

The GridWeaver Project: Summary and Conclusions

*Paul Anderson <dcspaul@inf.ed.ac.uk>
Peter Toft <peter.toft@hp.com>*

Abstract

This document is a final report for the GridWeaver project. It provides a summary of the work undertaken, and the conclusions from the main workpackages. It also draws some overall conclusions and presents a suggested roadmap for future research and developments in fabric management.

The logo for GridWeaver features the word "GridWeaver" in a blue, sans-serif font. A green, wavy line arches over the text, starting from the left, passing over the 'i' and 'd', and ending under the 'r'.

GridWeaver

Revision 1.0 – July 31, 2003

Chapter 1

The GridWeaver Project

The GridWeaver project is a collaboration between:

- The School of Informatics, Edinburgh University
- HP Laboratories, Bristol
- Edinburgh Parallel Computing Centre (EPCC), University of Edinburgh

The motivation for this project was based on the observation that:

An effective Grid infrastructure assumes the existence of reliable, and correctly configured fabrics, on which the middleware and applications can be hosted.

We suspected that current fabric management technology was unable to deliver these requirements and to provide the necessary foundations for the next generation of Grid applications. Furthermore, we believed that this was not simply a question of implementation; rather, we suspected that there would need to be a fundamental paradigm change in the way that most fabrics are currently managed, and in the design of the tools used to support this management.

Hence, we aimed to:

- Develop a good understanding of the current state-of-the-art in fabric management technology. In particular, to identify, classify and assess the different approaches, and to compare their strengths and weaknesses with the tools already developed by two of the project partners – SmartFrog from HP Labs, and LCFG from Informatics.
- Identify the management requirements of the next generation of Grid Fabrics, particularly those requirements which would require a fundamental change in approach, rather than a simple increase in scale.
- Investigate how the SmartFrog and LCFG tools might be used to address some of the problems identified above. In particular to investigate whether a combination of the complementary properties of these two tools might be used to demonstrate new approaches.

- Develop a proof-of-concept implementation (not a production system) to illustrate and evaluate some of the above approaches.
- Produce a “roadmap” identifying the major research directions and future developments which we believe necessary to deliver an effective fabric management infrastructure for the next generation of Grid applications.

The GridWeaver project has been extremely successful in meeting these aims, and this document is the final project report.

Chapter 2 provides a summary of the workpackages which cover the first four of the above aims; this includes a summary of the work undertaken, and the conclusions for each workpackage. More details are available in the full reports for each workpackage. Chapter 3 sets out the overall conclusions from the project, and proposes a roadmap for future developments. The appendices A and B list the project deliverables and describe the dissemination work.

Acknowledgements

This project would not have been so successful without the contributions of the whole project team, whose efforts were often well in excess of the original plan. This includes Paul Anderson, George Beckett, Kostas Kavoussanakis, Guillaume Mecheneau, Jim Paterson, and Peter Toft.

Patrick Goldsack from HP also provided expertise on the SmartFrog system and was an invaluable source of general wisdom on configuration issues. Carwyn Edwards from Informatics provided LCFG expertise and much useful work on the portability and exporting of the system. A number of other people also made helpful contributions, including Lex Holt and John Hawkins.

GridWeaver is part of the UK e-Science Core Programme, and we gratefully acknowledge the programme’s assistance with this activity. We would also like to thank the National eScience Centre for their help and encouragement.

For administrative purposes, GridWeaver is also known under the project name *HPFab-Man*.

Chapter 2

Project Summary

This chapter summarises the work and conclusions from each of the individual workpackages of the GridWeaver project. Further details are available in the separate reports which form part of the deliverables for each package.

2.1 Workpackage 1

The report *Technologies for Large-Scale Configuration Management* [ABK⁺02] describes a comprehensive survey of the current state-of-the-art in system configuration, including a critical examination of LCFG and SmartFrog. This report serves two purposes:

- It establishes a baseline for future research by understanding the approaches, technologies and best-practises that are available today.
- It captures this knowledge in a form that is useful to others wanting a general review paper to introduce them to the area.

As a pre-requisite to the technology survey, the report includes a definition of the scope and meaning of “system configuration” as addressed by the GridWeaver project. There is also a discussion of the different techniques and principles, which helps to classify aspects of the various tools, and a glossary.

The intention of the survey itself is to cover tools which are representative of all the various different techniques and approaches. In some cases, there are many tools which adopt a similar approach, and, in these cases, only one or two representative tools are discussed in detail.

The report identifies a hierarchy of approaches in system management tools, ranging from simple low-level installation (which is supported by numerous different products), through to more sophisticated approaches which attempt to specify and manage the evolving relationships between the nodes in a fabric. LCFG and SmartFrog are both identified as tools which are attempting to address some of these higher-level problems.

