



# Configuration Services at CERN

Ben Jones, HEPiX Fall 2014

# Agenda

- Infrastructure
- Security
- Continuous Integration
- User Training
- Current Issues
- Community

Quattor at CERN 2004 - ~~Oct~~ Nov 2014  
R.I.P.

# Infrastructure: puppet

<b>Puppet Masters</b>	49	(46 “batch” 3 “interactive”) all VMs
<b>Puppet Clients</b>	12,814	
<b>Puppet Version</b>	3.4.3-1, 3.4.3cern2.dirty-1	(masters are patched, more later)
<b>Catalog requests / min</b>	140	
<b>Average compilation time</b>	92 seconds	
<b>Modules</b>	249	
<b>Top-level hostgroups</b>	138	
<b>Environments</b>	141	2 “golden”: prod & qa

# Infrastructure: Foreman

- 7 Servers
  - 2x UI, 2x ENC, 2x Reports, 1x Rake tasks
- Running v1.3 + patches
  - Bugs/missing features preventing upgrade
- Use as ENC for hostgroup membership and some variables
- Inventory
- Report Visualization
- Kickstart generation
- BMC management
- Not used for modules (though perhaps some use cases)

# Infra: misc

- Git repository per hostgroup / module
- Homegrown tools “jens” to create manifest trees on puppet masters
- PuppetDB with two machines, postgres & puppetdb process
  - some issues with postgres replication
- Jira & jenkins used heavily in workflow
- mcollective kerberos plugin for PuppetDB & foreman (sent upstream)
- Looking at rundeck for some automation

# CERN Puppet user stats

	Module authors	Hostgroup authors
2013	152	85
2014	125	206
September 2014	39	122

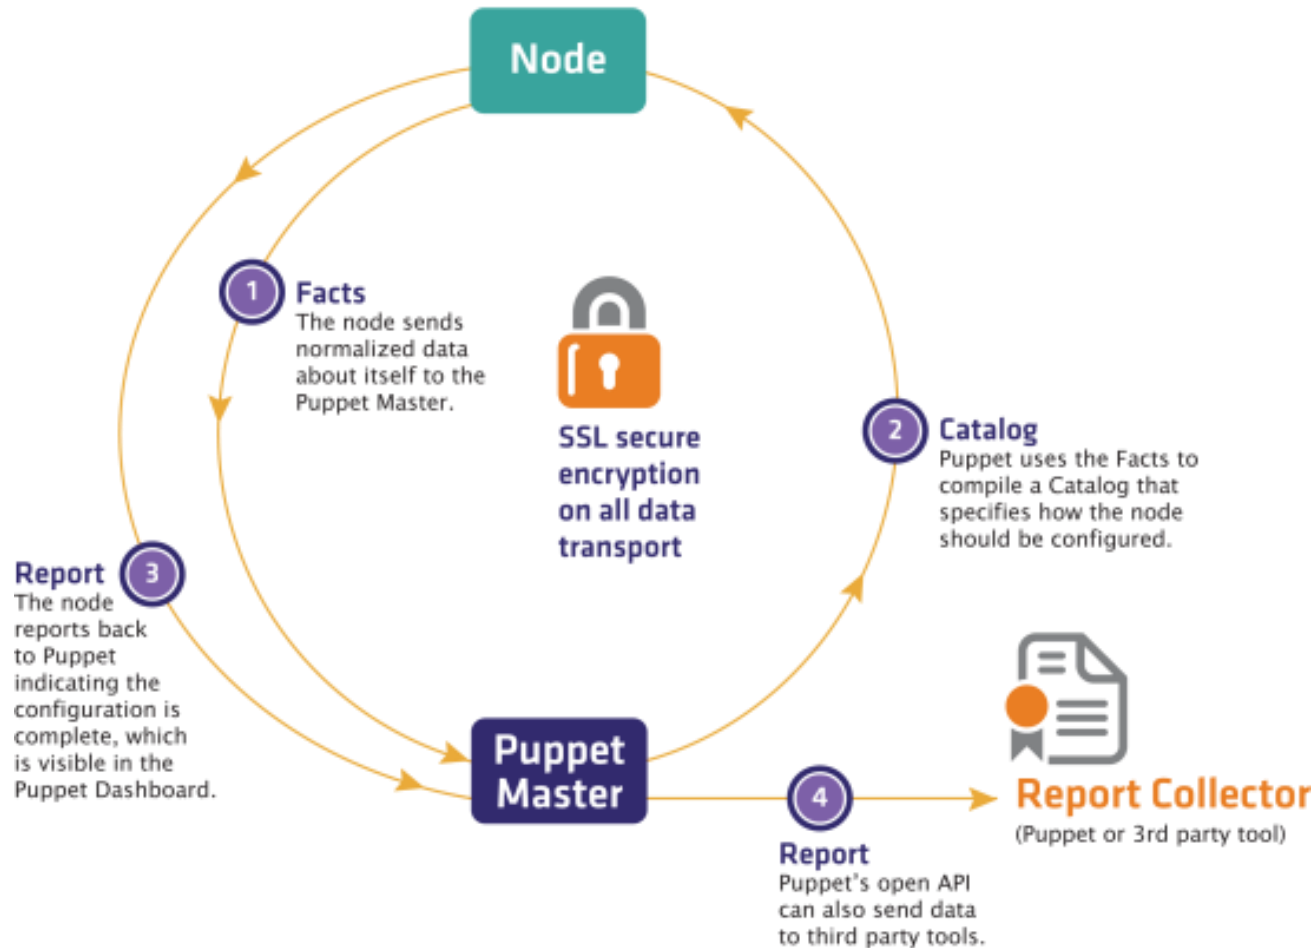
- Modules have between 400 and 700 production commmits a month
- Hostgroups 1500-2200 production commmits a month



