JISC DEVELOPMENT PROGRAMMES
Project Document Cover Sheet

Project Plan

Project

Project Acronym	GROWL	**Project ID**	None Assigned
Project Title	VRE Programming Toolkit and Applications (GROWL: Grid Resources on Workstation Library)		
Start Date	2005-02-01	**End Date**	2006-07-31
Lead Institution	University of Cambridge, e-Science Centre		
Project Director	Mark Hayes		
Project Manager & contact details	John Kewley, j.kewley@dl.ac.uk CCLRC e-Science Centre, Daresbury Laboratory, Keckwick Lane, Daresbury, Warrington, Cheshire, WA4 4AD tel: 01925 603513 fax: 01925 603634		
Partner Institutions	CCLRC Daresbury Laboratory, e-Science Centre University of Lancaster, e-Science Centre		
Project Web URL	http://www.growl.org.uk		
Programme Name	Virtual Research Environments (VRE) 5/04		
Programme Manager	Maia Dimitrova, JISC Programme Manager (VRE), 020 7848 2569, m.dimitrova@jisc.ac.uk		

Document

Document Title	GROWL: Project Plan		
Author(s) & Role	Mark Hayes (Project Director), 01223 765251 John Kewley (Project Manager), 01925 603513 Rob Allan (Associate Director), 01925 603207 Rob Crouchley (Associate Director), 01524 593161 Lorna Morris, 01223 765518 Ties van Ark, 01524 594818		
Date	2005-04-27	**Filename**	GROWL_ProjectPlanFinal-1.pdf
URL	www.growl.org.uk		
Access	☒ Project and JISC internal	☐ General dissemination	

Document History

Version	Date	Comments
Draft-1	2005-03-10	Draft submission to JISC
Final-1	2005-04-27	Submission to JISC for final comments

Contents

List of Figures

Overview of the Project

1 Background

Since the mid 1970s, computational scientists have seen an explosion in the power and functionality of the high-end systems to which they have had increasing access. Most developers of scientific software are now familiar with getting optimal performance by linking their bespoke applications to one or more specialised and highly tested libraries, for numerical algorithms, visualisation or data management. Several changes in paradigm for application development and execution have taken them from proprietary to open-source numerical libraries and to shared services. There is therefore a need for a lightweight client Grid toolkit to offer the scientists transparent access to the Grid.

Such a requirement was identified in late 2003 and resulted in the creation of GROWL (Grid Resources on Workstation Library) which utilises WS-I technology to provide a Web service layer to grid middleware as a programmatic VRE interface.

2 Aims and Objectives

Our aims in this project are to encourage the uptake of Grid-based computing and distributed data management, focusing on the issues which may hinder or facilitate end-user application development. We refer to the difficulties identified as the “client problem” and suggest a solution to build upon the existing prototype GROWL library to produce a truly lightweight extensible toolkit which complements other solutions.

“Lightweight” implies that the GROWL API should be minimally sufficient for the user requirements we will identify in three specific application areas. It should also be possible to install the GROWL library quickly on a variety of client workstations running Linux or a similar UNIX-like operating system with a minimum of additional software.

“Extensible” means that it should be possible to easily extend the GROWL library to provide interfaces to additional middleware services (e.g. CONDOR, Netsolve, SRB) or to use additional security mechanisms (e.g. Shibboleth).

The project intends to fulfil these aims by meeting the following objectives:

1. Generate user requirements for a lightweight grid application toolkit for the three target user communities (see Section 3).
2. Produce GROWL toolkit: it should be possible to install the GROWL library quickly on a variety of client workstations running Linux or a similar UNIX-like o/s with a minimum of additional software.
   (a) Provide client and wrappers to existing VRE resources and services developed in e-Science projects;
   (b) Integration of new common services (such as Condor, NetSolve and SRB) into GROWL;
   (c) Produce Clients for the National Grid Service;
3. Evaluate GROWL for Bioinformatics applications:
   (a) provide Web service wrappers to bioinformatics algorithms (e.g. from bioconductor.org) running on large computational resources, including the NGS and local institutional Condor pools;

4. Evaluate GROWL for Physics and Chemistry applications:
   (a) Provide Web service wrappers to chemistry applications using XML meta-data model from the e-CCP project for describing input and output;
   (b) Implement GROWL functionality in existing GUI interfaces, e.g. for DL_Visualize;
5. Evaluate GROWL for Social Science statistical analysis applications:
   (a) provide Web service wrappers to a selection of the functionality of the SABRE package (some of this work is already ongoing);
   (b) implement GROWL functionality in existing GUI interfaces for the social science community;
   (c) use GROWL client-side tools to make SABRE functionality available via the ReDRESS Web portal.