2.2 Workpackage 2

The report *Experiences and Challenges of Large-Scale System Configuration* [ABK⁺03b] attempts to identify the current and future problems with large-scale configuration management technologies. Two approaches are used:

- A number of case studies were conducted with managers of existing large fabrics of various different types. The managers were questioned about their current requirements and tool usage, as well as their thoughts on future directions.
- A number of “projected” case studies are described. These present what we consider to be “ideal” scenarios for managing typical future fabrics. The studies are designed to show the way in which system managers would like to control their fabric configuration, based on some of the current difficulties identified in the real case studies.

The report concludes that there is a significant gap between the capabilities of existing configuration tools and the requirements of fabric managers. Existing tools are clearly being used to manage existing fabrics, but they are barely adequate; there is a high chance of errors due to incorrect configurations, configurations are inflexible and unreliable, and there is a considerable dependence on manual intervention. Significant increases in fabric scale and (particularly) complexity are likely to intensify these problems.

A useful classification is made between the *specification*, *manipulation* and *deployment* of configurations. Fundamental problems are identified in each of these areas and a small number are selected for further investigation by the project.

2.3 Workpackage 3

The report *Large-Scale System Configuration with LCFG and SmartFrog* [ABK⁺03a] describes a deeper investigation into some of the above problems, using a combination of LCFG and SmartFrog to explore possible solutions. Two (largely independent) investigations were undertaken, reflecting the major areas of “specification” and “deployment” identified in the previous section:

- The richer specification language available in SmartFrog is used to represent some of the configuration idioms which have proven important and often awkward in practical use with LCFG. The use of SmartFrog is also explored to represent a number of typical areas of system configuration.
- A testbed is designed which allows SmartFrog and LCFG components (and to some extent, specifications) to be used interchangeably. This allows the construction of applications which can select the most appropriate technology for different aspects of the application. For example, existing LCFG components can be used to install a node from scratch and manage the low-level configuration. These components may then be controlled by SmartFrog components which bring the additional capability of peer-to-peer communication, and a richer specification language.

This report illustrates that the combination of technology from LCFG and SmartFrog enables the development of prototype configuration tools which go some way towards addressing identified problems in both the specification and deployment areas. In particular:

- The SmartFrog specification language provides a clearer way of expressing many of the practical system configurations found in typical LCFG installations. Particularly, it provides a way of describing higher-level entities and inter-node relationships.
- The use of peer-to-peer techniques from SmartFrog to control the practical LCFG components allows the construction of practical, autonomic systems with a high degree of fault tolerance.

2.4 Workpackage 4

The report *The GPrint Demonstrator* [BMP02] presents the design of a complete OGSA-enabled print service, built using the framework described in the previous section.

The entire cluster can be installed and configured from scratch using SmartFrog and LCFG configuration descriptions. This provides a multi-server, multi-printer print service with an OGSA portal which can accept print jobs from the Grid. The cluster is highly reliable and can reconfigure itself automatically to handle new or failing printers, or print servers. A video [Tof03b] demonstrating these features is available online.

The importance of this demonstrator is not in the printing service example itself, but in the ease with which fault-tolerant distributed services can be constructed, deployed, and configured using the principles developed from LCFG and SmartFrog.

Conclusions and Roadmap

This report itself [AT03] is part of Workpackage 4 and the following chapter presents some overall conclusions and a roadmap for future developments.

Chapter 3

Conclusions

This chapter presents a summary of the overall conclusions of the GridWeaver project: Firstly, it outlines a generally accepted configuration architecture that we believe is likely to form the basis of future configuration systems. It then summarises the outstanding problems that have been identified in current implementations of the architecture, and describes the explicit contribution of the GridWeaver project towards the solution of some of these problems.

3.1 A Configuration Architecture

The following overall architecture has been adopted by several practical configuration tools, including LCFG and others which are particularly focused on the management of entire fabrics.