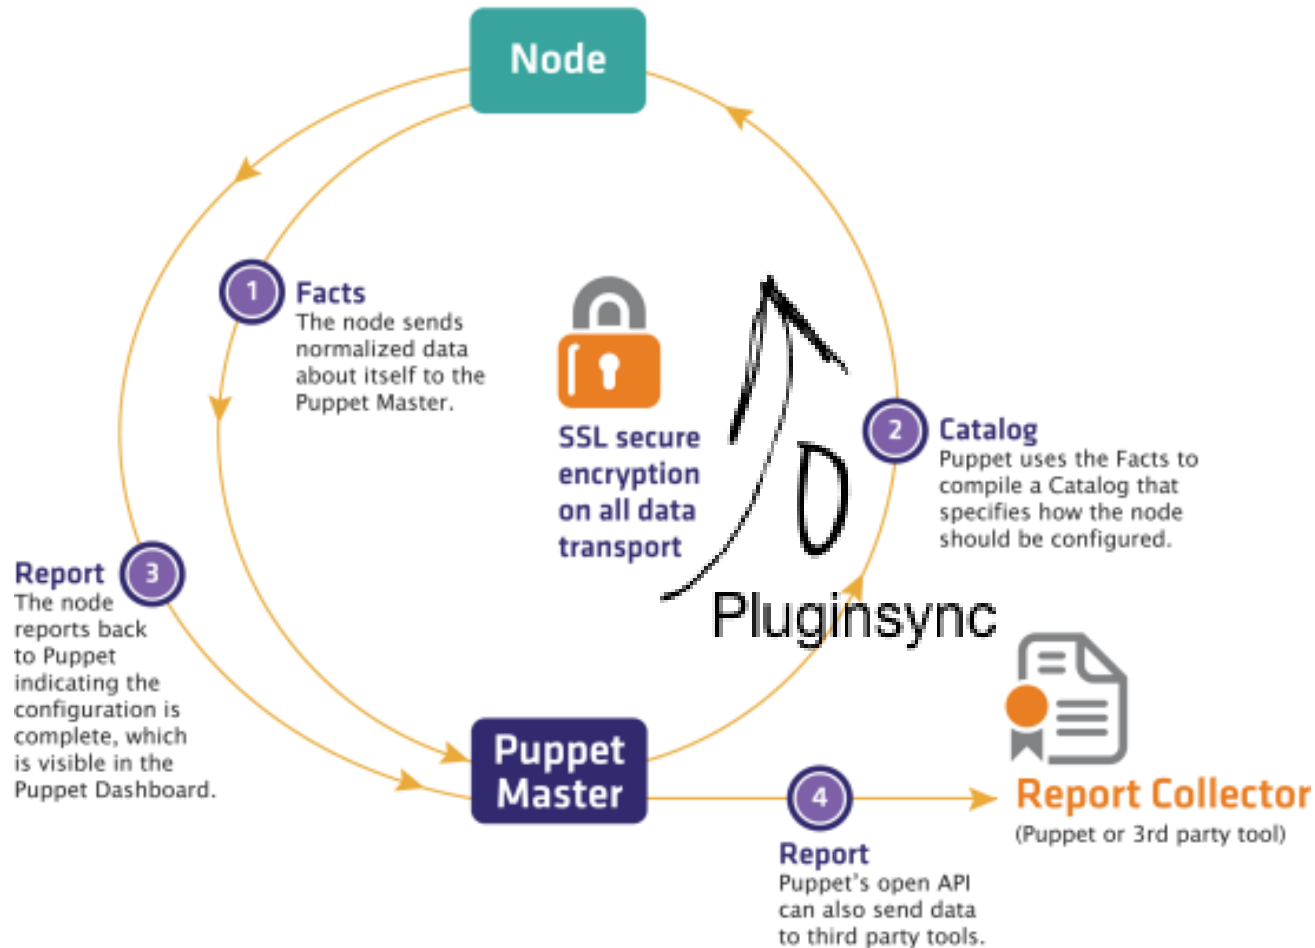
# Security

- CERN has puppet admins throughout IT department
- Even assuming trust across all groups one account compromise shouldn't affect unrelated services
- Our git repo per hg, module should be the limit of permissions
- Puppet's sweet spot may not be for this admin model

# Pluginsync



# Pluginsync



# Pluginsync

- Aside from configuration in the form of puppet manifests, puppet ships “libdir” code
  - facts
  - functions
  - types & providers
  - Augeas
- Pluginsync is process that ships libdir to clients to enable custom facts & providers to be present on client for runs
- All of libdir on modulepath of server is shipped to clients

Wait. All of libdir on modulepath of server is shipped to clients?

# Pluginsync problems

- Facter runs as root
  - So does everything else of course, but `_all_` facts are run `_always_`
- Pluginsync sends all facts (and other libdir) defined in all modules
- In our model that means any hostgroup (nevermind module) can run arbitrary code as root on all machines

## 3.4.3cern2.dirty-1

- Patched puppet to change behaviour of pluginsync
- Hiera variable for “pluginsync whitelist”
- Extra ENC call to determine hostgroup to determine hiera lookup
- Hiera lookup means we can have “base” modules in whitelist globally
- Yes, it leads to bad surprises occasionally

