



RAL Quattor Update

Quattor at the RAL Tier1 and Quattor toolkit developments

Ian Collier
hepsysman June 10th 2010
ian.collier@stfc.ac.uk

RAL Status & Progress

- Last autumn we'd just deployed a new SL5 batch service with the Quattor toolkit.
- 95% of our batch capacity now Quattor managed. Will be 100% when we turn off SL4 on August 1st.
- This years disk servers (both SL4 and SL5) deployed with Quattor.
- Also BDIIs, CE, VO boxes, nagios slaves, castor tape servers, SRMs, CIPs, repository servers, Quattor server itself.
- Have started contributing code fixes and documentation to community.

Outlook

- New Quattor server to tackle performance.
- On the point of deploying SINDES Secure INformation DElivery System - will allow certificates passwords etc to be managed by Quattor.
- OS and software management for Castor servers.
- Ready to start work on **Aquilon** - configuration database developed by Morgan-Stanley to manage 10s of 1000s of systems with Quattor.
 - By the end of the year all new systems deployed using Quattor.
- Will also be able to manage RHEL in the next few weeks.

Quattor - what is it?

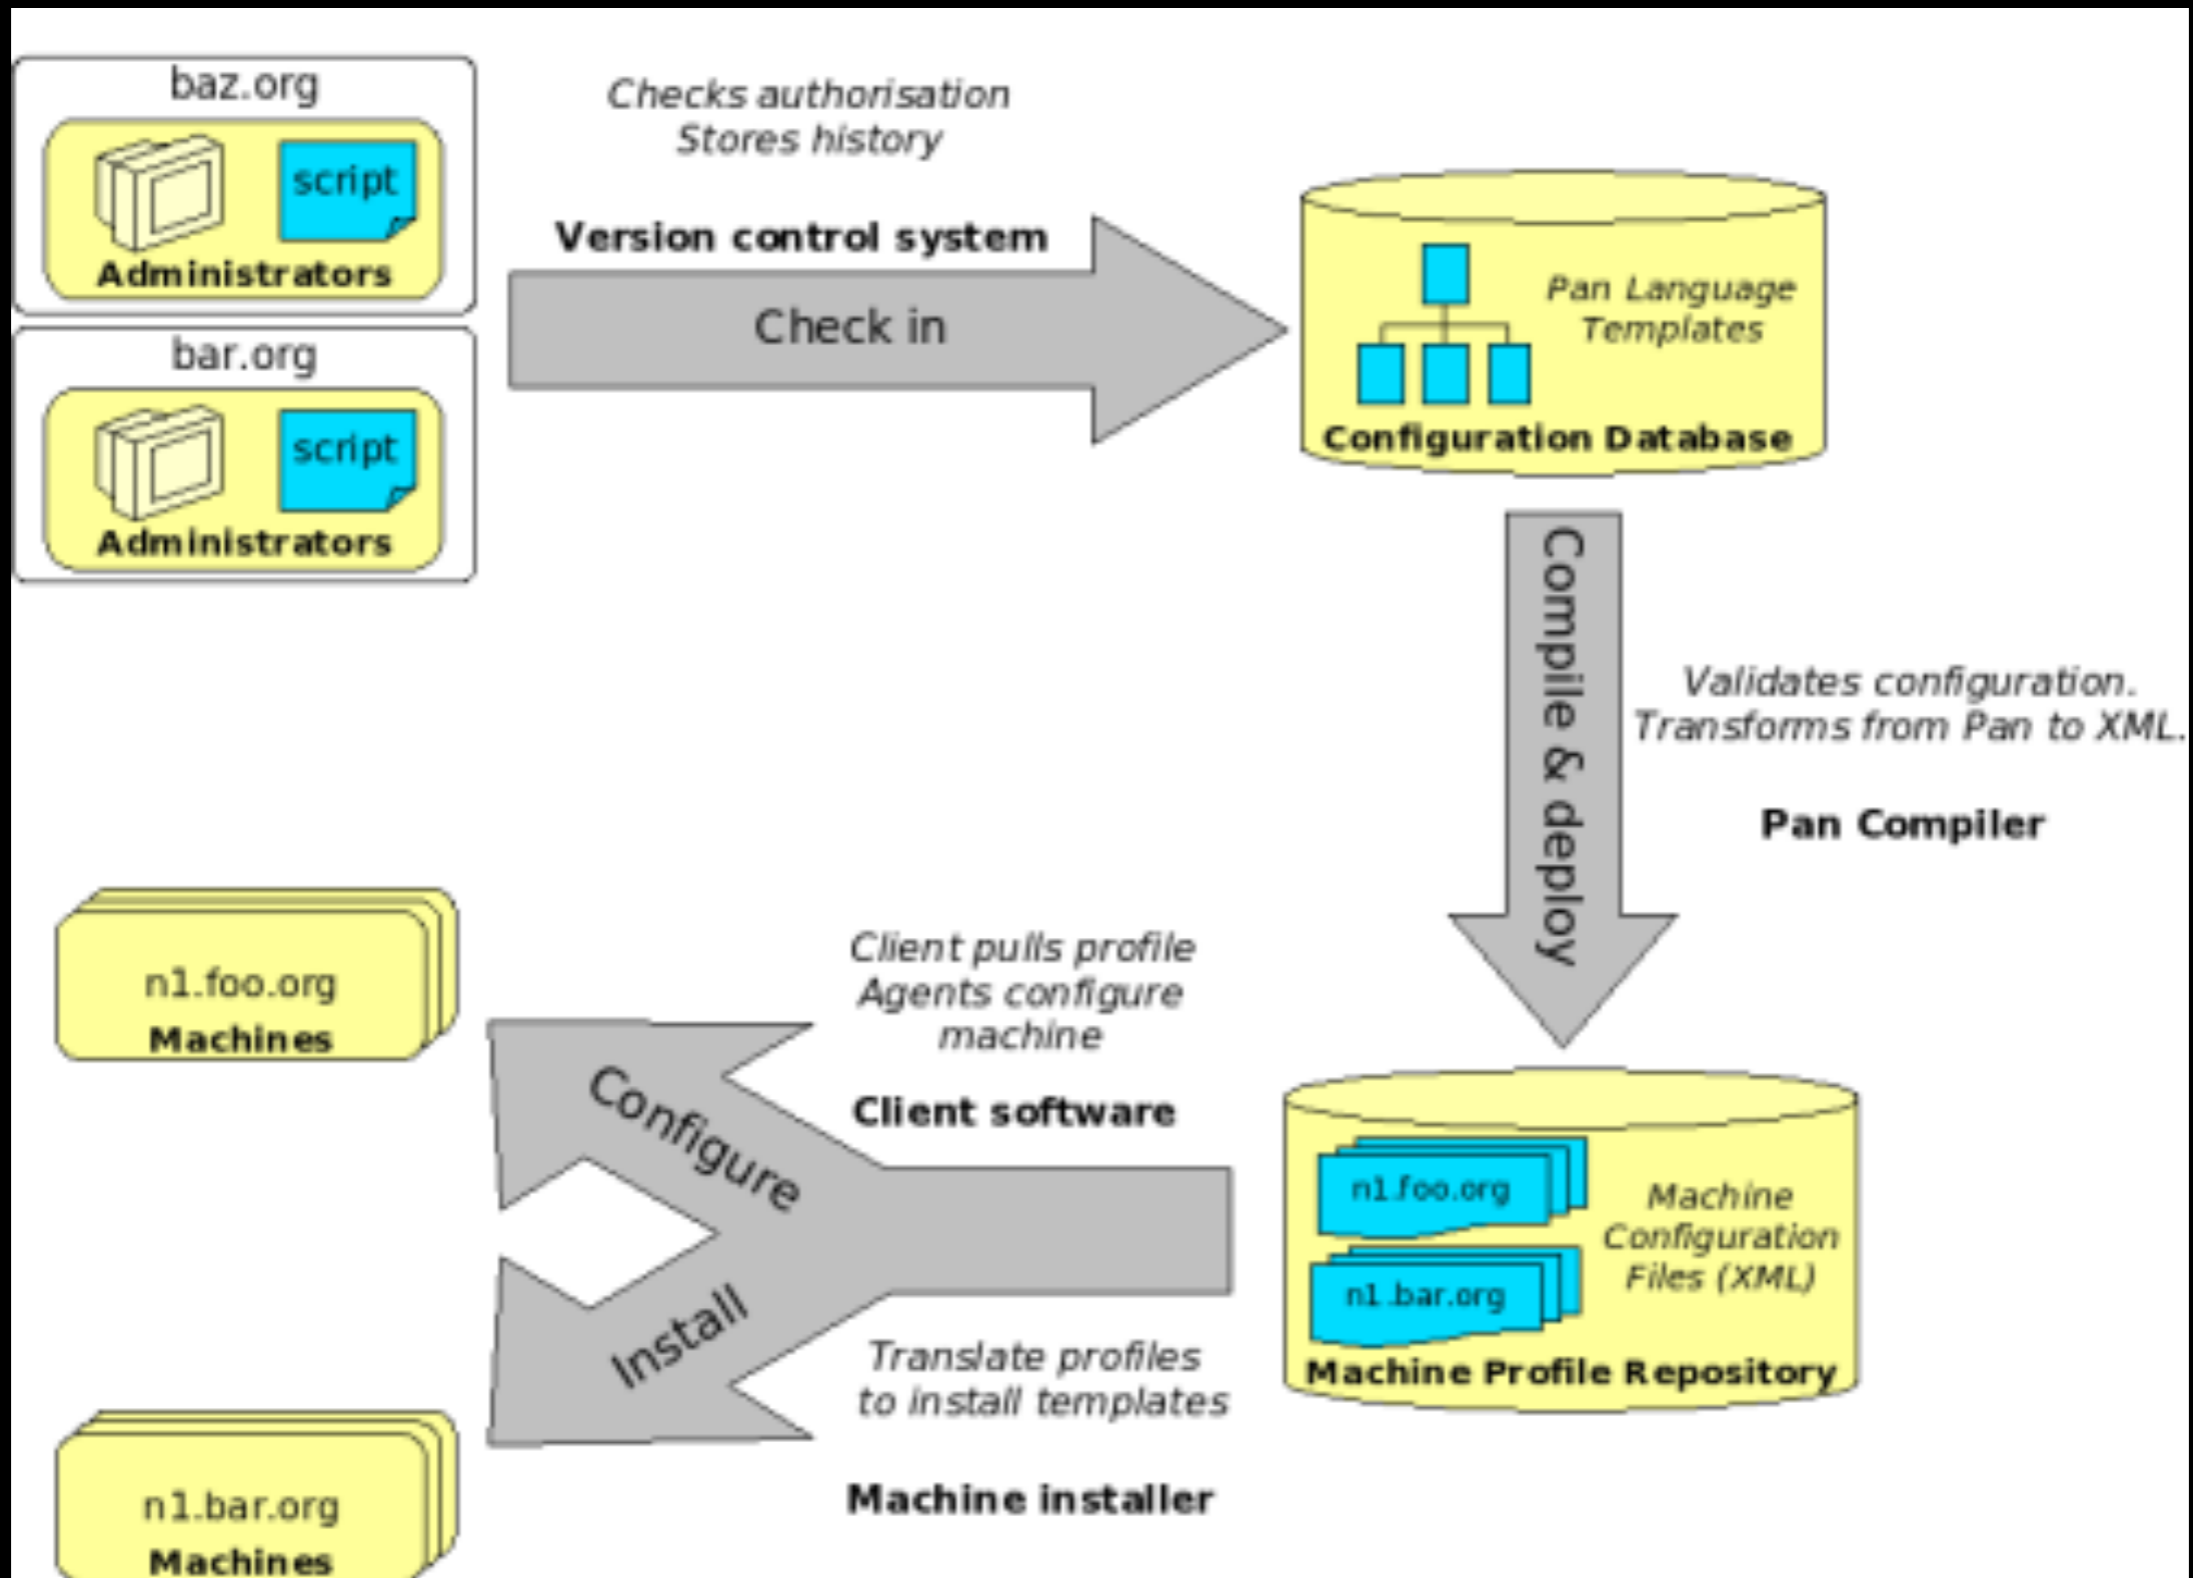
Those who aren't familiar

From quattor.org:

Quattor is a system administration toolkit providing a powerful, portable, and modular set of tools for the automated installation, configuration, and management of clusters and farms. It is developed as a community effort and provided as open-source software.

It is used at over 50 WLCG sites, and is managing 10s of 1000s of systems in commercial installations as well.

