quattor Update

Integrating Aquilon at a Tier 1

Ian Collier
with contributions from James Adams
RAL Tier 1
STFC Scientific Computing Department
HEPIX Beijing, 15 October 2012

Quattor Toolkit Update

- Configuration components moved to GitHub
- Maven build system for components
- Components being updated to meet new requirements
- Compiler Has been able to output profiles as JSON as well as XML
  - Opens up options for data warehousing
  - Client still needs updates to use json output (in progress)
- Option to use yum for package management in development

RAL Quattor Data warehouse

- Using the compiler to output all profiles as JSON
- Two (secondary/high school) summer students built indexing/search/display tool
- Allows search on any assigned attribute in profile
  - NOT free search on any strings
- Should be ready to share with other sites in coming months.

quattor

Distribution | Server Attribute

Path to a specifc value

system/cluster/name

Minimum profile count: 26

Analyse

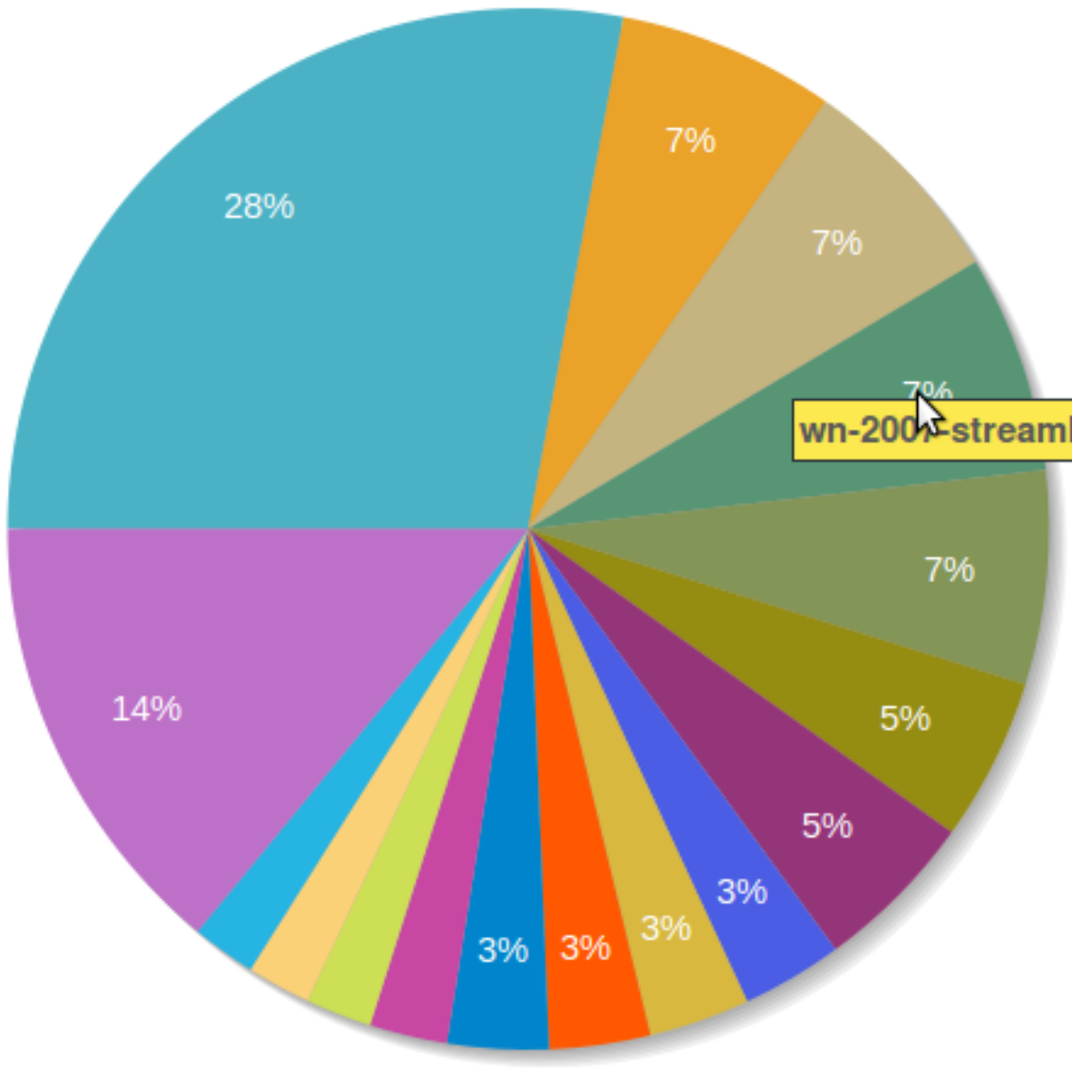

Value	Count
hadoop	22
RAL UMD 1	21
Cloud Hypervisor Nodes	16
RAL Tier1 Databases	13
RAL Hyper-V	12
lustre headnodes	10
CASTOR certification servers	10
CASTOR Facilities Tape Servers	10
Squid	7
wn-virtual-umd	6
wn-virtual	6
ties Instance	6
Web Servers	5
FDS CEDA StorageD servers	5
RAL gLite 3.1 Test	4
RAL EMI 2	4
CASTOR Virtual Certification Instance	4
CASTOR Repack disk servers	4
CASTOR Preproduction Instance	4
CASTOR ATLAS SRM Cluster	4
wn-legacy	3
SoftwareInstallNodes	3
DNS Servers	3

quattor

Distribution | Server Attribute

Attribute
hardware/cards/nic/eth0

Value
e1000

Submit Query

Show 10 entries Search:

