

# A not-so-bumpy road to Puppet 4 at CERN


nacho.barrientos@cern.ch  
CERN IT Department



On the French-Swiss border  
22 member states  
Collaborators from all around the world  
What do we do?








# Large Hadron Collider (LHC)





CMS

LHCb

ALICE

ATLAS






# Compact Muon Solenoid (CMS)



<https://home.cern>  
<https://visit.cern>















**STORAGETEK**









## Service Availability Overview

[Public screens view](#) | [Minimal view](#) | [FAQ](#)

Service Available Service Degraded Service Unavailable Informational Message No Information

### Batch Services

- Batch
- BOINC
- HPC

### Collaboration Services

- Conference Rooms
- E-Mail
- Eduroam
- Lync
- Sharepoint

### Computer Security Services

- Certificate Authority
- Single Sign On and Account Management

### Data Analytics Services

- HADOOP

### Database Services

- Accelerator Database
- Administration Database
- Database on Demand
- Database Replication
- Experiment Database
- General Purpose Database

### Desktop Services

- Windows Desktop

### Development Services

- Git
- JIRA
- SVN

### Document Management Services

- CDS

### Engineering Software Services

- Electronics Design Software
- Mathematics Software
- Mechanical Design Software

### GRID Services

- File Transfer
- GRID Compute Element
- GRID Development
- GRID Information
- GRID Infrastructure Monitoring
- MyProxy
- VOMS

### Infrastructure Application Services

- Indico Event Application Support

### Interactive Services

- LXPLUS
- Windows Terminal Servers

### IT Infrastructure Services

- ACRON
- Centralised Elasticsearch
- Configuration Management
- Linux Operating System
- Load Balancing
- Messaging
- Monitoring
- Server Provisioning

### Network Services

- Campus Network
- CIXP
- Datacenter Network
- Network Database and Registration
- Network for Projects and Experiments
- Technical Network
- WIFI
- WLCG Network

### Printing Services

- Printing and Copying

### Storage Services

- AFS
- Backup and Restore
- CASTOR
- Ceph
- CERNBox
- CVMFS
- DFS
- EOS
- FILER

### Telephone Services

- E-Fax
- Fixed Line Phone

### Text and Media Services

- Alerter
- MultiMedia
- Public Information Display

### Web Services

- AFS Web Hosting
- CERN Search
- Databases Applications
- Drupal
- IIS Web Hosting
- Java Web Hosting
- PaaS Web Application Hosting
- Twiki