3 Overall Approach

3.1 Strategy

The project is organised into the four work-packages described below.

3.1.1 WP1: Toolkit and Services

Lead partner: CCLRC Daresbury Laboratory

The toolkit to be used and extended is currently called GROWL: *Grid Resources On Workstation Library*. It is a working prototype and is being used by Daresbury and Lancaster in the SABRE-R e-Social Science pilot project. A number of other prototype lightweight Grid toolkits are being developed worldwide and one brief of this project would be to monitor this activity. In particular we cite *gLite* from the EU EGEE project and implementations of specifications developed in the GGF SAGA research group. This project will also be represented at SAGA workshops.

Prototype interfaces for GROWL services have already been created and can be called from C or R language programs and there is a prototype for creating C, Perl and R services. R is an interpreted functional language. It is a GNU open-source version of S, originally from Bell Labs, see http://www.r-project.org/. R is very much a vehicle for newly-developing methods of interactive data analysis and particularly use in statistical computing applications. R appears to be not only the environment of choice of many social scientists and bioinformaticians engaged in statistical analysis, but an ideal test-bed for evaluating a basic set of Grid services for the community. The aim of this work is to make it easy to call remote functions from an R script and to use R applications as Web services methods. The underlying gSOAP C library for Web services is used which can equally be called from Fortran applications.

Wrappers in GROWL will be designed to interface to a variety of services such as the workload management system Condor (http://www.cs.wisc.edu/condor/) and the data management system SRB (http://www.sdsc.edu/DICE/), in addition to Globus (http://www.globus.org/). Wrappers to services supported by the OMII will be addressed as these emerge and are taken up on the UK Grid.

3.1.2 WP2: Bioinformatics Applications

Lead partner: University of Cambridge

This work-package aims to enable the GROWL toolkit to be useful for bioinformaticians, working in the field of micro-array data analysis. Micro-array data analysis is notoriously resource intensive and requires the manipulation of large sets of data in a variety of different formats. Algorithms for the analysis of micro-array data make use of a plethora of statistical methods, and the R system is rapidly becoming the method of choice for this purpose. A collection of tools and extensions is being developed by statisticians and made available to the community via the Bioconductor project (http://www.bioconductor.org/) to encourage this.

There is a need for analysis to exploit multi-node computer systems to speed up processing, enable numerically intensive calculations to be performed and aid seamless automation of performing multiple sub-tasks. Our strategy is to enable ease of access to grid technology for biologists performing micro-array data analysis using a variety of software tools, written in different languages (in particular R and C) through the use of the GROWL toolkit.

We will identify specific application areas that would benefit from using GROWL. One such application area is the design of oligo-nucleotide probes for DNA micro-arrays. Identification of suitable probes requires calculation of the thermodynamic hybridization properties of candidate probes and assessing likely cross-reactions to non-targets from their nucleotide sequence. Present approaches are limited to traditional sequence similarity searches and rough approximations in calculating thermodynamic properties. Extensions made feasible by transparent access to distributed and specialized computing resources will allow DNA folding criteria to much earlier enter the scan for permissible probe candidates improving the quality of the obtained final set.

Three aspects play a major role in this context:

1. the availability of large computational resources,
2. library support in allocating resources to sub-tasks and collecting and integrating results,
3. accessibility for non-experts in the application sciences.

While shared resources can address aspect 1, the overhead required for developing and deploying solutions that adequately exploit heterogeneous shared resources for any individual laboratory in the applied sciences is prohibitive. By addressing each of the above issues in a centrally coordinated manner with the help of dedicated staff, the GROWL project can extend the scope of feasible analyses and functions as a facilitator of ambitious scientific applications.