Quattor schematically

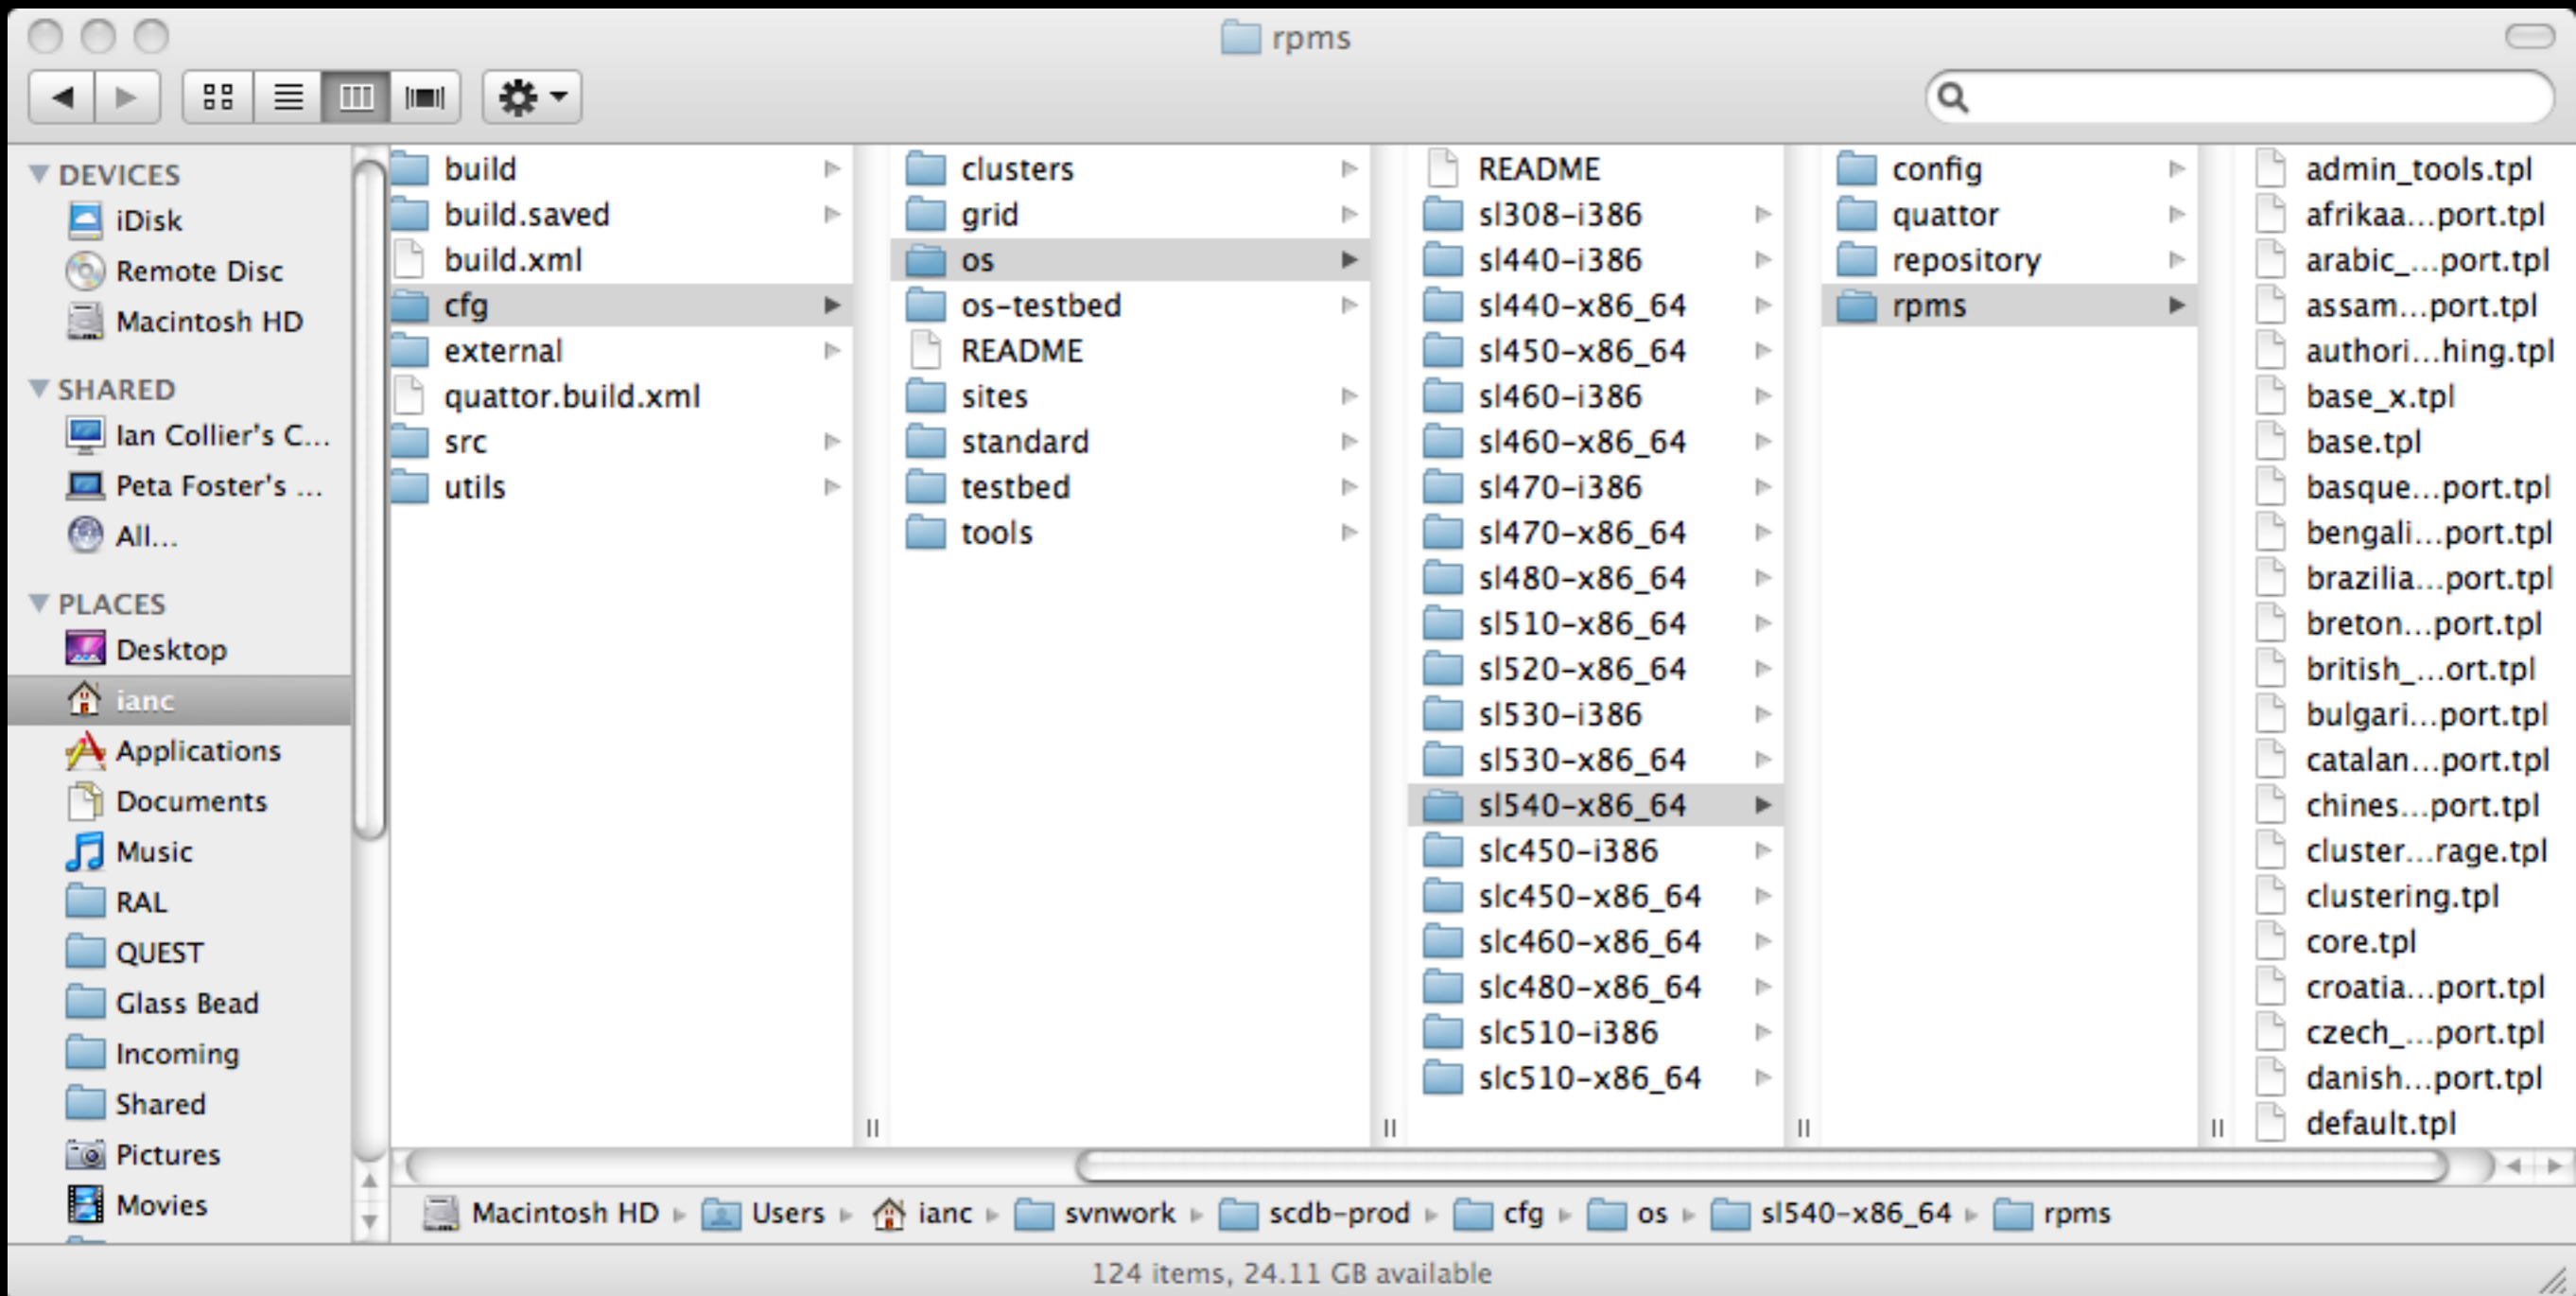


Working with Quattor

- Templates are written in the Pan Language.
 - Specific to Quattor - in fact almost defines 'using Quattor'
 - Used to specify machine configuration.
 - Flexible, hierarchical organization of information.
 - Allows arbitrary validation of configuration.
 - Declarative - says what config should be not how to get there - that is the job of the configuration components.
- Most shared among some or many systems.
 - Can get confusing - lots of places to configure things.
- We are using Quattor Working Group (QWG) framework.
 - Sharing gLite and OS configuration between ~50 other sites.

