

# Field Experience in the Agile Infrastructure

## HEPiX Spring 2014

Jérôme Belleman  
CERN IT  
May 2014

- 1 Comfortably Running Services
- 2 Challenges
- 3 Opportunities

## Section 1

### Comfortably Running Services

### Cloud infrastructure (OpenStack):

- Create VMs
- Probe, reboot VMs

### Server lifecycle management (Foreman):

- Organise nodes in groups
- Report on Puppet runs

### Configuration management (Puppet):

- Editing configuration files
- Software installation, daemon management

### Monitoring, Notifications (Elasticsearch/Kibana, GNI):

- Analytics and dashboards
- Notifications

Wrapper around `nova` making VMs agile:

```
% ai-bs-vm --nova-flavor m1.large  
-g bi/batch/blah  
-i "SLC6 CERN Server" foo
```

Puppetised + Foremanned + Certified + Alarmed

Wrapper around `nova` making VMs agile:

```
% for i in {00..99}; ai-bs-vm --nova-flavor m1.large  
-g bi/batch/blah  
-i "SLC6 CERN Server" foo$i
```

Puppetised + Foremanned + Certified + Alarmed

- Reliably getting console dumps
- Grep for *Kernel panic* in a console

```
% for h in foo42 foo43;  
    (nova console-log $h | grep 'Kernel panic') &&  
    nova reboot $h
```

- Reliably rebooting VMs
- Removing and creating VMs faster than fixing them

**FOREMAN** Jerome Belleman

Dashboard Hosts Reports Facts Audits Statistics Trends More

## Hosts

Filter ... Q Search New Host

<input type="checkbox"/>	Name	Operating system	Environment	Model	Host group	Last report	Owner	
<input type="checkbox"/>	batch00.cern.ch	SLC 6.5	production	OpenStack Nova	bi/batch/gridworker/ais...	about 1 hour ago	batch-3rd	<span>Edit</span>
<input type="checkbox"/>	batch01.cern.ch	SLC 6.5	qa	OpenStack Nova	bi/batch/gridworker/ais...	43 minutes ago	batch-3rd	<span>Edit</span>
<input type="checkbox"/>	batch02.cern.ch	SLC 6.5	qa	OpenStack Nova	bi/batch/gridworker/ais...	18 minutes ago	batch-3rd	<span>Edit</span>
<input type="checkbox"/>	batch03.cern.ch	SLC 6.5	production	OpenStack Nova	bi/batch/gridworker/ais...	about 1 hour ago	batch-3rd	<span>Edit</span>
<input type="checkbox"/>	batch04.cern.ch	SLC 6.5	production	OpenStack Nova	bi/batch/gridworker/ais...	21 minutes ago	batch-3rd	<span>Edit</span>
<input type="checkbox"/>	batch05.cern.ch	SLC 6.5	qa	OpenStack Nova	bi/batch/gridworker/ais...	27 minutes ago	batch-3rd	<span>Edit</span>
<input type="checkbox"/>	batch06.cern.ch	SLC 6.5	production	OpenStack Nova	bi/batch/gridworker/ais...	about 1 hour ago	batch-3rd	<span>Edit</span>
<input type="checkbox"/>	batch07.cern.ch	SLC 6.5	production	OpenStack Nova	bi/batch/gridworker/ais...	5 days ago	batch-3rd	<span>Edit</span>

Displaying all 8 entries - 0 selected



- REST API → `ai-foreman-cli`
- Shuffling resources
- Retiring machines

```
% ai-foreman-cli updatehost  
  --hostgroup bi/batch/gridworker/terminate foo
```

- Change environment

```
% ai-foreman-cli updatehost --environment qa foo
```

- REST API → ai-foreman-cli
- Shuffling resources
- Retiring machines

```
% for i in {00..99}; ai-foreman-cli updatehost  
  --hostgroup bi/batch/gridworker/terminate foo$i
```