Prototypes will be made available to the user community throughout the different stages of development for evaluation.

3.1.3 WP3: Chemistry and Physics Applications

Lead partner: CCLRC Daresbury Laboratory

This work-package is tasked with ensuring that the GROWL software is genuinely useful for grid-enabling chemistry and physics applications based on the requirements of the CCLRC quantum chemistry community and Physicists at CCLRC and the University of Lancaster. Many of their applications have been developed over a long time period and with considerable effort from many researchers in the UK, in particular those involved in the collaborative computational projects (http://www.cse.clrc.ac.uk/ccp/), for instance CCP1, CCP3, CCP4, CCP5, CCP6 and CCP9. These include large high-performance computer codes running on national facilities (CSAR, HPCx) and the JISC JCSR-funded National Grid Service clusters. Examples of the codes include: GAMESS-UK, MOLPRO, DL_POLY, Siesta, CRYSTAL and CASTEP. Some of these codes also have their own GUI front ends which are used to construct appropriate input, submit the jobs for execution and display and annotate output. Use of the GROWL

C-based client library will enable Grid functionality to be integrated directly into these "heritage" applications through appropriate "wrappers" to allow the codes to interact with one another, e.g. in simple work-flow scenarios. Indeed there is an associated e-Science project called WOSE, Work-flow Optimisation Services (see http://www.grids.ac.uk/WOSE/) which could contribute to this goal.

This work-package will thus address the use of GROWL to interface to and link with some of the leading chemistry and physics applications listed above.

3.1.4 WP4: Social Science Applications

Lead partner: University of Lancaster

The SABRE-R e-Social Science pilot project aims to port the open-source SABRE statistical analysis package to a parallel computing Grid-based environment. Middleware extensions to the functional R programming language are being added using Web services in the GROWL library and appropriate additional wrappers for R clients and services; this has already been demonstrated for the simplest cases.

This pilot project will play a key role in the development of some of the middleware components and services appropriate to e-Social Science.

In this project we are also applying the GROWL middleware to wrap the statistical modelling methods for analyzing work / life history data, and to make these developments available in the distributed environment as a componentised R library. The free-to-use R language and environment provides a wide variety of useful statistical and graphical techniques (linear and non-linear modelling, statistical tests, time series analysis, classification and clustering).

3.2 Scope

GROWL is a purely client-side Grid programming environment, it does not help you to create Web or Grid services, nor does it help you put your computer on the Grid. In support of the aim of providing a lightweight interface, the software supported is kept to the minimum needed to successfully demonstrate the applicability of GROWL to the research community. There are however associated services which GROWL uses and there is a procedure for enabling a particular service to be linked into the infrastructure, whether it be written in C, C++, Perl, Python, PHP, R or Java. We therefore focus on C and R-specific issues here. Installation and use of GROWL is described on the Web site. General philosophy and aims of the GROWL project are also explained separately and all documentation is available from the Web site http://www.growl.org.uk.

Note that currently no automatic publication of GROWL services is attempted and all coding is done by hand. This is a prototype – there is no GUI or IDE. The aim of the GROWL project is to keep the client side as simple as possible requiring only some C libraries and (optionally) Perl to be installed in addition to the application, which in some cases is the R package downloaded from the GNU Web site. Remote services are installed on Apache-enabled servers, and should not be the concern of the end users.

3.3 Critical Success Factors

- That the project achieves a robust, well documented and easily accessible lightweight toolkit which meets the requirements of its target users
- That this toolkit becomes widely accepted, well regarded and used in its targeted communities

- That the toolkit is being used and accepted outside its targeted communities.

4 Project Outputs

Reports and documentation

- Web page on JISC VRE programme site (http://www.jisc.ac.uk/index.cfm?name=vre_growl)
- Project Web site (http://www.growl.org.uk) with links from JISC and project partner's Web sites
- Project Plan
- Two 6-monthly Progress Reports (at 6 months and 12 months)
- User Requirements document with input from each of the three target user communities
- Architectural Design document
- User Documentation
- Final Project Report
- Completion Report

Technical outputs

- CVS repository containing:
  - The core GROWL C library
  - GROWL Wrapper for R (and potentially other languages, to be confirmed)
  - GROWL Wrappers for Services (final set to be decided) such as Netsolve and Condor
- An externally accessible GROWL demonstrator, accessible throughout the project with additional functionality added incrementally. This will provide a test of available functionality.