Machine Names
c2certdlf.ads.rl.ac.uk
c2certds6.gridpp.rl.ac.uk
c2certtape.ads.rl.ac.uk
castor200.ads.rl.ac.uk
castor201.ads.rl.ac.uk
castor208.ads.rl.ac.uk
ccse08.ads.rl.ac.uk
ccse09.ads.rl.ac.uk
ctsc12.ads.rl.ac.uk
ctsc20.ads.rl.ac.uk

Showing 1 to 10 of 434 entries First Previous 1 2 3 4 5 Next Last

Configuration Databases

- CERN using CDB
  - CVS based
- All but two sites use SCDB - the second generation of Quattor
  - SVN based
- Aquilon - Third generation Quattor configuration database
  - SQL + git

SCDB

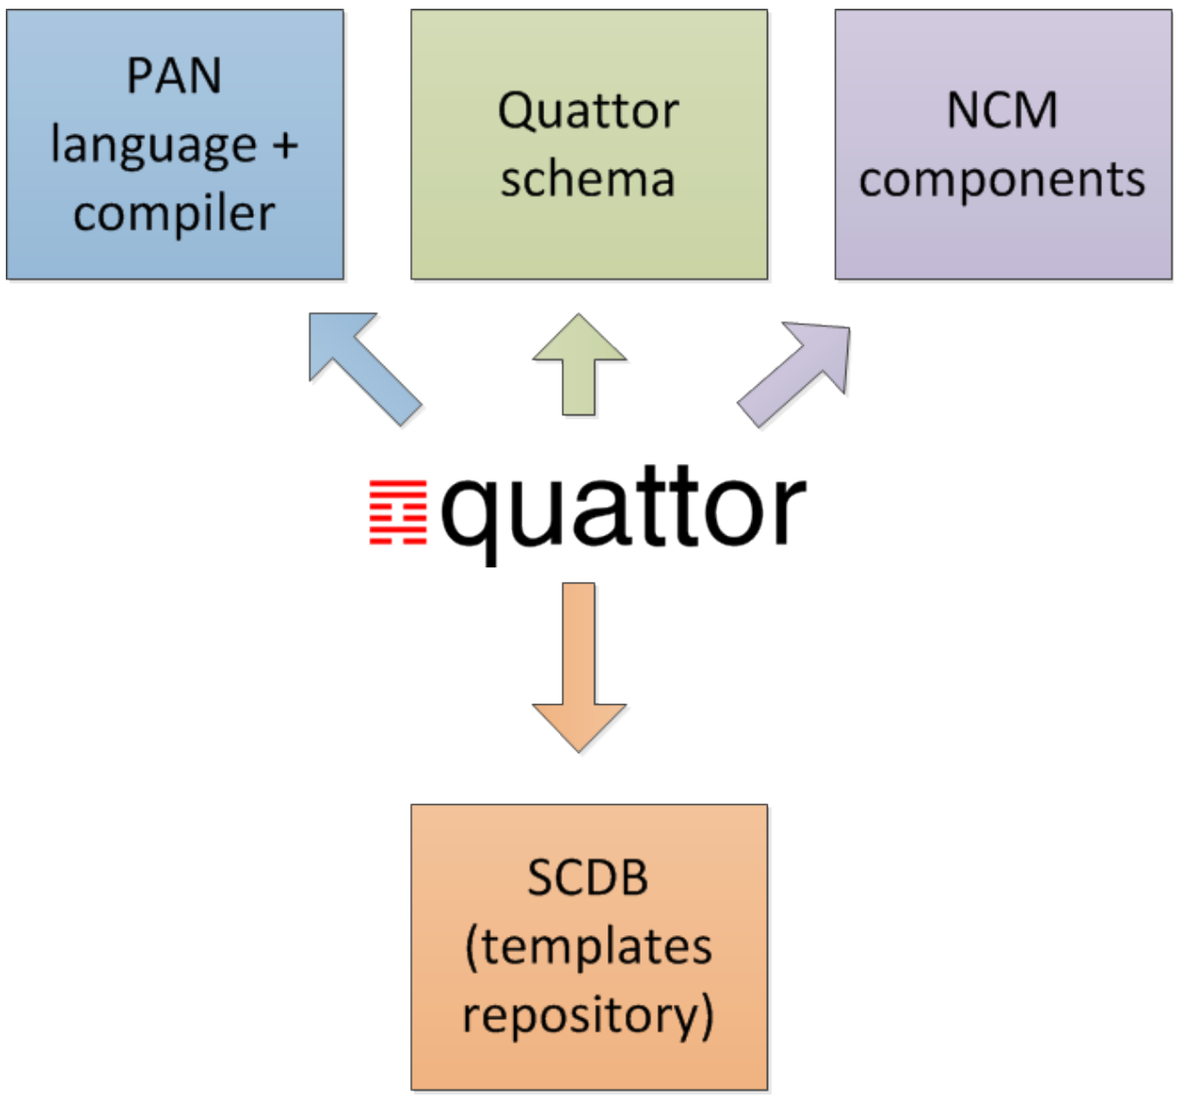

- It works
- It is certainly easier to work with and more flexible than CDB
- Together with Quattor Working Group framework it enables a lot of shared configuration among grid sites