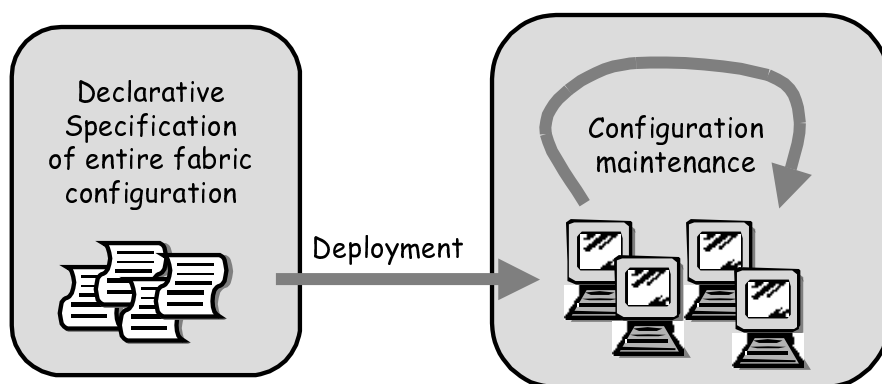


Figure 3.1: A Generic Configuration Architecture

- A **declarative description** is maintained for the configuration of the whole fabric. This describes the desired state of the entire fabric, including all configurable

devices, and the relationships between them. It does not describe procedures for implementing these configurations.

- Some tool is responsible for **deploying** this configuration by modifying the actual configuration of the nodes to match the specification.
- The same tool, or a different one, is responsible for responding to changes in the fabric (eg. failures) or the specification (eg. manual changes) and modifying the node configurations to **maintain the correspondence** between the actual and desired configuration.

3.2 Outstanding Problems

The above general model has many desirable properties which are discussed in more detail in the individual reports. The GridWeaver project has confirmed that this is an appropriate model for large-scale fabric configuration, but it has also shown that current instances of the model have fundamental problems in both the specification, and the deployment areas. A few of the most important of these problems are summarised below:

- Specification Descriptions
 - Current descriptions are too low-level.
 - Current descriptions are too explicit, allowing no room for autonomic fault recovery, or other dynamic changes.
 - Current descriptions do not represent the temporal information required to sequence changes reliably and automatically.
- Specification Use
 - Current descriptions do not adequately support the devolved management that is essential for large fabrics, either in terms of authorisation, or in terms of aspect composition.
 - Current languages are difficult to understand and use, making them error prone and unsuitable for unskilled staff.
- Deployment
 - Centralised control over all deployment does not scale.
 - Current systems have no uniform support for workflows which would allow a smooth change between configuration states.
 - Current systems do not support inter-node transactions on configuration changes, which would allow rollback on failure.
 - Current systems do not provide the ability for localised autonomic reconfiguration under a central policy control. This is necessary both for scaling, and for autonomic fault recovery.

The GridWeaver project has investigated some of the above issues in more detail, and provided prototypes for possible solutions. In particular:

- The use of the SmartFrog framework to provide a peer-to-peer layer above LCFG components has demonstrated a system which removes much of the central control and provides a scalable system with autonomic fault-tolerance, and central control over the policy [AGP03].
- The use of the SmartFrog language provides a way to describe LCFG configurations in a clear and high-level way.

However, the GridWeaver project has not attempted to address all of the above problems, and has not attempted to provide production-quality tools for those problems which it has addressed.

3.3 Recommendations

We believe that it would be possible to implement a new production fabric management system, based on the best of the currently available technology, and augmented by some of the techniques developed during this project. This should provide a significant improvement on existing tools. However, the development of production systems is extremely expensive, and we believe that some of the currently unsolved problems would still render such a system unable to fully meet the requirements of the next generation of Grid fabrics.

We would like to see further fundamental research into a number of areas before considering implementation of a new configuration tool. These areas include:

- Languages for configuration specification, including constraint specification, levels of representation, and support for devolved management of aspects.
- Integration of more peer-to-peer technology into the language and deployment architecture.
- Specification of temporal dependencies, and planning and deployment of the workflows necessary to effect smooth changes between configuration states.

Appendix A

Deliverables

D1.1 Report: Technologies for Large-Scale Configuration Management [[ABK⁺02](#)]

D2.1 Report: Experiences and Challenges of Large-Scale System Configuration [[ABK⁺03b](#)]

D3.1 Report: Large-Scale System Configuration with LCFG and SmartFrog [[ABK⁺03a](#)]

D4.1 Report: The GPrint Demonstrator [[BMP02](#)]

D4.2 Report: The GridWeaver Project: Summary and Conclusions [[AT03](#)]

Appendix B

Dissemination

- GridWeaver: exploring automated configuration and management for Grid computing fabrics “All Hands 2002” Poster [[MBH⁺02](#)].
- Automatic System Configuration: Tasks, Principles and Techniques. Presentation for the LISA 2002 Large-Scale System Configuration Workshop. [[And](#)]
- GridWeaver: Large-Scale, Adaptive Fabric Configuration for Grid Computing. Presentation for the GGF Production Grid Management Research Group [[Tof03a](#)]
- GridWeaver: Technology demo. Streaming video [[Tof03b](#)].
- The GridWeaver Project. Edinburgh University, Informatics Presentation [[AP03](#)].
- The Future of Fabric Management. Presentation for the 7th GridPP Collaboration Meeting. 30th June 2003.
- GridWeaver: Automatic, Adaptive, Large-Scale Fabric Configuration for Grid Computing. “All Hands 2003” Paper. [[APE⁺03](#)]
- SmartFrog meets LCFG - Autonomous Reconfiguration with Central Policy Control. LISA 2003 paper [[AGP03](#)].