**34k** active nodes in PuppetDB

Base catalog contains **1.2k** resources

**350** catalogs compiled per minute

**191** changes\* per day

**350** puppeteers





**CONFIG TEAM**






**CONFIG TEAM**

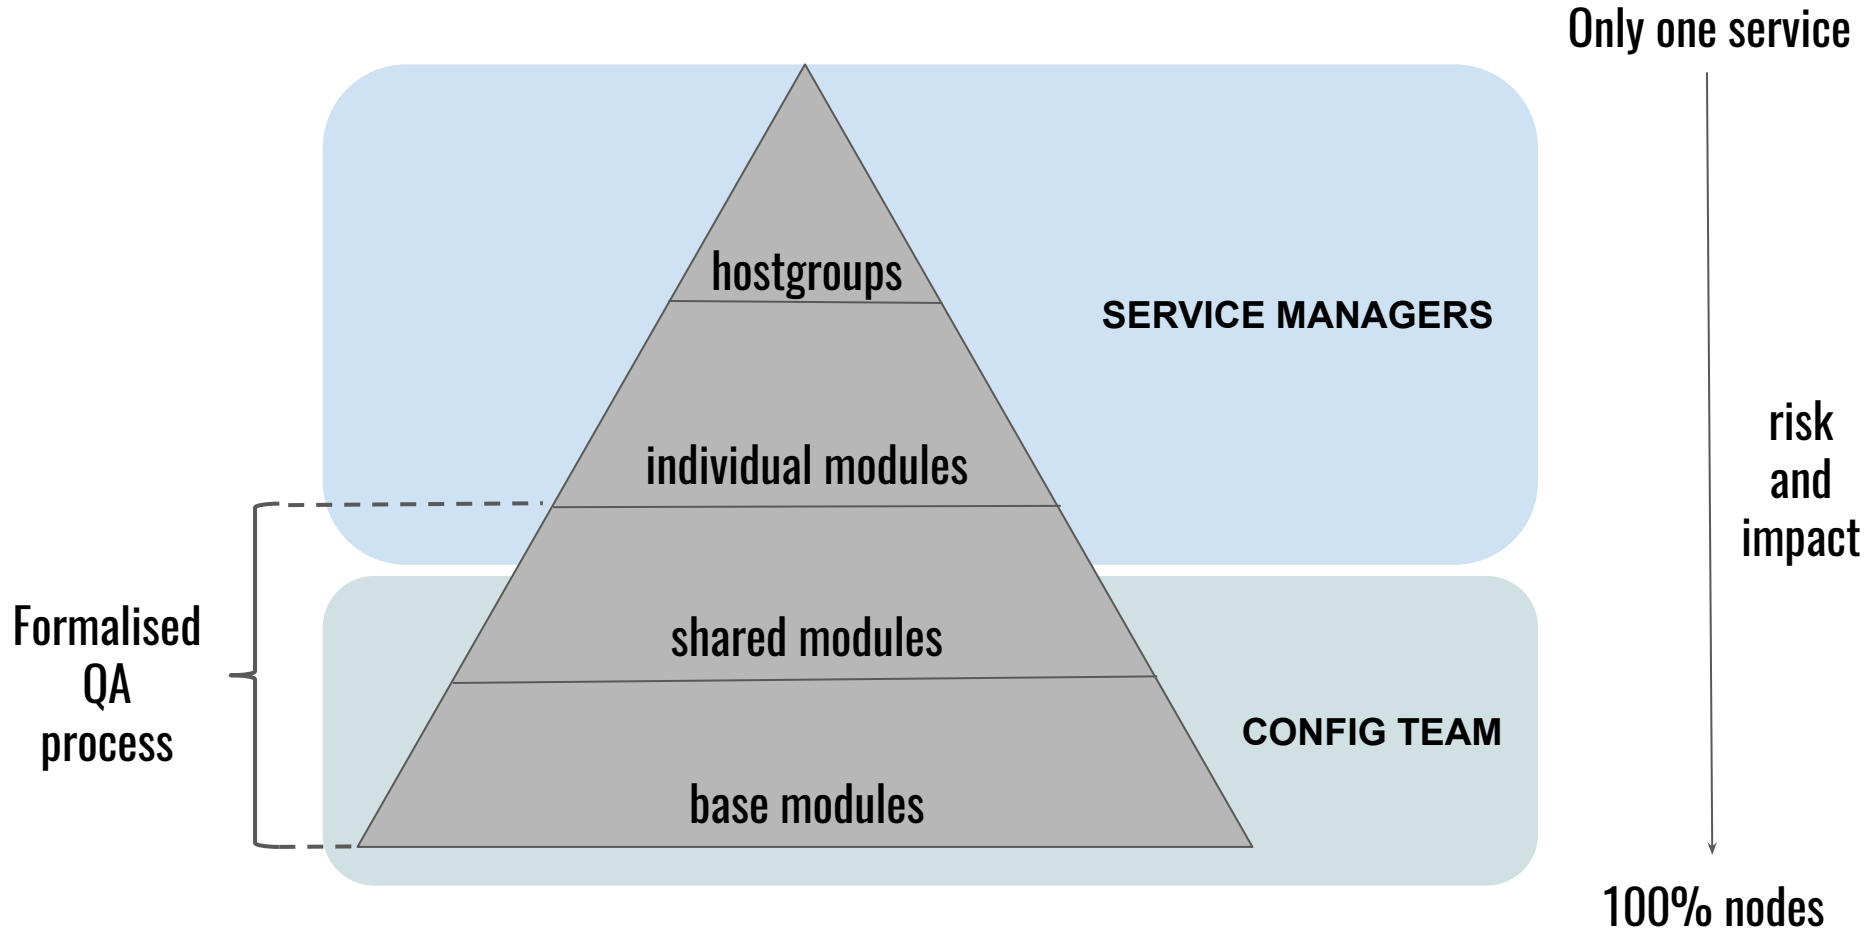


## CONFIG TEAM



## SERVICE MANAGERS






# Changes to QA are announced publicly

## Typical QA duration: 1 week

## All service managers can stop them

 Config Release Management / CRM-1943

### custom datacentre fact - new network ITS is meyrin.

EditCommentAssignMore

Merged into ProductionProblem found in QAReject

Export

#### Details

Type:Configuration Change

Status:III QA (View Workflow)

Resolution:Unresolved

Security Level:Internal Data (Only authenticated CERN users can see this issue)

Labels:custom

Change Type:New feature

Change scope:Manifest change

Feature Branch:[https://gitlab.cern.ch/ai/it-puppet-module-custom/merge\\_requests/8](https://gitlab.cern.ch/ai/it-puppet-module-custom/merge_requests/8)

Changelog:

- Trivial change, if the ITS landb network is encountered this is also in the Meyrin datacentre.

ITS = "Domain ITS, IT services (old GPN in data center)"  
<https://network.cern.ch/sc/fcgi/sc.fcgi?Action=Describe&Starpoint=513-C>

#### People

Assignee:Steve Traylen  
[Assign to me](#)

Reporter:Steve Traylen

Votes:1 [Vote for this issue](#)

Watchers:2 [Start watching this issue](#)

#### Dates

Created:1 week ago

Updated:5 hours ago

Proposed date for Production:08/Nov/16



*Fork me on GitHub*

<https://github.com/cernops/jens>

**7450 ESS-6**  
Ethernet Service Switch

CERN

Stat  
Pwr

Lnk  
Act

**Our trip from Puppet 3 to Puppet 4**



# The preps...



**Parallel server infrastructure**




**CI framework**



**Base configuration dual compatibility**



**Planning and communication**







1

Volunteer runs against Puppet 4 servers  
CI preaching

2

3

# When the ground started to shake

## Castings that became errors

```
2. if $::operatingsystemmajrelease <= 6 {  
3.   ...  
4. }
```

## Changed boolean interpretation ("")

```
3. if $variable {  
4.   ... you are running puppet 4.  
5. } else {  
6.   ... you are running puppet 3.  
7. }
```