SCDB - workflow

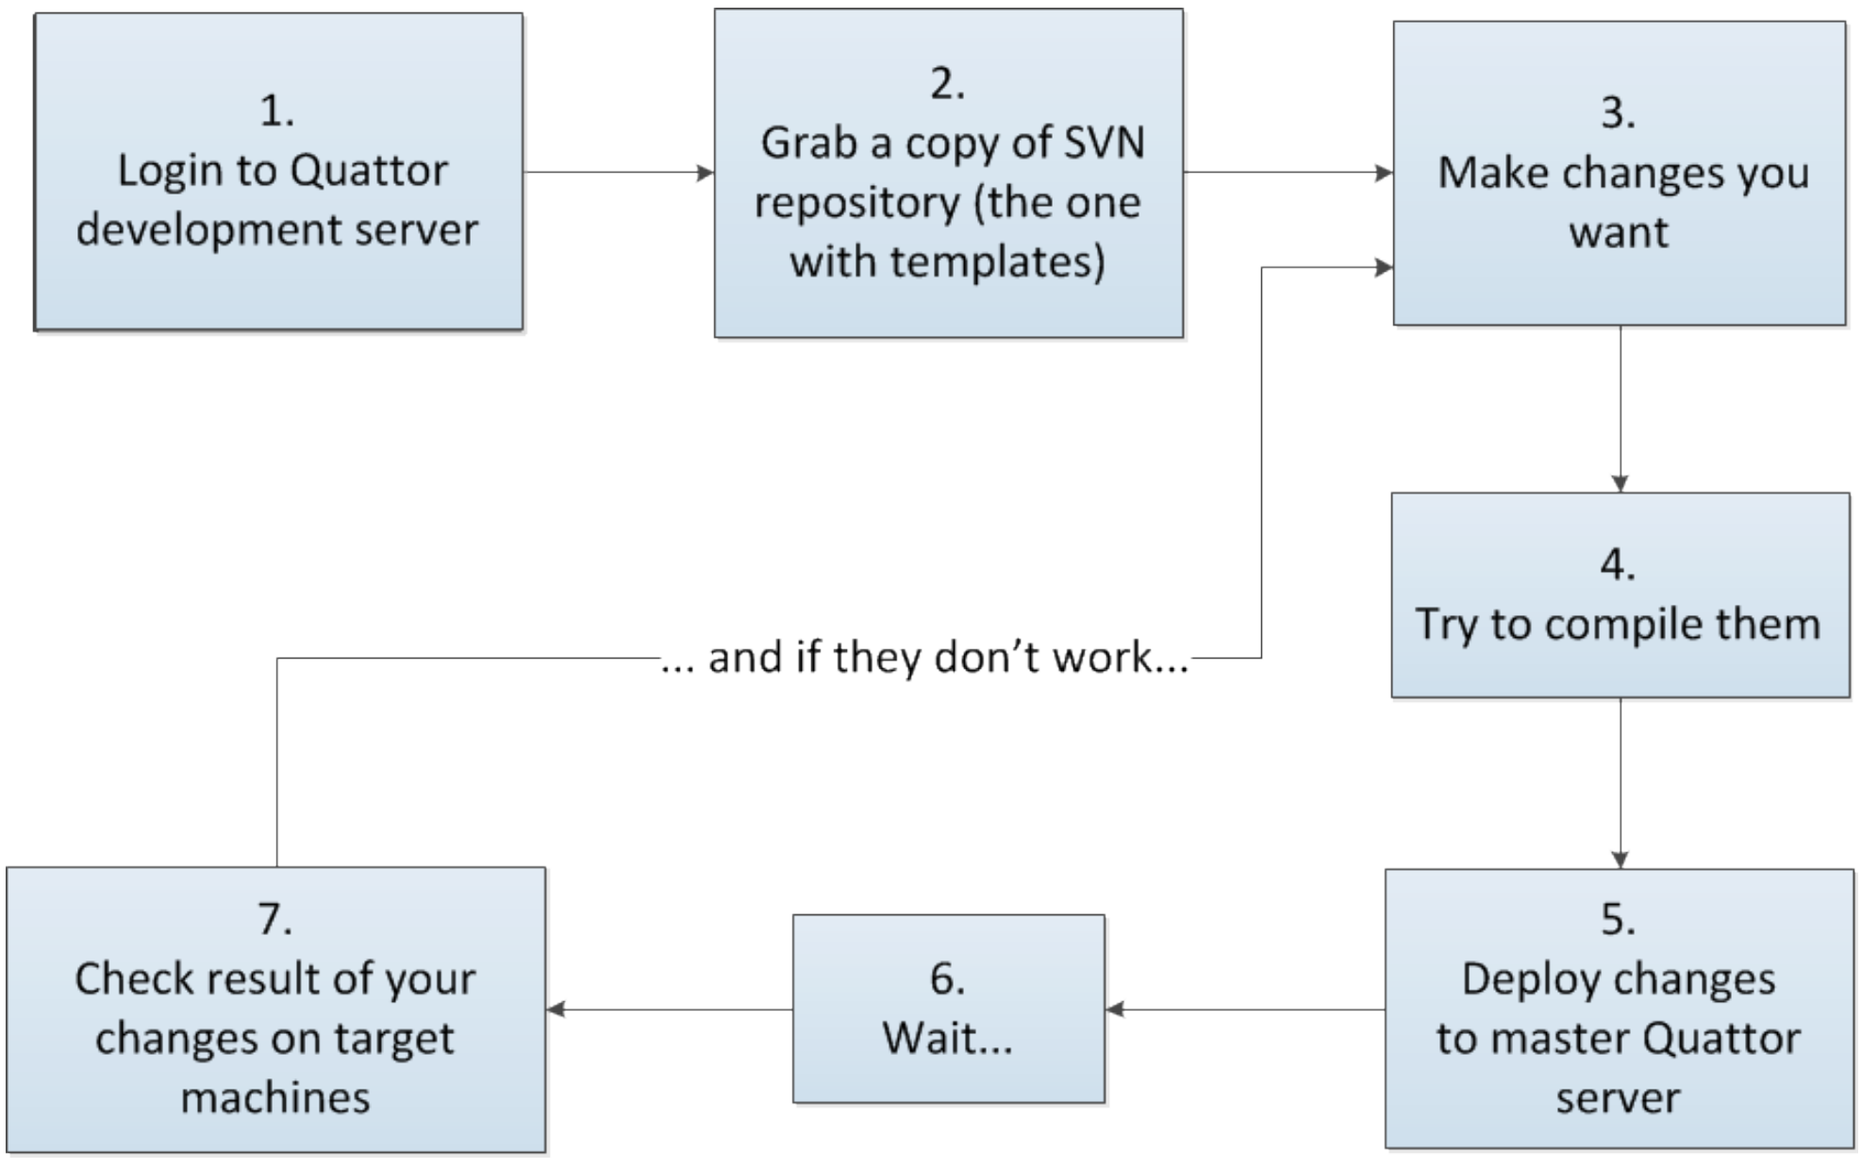

SCDB limitations

- Changes cannot be tested on one individual machine (in an easy way)
- User has to wait for other hosts' templates to compile even they do not care about them.
- Every one is a developer editing templates, using svn
- Structured data kept in flat files

```
variable DB_IP_EXT = nlist(
  escape("ads0cntr.cc.rl.ac.uk"), "130.246.183.165",
  escape("ads0pt01.cc.rl.ac.uk"), "130.246.183.157",
  escape("ads0pt02.ads.rl.ac.uk"), "130.246.183.142",
  escape("ads0sb01.cc.rl.ac.uk"), "130.246.183.143",
  escape("ads0sb02.cc.rl.ac.uk"), "130.246.183.155",
  escape("afs1.gridpp.rl.ac.uk"), "130.246.183.203",
  escape("afs2.gridpp.rl.ac.uk"), "130.246.183.204",
  escape("afs3.gridpp.rl.ac.uk"), "130.246.183.205",
  escape("amanda01.gridpp.rl.ac.uk"), "130.246.180.243",
  escape("amanda02.gridpp.rl.ac.uk"), "130.246.180.244",
  escape("amanda03.gridpp.rl.ac.uk"), "130.246.180.245",
...
```

So, what is Aquilon?

