



Testing as a Service with HammerCloud

Ramón Medrano Llamas
CERN, IT-SDC

15.10.2013



Agenda

- Introduction
- New use cases
- Architectural changes
- Tests performed
- Future work

HammerCloud

- The grid testing framework
- ATLAS, CMS, LHC*b*
- 50 M jobs/year
- 20 machines cluster hosted by CERN IT
 - *voboxes*

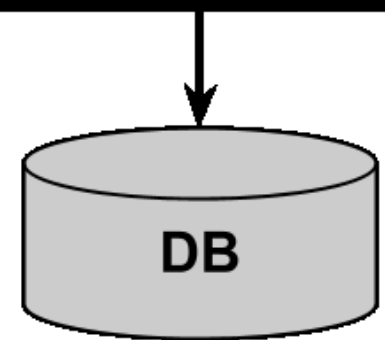
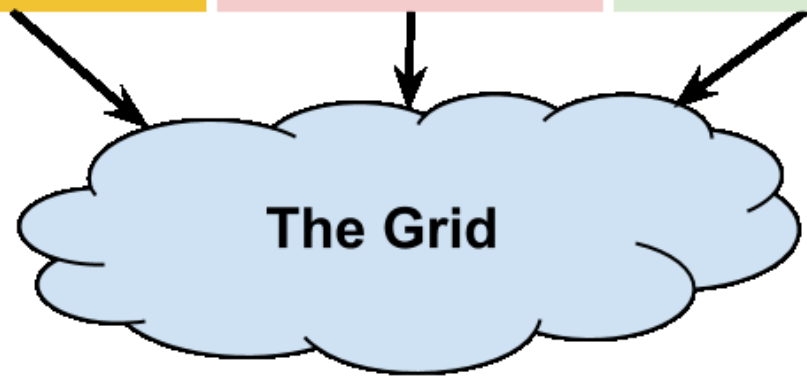
HC Testing Infrastructure