## Variable scope and nested classes

```
class outer {  
  $var = 'dynamic'  
  include inner  
}
```

```
class inner {  
  notice(inline_template('<%= @var %>'))  
}
```

## Hiera variable names (dots)

```
$stream_prune_log_max = hiera('snort.stream_prune_log_max', '0')
```



```
34 task test: [  
35     :metadata_lint,  
36     :syntax,  
37     :lint,  
38     :spec,  
39 ]
```

Status

Build ID

Name

✓ Tests

#201467

✓ passed

docker

puppet3

🕒 01:09

📅 36 minutes ago

#201468

✓ passed

docker

puppet4

🕒 01:12

📅 36 minutes ago

#201469

✗ failed

docker

allowed to fail

puppet4\_latest

🕒 00:51

📅 36 minutes ago

1

Volunteer runs against Puppet 4 servers  
CI preaching  
Bugging people!

2

3



# RECENT ACTIVITY IN PUPPET LAND

Bigger is more recent. Includes topic branches. Modules are represented in blueish, hostgroups in yellowish.

Countdown to Puppet 4 in Q4


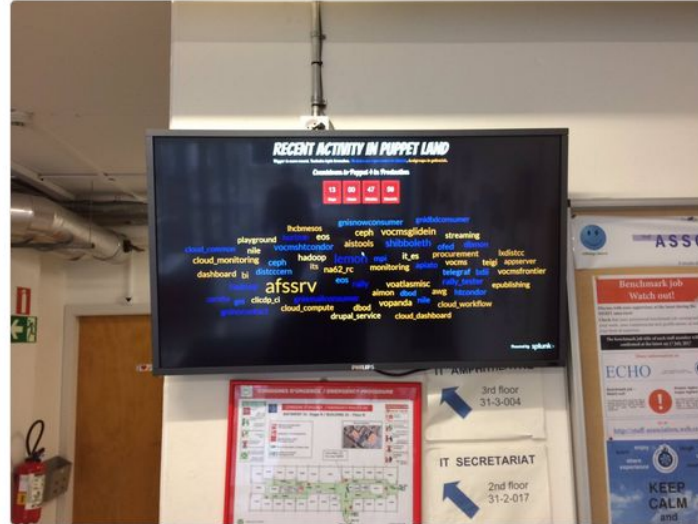




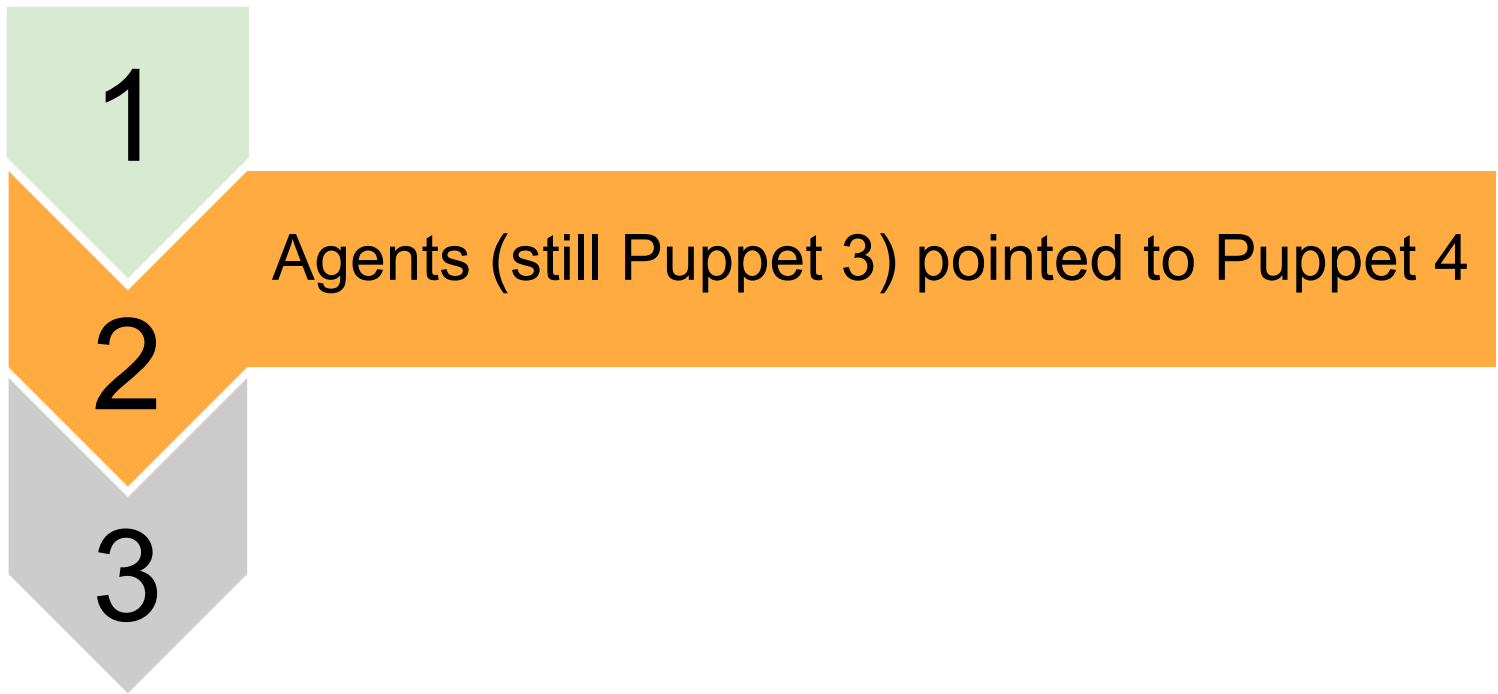
**Belmiro Moreira**  
@belmiromoreira

Following

Countdown for [#puppet](#) 4 in production at [#CERN](#)








(in the meantime)

## Preparing the agent upgrade...



(in the meantime)

# Preparing the agent upgrade...

Custom facts (Facter 2.x → Facter 3.x)

```
m/params.pp
1 class landb::params {
2
3   if versioncmp($::puppetversion, '4') >= 0 {
4     $default_landb_config = '/etc/puppetlabs/puppet/landb_config.yaml'
5   } else {
6     $default_landb_config = '/etc/puppet/landb_config.yaml'
7   }
8 }

```


```

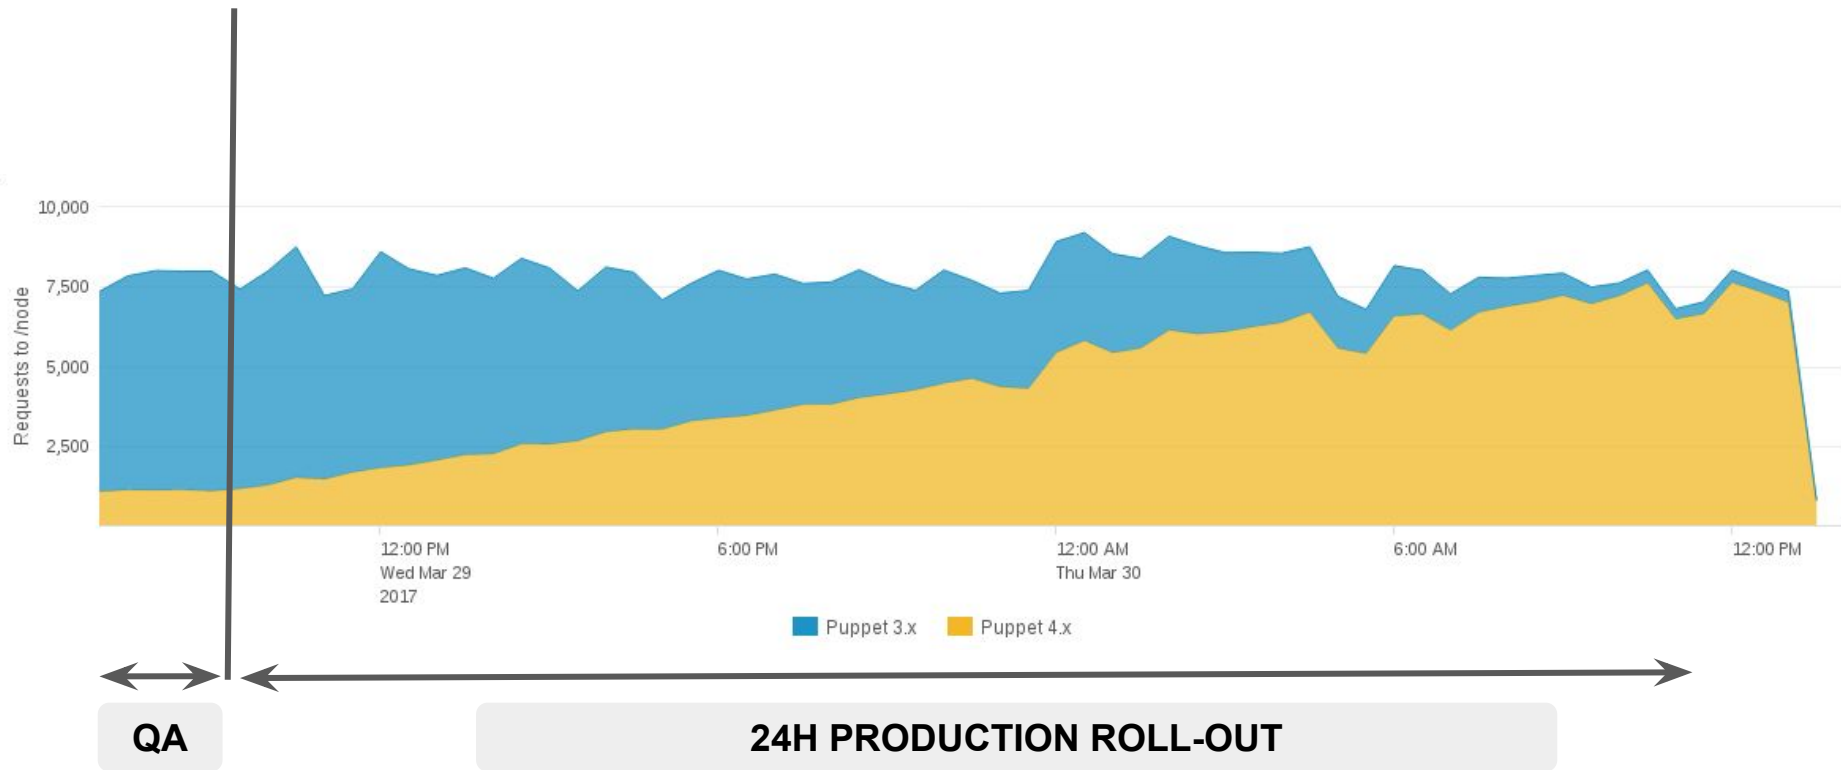
if Gem::Version.new(Facter[:puppetversion].value) >= Gem::Version.new('4')
  config_path = '/etc/puppetlabs/puppet/landb_config.yaml'
else
  config_path = '/etc/puppet/landb_config.yaml'
end

