

GridWeaver

www.gridweaver.org/

Automated Configuration Management for Grid Computing

The success of Grid technology in e-Science, e-Commerce, and e-Government assumes the availability of large-scale, reliable and correctly configured computing installations. GridWeaver focused on the core problems of effective system configuration. Here we discuss GPrint, the research prototype we developed to showcase to the Grid community key concepts and considerations.

The Configuration Testbed

A testbed was built combining the following complementary technologies (see Figure 1):

- LCFG, by the School of Informatics at the University of Edinburgh, is used for the basic installation of the cluster, setup of Linux and installation of essential software applications.
- SmartFrog, by HP Labs, is a framework that coordinates deployment and management of complex, pan-nodal, compound services and applications.

The GPrint demonstrator

Our Grid-enabled Printing Service features:

- installation and configuration of the base infrastructure;
- installation and configuration of the Globus Toolkit Version 3;
- deployment of our GPrint printing service onto the fabric;
- automatic maintenance of a reactive, event-driven service which is responsive to demands for load balancing and reliability;
- exposure of the OGSA-compliant printing service using Globus.

The printing system GPrint is presented in Figure 2. The key component is the Print Controller which:

- discovers printers and print servers on the network;
- establishes bindings between discovered printers, print servers and queues as per the functional specification and the load-balancing requirements;
- responds to printer/print server failure, maintaining an SLA;
- configures LPRng daemons and system files to reflect bindings, permissions, etc;
- catalogues user permissions related to printing.

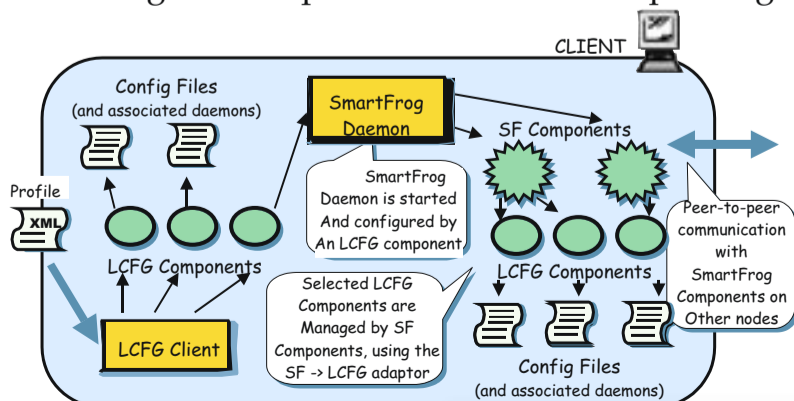


Figure 1: Integration of LCFG and SmartFrog in the testbed. The LCFG Client deploys the SmartFrog daemon, which then coordinates both LCFG and SF components.

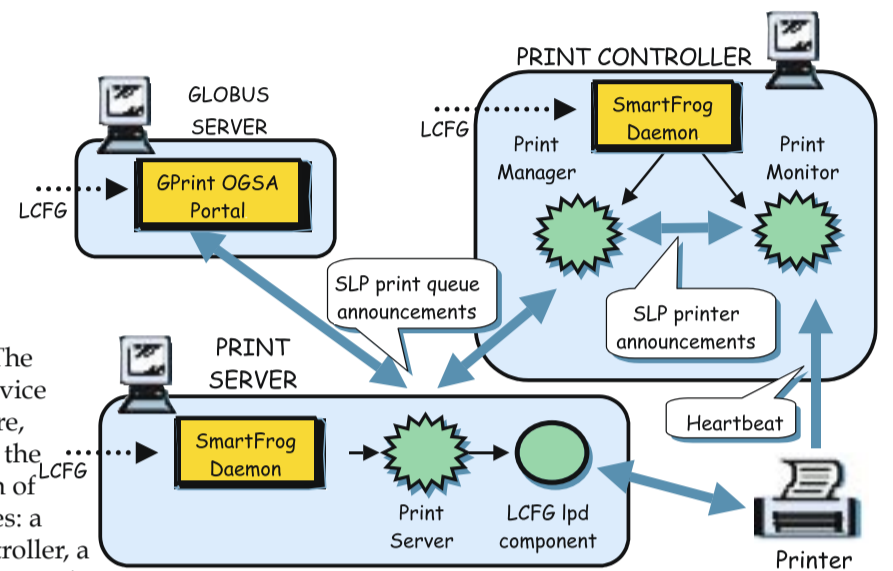


Figure 2: The GPrint service architecture, managing the interaction of three nodes: a Print Controller, a Print Server, and a Globus server.

Demonstrator achievements

The GPrint demonstrator illustrates that the following requirements for large-scale configuration can be satisfied:

- Multi-level representation of system configuration – abstract concepts (eg a queue-to-printer mapping) and real entities (eg a node) coexist; system views filter out irrelevant information.
- Layered deployment – LCFG for low-level nodal configuration, SmartFrog for pan-nodal service configuration.
- Ability to develop reusable component elements, independent from the particular solution.
- Configuration of external agents – proxy component allows new printers to be assimilated seamlessly.
- Exploitation of dynamic information from multiple sources and formats that span the fabric.

What next

We highlight the following areas for further research:

- constraint satisfaction – generation of a valid configuration which satisfies prescribed set of rules;
- investigation of peer election – for example, to eliminate Print Service Manager as a single point of failure;
- configuration of the service through the Grid – provide an OGSA portal to configure the fabric in a secure and reliable manner.

Contact details and acknowledgements

GridWeaver is co-funded by the UK e-Science Core Programme and HP.

The GridWeaver team consists of: Paul Anderson, Jim Paterson, Carwyn Edwards (School of Informatics, The University of Edinburgh), George Beckett, Kostas Kavoussanakis (EPCC, The University of Edinburgh), Guillaume Mecheneau, Patrick Goldsack, Peter Toft (HP Labs Bristol).

You can email the team at gridweaver@inf.ed.ac.uk.