



Paul Anderson, School of Informatics
Kostas Kavoussanakis, EPCC
Peter Toft, HP Laboratories

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

Contents

- Project goals
- Questions from the last review
- Progress update
- GridWeaver2 proposal

GridWeaver project goals

- Explore the challenges of large-scale system configuration, especially as applied to Grid computing fabrics
 - Identify the current state-of-the-art in configuration technologies and approaches
 - Identify the configuration requirements of current and future fabrics
 - Identify the key research challenges
- Prototype solutions for a subset of problems
 - Build on insights and technologies from Edinburgh LCFG system and HPL SmartFrog
 - Develop prototypes to test our ideas

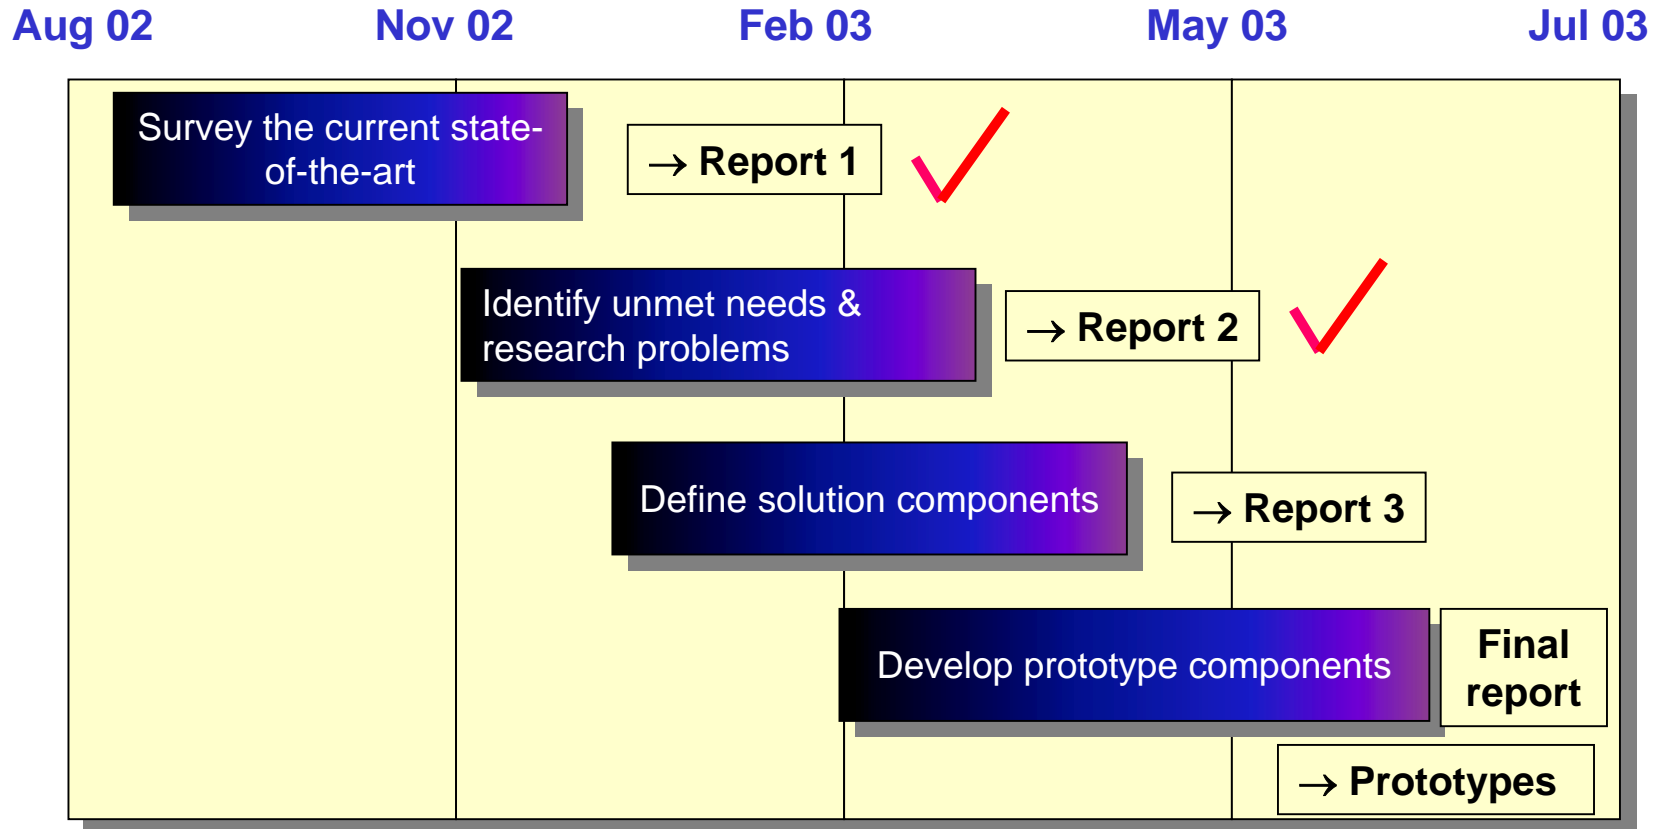
Questions from the last review

- **Connect with Les Robertson at CERN**
 - Paul Anderson works with CERN
 - ▶ LCFG used by EDG
 - ▶ CERN is a featured case-study in our second report
 - Peter Toft also discussed the project with David Foster, and with Sverre Jarpe of CERN
 - HP is part of CERN OpenLab – presents opportunities
- **Engage with GGF (e.g. a BoF at GGF7)**
 - GGF participation written into our contract for last 6 months of the project
 - GGF7 participation would have small returns
 - GGF8 participation planned

Questions from the last review

- **Relevance / connections between GridWeaver and other GCP activities**
 - Not developing production software that can be used right now
 - Configuration problem can be decoupled / is orthogonal to GCP activities
 - We are anticipating the fabric problems we will face when getting from 10-piece Grids to +1,000-piece Grids
 - ▶ Report 2, "Experiences and Challenges of Large-Scale System Configuration" lists these problems

Progress update



NOTES

- We are 30% behind schedule according to the original workplan
- We will produce all the specified deliverables by project conclusion
- We invested more time than originally anticipated in the reports 1 & 2
 - ▶ No regrets; necessary, excellent work
- Overlap of workpackages not captured in the original workplan

Progress update: deliverables




Report 1: "Technologies for Large-Scale Configuration Management"

December 2002

Report 2: "Experiences and Challenges of Large-Scale System Configuration"

March 2003



Progress update: plans

- **Research problems**

- Report 2 identifies a large set of research problems
- **Initial focus areas for GridWeaver chosen**
 - ▶ Configuration language – representing complex configurations