```


Dual configuration ([/etc/puppet, /etc/mcollective] ↔ /etc/puppetlabs/\*)







3.5 months






# And that 5%?


broken catalog, Puppet disabled, broken software updates...

please rebuild/reinstall 😊



# Unable! 🕶️

unpin software repositories,  
fix dependency problems,  
copy puppet.conf and run Puppet...



...usually helped

# What did go wrong?

## Double-check to make sure it's safe before purging cron resources

Previously, using `resources { 'cron': purge => true }` to purge `cron` resources would only purge jobs belonging to the current user performing the Puppet run (usually `root`). In Puppet 4, this action is more aggressive and causes **all** unmanaged cron jobs to be purged.

Make sure this is what you want. You might want to set `noop => true` on the purge resource to keep an eye on it.

## User crontabs removed from CVMFS Stratum 0 release manager machines

**Type:**

Service Incident

**Incident Status:**

Resolved

**Begin:**

Wed, Mar 29, 2017 13:48

**Service Impact:**

Degraded

### Sixtrack user cronjobs lost

**SE: BOINC Service, FE: BOINC Support,**

As a consequence of the change of Puppet behaviour to purge all non-puppet managed cron jobs (CRM-2188),

**Root Cause:** Puppet 3 to Puppet 4 upgrade, and change of default behaviour with respect to cron. (CRM-2188), not detected in QA,

### cron jobs removed for AIS

**SE: Oracle Web Hosting Service, FE: Application Middleware Infrastructure,**

As a side effect of Puppet 4 agent installation, crontab jobs have been removed from AIS main job scheduling machine.



# What could we have done better?

## Earlier CI


Module or Hostgroup	Master Brach	QA Branch	Modulesync Branch
it-puppet-hostgroup-acronsvc	build success	build success	build success
it-puppet-hostgroup-afssrv	build success	build success	build unknown
it-puppet-hostgroup-aiadm	build success	build success	build success
it-puppet-hostgroup-aimon	build success	build success	build failed
it-puppet-hostgroup-appserver	build success	build success	build success
it-puppet-hostgroup-backup	build success	build success	build success
it-puppet-hostgroup-batchmon	build unknown	build unknown	build failed

it-puppet-module-security	build success	build success	build unknown
it-puppet-module-selinux	build success	build success	build success
it-puppet-module-sendmail	build success	build success	build success
it-puppet-module-serf	build failed	build failed	build failed
it-puppet-module-shibboleth	build success	build success	build unknown
it-puppet-module-singularity	build success	build success	build success
it-puppet-module-slurm	build success	build success	build success

Catalog diffs (octocatalog-diff?)