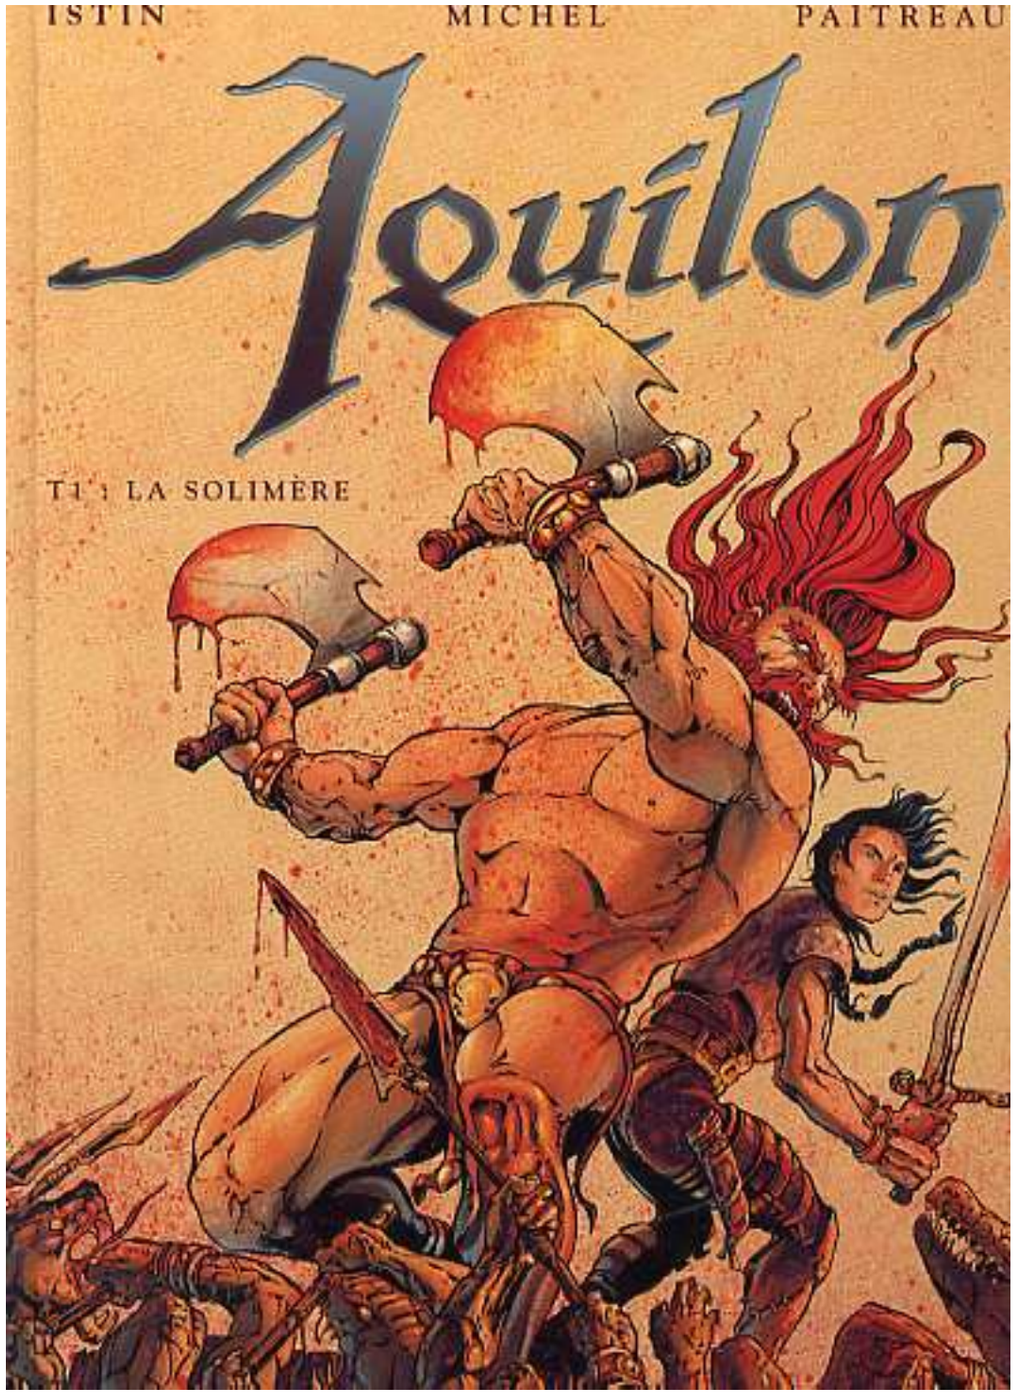

So, what is Aquilon?

So, what is Aquilon?