Knowledge and other outputs

- A better understanding of whether a toolkit such as GROWL can provide a sufficiently rich subset of Grid technology while remaining lightweight and easy to use.
- Demonstrations of the ease of use of GROWL so that it its use could be incorporated into Grid induction workshops (for NGS for instance).
- Presentations (and accompanying papers) at UK and international conferences, including the UK e-Science All-Hands Meeting.
- Article in relevant peer-reviewed journal
- A proposal for how the work on the project can be carried forward.

5 Project Outcomes

The GROWL project will provide a demonstrator quality, lightweight and user-friendly toolkit which will enable researchers to access the Grid from within their current applications in a transparent way. We think that there are three benefits to be had from this approach. The first benefit is that GROWL will allow the scientist already using the computational power of the Grid to create their applications more efficiently. The second benefit is that the GROWL toolkit will make it significantly easier for any scientist contemplating the use of the Grid to make the Grid step. Our tools will be designed to aid and assist these researchers and lower some of the key barriers to using the Grid, this will encourage more and more communities to see the possibilities of the Grid for their research. The third benefit we expect from the GROWL toolkit and its related demonstrators is that it will form part of core e-Research training. This will help students, researchers and scientist to become proficient users of the Grid.

We also expect that it is likely that the presence of the GROWL toolkit will lead to proposals for new tools/applications to be added to the toolkit and for calls for further development/refinements.

6 Stakeholder Analysis

Stakeholders in the project include:

Stakeholder	**Interest / Stake**	**Importance**
Associated User Communities	Will be provided with demonstrator-quality tools they can use to grid-enable their applications for more effective use of available computational resources.	High
Project Developers	Will learn new skills in Web and Grid service technologies and security.	Medium
UK e-Research community	Will be provided with a lightweight programming library demonstrator for linking “heritage” applications into a distributed Virtual Research Environment. It will provide “proof of concept” of the lightweight library approach to making grid computing more accessible by lowering the conceptual barrier that new grid users have to negotiate to grid-enable their applications.	Medium
Wider academic community	The wider academic community will be made aware of the GROWL toolkit. Since much of the core GROWL capability will be generic, many other user communities will obtain the advantages of the GROWL approach. In particular, we envisage the GROWL would be an ideal tool (even as a demonstrator) for undergraduate education about Grid technology.	Medium
JISC	Will be provided with feedback on using Web services as a framework for VRE development and delivery of remotely-hosted services.	Low
System Administrators	The use of GROWL will enable users to grid-enable their software without relying on the skills or passwords of their system administrators.	Medium
Service providers	Will be encouraged to publish their services via Web service interfaces for remote programmatic access.	Low

7 Risk Analysis

Risk	Probability (1–5)	Severity (1–5)	Score (P×S)	Action to Prevent/Manage Risk
Loss of staff	2	4	8	Attend to needs of staff and personal development. Cannot compete if staff apply for a job with higher pay or promotion. Ensure more than one person is involved, if we have the luxury to do so or it is a bigger project with more staff.
Poorly defined requirements or tasks mean project cannot be kept within scope	1	2	2	Ensure requirements are defined at start of project and developers understand them. Continue to update as functionality develops via regular technical project meetings with partners. Act quickly to contain functionality creep or redefine project scope if required.
Loss of software	1	3	3	Ensure software and paperwork relating is backed up using CVS or RAID / tape file stores as appropriate. Instruct staff how to do this and review regularly with assistance from SysAdmin to ensure procedures actually work.
Critical failure of project resources	2	2	4	Ensure backups are carried out as above. Use fail-over hardware if available. Ensure other hardware is on warranty or other maintenance contract, if appropriate, or purchase new (e.g. desktop or laptop PCs).
Firefighting	2	2	4	Other priorities within the busy portfolio of the partners mean that staff effort has to be diverted to unforeseen tasks. This should be avoided by correct portfolio planning.
Failure of dependencies	1	3	3	Dependencies, e.g. software such as Globus, should be evaluated and understood prior to undertaking the project. If they fail during use (e.g. are unreliable or do not function as described) this can lead to full deliverables not being achievable.
Failure of project partners	1	3	3	Other project partners may suffer consequences of any of the risks above leading to knock-on effects. This must be mitigated by regular management meetings with all project partners to forestall such events and adapt as necessary.