 - ▶ Configuration system architecture – initial architecture for the configuration deployment system
 - ▶ Prototype:
 - Configure parts of Globus Toolkit 3
 - » Show we can do it, rather than solve the problems
 - Provide a Print Webservice on it (optional)
 - » Build in resilience of fabric in face of HW failures

Progress update: plans

- **Showcase our work at GGF8**
 - Exact form of participation is to be confirmed
 - Options (probably all of):
 - ▶ Role in “Production Grid Management” Research Group
 - ▶ Paper(s) in the corresponding Workshop
 - ▶ Demo
- **The level of involvement that we can achieve in the RG depends on follow-on funding**
 - We have original, on-topic contributions to make
 - GridWeaver2 can fund participation in 4 more GGFs

GridWeaver2 proposal

- GridWeaver1 is primarily an exploratory activity
- It has exceeded our expectations in the:
 - Potential of the technical approach
 - Quality of partnership between collaborators
 - Opportunities in this space
- GridWeaver2 Proposal submitted for continuation
 - GW1: $15 + 6 + 12$ PMs = 33 PMs over 1 year
 - GW2: $44 + 30 + 24$ PMs = 98 PMs over 2 years

Backup

GridWeaver: large-scale configuration

- An effective Grid assumes the existence of correctly operating large-scale fabrics, on which the middleware and applications can be hosted.
 - What happens if one node in 500 has the wrong version of a math library?
 - What happens if we are unable to deploy a security patch quickly enough?
- We define *configuration* to include:
 - Initial software installation
 - Initial customisation of the nodes
 - Node reconfiguration to implement changes in the specified configuration
 - Monitoring and correction of unwanted changes in the actual node configuration
- There are many unsolved problems as fabrics become larger in scale, more heterogeneous, more dynamic and more complex

GridWeaver approach

- We believe the way forward for very large scale configuration requires:
 - Models of the systems being configured
 - Languages to describe the desired configuration of the models
 - Systems to deploy the specified configurations on the target hardware
 - Systems to monitor and correct discrepancies between the specified configuration and the actual configuration
- Our technology assets
 - LCFG from Edinburgh School of Informatics for per-node configuration
 - *SmartFrog* from HP Labs includes a specification language and a deployment engine for distributed applications
- The aim of the project is to use the combined experience of LCFG and SmartFrog to investigate solutions to some of the above problems.