HC Web Interface

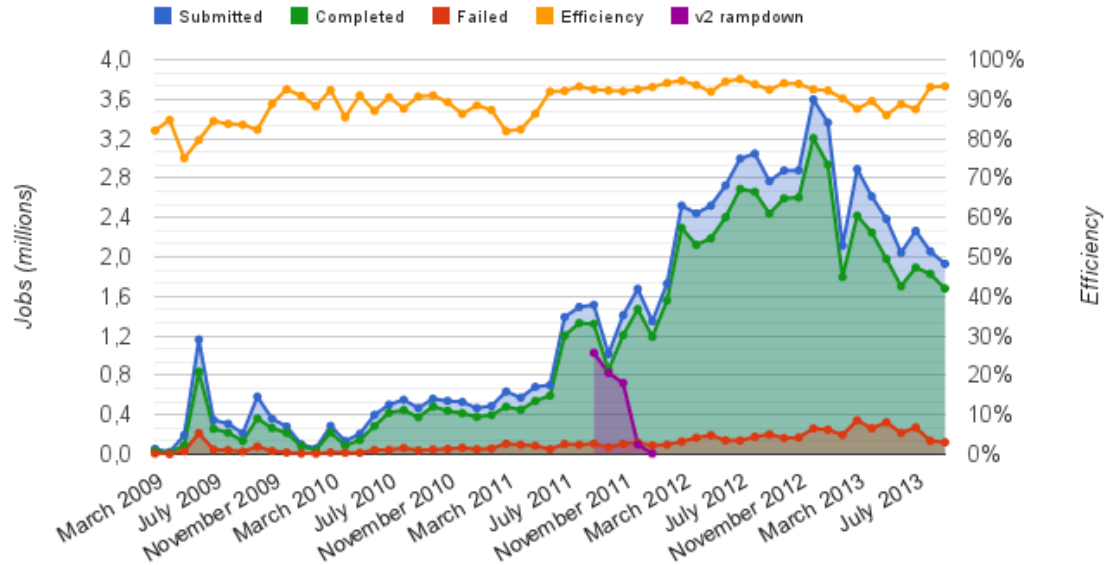


django


PanDA



HammerCloud for ATLAS job