- Change environment

```
% ai-foreman-cli updatehost --environment qa foo
```

Feature branch → QA → Production

- Several safety steps
- Better communication
- Incentive to try

- Grep around for inspiration
- `git push`  $\xrightarrow{1 \text{ min}}$  Puppet run  $\xrightarrow{3 \text{ min to } 1 \text{ h}}$  Done
- Security fixes

- `git rebase -i`
- `git diff/git log -p`
- ```
% git log --graph --oneline
* 0dad81e Add Condor config file
* cac33eb add lemon monitoring
* 4c0dce7 move LSF instance definition one down
* 00a1cd7 Merge branch 'qa'
| \
| * 5189e63 change masters for batchint
* | 0496b36 Merge branch 'qa'
| \ \
| | /
| * e579c7e batchint setup
* | 1474105 Merge branch 'qa'
| \ \
| | /
| * b15e545 GGUS #102823 SLC5 only
| * 0d1085f Bump osdate
```

**FOREMAN** Jerome Bellemen More

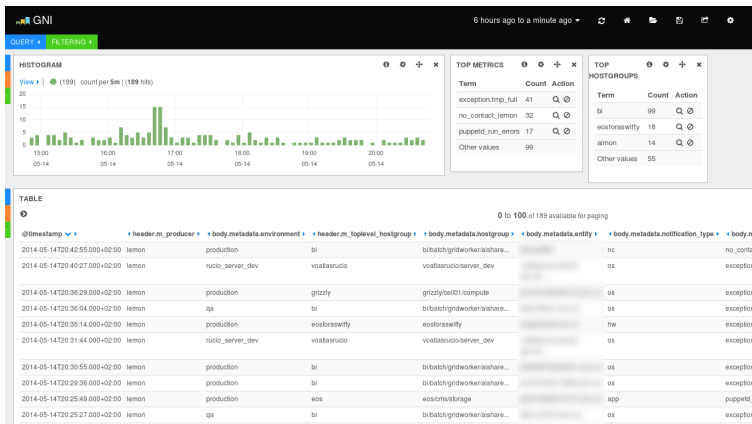
Dashboard Hosts Reports Facts Audits Statistics Trends

