Scaling Agile Infrastructure

Development, Change Management, Security

Ben Jones IT-PES-PS

CHEP 2015 Okinawa, Japan

Definitions

- Admins: Users who are able to commit to git repositories used to create puppet code tree
- Hostgroups: machines grouped by common configuration, typically providing a related service and having common admins. In terms of puppet code, pulls together modules and uses data to define a service.
- Modules: common reusable code. When used by more than one hostgroup is “shared” and is subject to change control
- Puppet Masters: machines compiling puppet code on demand, serving around 200 catalog compilations / second
- Client Machines: Windows, Linux and dumb devices being configured by puppet

Development

- There is only one code tree which is applied to all puppet masters.
- Each module and hostgroup is a separate git repository. The code tree is compiled from a list of all modules and hostgroups
- Each repository must contain a production and a QA branch, from which respective production and QA environments are created.
- Other environments can be created by accepting production or QA as the base, but having feature branches of individual modules or hostgroups as overlays.
- Different services move at different speeds. Whereas the configuration tree is continuously delivered, deployment can be staggered. This is achieved by snapshotting - creating an environment from a point in the git history of a module.

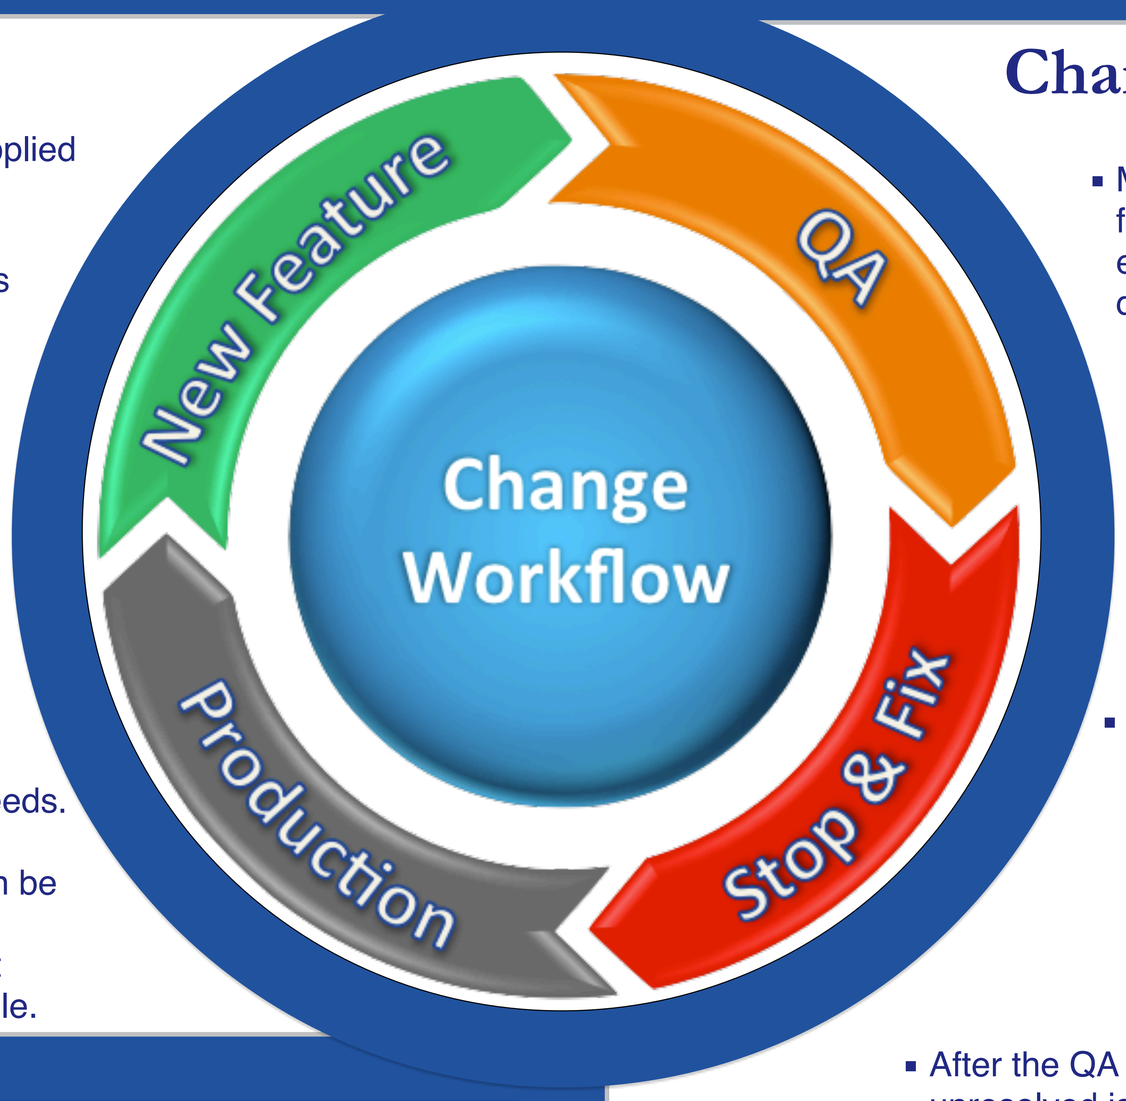

Change Workflow

- Module changes are tested in a git feature branch, mapped to a puppet environment. Test machines are deployed to the puppet environment.
  - Puppet’s data warehouse (puppetdb) is used to determine if the puppet module is shared, and thus requires change management.
