
CERN

Update on Configuration Management at CERN

Alberto Rodríguez Peón, HEPiX Autumn 2015

Outline

- **Agile Infrastructure**
  - Current status and numbers
  - New CC7 Puppet Masters
  - Hiera GPG phase-out
- **Moving to new Apache module**
  - How we used environments to migrate a non-backwards compatible module
- **A new hostgroup ownership model**
- **Package Inventory**

Some Puppet numbers

- Last HEPiX meeting, Spring 2015

Puppet Clients	15,800
Catalog requests / min	193
Catalog compilation time	121 seconds
Modules	276
Top-level hostgroups	159
Environments	196

Some Puppet numbers

- Last HEPiX meeting, Spring 2015

Puppet Clients	~~15,800~~ **18,577**		**+ 17 %**
Catalog requests / min	~~193~~ **199**		**+ 9 %**
Catalog compilation time	~~121~~ **38** seconds	**(!!!)**	**- 68 %**	**(!!!)**
Modules	~~276~~ **289**		**+ 5 %**
Top-level hostgroups	~~159~~ **154**		**- 3 %**
Environments	~~196~~ **211**		**+ 7 %**

CC7 Puppet Masters

- **Catalog compilation time drastically reduced** by upgrading our Puppet Masters

  OS: **SLC6** -> **CC7**

  Puppet: **3.7.4** -> **3.7.5**

  Ruby: **1.8.7** -> **2.0.0** **(not backwards compatible!)**

- The upgrade involved chasing people to **adapt their code to Ruby 2.0**
  - Only **affected Ruby code** (i.e. server-side functions, ERB templates, **YAML files**) but no Puppet manifests
  - Service managers were given the chance to **test their configuration code** before the upgrade