submission evolution



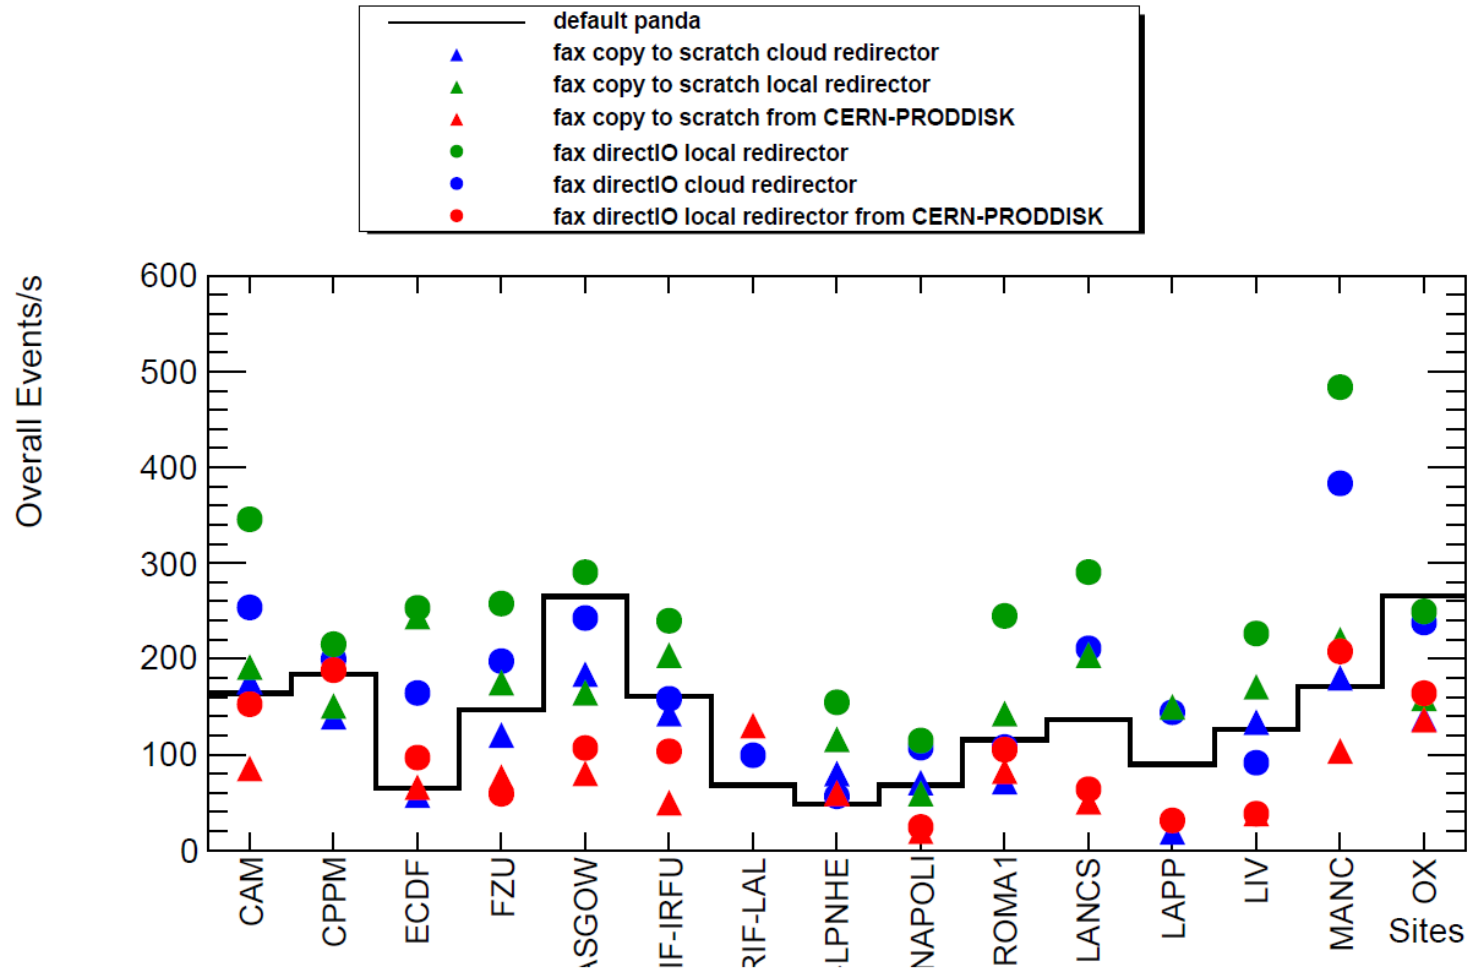
HammerCloud for CMS job submission evolution