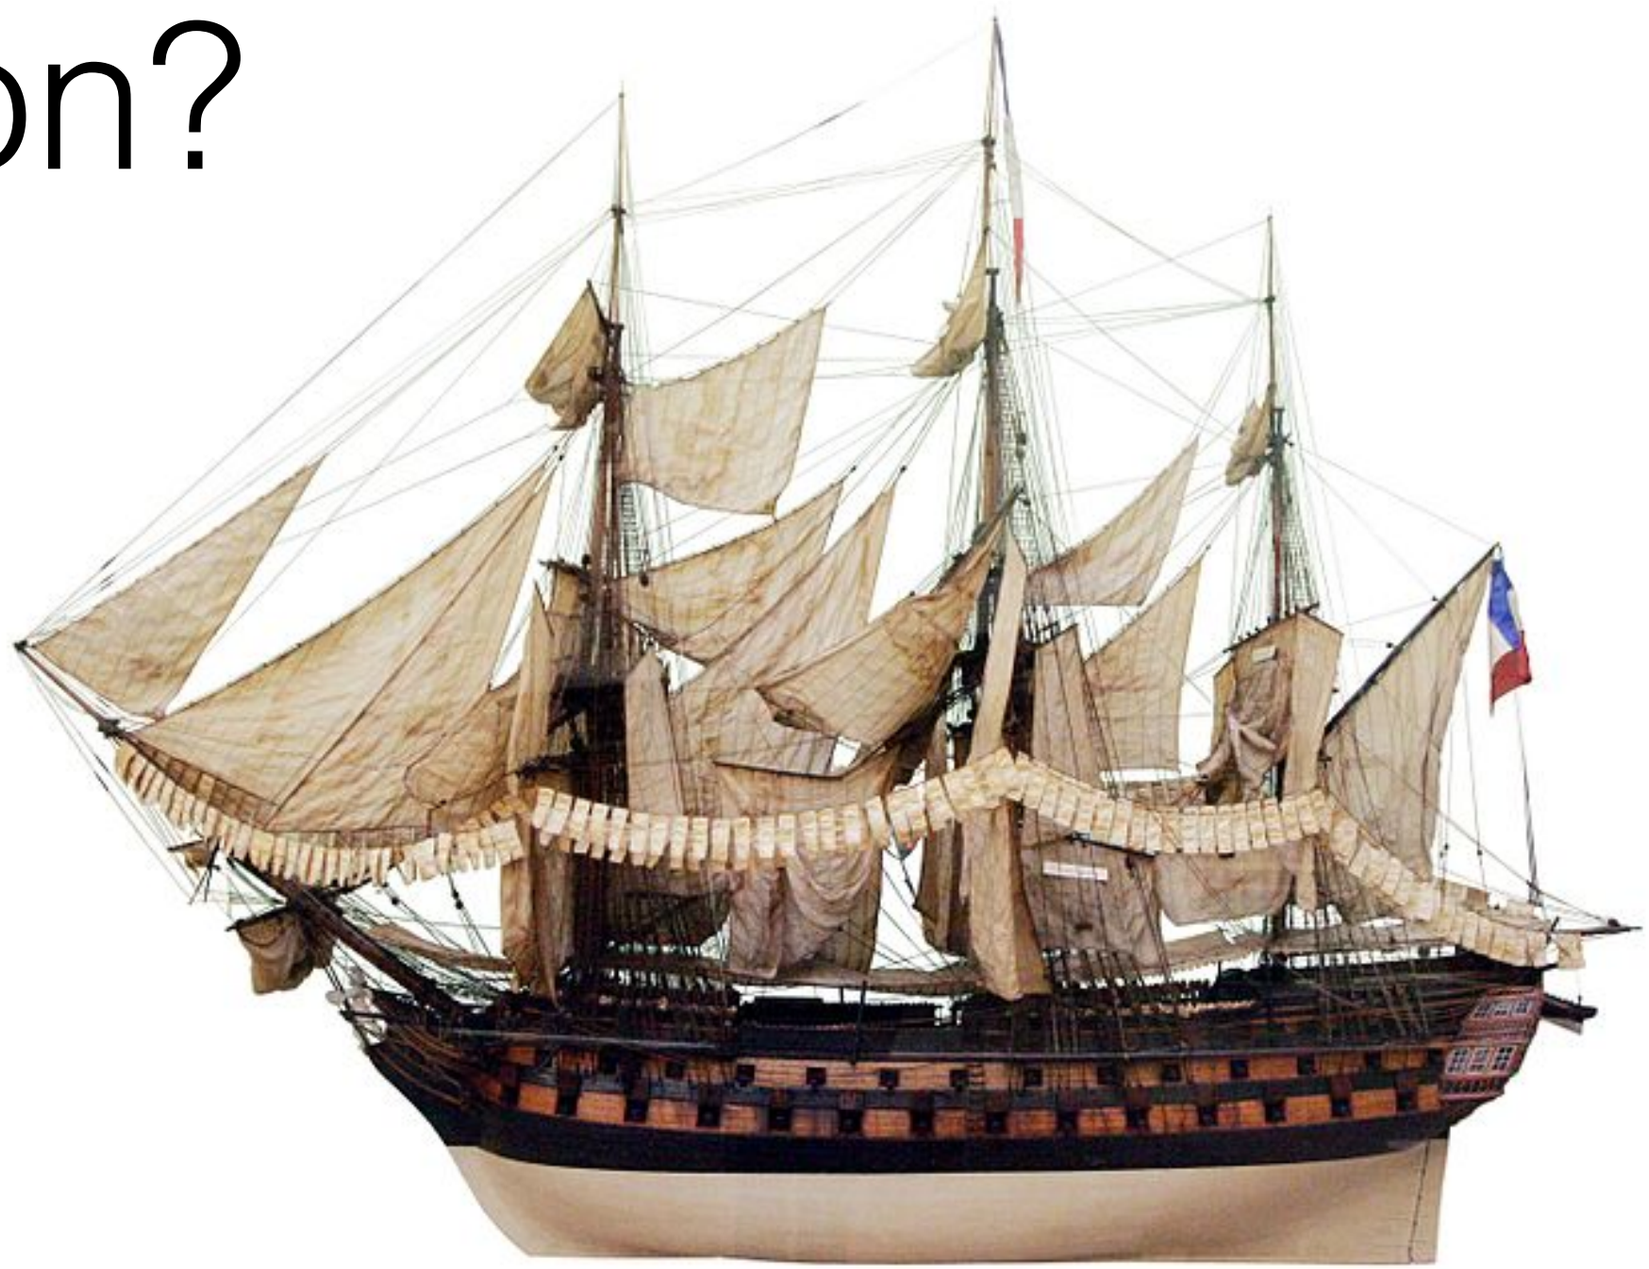

So, what about Aquilon?

- Replaces SCDB (or CDB)
- SQL database plus command line broker
- Templates in git
- Developed by Morgan Stanley
  - Manage ~30K machines (plus VMs, switches, filers etc)
- Now working at RAL - first site outside MS

