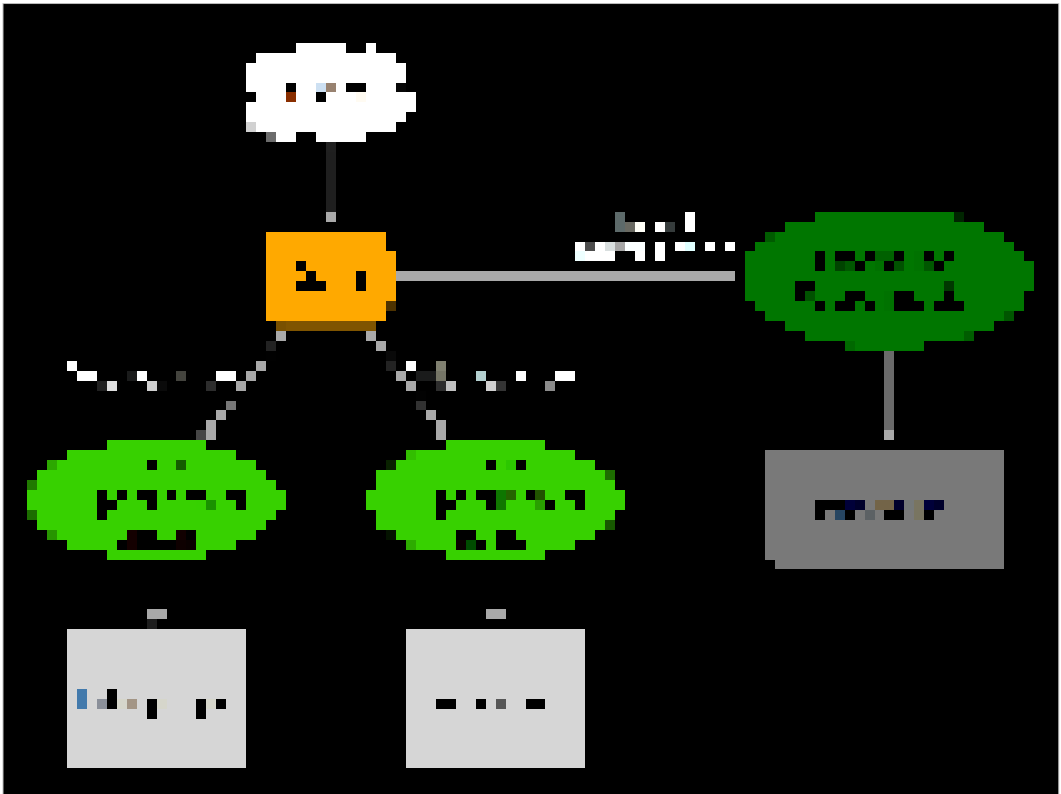


















Ryan Lane  
Wikimedia Foundation  
[ryan@wikimedia.org](mailto:ryan@wikimedia.org)

IRC: Ryan\_Lane on  
Freenode; channels:

[#mediawiki](#)  
[#wikimedia-labs](#)  
[#wikimedia-tech](#)  
[#openstack](#)  
[#openstack-dev](#)



Please don't hesitate to contact me. I'm very active on IRC, and would love to talk to you.