Catalog compilation time comparison

Average per fixed host

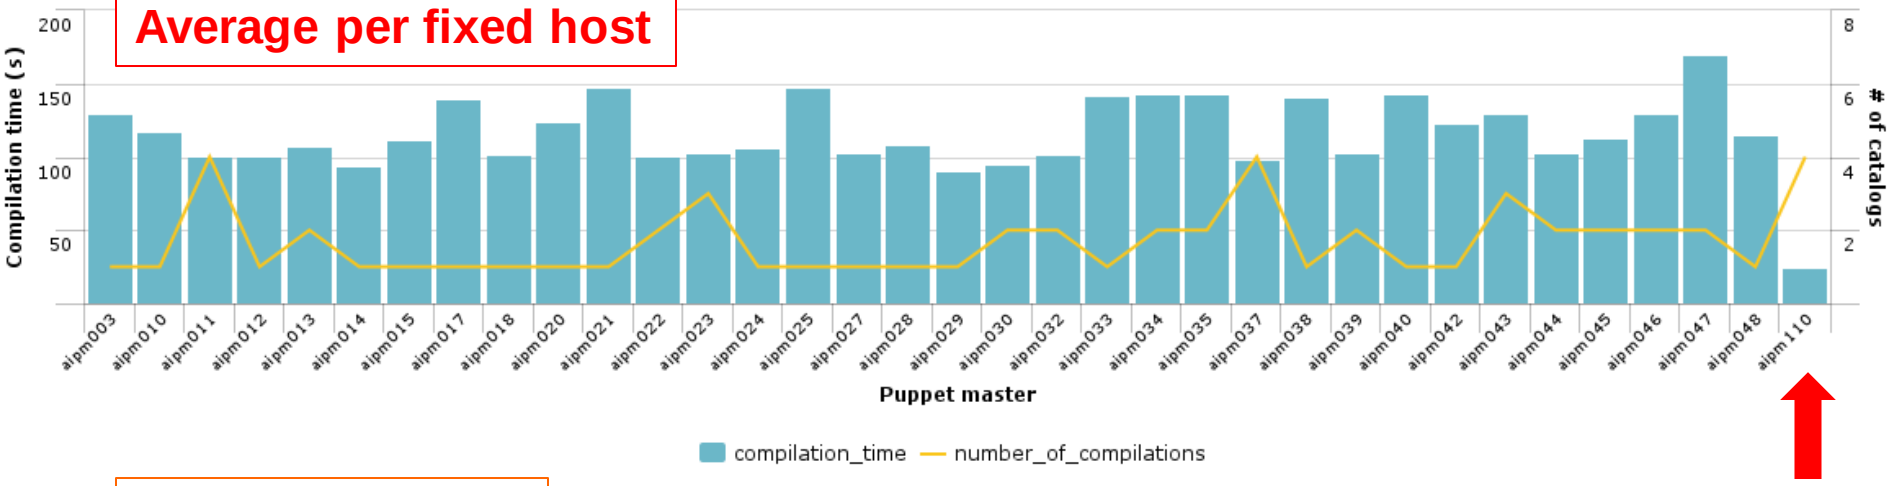

Average per pool

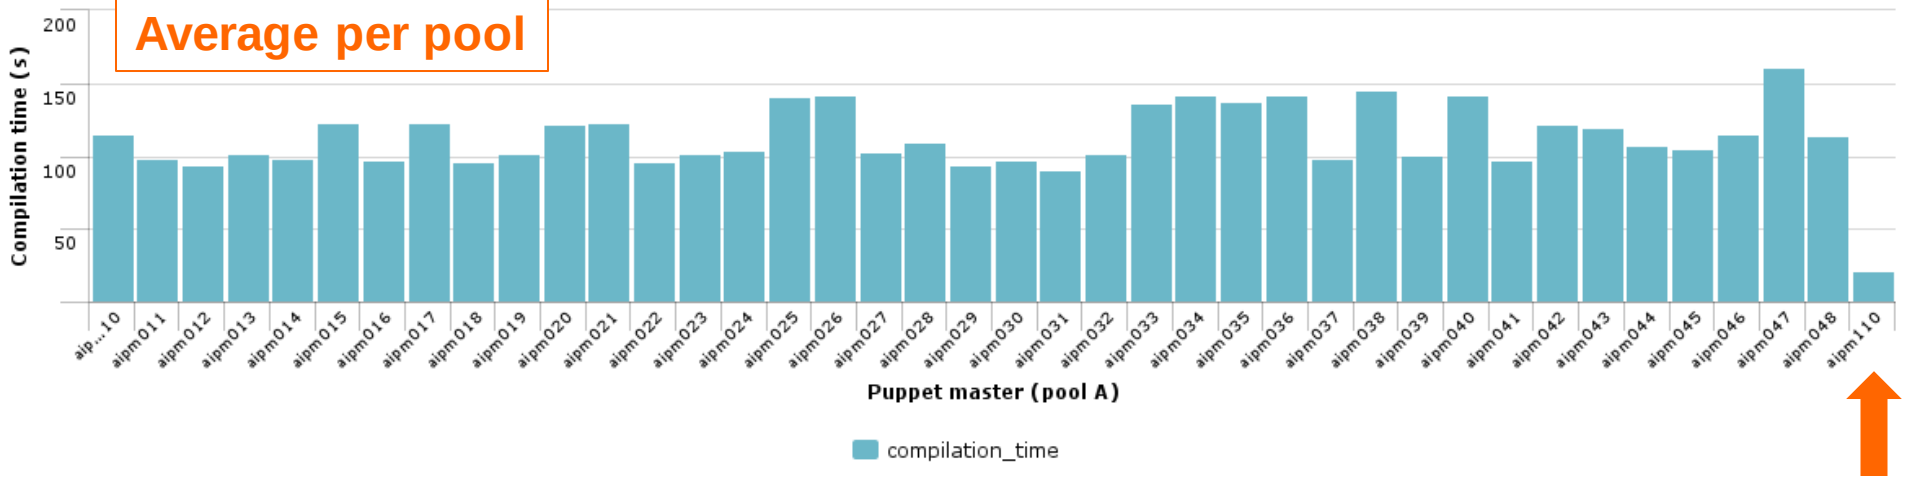

CC7 Puppet Masters

- **CPU usage** on the PMs also **dropped**
  - From ~**80%** to **~25%**
  - Now we can **accommodate the same number of** catalog **requests** with less than **half** of the hardware