# Ongoing activities



## SERVICE MANAGERS





# Multi-tenancy

Packing problem  
Auto-scaling  
Auto-creation  
Low maintenance



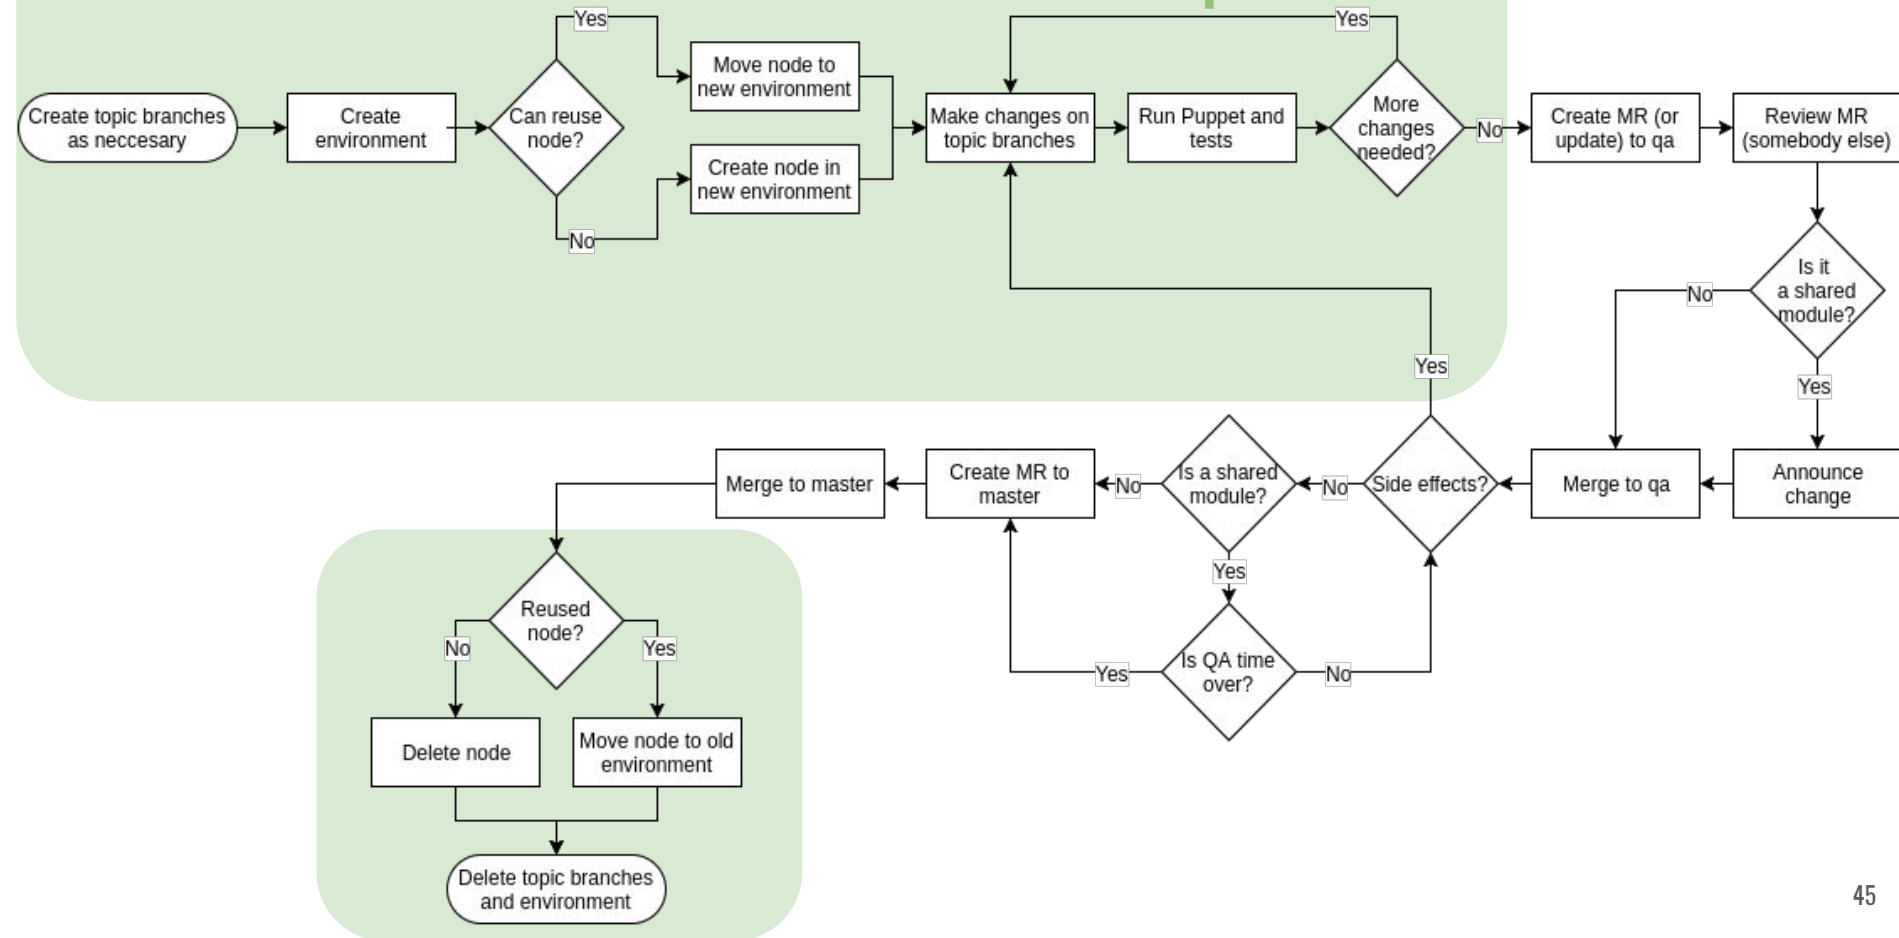


[www.cern.ch](http://www.cern.ch)

