

GridWeaver

Large-Scale, Adaptive Fabric Configuration for Grid Computing

Kostas Kavoussanakis
George Beckett
EPCC

All-hands meeting, 2-4 September 2003

The GridWeaver Project

- **A collaboration between**
 - Serrano Project, HP Labs Bristol
 - School of Informatics, University of Edinburgh
 - EPCC, University of Edinburgh
- **Combining research interests and technologies from HP Labs (SmartFrog) and Edinburgh (LCFG)**
- **Funded by the UK e-Science Grid Core Programme, and by HP**
 - A 1 year project, to July '03
 - Research-oriented, not building production systems

The Challenge

- An effective Grid assumes the existence of correctly operating, large-scale fabrics
- Every aspect of the fabric must be correctly configured
 - From base OS on a single fabric element, up to complex, coupled, distributed services
- Challenging problems:
 - Scale
 - Diversity
 - Complexity
 - Dynamism
- Aim to:
 - Describe whole-fabric configuration
 - Deliver a correctly-configured fabric
 - Showcase automatic adaptive behaviour

Approach

- **4 Steps:**
 - Review of existing technologies
 - Identification of next-generation requirements
 - Focus on a subset we can investigate further
 - Development of proof-of-concept prototype

Weaponry

- **Five main categories of tools:**
 - **Low-level, single node start-up installation tools**
 - ▶ No support for diversity and change; usually tied to an OS
 - **Multiple node, central database tools**
 - ▶ No support for diverse requirements
 - **Managing change through central database**
 - ▶ No support for dynamic parameters
 - **Integrated packages**
 - ▶ Limited diversity, low-level descriptions
 - **Automated distributed applications**
 - ▶ Right track, none complete

<http://www.gridweaver.org/WP1/report1.pdf>

Status Quo

- **Large Installations depend on:**
 - Home-grown or adapted, piecemeal-developed tools
 - Heroics from staff
- **What the managers want and the fabrics need:**
 - Abstract representations of resources and relationships
 - On-the-fly validation of descriptions and dependencies
 - Devolved, self-resolving configuration descriptions
 - Support for mobile users
 - Automated (if not automatic) management methods

<http://www.gridweaver.org/WP2/report2.pdf>

Technologies: LCFG and SmartFrog

- **A common philosophy...**
 - Declarative language-based approaches for expressing whole-system configuration
 - Frameworks and extensible component sets for realising system configurations
 - “Asymptotic” configuration to deal with scale
- **...with complementary foci**
 - LCFG focuses on configuring and managing individual nodes in a fabric
 - SmartFrog focuses on configuring and orchestrating distributed applications running across nodes

Combining LCFG and SmartFrog

- **Separation of concerns**
 - LCFG: per-node configuration, bare-metal upwards
 - SmartFrog: higher-level, distributed, adaptive services
- **Integrated configuration infrastructure**
 - LCFG configures, starts and manages the SmartFrog daemon on each node
 - SmartFrog controls LCFG components using a generic SmartFrog / LCFG adaptor
- **Planned (but not yet done): unified description language (using SmartFrog language v2)**
 - Complete fabric description using one, powerful representation
 - “Compilation” results in LCFG node configuration profiles plus deployable SmartFrog descriptions

http://www.gridweaver.org/WP3/report3_1.pdf

GridWeaver Prototype

- Bare-metal OS installation, configuration and ignition
 - Basic fabric management
- Automatic installation and configuration of a Globus GT3 (OGSI) infrastructure
 - Grid-enabled fabric
- GPrint: an example adaptive, cross-fabric print service
 - Fabric + adaptive service
- Exposure of GPrint as a Grid Service via Globus
 - Grid-enabled service

http://www.gridweaver.org/WP4/report4_1.pdf



Video Sequence

<http://www.gridweaver.org/docs>



Future Work

- **From a research prototype to a production system**
 - LCFG is open source (GPL)
 - SmartFrog release to open source (LGPL) in '03
 - Work needed to create a portable, production-quality system
- **Grid Configuration Service Interface**
 - A Grid Service interface for fine-grained control over aspects of fabric configuration
- **Many other interesting research problems**
 - Representing time-based properties in configurations
 - Devolving control of different configuration aspects, securely

http://www.gridweaver.org/WP4/report4_2.pdf

More Information

• Contacts

- GridWeaver List: <gridweaver@inf.ed.ac.uk>
- HP: Peter Toft <peter.toft@hp.com>
- University of Edinburgh: Paul Anderson <dcspaul@inf.ed.ac.uk>

• Web pages

- GridWeaver: www.gridweaver.org
- SmartFrog: www-uk.hpl.hp.com/smartfrog
www.smartfrog.org
- LCFG: www.lcfg.org