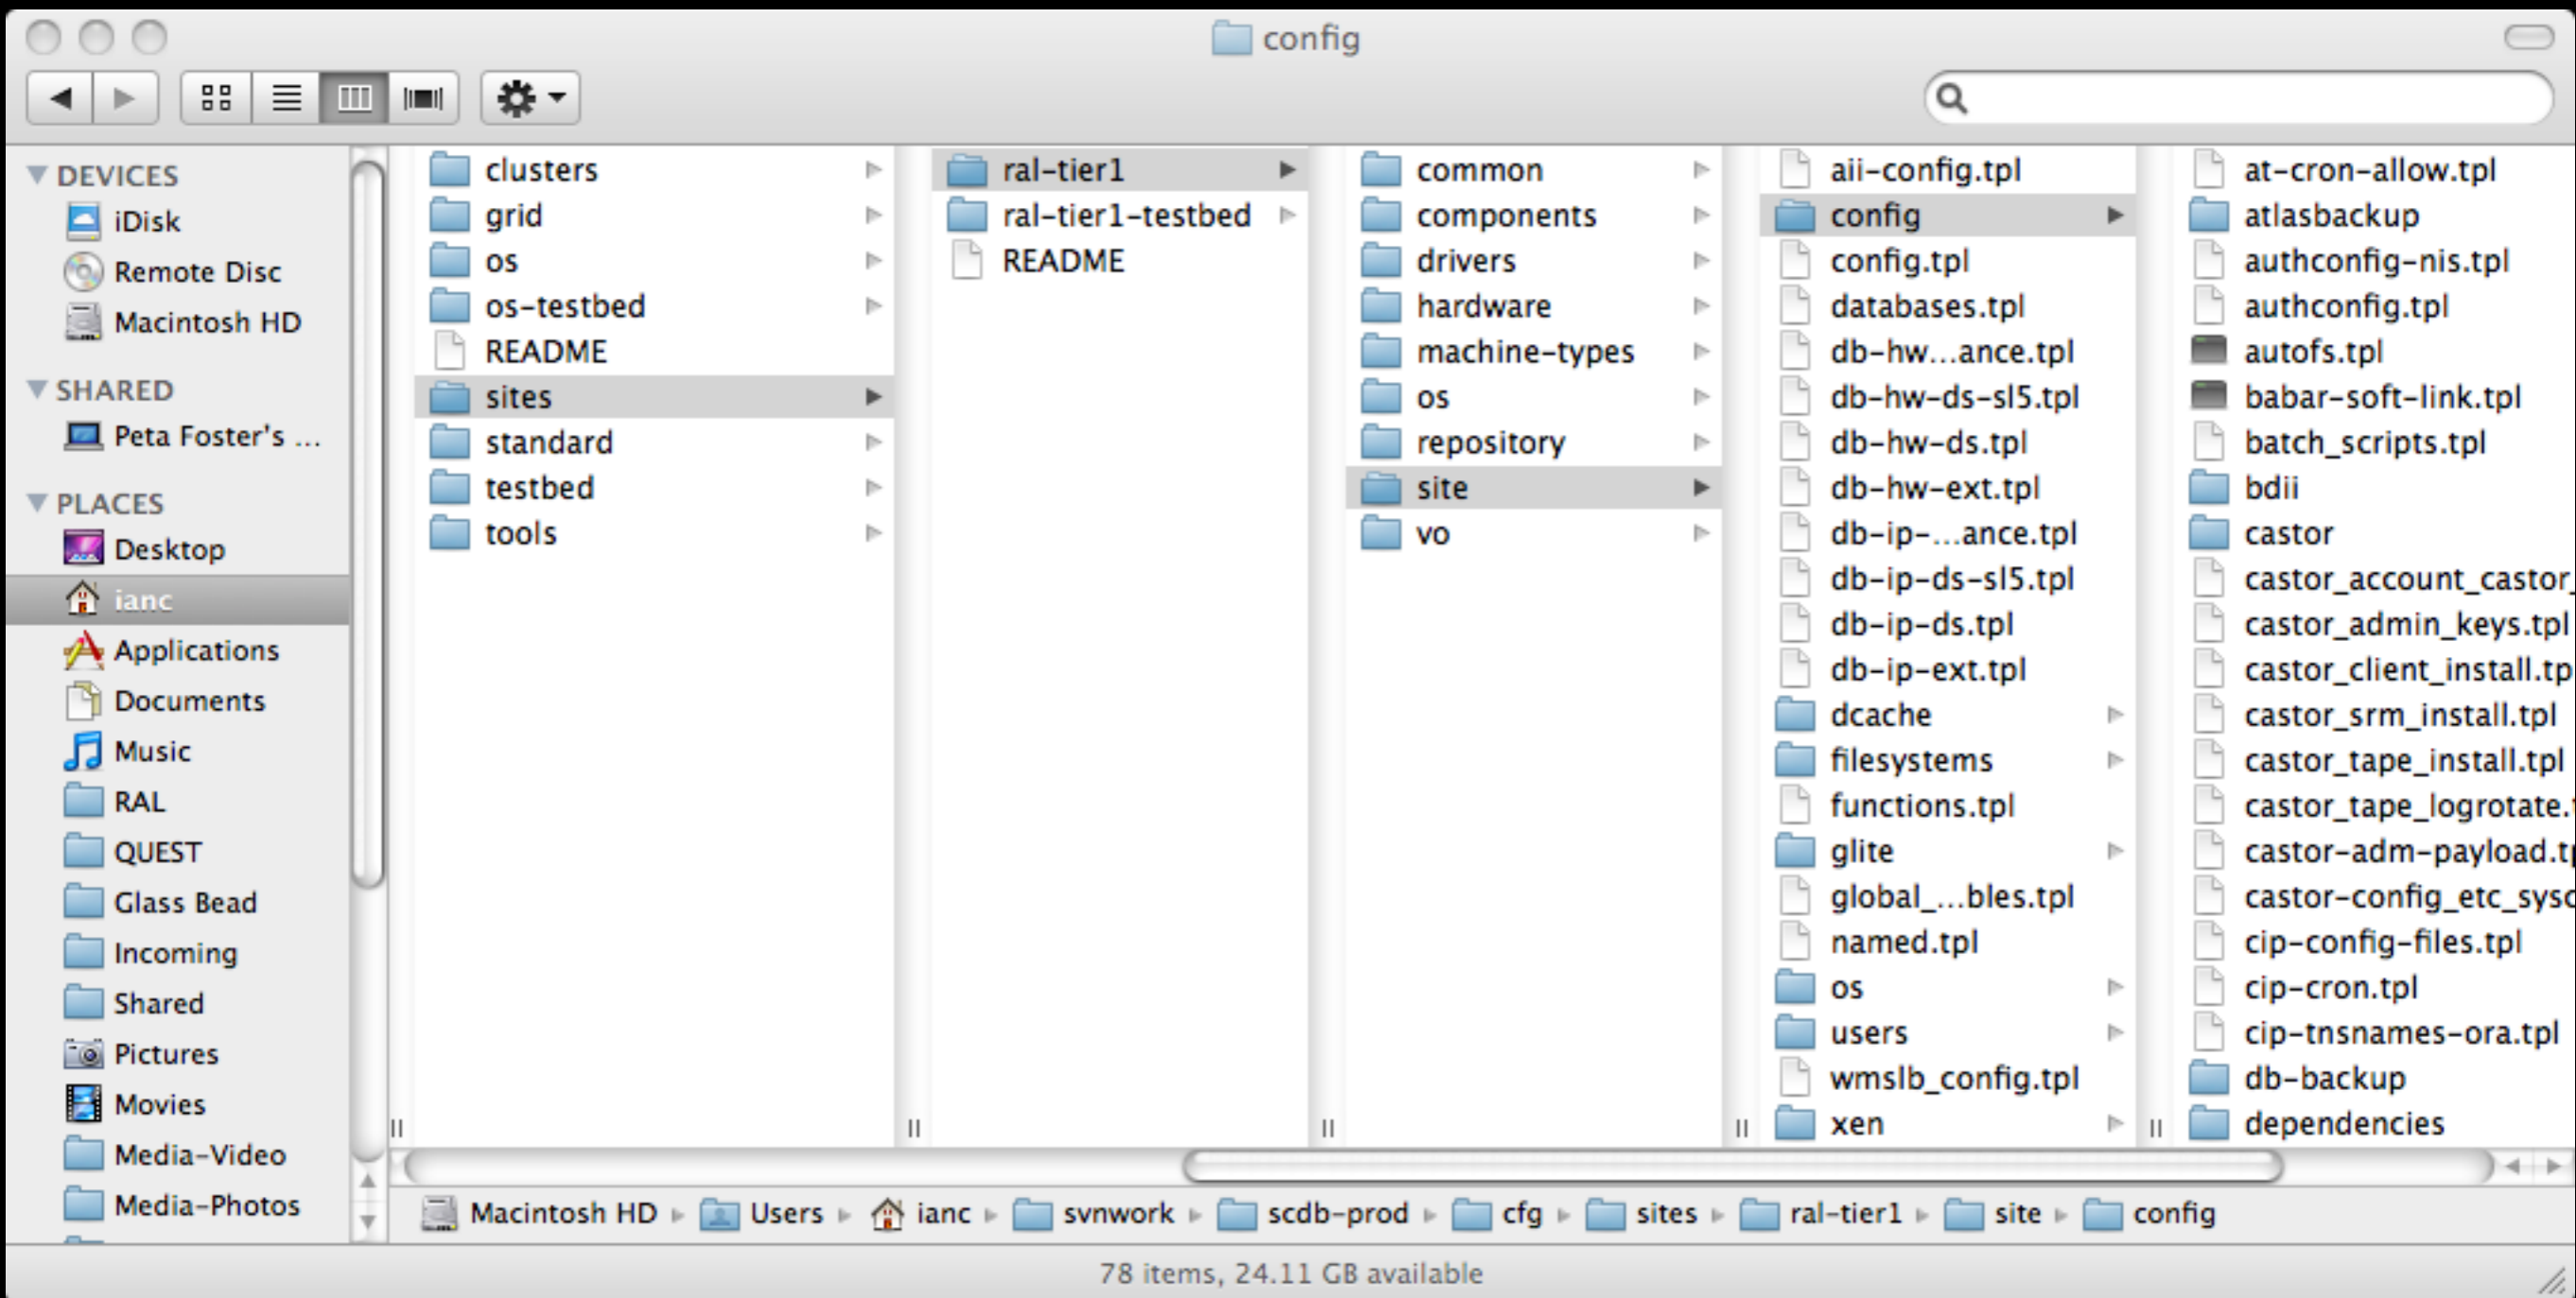
Working with Quattor

OS templates:



Working with Quattor

Site specific configuration templates:



Working with Quattor - Testbed

- Test bed developing well - very easy to switch from one machine "personality" to another.
 - A dozen or so machines at present - will be many more when we retire SL4 batch service.
 - allow us to test changes using Quattor before deploying them to production.
 - Becoming vital when bringing in updates from QWG or OS errata.
 - Less nasty surprises.

Working with Quattor - Testbed 2

Lets say we want to test a new batch server and WN.

A standard torque server template looks something like this:

```
variable FILESYSTEM_LAYOUT_CONFIG_SITE = 'site/filesystems/ral-  
batch-sw_raid';
```

```
include { 'machine-types/ral/torque_server' };
```

```
include { 'machine-types/ral/blahparser' };
```

```
include { 'site/config/atlasbackup/lcgbatch01' };
```

Working with Quattor - Testbed 3

And for the testbed the torque server template has just a couple of extra lines:

```