New use cases

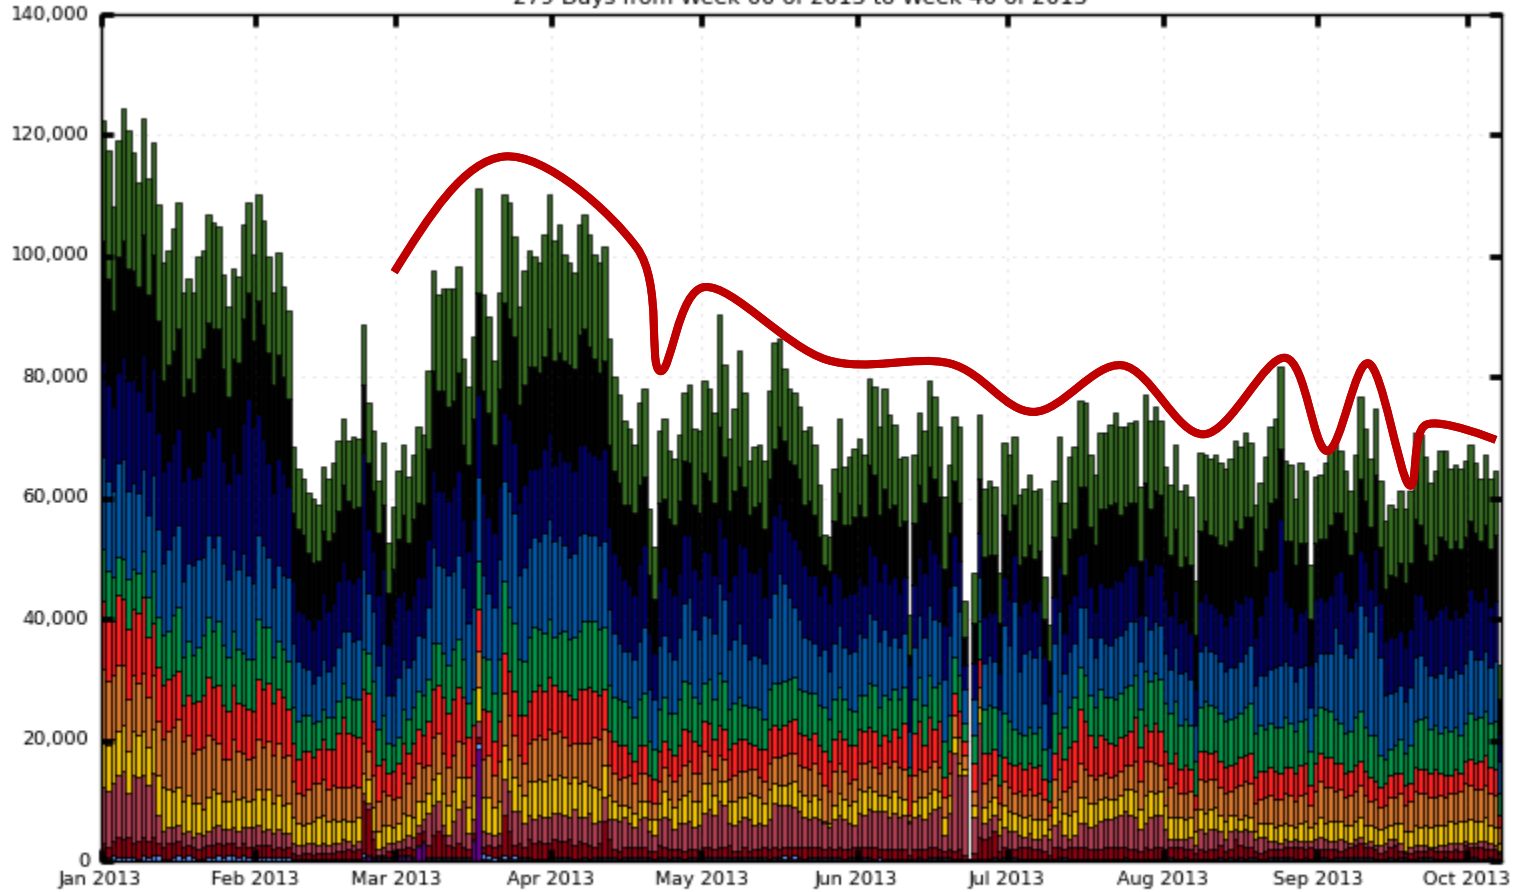
- Stress testing of sites
 - Functional testing of sites
 - AFT/PFT testing suite
- } **12,000 test/year**
- *Benchmarking testing* **NEW!**
 - *Cloud resource validation* **NEW!**
 - *Athena nightly build system* **NEW!**
 - *XRootD federation (FAX)* **NEW!**
 - *ROOT I/O and WAN tests* **NEW!**