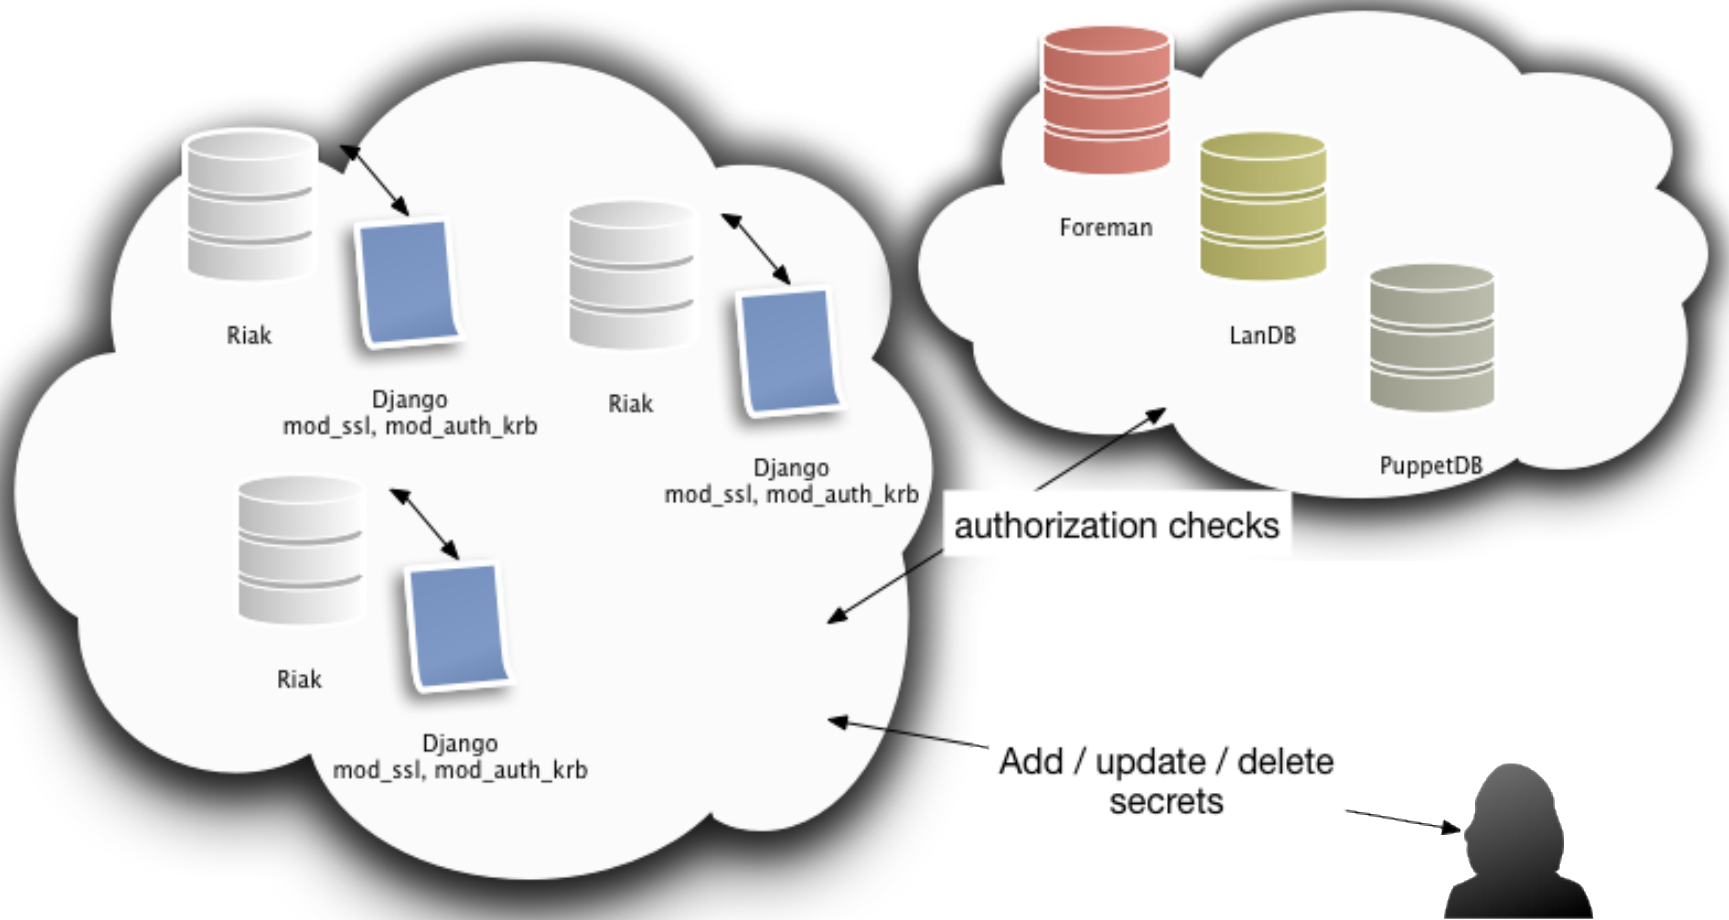
# Secrets

- Problem of secrets has been discussed at previous HEPiX
- Solutions such as hiera-gpg mean that puppet master must decode secrets to add to manifests
- Many methods to get data from puppet masters
  - hiera calls, functions, erb files
- Resources are stored in PuppetDB



# secrets at CERN

- External datastore using hostgroup and host authorization
- Types and providers to allow for files and templates to include keys for datastore
- Host authentication is used to download secrets (x509 or keytab)
- Anyone who can get root on a machine can update/see secrets



Secrets downloaded by key



← Catalog contains secret keys —



Puppet Catalog

## Teigi architecture

# Continuous Integration

- We want to enable regular changes of configuration
  - small changes safer than big changes
  - drift and freezing causes other problems
- Requires rigorous testing and QA to be possible
- Testing should be frictionless
- Normal QA flow: change from topic to QA, wait a week without complaints, QA to Prod
- Automated Pipeline with tests vs forgetful busy humans

Project\*  Config Release Management ▾

Issue Type\*  Configuration Change ▾ 


Some issue types are unavailable due to incompatible field configuration and/or workflow associations.

