Middleware development: fabric management

GridPP 8th Collaboration Meeting
23 September 2003

Alexander Holt
School of Informatics
University of Edinburgh
lex@fixedpoint.org

Overview

- EDG 2.0: LCFGng
- EDG 2.1: no major changes planned
- The new EDG/CERN tools: quattor
  - configuration
  - package management
  - installation
- UK plans for 2004: novel configuration mechanisms
- GridPP-2/EGEE: what configuration system?

EDG 2.0: LCFGng

- RH 7.3
- Server install recipe (Steve Traylen) on GridPP web site
- PXE and floppy installs
- Component changes: `fstab`, Perl components, . . .
- Lots of small improvements

EDG 2.1: no major changes planned

The new EDG/CERN tools: quattor

Retain basic LCFG architecture for configuration, but:

- New language, new components
- Less 'intrusive' package management, Solaris PKG as well as RPM
- Vendor-specific install tools

EDG production use not planned within project lifespan, but already partly in production at CERN, and pilot deployment intended for LCG-1.

`http://cern.ch/quattor/`

quattor: configuration

- New high level language: Pan. More expressive: named templates, types, validation expressions, powerful built-in functions
- Pan files held in database (CDB): akin to API wrapper round CVS repository: graphical tool available (Panguin)
- XML profiles much the same, but more deeply nested
- XML profiles $\Rightarrow$ SQL database convertor
- Client-side cache & notification machinery reimplemented
- Components: Perl, new library (NVA-API), single `Configure()` method

LCFG and Pan example, part 1

LCFG (`disk.h`):

```
disk.deftype     ext2
disk.parts       hda1 hda2
disk.size_hda1   5000
disk.type_hda1   <%disk.deftype%>
disk.mount_hda1  /
disk.size_hda2   512
disk.type_hda2   swap
```

Pan (`disk.tpl`):

```
template disk;
"/disk/deftype"    = "ext2";
"/disk/parts/hda1" = nlist(
  "size",   5000,
  "type",   value("/disk/deftype"),
  "mount",  "/");
"/disk/parts/hda2" = nlist(
  "size", 512,
  "type", "swap");
```

LCFG and Pan example, part 2

LCFG (mysys):

```
#include "disk.h"
!disk.parts     mADD(hda3)
!disk.size_hda1 mSET(7000)
disk.size_hda3  3000
disk.type_hda3  <%disk.deftype%>
disk.mount_hda3 /opt
```

Pan (mysys.tpl):

```
object template mysys;
include disk;
"/disk/parts/hda1/size"  = 7000;
"/disk/parts/hda3" = nlist(
  "size",  3000,
  "type",  value("/disk/deftype"),
  "mount", "/opt");
```

quattor: package management

- Support for RPM and Solaris PKG
- Common management agent on client that feeds packages to `rpmt` or `pkgt` as appropriate (replacement for `updaterpms`)
- Distribution by HTTP, FTP, NFS/AFS
- Local per-node caches, as well as standard HTTP caching
- Can manage subset of all packages (add-on installations, desktops)

quattor: installation

- Unlike LCFG, not tightly integrated with config system
- Relies on vendor install tools: Kickstart & Jumpstart (translators from configuration DB)
- New programs for management of boot servers (PXE, DHCP)

Edinburgh plans for 2004: novel configuration mechanisms

Experiment with:

- 'Autonomous' configuration: harnessing peer-to-peer protocols like SLP (`http://www.srvloc.org/`) and group membership/election schemes within prescribed policy
- More abstract specifications: 'I want a 3-tier web/script/database service with these parameters . . . '
- Delegated management: different aspects of a local site managed independently and combined intelligently through, e.g., constraint satisfaction (cf. merging MON and UI source files in LCFGng)

GridPP-2/EGEE: what configuration system?

- EDG LCFGng trails Edinburgh developments. Future support path unclear
- LCFG: both a production system (1000 machines at Edinburgh) and a vehicle for research & experimentation. Slowly spreading at Edinburgh beyond Informatics; also some US interest. No explicit external support resources, but interested in helping users of current version
- quattor: production requirements dominate; driven by immediate needs at CERN and other EDG sites. Under active development. CERN have made strong commitment $\Rightarrow$ support (& training)
- Other tools. Investment would be required to equal current levels of LCFG-based support for testbed sites.