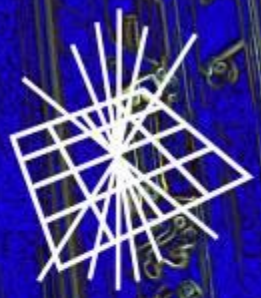

GridPP

UK Computing for Particle Physics

RAL Tier1 Quattor experience and Quattor outlook

Ian Collier

RAL Tier1 Fabric Team

Hepix, Lisbon April 2010

Science & Technology Facilities Council

Contents

- Quattor Introduction
- Quattor at RAL Tier 1
  - History
  - Status
  - Benefits
  - Issues
- Quattor Toolkit /QWG Report
- Key Developments
- Outlook

Quattor Toolkit introduction

- Overview:
  - Systems described by one or more templates (pan language)
  - Compiled to produce XML profile for each machine
  - Basic system + Quattor client installed by minimal (auto-generated) kickstart
    - AII takes care of DHCP and pxeboot – just need to boot/reboot target system
  - Client then downloads XML profile and carries out rest of software install and system configuration (client configuration modules written in perl)
  - When configurations are changed fresh XML profile is downloaded and changes carried out by Quattor client
- Hierarchical template structure
  - Configuration elements shared between machines, machine-types, clusters and sites as appropriate
- Modular – can pick and choose components
- Quattor Working Group – last 4 years
  - community supported framework sharing gLite (and OS) configuration
  - Used and contributed to by ~50 sites of varying sizes