Bibliography

- [ABK⁺02] Paul Anderson, George Beckett, Kostas Kavoussanakis, Guillaume Meche-
neau, and Peter Toft. Technologies for large-scale configuration manage-
ment. Technical report, The GridWeaver Project, December 2002.
<http://www.gridweaver.org/WP1/report1.pdf>.
- [ABK⁺03a] Paul Anderson, George Beckett, Kostas Kavoussanakis, Guillaume Meche-
neau, Jim Paterson, and Peter Toft. Large-scale system configuration with
LCFG and SmartFrog. Technical report, The GridWeaver Project, July
2003.
<http://www.gridweaver.org/WP3/report3.1.pdf>.
- [ABK⁺03b] Paul Anderson, George Beckett, Kostas Kavoussanakis, Guillaume Meche-
neau, Peter Toft, and Jim Paterson. Experiences and challenges of large-
scale system configuration. Technical report, The GridWeaver Project,
March 2003.
<http://www.gridweaver.org/WP2/report2.pdf>.
- [AGP03] Paul Anderson, Patrick Goldsack, and Jim Paterson. SmartFrog meets LCFG
- autonomous reconfiguration with central policy control. In *Proceedings
of the 2002 Large Installations Systems Administration (LISA) Conference*,
Berkeley, CA, 2003. Usenix. To be published.
- [And] Paul Anderson. Automatic system configuration: Tasks, principles and
techniques. Presentation at LISA 2003 Large-Scale System Configuration
Workshop.
[http://homepages.informatics.ed.ac.uk/group/ ...
lssconf/config2002/slides/anderson.pdf](http://homepages.informatics.ed.ac.uk/group/...
lssconf/config2002/slides/anderson.pdf).
- [AP03] Paul Anderson and Jim Paterson. The Gridweaver Project. Informatics
Talk, June 2003.
[http://www.epcc.ed.ac.uk/gridweaver/docs/ ...
Gridweaver_3_7_03.pdf](http://www.epcc.ed.ac.uk/gridweaver/docs/...
Gridweaver_3_7_03.pdf).
- [APE⁺03] Paul Anderson, Jim Paterson, Carwyn Edwards, George Beckett, Kostas
Kavoussanakis, and Peter Toft Guillaume Meche-
neau, Patrick Goldsack. Gridweaver: Automatic, adaptive, large-scale fabric configuration for grid
computing. All Hands Conference, September 2003. To be published.

- [AT03] Paul Anderson and Peter Toft. Conclusions from the GridWeaver project. Technical report, The GridWeaver Project, July 2003.
http://www.gridweaver.org/WP4/report4_2.pdf.
- [BMP02] George Beckett, Guillaume Mecheneau, and Jim Paterson. The GPrint Demonstrator. Technical report, The GridWeaver Project, December 2002.
http://www.gridweaver.org/WP4/report4_1.pdf.
- [MBH⁺02] Guillaume Mecheneau, George Beckett, Alexander Holt, Peter Toft, and Kostas Kavoussanakis. Gridweaver: exploring automated configuration and management for grid computing fabrics. Poster for "All Hands 2002" Conference, 2002.
[http://www.epcc.ed.ac.uk/gridweaver/docs/ ...
gridweaver_poster.pdf](http://www.epcc.ed.ac.uk/gridweaver/docs/...gridweaver_poster.pdf).
- [Tof03a] Peter Toft. Gridweaver: Large-scale, adaptive fabric configuration for grid computing. Presentation for the GGF Production Grid Management Research Group, June 2003.
[http://www.epcc.ed.ac.uk/gridweaver/docs/ ...
GridWeaver-GGF8-PGM-RG-Workshop.pdf](http://www.epcc.ed.ac.uk/gridweaver/docs/...GridWeaver-GGF8-PGM-RG-Workshop.pdf).
- [Tof03b] Peter Toft. Gridweaver technology demo. Realplayer Video, 2003. A Realplayer v8 version of the video is also available in the file `gridweaver-v8.rm`.
[http://boombox.ucs.ed.ac.uk/ramgen/ ...
informatics/gridweaver.rm](http://boombox.ucs.ed.ac.uk/ramgen/...informatics/gridweaver.rm).