Risk	Probability (1–5)	Severity (1–5)	Score (P×S)	Action to Prevent/Manage Risk
Failure for partners and JISC to reach an agreement for a software license under which GROWL will be distributed.	1	4	4	Work is already underway to obtain this is as soon as possible.
Distributed development team (3 people at 3 sites) could lead to Lack of communication and co-ordination	3	2	6	Regular meetings over Access Grid (see Section 12) will be used for essential face-to-face communication.
Use of GROWL library poses a security risk	1	4	4	We will endeavour to use techniques such as “tainted data” (in Perl) to prevent, or at least minimise the risk of certain known attacks.

8 Standards

The success of a lightweight toolkit such as GROWL relies on its ability to inter-work with a variety of software and tools. To this end, it is vital that standards are followed. Moreover, as the Grid and Web arenas are continuously and rapidly changing, it is also important that emerging standards are tracked, exact versions of some of these will not be known until the project is considerably advanced.

The following standards are considered both relevent and important for GROWL:

- Web Services standards: WS-I, WSDL, SOAP and UDDI
- Globus toolkit: initially GT2.4, but see below
- Authentication: X.509, GSI and TLS
- Documentation standards: combination of JISC and institution’s own guidelines
- Web pages: XHTML 1.0, CSS, paying attention to latest accessibility guidelines
- Coding: for a project of this size, it is not feasible to generate standards for all the languages we will be coding in. Our build system will however utilise compiler warning options (Perl) and flags (C/C++) to provide automatic checking.

The following projects will be tracked for possible future inclusion.

- Globus toolkit: WSRF, GT4
- Authentication: Shibboleth
- Authorisation: PERMIS
- Application Programming Interface: GGF SAGA work (also DRMAA)

9 Technical Development

Following best practice in the field of distributed development, we will make use of CVS for configuration management of software, documentation and Web pages. A CVS server will be set up on a reliable server.

Compilers for several of the languages we will be using (C, C++ and perl for instance) have the option of requesting warnings about unusual constructions and dubious practices. We will ensure that all compilation scripts utilise these compiler flags to reduce the risk of unforeseen problems in the code. An extra check will be done for C and C++ code to check for memory leaks and memory corruptions. It is recommended that the open-source memory debugger `valgrind` is used for this.

In addition, we will endeavour to use the latest stable releases of compilers and, where possible, use ISO/ANSI/POSIX versions of these (enforced with compiler/linker flags if necessary).

We will also use stable releases of 3rd party software that we depend on. Where feasible, these will be the most up-to-date versions of this software, but this may depend on the timing of such releases. Software in this category would include gSOAP, Globus, R and SOAP::Lite.

10 Intellectual Property Rights

Management of IPR between collaborators
Details of the Intellectual Property Rights of outputs produced by the project in collaboration or individually will be as defined in the project Collaboration agreement(s).

How we will respect IPR of 3rd party dependent components / suppliers
It is anticipated that dependent software components will be treated in different ways, depending on the prevailing rights concerning their use:

1. Some 3rd party software may be bundled with our project outputs; this will be restricted to those components for which such use is either clearly permitted or for which permission has been sought and given.

2. For other software it might not be clear on whether it can bundled or not. In this case, we would bundle a link to the latest compatible version. Where possible, the download and installation of these products will be handled by our build system.

3. A final category would for software which is definitely not free to the community at large, for instance software libraries for our three target application areas. For this case, use of the software would be optional so we will provide instructions for obtaining the software as appropriate.

A register of Software Licenses used and needed for the GROWL toolkit will be maintained for inclusion as appropriate in User and other documentation.

How IPR is managed with the Users
Any GROWL software release will be subject to a license. The license details will be agreed between our respective commercial departments and will ensure that GROWL software is freely available to UK HE and FE institutions in accordance with Grant Conditions (clause 9). The license will also contain statements on Warranty, Liability, Reselling, Commercial Use and Copyright.