Example of FAX system testing



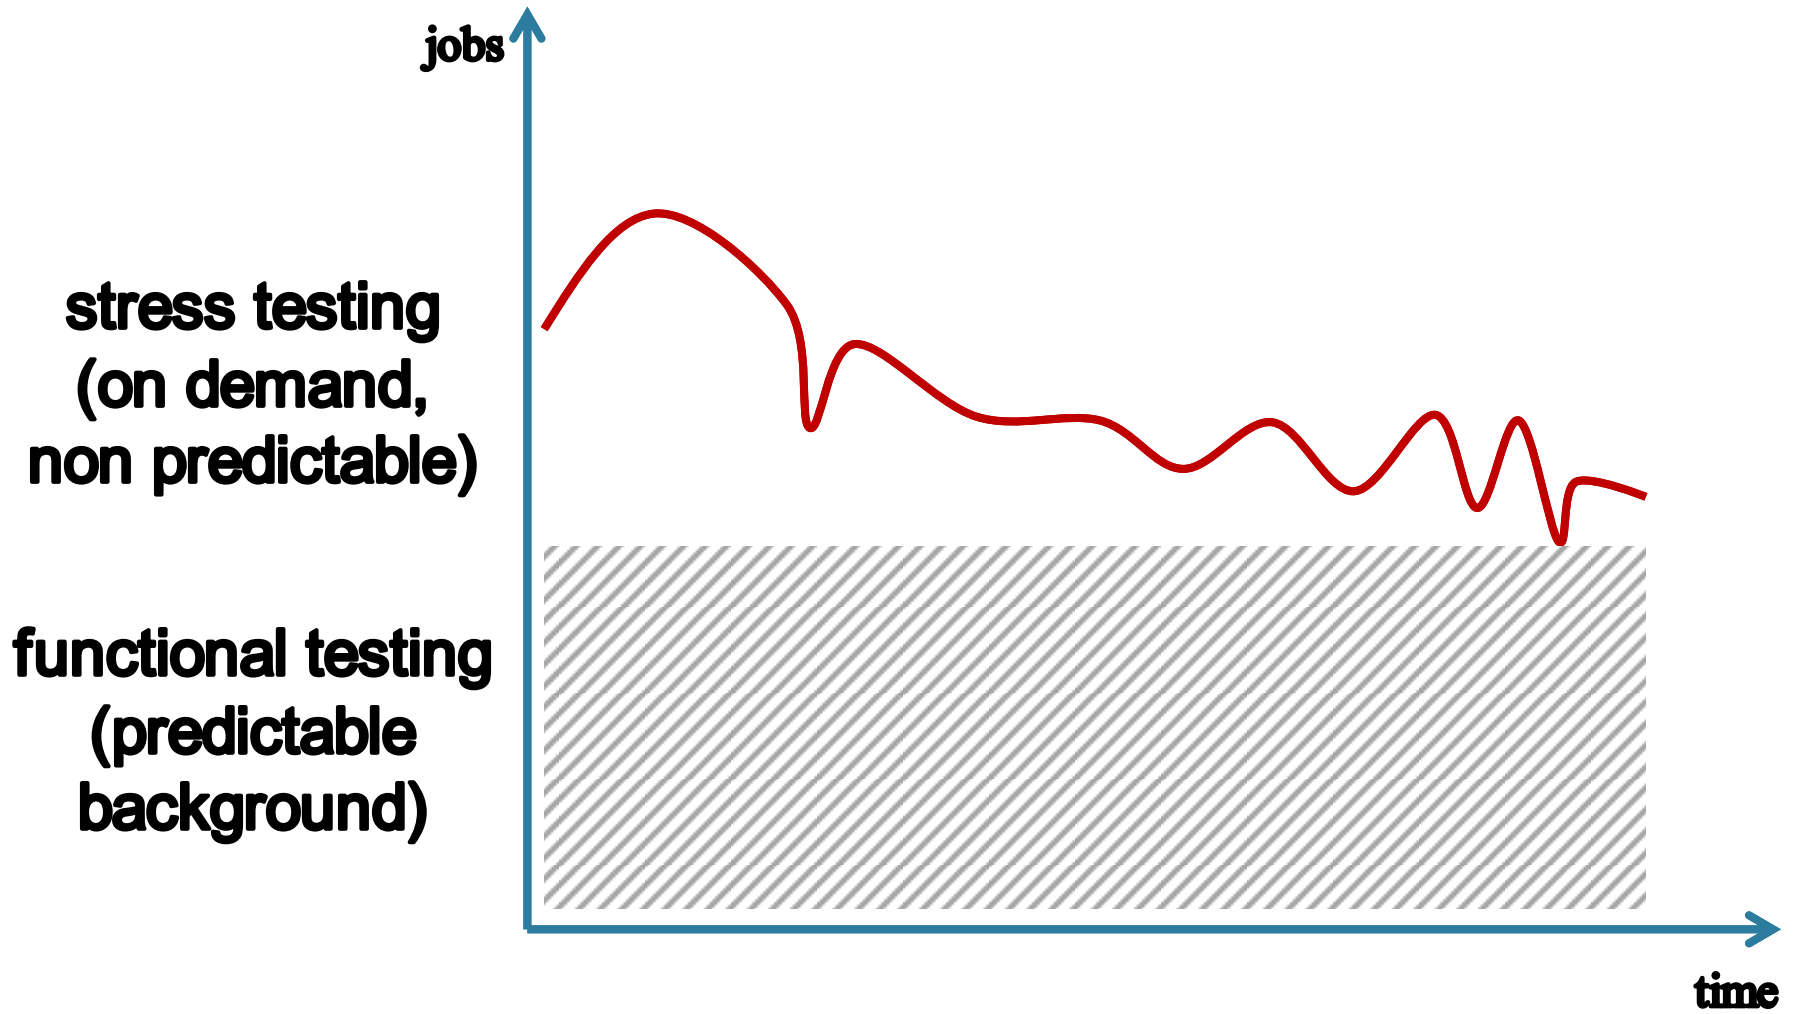
Submitted jobs

279 Days from Week 00 of 2013 to Week 40 of 2013



- UK
- DE
- US
- FR
- IT
- CA
- NL
- ES
- CERN
- ND
- RU
- Deactivated
- TW

Maximum: 124,509 , Minimum: 289.00 , Average: 76,466 , Current: 32,271



This behavior needs some *elasticity*

Testing as a Service

- Need to cope with increasing demand,
- Requested by users and tools on demand

- *Elastic* testing infrastructure
- Testing anything that allows:
 - `submit()`
 - `status()`
 - `cancel()`

Current infrastructure

- 20 machines → manually managed
- Capped scalability
- Long time to provision
- Difficult lifetime management
- Security and updates
- Challenging software deployment

