

GridWeaver

Large-scale, adaptive fabric configuration for Grid Computing

www.gridweaver.org

Introduction

The success of Grid technology in e-Science, e-Commerce, and e-Government depends upon the availability of reliable and correctly configured computing installations.

Typical characteristics of such installations include:

- A large number of component nodes (up to tens of thousands);
- Extreme levels of hardware and software diversity distributed across multiple sites or virtual organisations;
- Complex dependencies between software and hardware elements;
- Massive levels of dynamism with a perpetual flow of configuration;
- Devolved control and responsibility between personnel for aspects of the infrastructure.

The GridWeaver project has conducted research into the problems which are at the heart of an effective configuration system. The project has achieved the following goals:

- Surveyed a representative set of current configuration systems and tools;
- Conducted case study investigations of current and projected large-scale computing installations;
- Explored the language/tools of configuration specification, highlighting key linguistic/modelling constructions and generating simple but illustrative case study models;
- Integrated two existing complementary technologies, contributed by the partners (LCFG and SmartFrog);
- Constructed a research prototype to test and demonstrate key concepts;
- Proposed important research targets to advance further the community's understanding and capabilities in the field.

Configuration technology today

The GridWeaver team has conducted a survey of the state-of-the-art in configuration technology. We identified a representative subset of existing applications and appraised the scope of functionality.

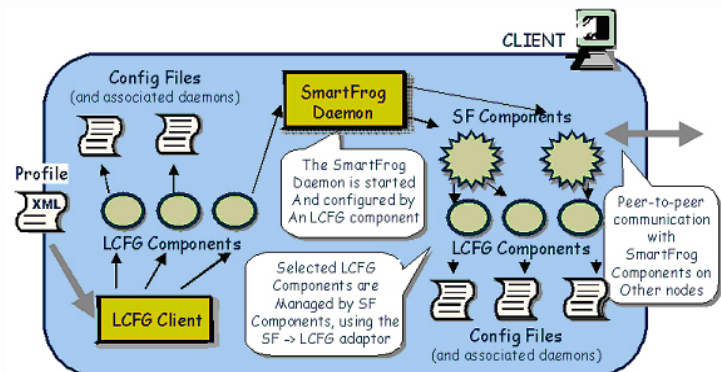
<http://www.gridweaver.org/WP1/report1.pdf>

The team also interviewed key personnel involved with real large installations and additionally generated realistic "Projected Use Case" scenarios.

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

Based on these investigations, we conclude that the existing tools and the underlying approaches to configuration are inadequate, because:

- Many tools perform only basic installation tasks, focusing on client nodes;
- Too much manual intervention is required, with high dependence on technical expertise of staff;
- Automatic, on-going configuration maintenance is rarely considered;
- There is no tool offering a mature configuration language which supports multiple levels of abstraction in the model;
- There is little consideration of pan-nodal relationships and interactions;
- Any significant increase in scale or complexity would necessitate a migration to a more autonomic environment.



A configuration architecture

The team has defined key aspects of a configuration description and illustrated them using three case study models: a Globus installation supporting OGSA services; a network printing system; and a DHCP/DNS naming service.

We propose that a configuration environment manage the whole fabric (Servers and Clients), using a Peer-To-Peer philosophy to avoid scaling constraints. We assert that system description should be a *Declarative Specification*, expressing pan-nodal relationships, not grounded in node-centric expression of configuration.

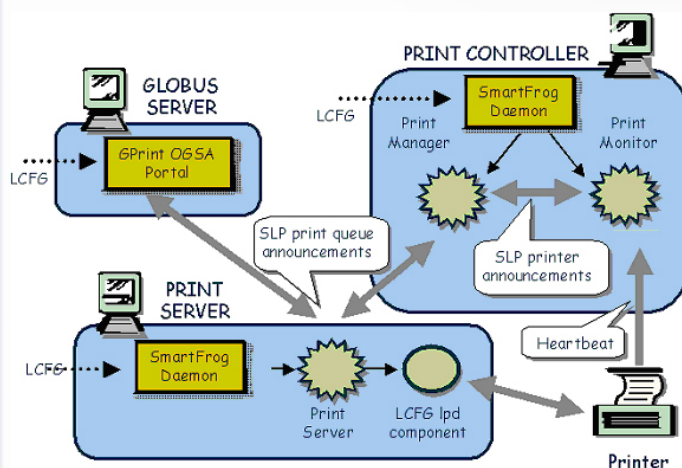
We also identify the need for autonomic maintenance of system integrity. The fabric reconfigures in response to:

- Changes to a high-level specification;
- Events like component failure or appearance of new resources.

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

We combined two complementary technologies to produce a testbed on which to experiment:

- 1) LCFG—Informatics' system for managing configuration of a diverse cluster of workstations;
- 2) SmartFrog—HP's experimental framework for configuring, deploying, and orchestrating distributed applications.



GPrint prototype

The team has developed GPrint, a research prototype which demonstrates and exercises the achievements of the GridWeaver project. The prototype has the following key features:

- A declarative configuration is derived from a succinct, modular, multi-level model, expressed in the languages of SmartFrog and LCFG;
- It allows "bare metal" OS installation, configuration and ignition;
- It automatically configures and instantiates an OGSA-enabled environment (Globus Toolkit 3);
- It installs and deploys a demonstration, OGSA-compliant service (Gprint-an adaptive, pan-nodal printing system);
- It automatically responds to events such as the appearance of new resources or the failure of existing elements, maintaining a service level agreement.

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

A video sequence of GPrint may be viewed at:

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

Further work

Based on the outcomes of the project, the consortium is well placed and keen to:

- Propose a unified configuration description language;
- Incorporate the representation of temporal aspects into the configuration language;
- Investigate security issues;
- Address automatic constraint satisfaction.

The GridWeaver team

GridWeaver is a collaboration between:

The School of Informatics, the University of Edinburgh;
EPCC, the University of Edinburgh;
HP Laboratories (Bristol, UK).

The project was funded by the UK e-Science Grid CoreProgram and HP, had a duration of one year and finished in July 2003.