Quattor Toolkit Introduction

- Complete fabric automation toolkit

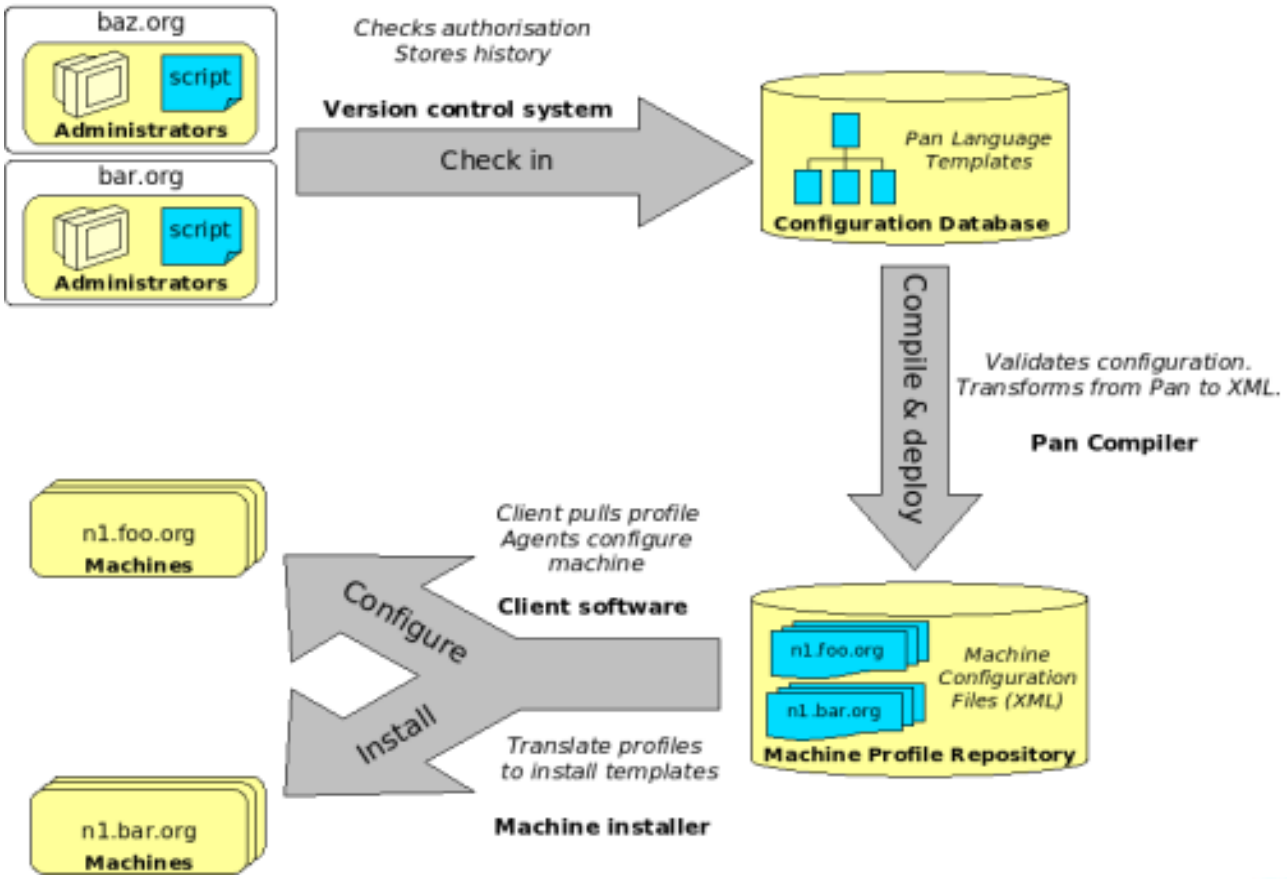

Specification and management of system (OS and payload) and site configuration

Provisioning of new machines

Ongoing maintenance of software and services

Quattor at RAL Tier 1

- History
  - Traditionally, using kickstart + parallel ssh + manual methods
    - Extensive use of Puppet for Castor server config
    - Reaching limits of scalability
  - In 2009 began search for management tools for (growing) Tier 1 fabric
  - Considered many options – including cfengine, wider use of puppet, Platform Manager etc. and Quattor
  - Recommendation to try Quattor – seemed to meets all requirements
    - Must be proven to scale to >1500 systems
    - Must automate of deployment **and** ongoing config
      - Single language to describe both stages a benefit
    - Support for gLite (via Quattor Working Group)
    - Open Source strongly preferred
  - Began with deployment of new SL5 batch service last August/September - very successful

Quattor at RAL Tier 1

- Status
  - 680 SL5 WNs managed by Quattor
  - 130 Castor disk servers being deployed now
  - This years new deployments:
    - 160+ WNs and 130 disk servers deployed from scratch with Quattor
  - gLite servers
    - Batch server, bdiis, vobox, ui, wms, lfc
    - Others in progress
  - Non-gLite:
    - nis servers, nagios servers, repository servers, license servers…
  - Starting work on Castor core servers
    - Castor information provider (CIP), SRM and tape servers currently testing
  - Sindes (secure information distribution) ready to deploy
  - 'Fed' by hardware tracking database

Quattor at RAL Tier 1

- Benefits
  - Consistency good
  - Security and other updates easier
  - Biggest single saving so is time managing (increasing numbers of) WNs
  - Repository of XML profiles tell us how systems are configured

- Used for deploying latest hardware to acceptance testing (just complete)
  - Much quicker and easier than traditional (kickstart based) methods
  - Currently repurposing machines for production – before we would have re-kickstarted.

Quattor at RAL Tier 1

- Issues
  - Quattor Working Group framework sometimes needs adapting
    - We are (currently) only Tier 1 using QWG
  - Castor systems currently managed by puppet
    - Interface between methods will need some care
  - It is still significant work to get new (to us) machine types set up
    - But OS changes easier
  - Of course some admins find the different approach challenging
    - Even so we have 9 active admins + 5 or 6 occasional

Quattor Toolkit Status

- There is a positive relationship between sites using Quattor and their availability – based on LCG statistics