Project Resources

11 Project Partners

Cambridge eScience Centre (Lead Institution)

Responsibility: WP2

Contact: Mark Hayes; mah1002@cam.ac.uk; Cambridge eScience Centre, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge, CB3 0WA;
tel: +44 (0)1223 765251; fax: +44 (0)1223 765900

Centre for eScience - Lancaster University

Responsibility: WP4

Contact: Rob Crouchley; r.crouchley@lancs.ac.uk; Centre for eScience, C Floor, Bowland Annexe, Lancaster University, Lancaster, LA1 4YT;
tel: +44 (0)1524 593161; fax: +44 (0)1524 594459

CCLRC e-Science Centre

Responsibility: WP1, WP3 and Project Management

Contact: Rob Allan, r.j.allan@dl.ac.uk; CCLRC e-Science Centre, Daresbury Laboratory, Keckwick Lane, Daresbury, Warrington, Cheshire, WA4 4AD;
tel +44 (0)1925 603207; fax: +44 (0)1925 603634

12 Project Management

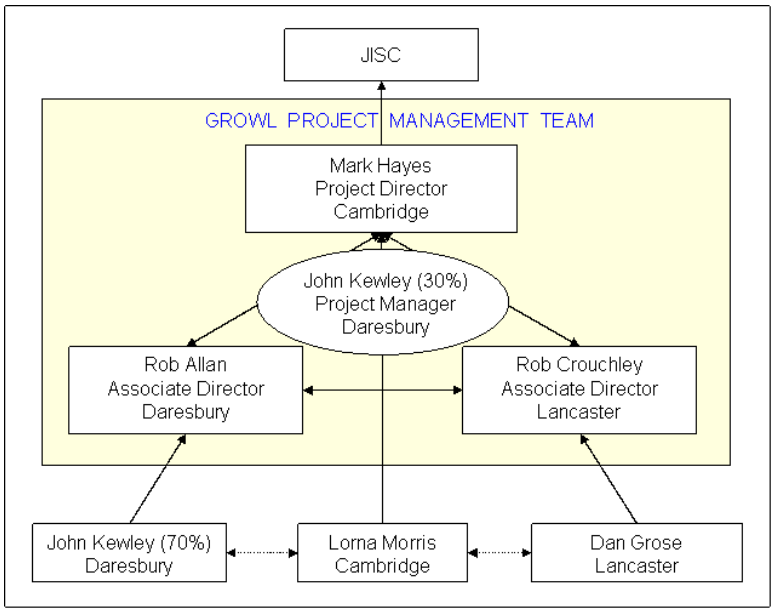

Figure 1: GROWL Management Structure

The organisation of the GROWL project is as follows:

Mark Hayes (MAH): Project Director, Cambridge eScience Centre, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge, CB3 0WA;

John Kewley (JK): Project Manager (30%) and Software Engineer (70%) (reporting to Mark Hayes and Rob Allan), CCLRC e-Science Centre, Daresbury Laboratory, Keckwick Lane, Daresbury, Warrington, Cheshire, WA4 4AD;

Rob Allan (RJA): Associate Director, CCLRC e-Science Centre, Daresbury Laboratory, Keckwick Lane, Daresbury, Warrington, Cheshire, WA4 4AD;

Rob Crouchley (RC): Associate Director, Centre for eScience, C Floor, Bowland Annexe, Lancaster University, Lancaster, LA1 4YT;

Ties van Ark (TvA): Local Project Manager, Centre for eScience, C Floor, Bowland Annexe, Lancaster University, Lancaster, LA1 4YT;

Lorna Morris (LM): Software Engineer (reporting to Mark Hayes), Cambridge eScience Centre, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge, CB3 0WA

Dan Grose (DG): Software Engineer (reporting to Rob Crouchley, to start 2005-05-16), Centre for eScience, C Floor, Bowland Annexe, Lancaster University, Lancaster, LA1 4YT

The Project Management Team (PMT) has a collective responsibility for achieving the project deliverables according to the Quality Plan in a controlled and coordinated manner and using the resources made available by the JISC to do so. The team will aim to make any decisions by consensus and will meet at least once every two months.

