

Managing distributed grid sites with quattor

Stephen Childs (TCD), Marco Emilio Poggi (INFN-CNAF), Andrea Chierici (INFN-CNAF), Michel Jouvin (CNRS-LAL), Charles Loomis (CNRS-LAL), Luis Fernando Muñoz (UAM), Ronald Starink (NIKHEF), Stijn de Weirdt (UGENT)

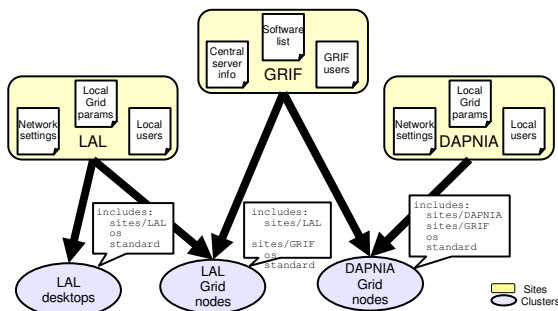
Grids increasingly organise their Tier-2 sites using a distributed model, where resources at multiple physical locations appear to the grid as a single logical site. For example, the GRIF site aggregates resources from six institutions in the Paris region. This arrangement allows the effort of administration to be spread across sites.

The Quattor Working Group (QWG) templates are developed by a consortium of grid sites. They provide an integrated "configuration distribution" for coordinating multiple collaborating sites running gLite middleware. They provide support for **sharing configuration** and **local customisation**.

Sharing configuration

QWG templates are developed using a collaborative model. Contributors check in to a central repository, and periodically releases are tested and tagged. Collaborating sites share the responsibility for developing and updating the configuration of OS and gLite middleware templates.

The configuration for nodes at a particular site is generated by including standard templates for the OS type and gLite node type. Local customisations may be applied without modifying the core grid or OS templates. In the figure below, two physical sites (LAL and DAPNIA) share common configuration information from the GRIF super-site while retaining local autonomy.



Pan configuration language

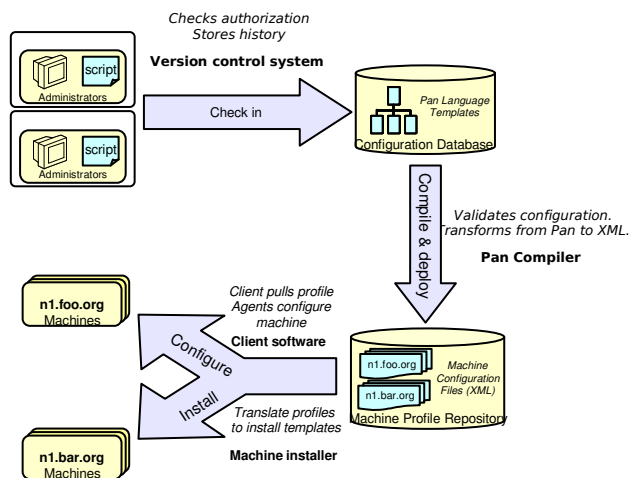
Machine configurations are written in Pan, a high-level, declarative language featuring a simple syntax and the ability to define simultaneously configuration values and a schema. Pan allows for easy development of complex configurations as it supports inclusion of existing configuration modules (e.g. gLite node types), iterating over lists (e.g. to create lists of work nodes), validation of typed variables (e.g. to check that IP addresses are validly formed), and cross-validation between machines (e.g. to ensure coherent configuration of clients and servers).

Current grid deployments

	Distributed				Single-site			Total
	BE Grid	Grid-Ireland	GRIF	CNAF	CERN	Nikhef	UAM	
Nodes	260	410	575	800	8000	300	230	10895
Admins	8	11	20	10	100	4	4	161
Sites	6	18	6	2	1	1	1	34

Full lifecycle management

Quattor supports distributed management over the full lifecycle: a machine's profile is first used to generate Kickstart installation scripts, then the same profile is used for ongoing configuration and maintenance of running systems. Over 200 agents are available for managing various grid and OS services on client nodes. Configuration for monitoring systems such as Nagios and LEMON is also generated automatically from the system information contained in profiles.



Supported services

gLite	BDII, CE, LB, LFC, MONBOX, PX, DPM, DCACHE, UI, VOBOX, VOMX, WMSLB, WMS, WN
OS	SL 3 (32/64), SL 4 (32/64), SL 5 (32/64)
Monitoring	LEMON, Nagios
Standard	Large range of standard Linux services
Virtualisation	Xen (host and guest)

More information

<http://quattor.org> (website) <http://quattor.org/QWG> (QWG wiki)
 "Devolved management of distributed infrastructures with Quattor"
 USENIX LISA '08 (to appear)