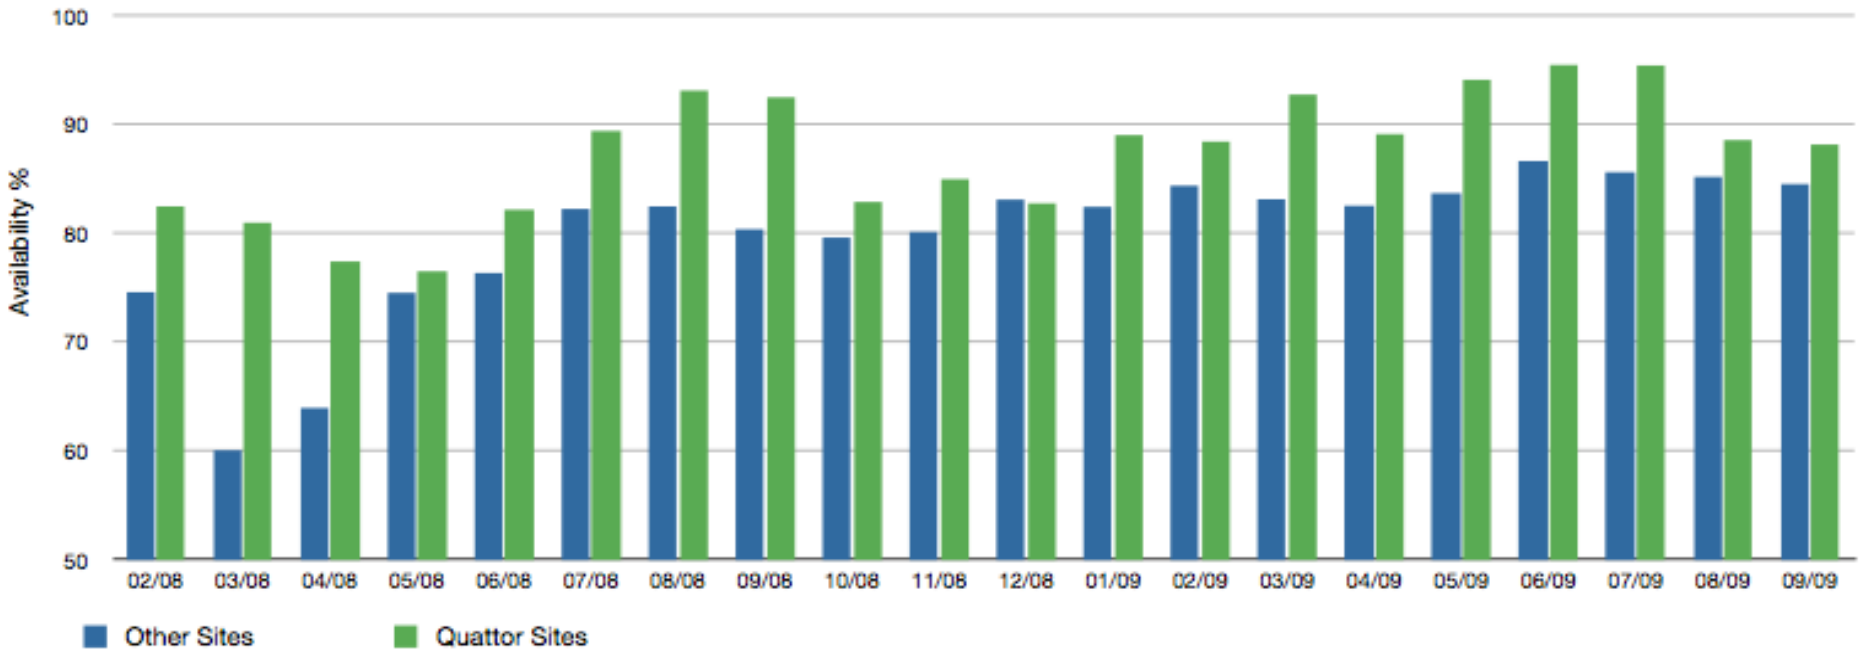

Quattor Toolkit Status

- The percentage gain in availability:

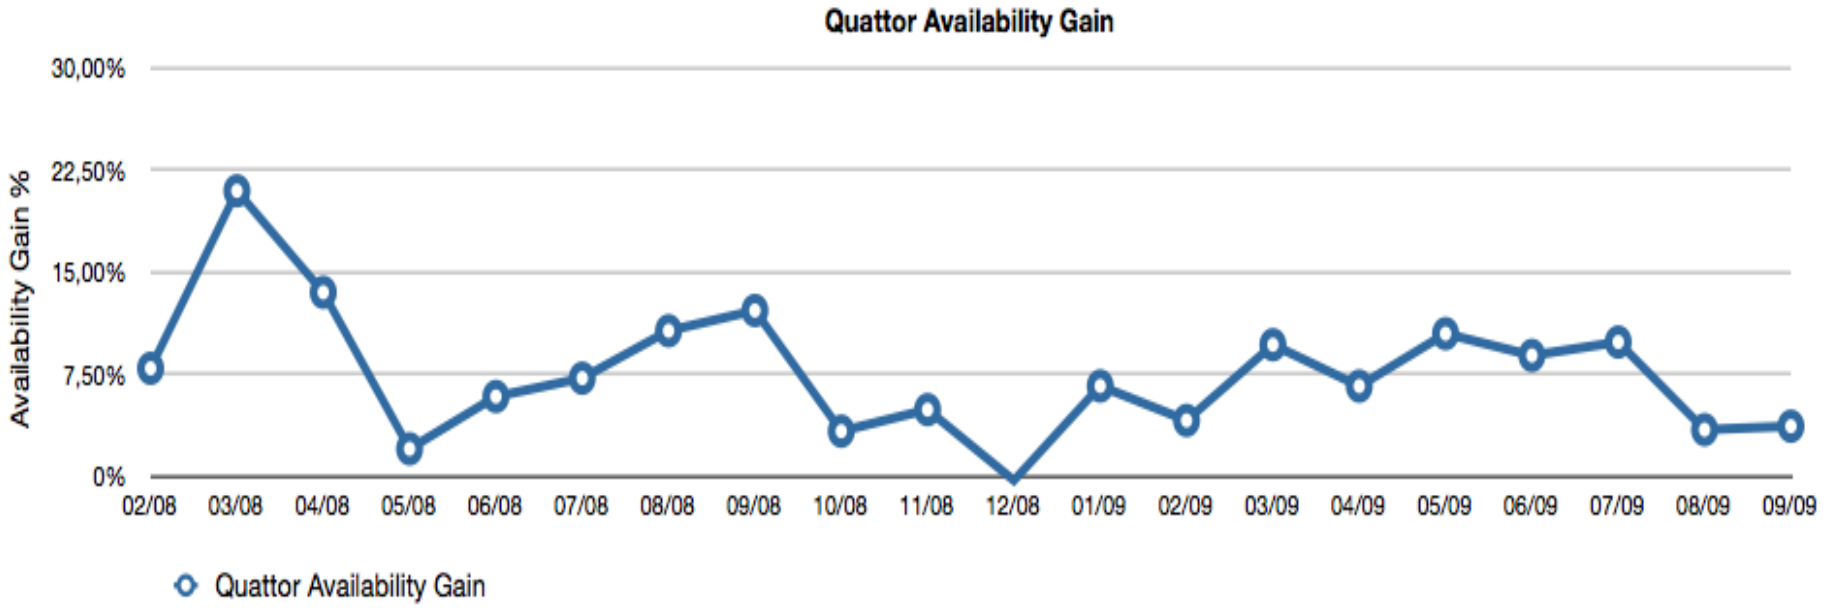

Quattor Toolkit Status

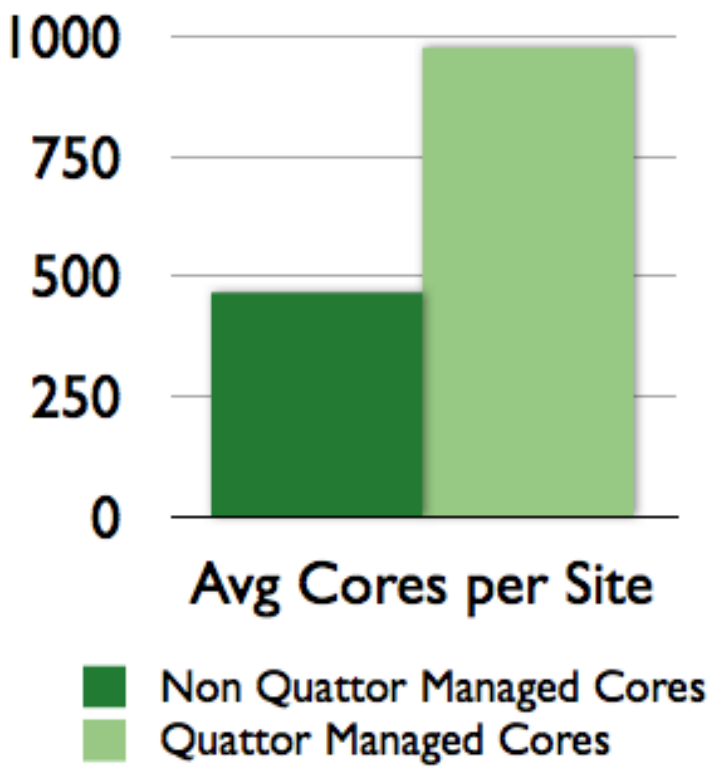

- Quattor managed sites tend to be larger
- In fact Quattor managed sites accounting for 15% of the cores in WLCG

Quattor Toolkit Status

- Move to sourceforge more or less complete
  - But documentation consolidated at LAL – awaiting trac features on SF
- Deployments growing
  - One commercial site manages over 20000 servers (using Aquilon configuration database)
    - Plus 1200 ESX servers with 1700 ESX nodes
  - RAL already one of larger QWG/SCDB sites (CERN of course larger again using CDB)
- FP7 funding bid was unsuccessful
  - Preparation did help set development proprieties
  - Reflected in actions

Quattor Developments

- Automated monitoring configuration (nagios) active area
  - Several different approaches currently
- OS Errata structure in QWG now very granular – can ‘pin’ errata to specific date/versions and specify kernels on a system or cluster basis
  - Modular so can easily be used by non-QWG sites
- StratusLab project will use Quattor for managing virtualised cloud resources
- Action highlights
  - Monthly development meetings
  - Convergence of automated nagios configurations
  - RAL about to start ‘porting’ Aquilon to non-MS environment
  - Set up automated build and test infrastructure
  - IPV6
  - Investigate non-RPM package management

Conclusions and Outlook - RAL

- RAL Tier 1 now fully committed to Quattor
  - (And mostly happy about it!)
- It is still true that getting started with Quattor takes effort
  - Community support was vital at RAL
- At RAL we are seeing real benefits:
  - In consistency
  - In speed of deployment and updates
- Sindes will allow all sensitive info in quattor profile
- Aquilon configuration database will aid scalability
- Roadmap and monthly meetings will focus community effort

Outlook - Quattor

- Roadmap and monthly meetings will focus community effort
- Increased sharing of e.g OS templates
- Automated nagios configuration would be big benefit – again not just in QWG
- Aquilon configuration database option should aid large sites