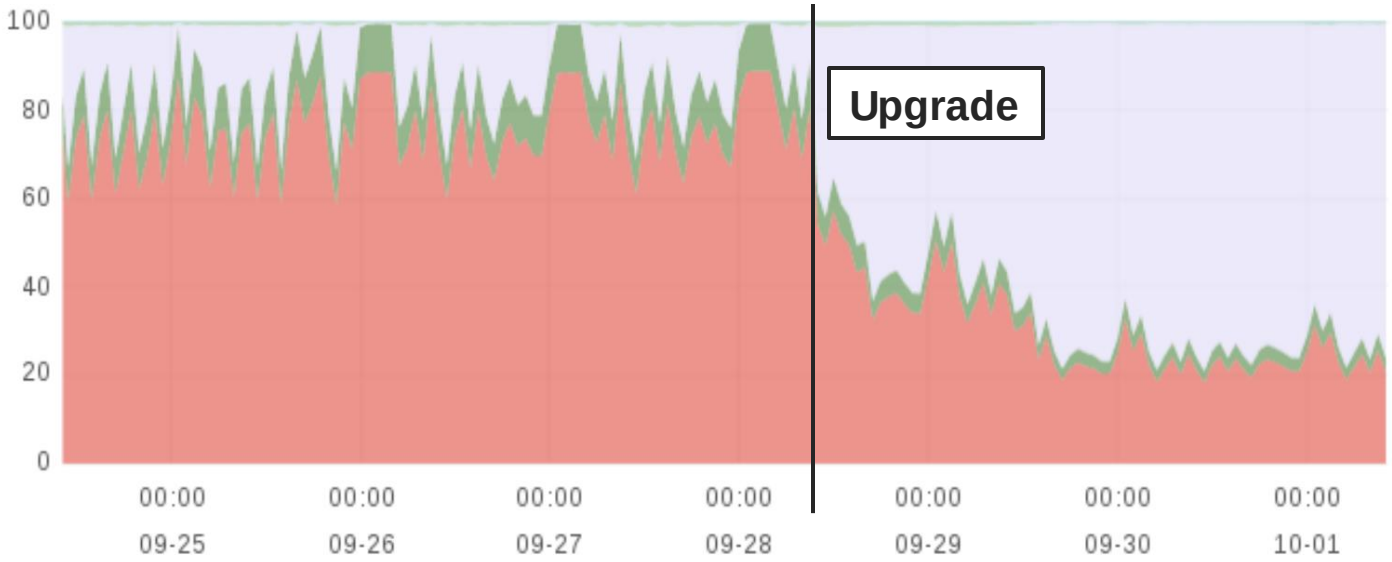

Next steps for the infrastructure

- Move to **puppet-server**
  - PM re-implementation in *Clojure* + *JRuby* and run on the JVM
  - Supposedly much faster

- Start thinking on **Puppet 4**
  - **Not backwards compatible (again!)** with **Puppet 3**
  - Making the **puppet future parser available and optional** could ease the transition for all the users

Hiera GPG phase-out

- Although **secrets** are **GPG encrypted**, they are evaluated in the **puppet masters**
  - Therefore **readable by anyone** with access to them

- At CERN, we store secrets using an in-house service that **evaluates** them **directly in the client**
  - More details in our slides from **HEPiX Fall 2014**!

- **No need to support Hiera GPG anymore**
  - Decrease in the **catalog compilation time**

Upgrading Apache puppet module

- Moving our Apache module from **0.4.0** to **1.2.0** was tricky
  - https://forge.puppetlabs.com/puppetlabs/apache
  - Newer releases are **not always compatible** with previous code

2013-09-06 Release 0.9.0

Summary:

This release adds more parameters to the base apache class and apache defined resource to make the module more flexible. It also adds or enhances SuPHP, WSGI, and Passenger mod support, and support for the ITK mpm module.

Backwards-incompatible Changes:

- Remove many default mods that are not normally needed.
- Remove `rewrite_base` `apache::vhost` parameter; did not work anyway.
- Specify dependencies on stdlib >=2.4.0 (this was already the case, but making explicit)
- Deprecate `a2mod` in favor of the `apache::mod::*` classes and `apache::mod` defined resource.

Apache upgrade: environments

- Apache is **widely used** in our infrastructure
  - **385 servers** across **40 different services**

- How to guarantee **a smooth transition**?
  - Using **puppet environments!**
  - The idea is to maintain **both versions** of the module **available** in different environments for a while
  - Existing users can migrate **at their own pace**!