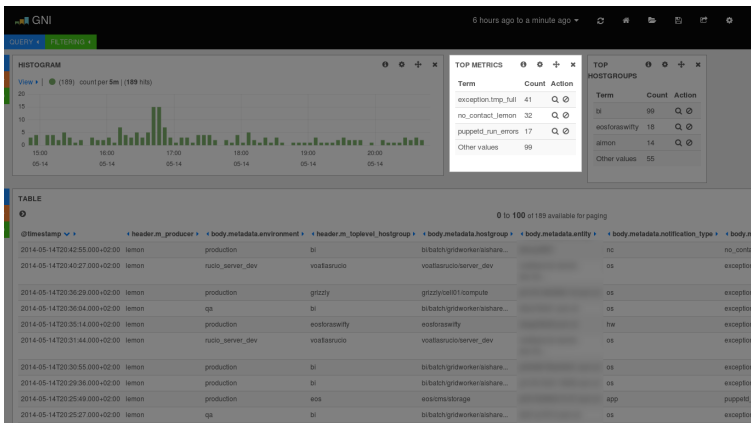
## Hosts

Filter ... Q Search New Host











| <input type="checkbox"/> | Name            | Operating system | Environment | Model          | Host group                 | Last report      | Owner     |                   |
|--------------------------|-----------------|------------------|-------------|----------------|----------------------------|------------------|-----------|-------------------|
| <input type="checkbox"/> | batch00.cern.ch | SLC 6.5          | production  | OpenStack Nova | bi/batch/gridworker/ais... | about 1 hour ago | batch-3rd | <span>Edit</span> |
| <input type="checkbox"/> | batch01.cern.ch | SLC 6.5          | qa          | OpenStack Nova | bi/batch/gridworker/ais... | 43 minutes ago   | batch-3rd | <span>Edit</span> |
| <input type="checkbox"/> | batch02.cern.ch | SLC 6.5          | qa          | OpenStack Nova | bi/batch/gridworker/ais... | 18 minutes ago   | batch-3rd | <span>Edit</span> |
| <input type="checkbox"/> | batch03.cern.ch | SLC 6.5          | production  | OpenStack Nova | bi/batch/gridworker/ais... | about 1 hour ago | batch-3rd | <span>Edit</span> |
| <input type="checkbox"/> | batch04.cern.ch | SLC 6.5          | production  | OpenStack Nova | bi/batch/gridworker/ais... | 21 minutes ago   | batch-3rd | <span>Edit</span> |
| <input type="checkbox"/> | batch05.cern.ch | SLC 6.5          | qa          | OpenStack Nova | bi/batch/gridworker/ais... | 27 minutes ago   | batch-3rd | <span>Edit</span> |
| <input type="checkbox"/> | batch06.cern.ch | SLC 6.5          | production  | OpenStack Nova | bi/batch/gridworker/ais... | about 1 hour ago | batch-3rd | <span>Edit</span> |
| <input type="checkbox"/> | batch07.cern.ch | SLC 6.5          | production  | OpenStack Nova | bi/batch/gridworker/ais... | 5 days ago       | batch-3rd | <span>Edit</span> |

Displaying all 8 entries - 0 selected







| TOP METRICS        |       |                                                                                                                                                                         |  |  |  |  |
|--------------------|-------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| Term               | Count | Action                                                                                                                                                                  |                                                                                   |                                                                                   |                                                                                     |                                                                                     |
| exception.tmp_full | 41    |   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |
| no_contact_lemon   | 32    |   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |
| puppetd_run_errors | 17    |   |                                                                                   |                                                                                   |                                                                                     |                                                                                     |
| Other values       | 99    |                                                                                                                                                                         |                                                                                   |                                                                                   |                                                                                     |                                                                                     |

| TOP METRICS        |       |        |  |  |  |  |
|--------------------|-------|--------|--|--|--|--|
| Term               | Count | Action |  |  |  |  |
| exception.tmp_full | 41    |        |  |  |  |  |
| no_contact_lemon   | 32    |        |  |  |  |  |
| puppetd_run_errors | 17    |        |  |  |  |  |
| Other values       | 99    |        |  |  |  |  |

GNI

6 hours ago to a minute ago

QUERY FILTERING

GRAPHS

TABLE

0 to 26 of 26 available for paging

| @timestamp                    | header_m_producer | body.metadata.environment     | header_m_toplevel_hostgroup | body.metadata.hostgroup        | body.metadata.entity | body.metadata.notification_type | body.m   |
|-------------------------------|-------------------|-------------------------------|-----------------------------|--------------------------------|----------------------|---------------------------------|----------|
| 2014-05-14T20:51:59.000+02:00 | lemon             | production                    | bi                          | bi/batch/gridworker/a/share... |                      | app                             | puppetd_ |
| 2014-05-14T20:49:08.000+02:00 | lemon             | ebonfil                       | playground                  | playground/ebonfil             |                      | app                             | puppetd_ |
| 2014-05-14T20:25:40.000+02:00 | lemon             | production                    | eos                         | eos/cms/storage                |                      | app                             | puppetd_ |
| 2014-05-14T20:14:01.000+02:00 | lemon             | \$erzero_control_dev_batch... | voafastz                    | voafastz/control_dev_batch...  |                      | app                             | puppetd_ |
| 2014-05-14T19:18:22.000+02:00 | lemon             | rtb3                          | hadoop                      | hadoop/test/datanode           |                      | app                             | puppetd_ |
| 2014-05-14T18:46:15.000+02:00 | lemon             | nucio_monitor                 | voafasrucio                 | voafasrucio/monitor            |                      | app                             | puppetd_ |
| 2014-05-14T18:33:59.000+02:00 | lemon             | production                    | bi                          | bi/batch/gridworker/a/share... |                      | app                             | puppetd_ |
| 2014-05-14T18:12:14.000+02:00 | lemon             | production                    | indico                      | indico/prod/indico/worker/s... |                      | app                             | puppetd_ |
| 2014-05-14T18:09:10.000+02:00 | lemon             | production                    | indico                      | indico/prod/indico/worker      |                      | app                             | puppetd_ |
| 2014-05-14T18:02:26.000+02:00 | lemon             | qa                            | indico                      | indico/prod/indico/worker      |                      | app                             | puppetd_ |
| 2014-05-14T17:59:58.000+02:00 | lemon             | qa                            | bi                          | bi/batch/gridworker/a/share... |                      | app                             | puppetd_ |
| 2014-05-14T17:59:30.000+02:00 | lemon             | production                    | bi                          | bi/batch/gridworker/a/share... |                      | app                             | puppetd_ |
| 2014-05-14T17:35:27.000+02:00 | lemon             | nucio_server_dev              | voafasrucio                 | voafasrucio/server_dev         |                      | app                             | puppetd_ |
| 2014-05-14T17:32:51.000+02:00 | lemon             | nucio_collector_prod          | voafasrucio                 | voafasrucio/collector_prod     |                      | app                             | puppetd_ |
| 2014-05-14T17:31:12.000+02:00 | lemon             | production                    | bi                          | bi/batch/gridworker/a/share... |                      | app                             | puppetd_ |
| 2014-05-14T17:26:44.000+02:00 | lemon             | production                    | cloud_infrastructure        | cloud_infrastructure/monits... |                      | app                             | puppetd_ |