The Agile Infrastructure

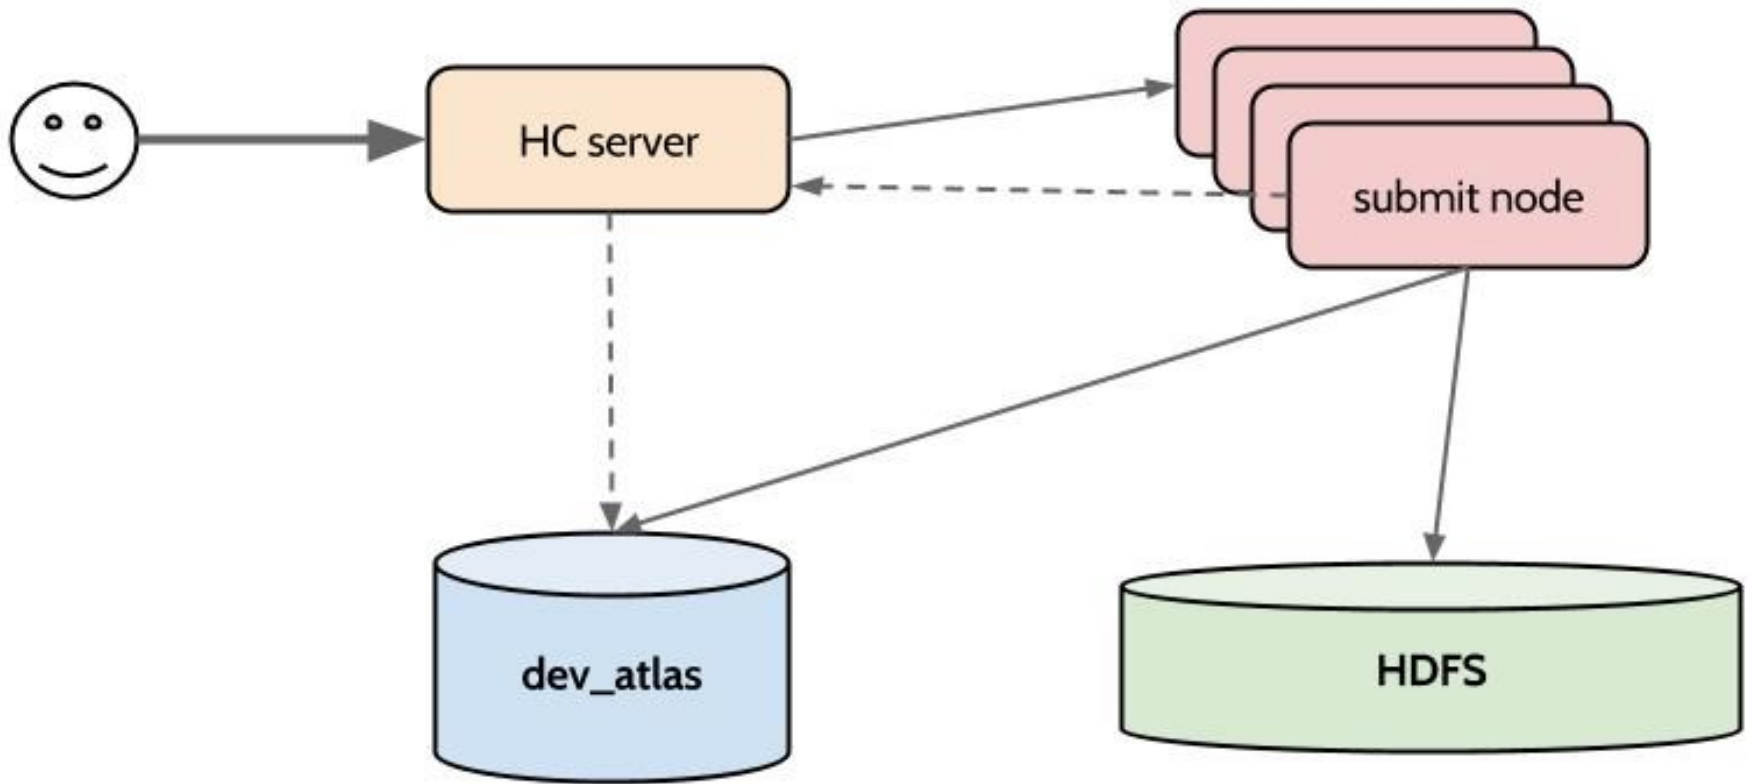
- Private IaaS cloud
- OpenStack based
- *Federates* Meyrin and Wigner
- 15,000 hypervisors by 2015
- 300,000 VMs by 2015
- Configuration management tools
 - Puppet, Foreman

Requirements

1. Elastic infrastructure (OpenStack)
2. Cloud orchestrator
3. Code sanitation (Gerrit)
4. Configuration Management (Puppet)
5. Deployment procedures

Dynamic testing scheduling

1. Test is requested (user, API, cron)
2. Creation of the VMs on demand
 - Isolation
 - Reliability
 - Elasticity
3. Configuration and startup of the VM
4. ...test runs...
5. Cleanup and destroy