<https://cern.ch/nacho/puppetconf2017>

# Backup slides

# but ... what about development?







```
ibarrien@aiadm13: (master)
~/it-puppet-environments $ cat ai4632.yaml
```

```
---
default: qa
notifications: nacho.barrientos@cern.ch
overrides:
  modules:
    flume: monit564
  hostgroups:
    punch: ai4632
```

```
ibarrien@aiadm13: (master)
~/it-puppet-environments $ ai-foreman -l punch/aijens showhost
```

Name	Hostgroup	Environment	OS
-dev01.cern.ch	punch/aijens/b/live/secondary	ai4632	CentOS 7.2
101.cern.ch	punch/aijens/a/live/primary	production	CentOS 7.2
103.cern.ch	punch/aijens/a/live/secondary	production	CentOS 7.2
112.cern.ch	punch/aijens/b/live/primary	production	CentOS 7.2
113.cern.ch	punch/aijens/b/live/secondary	qa	CentOS 7.2

```
ibarrien@aiadm13: (master)
~/it-puppet-environments $ █
```

```
overrides:
```

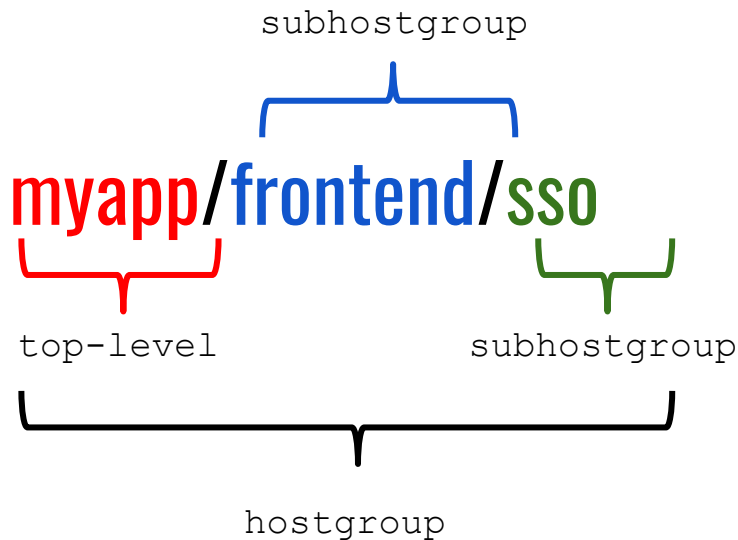
```
hostgroups:
```

```
database: master
```

```
modules:
```

20,1

# A node belongs to a hostgroup





backup/server/tape/lib0  
punch/fore/balance/external  
cloud\_compute/level2/kvm/gva\_project\_002

it-puppet-hostgroup-myapp

code

*manifests*

**init.pp**

frontend.pp

*frontend*

sso.pp

data

*hostgroup*

**myapp.yaml**

*myapp*

frontend.yaml

*frontend*

sso.yaml

**myapp**

The diagram illustrates a mapping from a hierarchical file structure to a target. On the left, a box contains a tree-like structure of files and directories. Two red arrows originate from this structure: one from 'init.pp' and another from 'myapp.yaml'. Both arrows curve and point towards the word 'myapp' on the right, which is rendered in a large, bold, red font.

it-puppet-hostgroup-myapp

code

manifests

init.pp

frontend.pp

frontend

sso.pp

data

hostgroup

myapp.yaml

myapp

frontend.yaml

frontend

sso.yaml

myapp/frontend

The diagram illustrates the mapping of files from a Puppet manifest to a directory structure. On the left, a box contains the manifest 'it-puppet-hostgroup-myapp' with sections 'code' and 'data'. The 'code' section lists 'manifests/init.pp', 'manifests/frontend.pp', 'manifests/frontend', and 'manifests/sso.pp'. The 'data' section lists 'data/hostgroup/myapp.yaml', 'data/hostgroup/myapp', 'data/hostgroup/frontend.yaml', 'data/hostgroup/frontend', and 'data/hostgroup/sso.yaml'. On the right, the directory structure 'myapp/frontend' is shown. Red dashed arrows indicate that 'init.pp' maps to 'myapp' and 'myapp.yaml' maps to 'myapp'. Blue solid arrows indicate that 'frontend.pp', 'frontend', and 'frontend.yaml' all map to 'frontend', while 'sso.pp' and 'sso.yaml' map to 'sso'.

it-puppet-hostgroup-myapp

code

manifests

init.pp

frontend.pp

frontend

sso.pp

data

hostgroup

myapp.yaml

myapp

frontend.yaml

frontend

sso.yaml

myapp/frontend/sso

```
graph LR; subgraph Puppet_Structure [it-puppet-hostgroup-myapp]; direction TB; C[code]; C --- M[manifests]; M --- IP[init.pp]; M --- FP[frontend.pp]; M --- F[frontend]; M --- SP[sso.pp]; D[data]; D --- HG[hostgroup]; HG --- MYA[myapp.yaml]; HG --- MY[myapp]; MY --- FEA[frontend.yaml]; MY --- FR[frontend]; MY --- SSA[sso.yaml]; end; IP -.-> myapp; FP -.-> frontend; SP -.-> sso; MYA -.-> myapp; FEA -.-> frontend; SSA -.-> sso; C -.-> Path[myapp/frontend/sso];
```