GNI

6 hours ago to a minute ago

QUERY FILTERING














































GRAPHS

TABLE

0 to 26 of 26 available for paging

| @timestamp                    | header_m_producer | body.metadata.environment     | header_m_toplevel_hostgroup | body.metadata.hostgroup        | body.metadata.entity | body.metadata.notification_type | body.m     |
|-------------------------------|-------------------|-------------------------------|-----------------------------|--------------------------------|----------------------|---------------------------------|------------|
| 2014-05-14T20:51:59.000+02:00 | lemon             | production                    | bi                          | bi/batch/gridworker/ishare...  |                      | app                             | puppetd... |
| 2014-05-14T20:49:08.000+02:00 | lemon             | ebonfil                       | playground                  | playground/ebonfil             |                      | app                             | puppetd... |
| 2014-05-14T20:25:49.000+02:00 | lemon             | production                    | eos                         | eos/cms/storage                |                      | app                             | puppetd... |
| 2014-05-14T20:14:01.000+02:00 | lemon             | tierzero_control_dev_batch... | voastaz                     | voastaz/control_dev_batch...   |                      | app                             | puppetd... |
| 2014-05-14T19:18:22.000+02:00 | lemon             | rb3                           | hadoop                      | hadoop/testdatanode            |                      | app                             | puppetd... |
| 2014-05-14T18:46:15.000+02:00 | lemon             | nucio_monitor                 | voastasnucio                | voastasnucio/monitor           |                      | app                             | puppetd... |
| 2014-05-14T18:33:59.000+02:00 | lemon             | production                    | bi                          | bi/batch/gridworker/ishare...  |                      | app                             | puppetd... |
| 2014-05-14T18:12:14.000+02:00 | lemon             | production                    | indico                      | indico/prod/indico/worker/s... |                      | app                             | puppetd... |
| 2014-05-14T18:09:10.000+02:00 | lemon             | production                    | indico                      | indico/prod/indico/worker      |                      | app                             | puppetd... |
| 2014-05-14T18:02:26.000+02:00 | lemon             | qa                            | indico                      | indico/prod/indico/worker      |                      | app                             | puppetd... |
| 2014-05-14T17:59:58.000+02:00 | lemon             | qa                            | bi                          | bi/batch/gridworker/ishare...  |                      | app                             | puppetd... |
| 2014-05-14T17:59:39.000+02:00 | lemon             | production                    | bi                          | bi/batch/gridworker/ishare...  |                      | app                             | puppetd... |
| 2014-05-14T17:35:27.000+02:00 | lemon             | nucio_server_dev              | voastasnucio                | voastasnucio/server_dev        |                      | app                             | puppetd... |
| 2014-05-14T17:32:51.000+02:00 | lemon             | nucio_collector_prod          | voastasnucio                | voastasnucio/collector_prod    |                      | app                             | puppetd... |
| 2014-05-14T17:31:12.000+02:00 | lemon             | production                    | bi                          | bi/batch/gridworker/ishare...  |                      | app                             | puppetd... |
| 2014-05-14T17:26:44.000+02:00 | lemon             | production                    | cloud_infrastructure        | cloud_infrastructure/monit...  |                      | app                             | puppetd... |