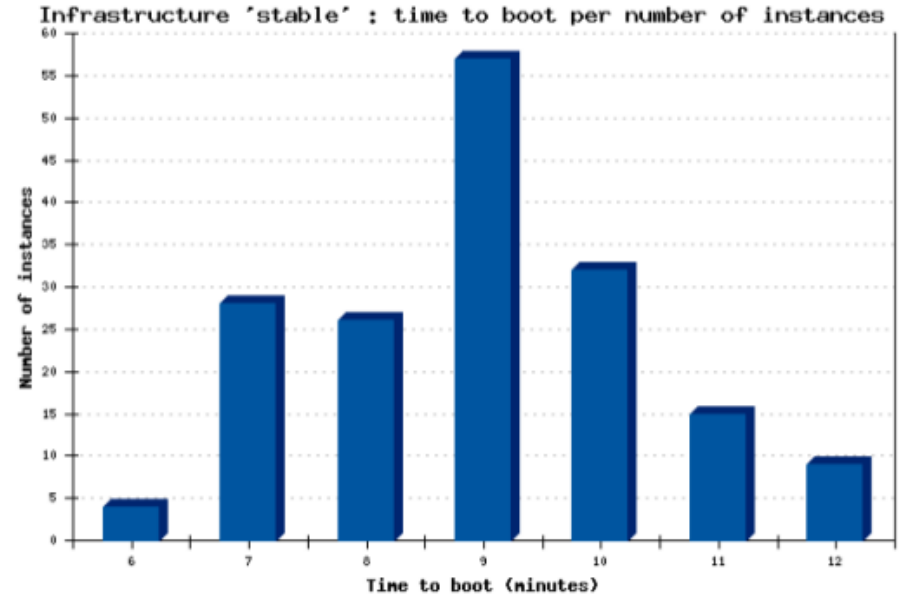
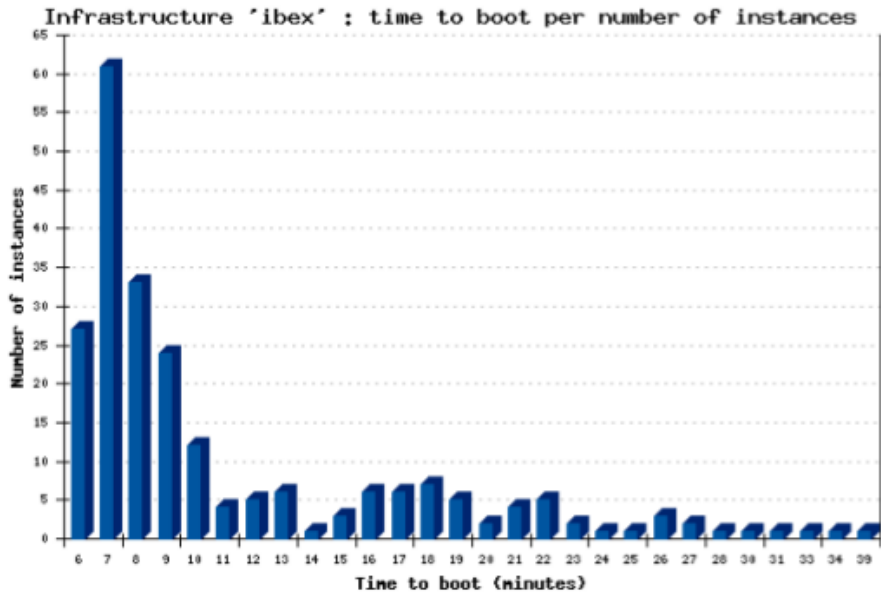


Dynamic testing scheduling

- Scheduling on time is hard
- IaaS predictable?
 - Provisioning times
 - Boot times
- HC will learn how is performing
- Tool to estimate boot performance

Provisioning latency

- Big improvement on OpenStack Grizzly
- Dynamic adaption of this time at runtime



Shared storage

- Needed a way to store all the logs
- Ideally accessible to others
- NFS for now,

- HDFS, elasticsearch from the Agile Monitoring?

Testing API

- Tools will need an API to request tests
 - And get information.
- Already done with LHC*b* RSS
- Athena Nightlies API ongoing


Current status

- Orchestration framework developed
- Currently deploying on OpenStack
 - OpenStack is now production ready
- Thinking about configuration frameworks

Conclusions

- New kind of testing requirements
- Big infrastructure changes
- Shift to a new *cloudy* paradigm

- Needs a transition period



Questions?