Summary\*

Change Type\* New feature ▾

Change scope\* Manifest change ▾

Whether the change is a manifest or software update or both

Security Level Internal Data ▾ 

Origin Repo/Branch\*

The git repo and origin branch for the change

Reporter\*  Ben Jones

Automatic Change  Use the automatic CI pipeline for this change

[TEST] This change will be driven using the automatic CI pipeline

# CI Tests

- Daily tests of base modules run on freshly installed VM
- Outside “base” tests are run based on change management ticket
- Functional tests can be defined for each module
- Jenkins test nodes spawned from openstack, with specific configuration as needed
  - for example: “SSO” app webserver

# CI radiator

## QA-continuous-integration-status



**CC7**

**test-afs**

#44 1h 32m 4s

#312 7m 7s

**test-base**

**test-certmgr**

#163 - aici-test-crm612-b01b030f91.cern.ch 1m 13s

#150 2s

**test-firewall**

**test-landb**

#156 0s

#29 - aici-test-crm624-2ea4fdd04f.cern.ch 4s

**test-lbd**

**test-mysql**

#12 1s

#20 22s

**test-osrepos**

**test-pgsql**

#131 1m 6s

#20 23s

**test-puppet**

**test-shibboleth**

#2 1m 41s

#616 27s

**test-sssd**

**windows-QA**

#358 13m 45s

#186 40s

# Change Management test



**Build CRM-724 (#56) (Oct 13, 2014 4:09:07**

Progress:  

**PM)**



No changes.



Started by Naginator

[Phase] - Create

 [create-generic](#)

[build #169](#) ( 20 min )

[parameters](#)

# test-shibboleth


test-shibboleth

ENABLE AUTO REFRESH

trend

## Build Time Trend

Build	Duration	Slave
<a href="#">#619</a>	24 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#618</a>	24 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#617</a>	24 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#616</a>	25 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#615</a>	33 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#614</a>	24 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#613</a>	25 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#612</a>	29 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#611</a>	50 min	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#610</a>	28 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#609</a>	25 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#608</a>	25 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#607</a>	29 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#606</a>	22 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#605</a>	25 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#604</a>	25 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#603</a>	37 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#602</a>	26 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#601</a>	25 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#600</a>	24 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#599</a>	21 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#598</a>	20 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#597</a>	23 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#596</a>	21 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#595</a>	0.66 sec	<a href="#">jenkins-slave3.cern.ch</a>
<a href="#">#594</a>	2.6 sec	<a href="#">jenkins-slave3.cern.ch</a>





# Training

- We initially thought we'd need basic Puppet training
  - ...and setup a course with external trainers
- What people actually prefer is the details for your site
  - Local conventions, local tools and practices
  - We found the rest you can get from books, Puppet documentation, and just from getting on with it

# Local training


- At CERN we now provide a general training for the new infrastructure
  - Cloud
  - Puppet and Configuration
  - Monitoring
- 1 day training, approx 1/month presently
- Documentation contains training material, so people can work on it in their own time

# Current issues

- Foreman permission model.
  - 1.3 we have patches to reduce hostgroup filter to hostgroups the user is permissioned to
  - 1.5 has regression of this behaviour
  - Foreman is perhaps on critical path for puppet upgrades
- Reacting to events such as Heartbleed & ShellShock could be improved
  - should be noted that mcollective has been useful
- PuppetDB performance issues more or less fixed
  - issues separating API use from compilation critical path
  - issues getting postgres replication to keep up

# Postgres slave 1 day behind...

lagging puppetdb, slave running postgres 9.2.9



# Community

- 9 CERN modules now on puppet forge
  - more on [github.com/cernops](https://github.com/cernops)
- Interest in finding gaps where YAIM currently relied upon
- Will be sending a survey to discover top modules required to migrate from YAIM

# WHAT??



KREATHOOD

KREATHOOD

KREATHOOD

KREATHOOD