  - Jira is used to create Change Management Requests, detailing the changeset and the originating feature branch.
- Changes reside in QA for a minimum of one week, during which time a Service Manager can press the “STOP” button which prevents the code being merged into production. The module maintainer can abandon the change request, or fix the issue, and then return the change to QA.
- After the QA period has elapsed, and if there are no unresolved issues, the code is merged to production.

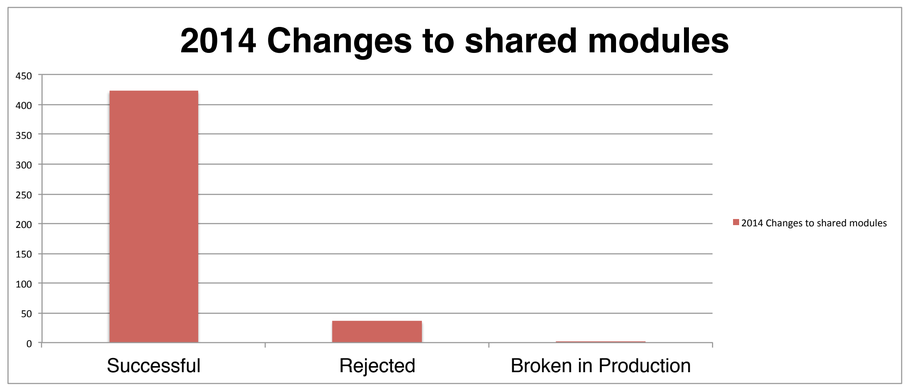

Test Pipeline

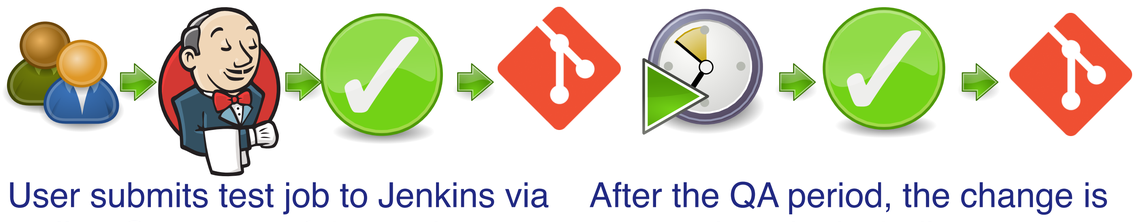

User submits test job to Jenkins via Jira. On successful test, the feature branch is merged and pushed to the QA branch

After the QA period, the change is tested again, including to ensure no service manager objections, then merged to production

Security

The security model of puppet relies on every user who has access to the puppet master to be trusted. Where we have more than 200 admins, this isn’t possible to do across the different groups, or allowing for account compromise.

The puppet pluginsync process takes ruby code in modules and ships them to the client. In the case of “facts” this means running code as root. Using a whitelist, we allow service managers to control which modules they accept such code from.

Secrets that are accessible from the puppetmaster cannot be trusted, as they must necessarily be readable by any user of that puppetmaster process. Instead host credentials are used to download secrets that are only referenced in the host manifest.

Info

Authors:

- Ben.Jones@cern.ch
- Steve.Traylen@cern.ch
- Ignacio.Barrientos.Arias@cern.ch
- Gavin.Mccance@cern.ch

Web:

- https://cern.ch/config

Statistics used are from data for 2014