Aquilon

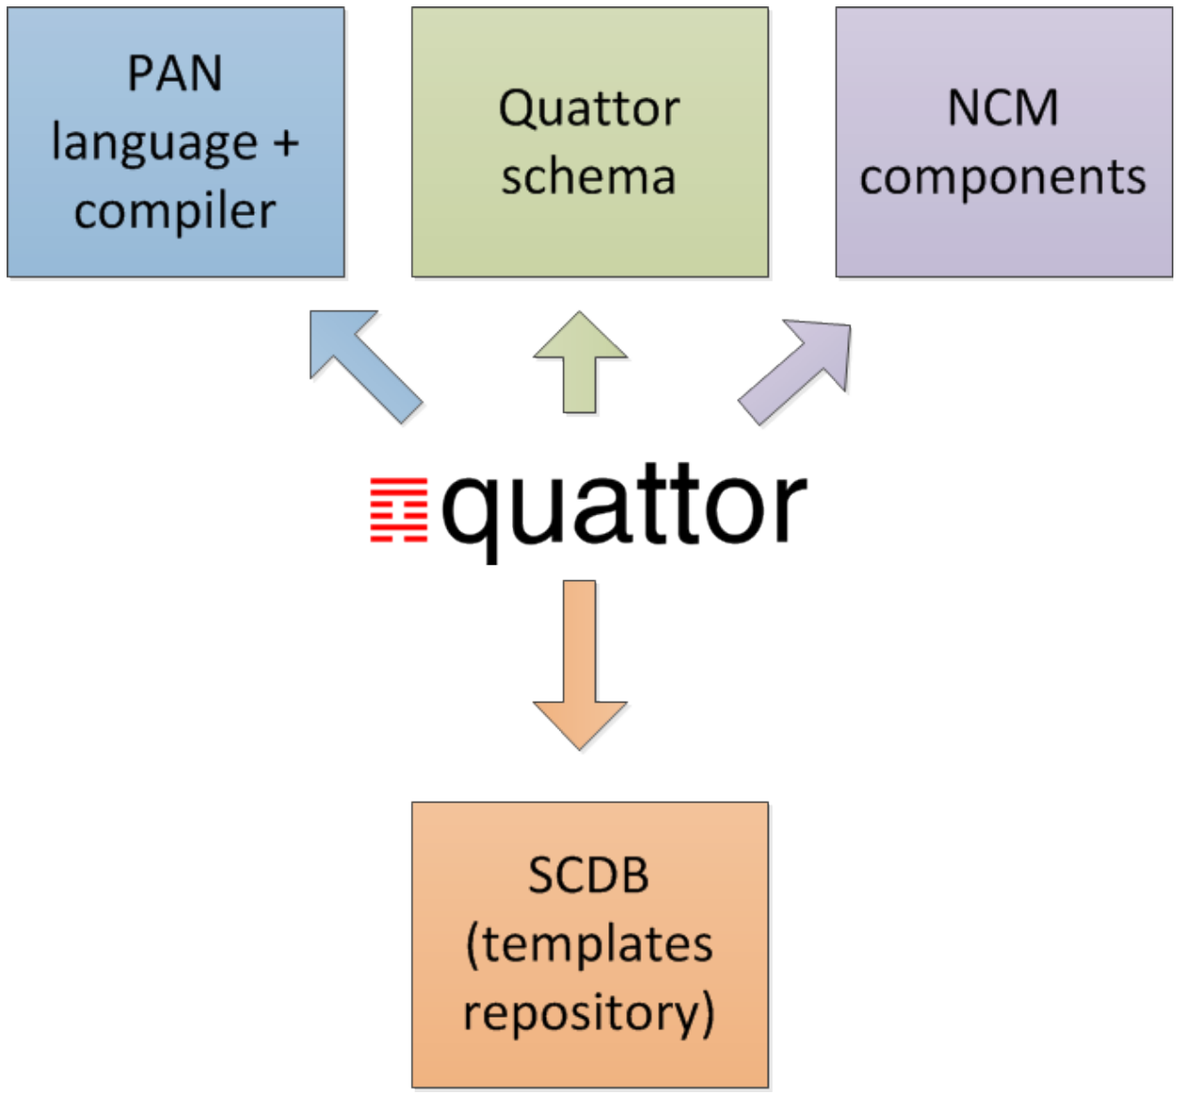

- Most information now in a relational database.
- No more flat files.
- PAN templates when needed for that which cannot be handled by database & broker.