### it-puppet-hostgroup-myapp

*code*

*manifests*

*init.pp*

***frontend.pp***

*frontend*

*sso.pp*

*data*

*hostgroup*

*myapp.yaml*

*myapp*

*frontend.yaml*

*frontend*

*sso.yaml*

### it-puppet-module-apache

*files*

*lib*

*manifests*

*init.pp*

*vhost.pp*

*...*

*templates*

### frontend.pp

```
class hg_myapp::frontend {  
  ...  
  class {'apache':  
    server_signature => 'off'  
    ...  
  }  
  apache::vhost {  
    ...  
  }  
  ...  
}
```



## it-puppet-hostgroup-myapp

*code*

*manifests*

*init.pp*

*frontend.pp*

*frontend*

**sso.pp**

*data*

*hostgroup*

*myapp.yaml*

*myapp*

*frontend.yaml*

*frontend*

**sso.yaml**

*fqdns*

*appc003.cern.ch.yaml*

## sso.pp

```
class hg_myapp::frontend::sso {  
  ...  
    include shibboleth  
  ...  
}
```

## sso.yaml

```
shibboleth::session_storage: 'memcache'  
shibboleth::memcached_servers:  
  - foomc01.cern.ch  
  - foomc02.cern.ch  
shibboleth::memcached_prefix: "MY_PREFIX:"  
shibboleth::remote_user: 'ADFS_LOGIN'  
shibboleth::application_name: foo_application  
shibboleth::site_name: foo.bar.cern.ch
```

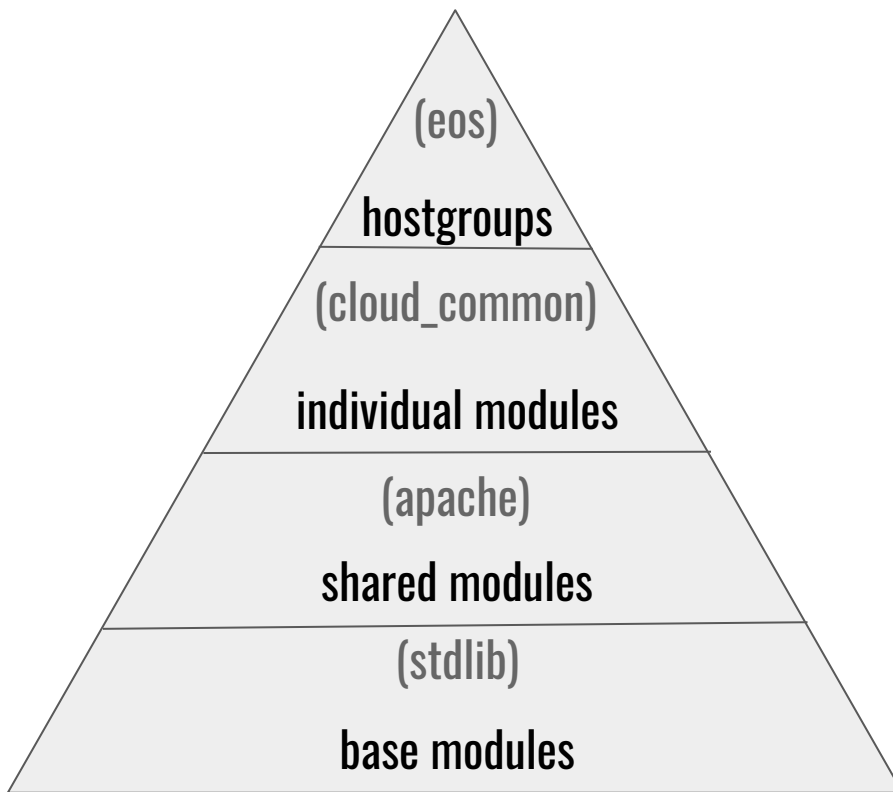
**393 modules (upstream and homegrown)**

**206 top-level hostgroups (services)**

Only one service

risk  
and  
impact

100% nodes



```
[root@nachodev06 ~]# yum clean all
Loaded plugins: kernel-module, priorities, rpm-warm-cache, security
Cleaning repos: ai6-qa ai6-stable epel epel-testing lemon6-qa lemon6-testing
               : slc6-testing slc6-updates
Cleaning up Everything
[root@nachodev06 ~]# puppet --version
3.8.7
[root@nachodev06 ~]# /usr/local/sbin/distro_sync.sh
[root@nachodev06 ~]# puppet
-bash: /usr/bin/puppet: No such file or directory
[root@nachodev06 ~]# source /etc/profile.d/puppet-agent.sh
[root@nachodev06 ~]# puppet --version
4.8.1
[root@nachodev06 ~]#
```

Or just log in and out!