variable FILESYSTEM_LAYOUT_CONFIG_SITE = 'site/filesystems/ral-  
batch-sw_raid';  
  
variable TORQUE_SERVER_HOST = "lcg0628.gridpp.rl.ac.uk";  
  
variable WORKER_NODES = list (  
    "wn.example.com",  
    "lcg0625.gridpp.rl.ac.uk",  
);  
  
include { 'machine-types/ral/torque_server' };  
  
include { 'machine-types/ral/blahparser' };  
  
include { 'site/config/atlasbackup/lcgbatch01' };
```

Working with Quattor - Testbed 4

And for one of our testbed WNs we simply add one line:

```
include { 'machine-types/ral/wn' };
```

```
variable TORQUE_SERVER_HOST = 'lcg0628.gridpp.rl.ac.uk';
```

```
include { 'machine-types/ral/wn' };
```

Working with Quattor - Testbed 5

- Hardware, OS, payload all abstracted.
- For a given test, we just assign the 'personality' or machine-type, and possibly move it to a new 'cluster', the rest is organised in the template structure.
- Much easier than taking the parts from one kickstart for given hardware, then getting the payload part from another and hoping they work together.....

Ch-ch-changes...

- We are saving approximately one week per month in effort looking after the batch farm.
- Disk server deployment streamlined - more database driven.
- Switching disk servers between instances easier - as seen in a recent middle of the night call out.
- A lot to learn, admins a getting to grips with very different workflows - this can be a challenge...

The Quattor Toolkit

- Last Quattor workshop led to clear work plan and monthly development meetings.
- Small but helpful changes like the unification of all documentation in one place making it (at least a bit) less confusing.
- Effort from several sites to use knowledge of the site embodied in the Quattor repository to automatically configure nagios - a number of strands to be unified here.
- Commercial users also continue to return significant developments to community.
- Quattor Working Group extending support for gLite 3.2 - the shared effort to support gLite and OS updates etc. saves a lot of work locally.
- 10th Quattor workshop will be hosted here at RAL in October.

Questions?

<http://quattor.org>

10th Quattor Workshop
11th-13th October 2010
RAL