Apache upgrade: action plan

1. Create a **new custom branch** on the module with the **old Apache 0.4.0** code
2. Add two new **golden environments defaulting to production and qa** but with apache (and shibboleth) **overridden** to the new custom branch
3. Identify **who's using the module** (by *tagging* the code)
   - https://docs.puppetlabs.com/puppet/latest/reference/lang_tags.html

```
---
default: master
notifications: ai-config-team@cern.ch
overrides:
  modules:
    apache: gold_apache040
    shibboleth: gold_apache040
```

4. **Move everyone** using Apache to these new environments
5. Upgrade **the Apache module (*master* and *qa)*** to **1.2.0**
6. Service managers can go back to *production* and *qa* once **their code is adapted to the new module**

A new ownership management model

- In **CERN's Agile Infrastructure** we use ***hostgroups*** as the administrative unit of management
  - A ***hostgroup*** is a collection of puppet code, foreman directives and a set of servers

Hostgroup ownership

- A ***hostgroup*** **responsible** must have ownership in all of the bits of our infrastructure
  - Mainly **Foreman** and **Git**

A single *source of truth* for ownership

- The ***source of truth*** for hostgroup ownership is stored in a key-value database (**Riak**)
  - Easy to **query** and **update**
  - Foreman and Git ACLs are periodically refreshed with the latest changes

Key

`/hostgroup/bi`

Value


```
{
    "hostgroup": "bi",
    "options": {},
    "owners": [
        "it-dep-pes-ps-support"
    ],
    "update_time": "1438007448",
    "updated_by": "mccance"
}
```

Package Inventory

- Solution to **detect package drift**
  - **Developed in-house** and with a stack formed by *elasticsearch*, *Flume* and *Kibana*

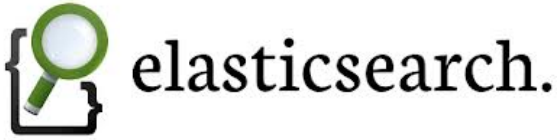

- Every server **reports the list of its installed packages** to a centralized *elasticsearch* cluster
  - **Async –** with a YUM plugin
  - **Sync –** with a daily cronjob

Package Inventory

- *PackageInventory* also ships a CLI that **can compare the package drift** of two or more servers

```
$ pkginv -m 'b6b3d71dfa b6b1576a51' compare
+-----------------------------------------+-------+------------+-------------+
|                 Package                 | Field | b6b3d71dfa |  b6b1576a51 |
+-----------------------------------------+-------+------------+-------------+
|                  httpd                  |       |  Present   | Not present |
| kernel-module-openafs-2.6.32-504.16.2.el6 |       |  Present   | Not present |
|               httpd-tools               |       |  Present   | Not present |
|              apr-util-ldap              |       |  Present   | Not present |
|                 gridsite                |       |  Present   | Not present |
|                 mod_ssl                 |       |  Present   | Not present |
|               castor-lib                | epoch |   8.slc6   |    9.slc6   |
|               castor-lib                |  arch |    i686    |    x86_64   |
|           castor-rfio-client            | epoch |   8.slc6   |    9.slc6   |
|            castor-ns-client             | epoch |   8.slc6   |    9.slc6   |
|          castor-stager-client           | epoch |   8.slc6   |    9.slc6   |
|              castor-devel               | epoch |   8.slc6   |    9.slc6   |
|              castor-devel               |  arch |    i686    |    x86_64   |
+-----------------------------------------+-------+------------+-------------+
```

More ongoing work

Moving all the Puppet code **to GitLab**

A better “merge request-based” workflow

Maturing our **Continuous Integration workflow**

On top of Jira and integrated with GitLab

Automating procedures with **Rundeck**

More granularity on complex operations

Easiness to **delegate actions** to shifters and Service Desk

More ongoing work

"*tellme*" : A **new root password generator** service

- Generated password **expire after some time**
- Useful for **on-call operators**

Foreman **upgrade to 1.9** and CC7

- Adapting our tools to changes on the API

Puppet balancers to **HAProxy** and CC7

- Make client **sticky to a specific puppet master**

www.cern.ch