|              |                                |  |
|--------------|--------------------------------|--|
| voatlasrucio | voatlasrucio/monitor           |  |
| bi           | bi/batch/gridworker/aishare... |  |
| indico       | indico/prod/indico/worker/s... |  |
| indico       | indico/prod/indico/worker      |  |

| 2014-05-14T18:33:59.000+02:00                                            | lemon                                                                                                                                                                                                                                                 | production                         | bi | bi/batch/gridworker/share... | app | puppetd... |
|--------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|----|------------------------------|-----|------------|
| View: <a href="#">Table</a> / <a href="#">JSON</a> / <a href="#">Raw</a> |                                                                                                                                                                                                                                                       |                                    |    |                              |     |            |
| Field                                                                    | Action                                                                                                                                                                                                                                                | Value                              |    |                              |     |            |
| @timestamp                                                               |    | 2014-05-14T18:33:59.000Z           |    |                              |     |            |
| body.metadata.entity                                                     |    |                                    |    |                              |     |            |
| body.metadata.environment                                                |    | production                         |    |                              |     |            |
| body.metadata.le_name                                                    |    | LXBATCH                            |    |                              |     |            |
| body.metadata.hostgroup                                                  |    | bi/batch/gridworker/share/share    |    |                              |     |            |
| body.metadata.metric_id                                                  |    | 33446                              |    |                              |     |            |
| body.metadata.metric_name                                                |    | puppetd_run_errors                 |    |                              |     |            |
| body.metadata.notification_type                                          |    | app                                |    |                              |     |            |
| body.metadata.state                                                      |    | open                               |    |                              |     |            |
| body.metadata.timestamp                                                  |    | 1400085239                         |    |                              |     |            |
| body.metadata.uuid                                                       |    | 8982238-d85-11e3-bd41-02163e00c6b6 |    |                              |     |            |
| body.metadata.validity                                                   |    | 24                                 |    |                              |     |            |
| body.payload.entity                                                      |    |                                    |    |                              |     |            |
| body.payload.metric_id                                                   |    | 33446                              |    |                              |     |            |
| body.payload.timestamp                                                   |    | 1400085239                         |    |                              |     |            |

|                           |        |                                   |
|---------------------------|--------|-----------------------------------|
| body.metadata.environment | Q Ø ■■ | production                        |
| body.metadata.fe_name     | Q Ø ■■ | LX BATCH                          |
| body.metadata.hostgroup   | Q Ø ■■ | bi/batch/gridworker/aishare/share |
| body.metadata.metric_id   | Q Ø ■■ | 33446                             |
| body.metadata.metric_name | Q Ø ■■ | puppetd_run_errors                |

|                           |        |                                   |
|---------------------------|--------|-----------------------------------|
| body.metadata.environment | Q Ø ■■ | production                        |
| body.metadata.fe_name     | Q Ø ■■ | LX BATCH                          |
| body.metadata.hostgroup   | Q Ø ■■ | bi/batch/gridworker/aishare/share |
| body.metadata.metric_id   | Q Ø ■■ | 33446                             |
| body.metadata.metric_name | Q Ø ■■ | puppetd_run_errors                |