Aquilon

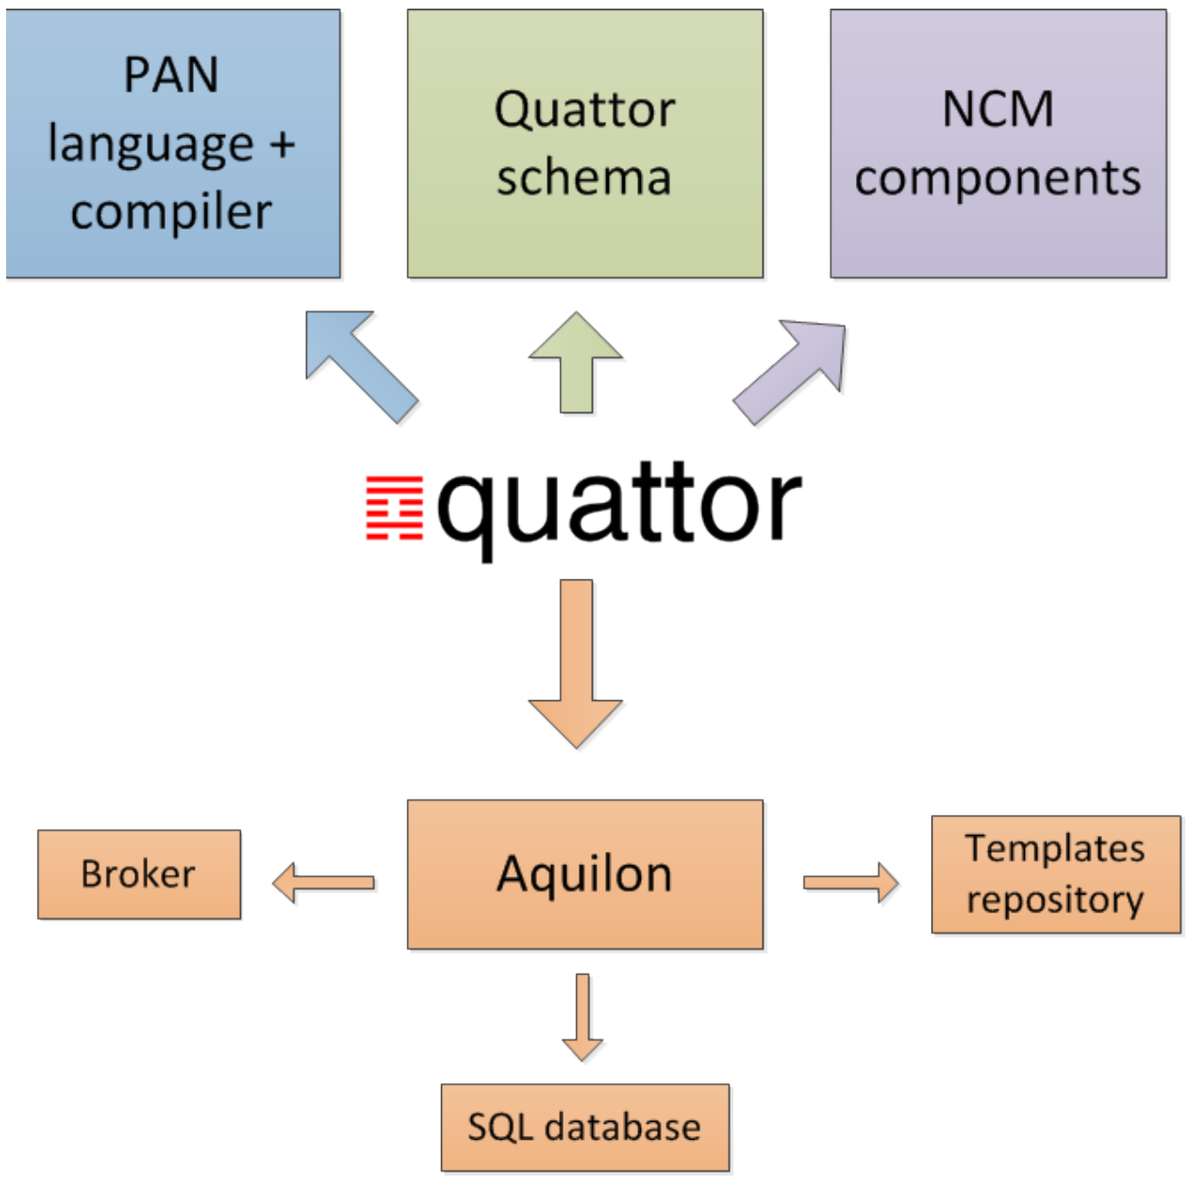

- Most information now in a relational database.
- No more flat files.
- PAN templates when needed for that which cannot be handled by database & broker.

Aquilon Concepts

- Personalities - which are made up of...
- Features
- Services (later slide)
- Command line broker for administering system
  - Setting up hardware, rooms, racks, network etc
  - Binding features to personalities & personalities to hosts

Aquilon Concepts

- Domains and sandboxes

- Domains are groupings of systems with some commonality, but some significant differences

  - Main 'prod' production domain plus however many you need, perhaps used to distinguish main facility from cloud platform, or linux machines from switches & routers

- Sandboxes

  - Essentially git clone where you develop configs & changes. Do not make changes to 'prod' domain, merge in tested changes from sandboxes.

  - You can 'bind' host(s) to a sandbox, only those hosts are compiled or affected by your changes. Development becomes much faster & more flexible.

Aquilon Broker

```
aq show model --all

          aq add_cpu --cpu "xeon_e5520" --vendor "intel"

 aq add rack --rackid r89rack96 --column 13 --row 4 --room ups

  aq search host --ip 130.246.180.88

aq add host --hostname lcg-erasmustest.gridpp.rl.ac.uk

          aq delete continent --continent europe
```

- Deploying 'repeat' systems largely a broker operation.
- Very easy to automate operations
- Pan coding reserved for 'experts'

Aquilon Broker

- Has a new (and so far rather simple) web interface
  - command reference
  - execution of read-only commands
  - investigating web authentication to allow commands which make changes

aquilon

quattor

Aquilon Objects

address	ADD DEL SHOW
alias	ADD DEL SHOW UPDATE
allowed_personality	ADD DEL
application	ADD DEL SHOW
archetype	ADD DEL SHOW UPDATE
auxiliary	ADD DEL SHOW
building	ADD DEL SHOW UPDATE
campus	ADD DEL SHOW
chassis	ADD DEL SHOW UPDATE
city	ADD DEL SHOW UPDATE
cluster	ADD DEL SHOW SEARCH UPDATE
cluster_aligned_service	ADD DEL
continent	ADD DEL SHOW
country	ADD DEL SHOW
cpu	ADD DEL SHOW
disk	ADD DEL