GNI 6 hours ago to a few seconds ago

QUERY FILTERING

GRAPHS

TABLE

0 to 9 of 9 available for paging

| @timestamp                    | header_m_producer | body.metadata.environment | header_m_toplevel_hostgroup | body.metadata.hostgroup        | body.metadata.entity | body.metadata.notification_type | body.me   |
|-------------------------------|-------------------|---------------------------|-----------------------------|--------------------------------|----------------------|---------------------------------|-----------|
| 2014-05-14T20:51:59.000+02:00 | lemon             | production                | bi                          | bi/batch/gridworker/aishare... |                      | app                             | puppetd_n |
| 2014-05-14T18:33:59.000+02:00 | lemon             | production                | bi                          | bi/batch/gridworker/aishare... |                      | app                             | puppetd_n |
| 2014-05-14T17:59:58.000+02:00 | lemon             | qa                        | bi                          | bi/batch/gridworker/aishare... |                      | app                             | puppetd_n |
| 2014-05-14T17:59:39.000+02:00 | lemon             | production                | bi                          | bi/batch/gridworker/aishare... |                      | app                             | puppetd_n |
| 2014-05-14T17:31:12.000+02:00 | lemon             | production                | bi                          | bi/batch/gridworker/aishare... |                      | app                             | puppetd_n |
| 2014-05-14T17:21:15.000+02:00 | lemon             | production                | bi                          | bi/batch/gridworker/aishare... |                      | app                             | puppetd_n |
| 2014-05-14T16:00:05.000+02:00 | lemon             | qa                        | bi                          | bi/batch/gridworker/aishare... |                      | app                             | puppetd_n |
| 2014-05-14T15:19:40.000+02:00 | lemon             | production                | bi                          | bi/batch/gridworker/aishare... |                      | app                             | puppetd_n |
| 2014-05-14T15:14:25.000+02:00 | lemon             | production                | bi                          | bi/batch/gridworker/aishare... |                      | app                             | puppetd_n |

0 to 9 of 9 available for paging

ADD A ROW

- Listen to a message broker for specific notifications
- Catch notifications for unreachable machines
- 

```
(nova console-log | grep 'Kernel panic') &&  
nova reboot
```

## Migrating a Squid service: 1/2 day

- Argus
- BDII
- Batch
- CEs (CREAM)
- CVMFS
- FTS
- Interactive services
- LFC
- MyProxy
- VOMS

## Section 2

### Challenges

- New twins aren't always identical
- Failed Puppet runs
- Corrupted RPM DBs

→ Just create a new one

wassh, parallel-ssh, MCollective, Puppet

```
% parallel-ssh -H'b6249f0276 b65920827e b61e68bda3
b668c3d4b4 b6f3bf325a b6627e2675 b6b72a3ecf' 'uptime'
[1] 21:20:06 [SUCCESS] b61e68bda3
[2] 21:20:06 [SUCCESS] b6627e2675
[3] 21:20:06 [SUCCESS] b6249f0276
[4] 21:20:07 [SUCCESS] b6f3bf325a
[5] 21:21:05 [FAILURE] b65920827e Timeout, Killed by signal 9
[6] 21:21:05 [FAILURE] b668c3d4b4 Timeout, Killed by signal 9
[7] 21:21:05 [FAILURE] b6b72a3ecf Timeout, Killed by signal 9
```

- Crashes
- SSH daemon down
- Heavy load

→ Failover method, automatic reboots, replace nodes

Need for automation, need for flexibility, need for CLIs:

- No transparent OpenStack tenancy → Facts, PuppetDB
- Heavy OpenStack queries ∼ Multithreading
- Inconvenient Foreman queries → Next version

Need for automation, need for flexibility, need for CLIs:

- No transparent OpenStack tenancy → Facts, PuppetDB
- Heavy OpenStack queries → Facts, PuppetDB
- Inconvenient Foreman queries → Next version



- Capped VM creation rate → Workflows
- New, Puppetised VM: 30 min → Parallelism
- I/O latencies in VMs → Different scheduler

## Section 3

# Opportunities

## Logging:

- Centralised for all AI components
- Entities which aren't there anymore
- Access control

## Documentation:

- Organise documentation
- Index documentation
- Load speed

- Fill up OpenStack with VMs
- Automatic reboots
- Draining
- Automatic removals

→ Procedures, automation

- Already happy: tools mostly mature enough for services
- Improvements needed to make it all really comfortable
- Agile Infrastructure: a source of inspiration



# Questions?