dns_domain	ADD DEL SHOW UPDATE
dns_environment	ADD DEL SHOW
domain	ADD DEL SHOW SEARCH UPDATE

aquilon

quattor

show_machine

Show the hardware configuration of the named machine.

show_machine_opts

show_machine_opts: *

all

show_machine

aquilon

quattor

show_machine

```
system1000
system1001
system1002
system1003
system1004
system1005
system1006
system1007
system1008
system1009
system1010
system1011
system1012
system1013
system1014
system1015
system1016
system1017
system1018
system1019
system1020
```

Aquilon Services

- Aquilon can understand about relationships between servers and clients.
- Defined by broker & stored in database.
- Add a client and Aquilon updates the server config.
- Change a server and client configs are updated.
- Can have rules binding (groups of) clients to specific servers. (Based on locations, functions, etc, etc.)

Other Aquilon tools

- Quattor remote configure
  - Allows description of switches, power controllers, vmware in Quattor, and delivery of configuration without installed client. In use at MS, potentially portable to others using Aquilon.
- Deployment windows
  - Define acceptable time window for e.g. kernel updates
  - Can make global changes knowing that disruptive changes will only take place at planned times.
  - Again in use at MS

Aquilon At RAL

- Since we first looked at Quattor - over 3 years ago - we knew we were interested in Aquilon.
  - Last year layed groundwork with MagDB - site system & network database with schema compatible with Aquilon.
- At RAL we are just beginning to understand workflows and concepts.
- Still need to integrate QWG framework.
  - Concepts different. No Machine types. No host object templates - forces some discipline.
  - To fully benefit will need some re-organisation.
- Need to establish processes for feeding developments back.
  - until now MS have been sole maintainers.
- Next site should be easier.

Summary

- Community active
  - 14th Quattor workshop in Ghent at end of month
- Toolset is healthy & developing
- Aquilon finally being deployed at additional sites
- Possibilities of the new generation of Quattor fabric management are exciting

Questions?