The Project Director is responsible for communication with the JISC and carries end-responsibility for achieving the deliverables. The Project Director is empowered by the collaboration agreement signed between the partners to conclude a partnership. The Project Director is ultimately responsible for technical issues.

The Project Manager is responsible for planning, coordination and monitoring progress and maintains effective communication with all parties. Most importantly, he will ensure that the Project Management team are updated with sufficient information to ensure that necessary decisions are taken. The Project Manager is available for 1.5 days per week, averaged over the duration of the project.

The Project Director and Associate Directors are responsible for both the work-packages assigned to their sites, and any associated staffing and personnel issues. They are also responsible for the resolution of any non-conformance with respect to the Quality Plan.

All staff involved in the project have the right and duty to maintain effective communication with each other member of staff to ensure that the deliverables are produced in the most efficient and effective way, and any problems or other issues arising are dealt with as swiftly as possible so as not to hinder progress.

For internal communication, sharing of documents and record keeping, the project will use a SAKAI portal provided by Lancaster University and funded by JISC as part of the Virtual Research Environments (VRE) programme.

For the majority of meetings involving at least two sites, Access Grid technology will be used. The remainder will be by phone or face-to-face.

13 Programme Support

JISC have already produced model licences[1]. It would be good if there was a licence written in UK legal english which described exactly the requirements that are described in the project offer letter and Terms and Conditions. This is currently under discussion with Alan Robiette and others; one suggestion was linking with the work done in the Creative Commons framework.

We suggest therefore that work is lead by JISC to produce this model licence for use in all JISC projects. Note that it might be neccessary to have different variants for programs as well as libraries such as GROWL.

14 Budget

See Appendix A

[1] See http://www.jisc.ac.uk/index.cfm?name=coll_guide_jiscmodel for further details

Detailed Project Planning

15 Work-packages

See Appendix B for work-package breakdown.

15.1 Project Milestones

1. Release of Final Project Plan
2. Site visit and Progress Report in July 2005, also JISC VRE Programme Meeting at the beginning of July.
3. Dissemination at All-Hands Meeting in September 2005
4. Second Progress Report in January 2006
5. We also expect to have a milestone in April or May 2006; its form is not yet decided.
6. Release of Final Report and Completion Report in July 2006

16 Evaluation plan

Timing	Factor to Evaluate	Questions to Address	Method(s)	Measure of Success
2005-05	Coherency of user requirements	Is there a common derived set of User Requirements which are satisfied by the GROWL Architecture.	Review by PMT.	Agreement on an architectural document as basis for development work
2005-08	Feasibility of initial GROWL Demonstrator	Can I submit a simple test job (one from each application area) to a remote Grid resource via GROWL? Is this significantly easier than alternative methods? (e.g. `globus-job-submit`)	Testing the demonstrator	Does the job run successfully and return sensible results? Is the technical knowledge required to use it significantly less than alternative methods for job submission.
Ongoing	Dissemination to the wider community	What level of impact has the project made?	Follow up contacts after dissemination activities, Web site logs, number of downloads, citations from other sources	Contacts made, uptake of toolkit outside initial target user communities

Timing	Factor to Evaluate	Questions to Address	Method(s)	Measure of Success
Ongoing	Project management	Is development work on schedule? Are stakeholders (especially users) engaged with the project?	Review by project board and request of comments from users	Milestones achieved, working code in use
Project end	Whether the project meets its Aims and Objectives	To what degree does the project demonstrate that a lightweight toolkit could satisfy the needs of the target user communities.	Review by PMT and users.	A measurable improvement in user productivity
Project end	Further development	On the basis of project outputs and user feedback, should funding be sought to continue the work.	Review by PMT, interview with stakeholders.	Increased confidence in and demand for GROWL technology
Post-project	Deployment of toolkit in the targeted user communities	Does the toolkit meet user requirements?	User interviews, external peer review?	The demonstrator toolkit is being used in research projects, enabling users to grid-enable their programs more readily.

17 Quality Assurance Plan

Timing	Quality Criteria	QA Method(s)	Evidence of Compliance	Responsibility	Quality Tools)
Output	**Required JISC Documentation: Project Plan, Progress Reports, Completion Report and Final Report**				
according to plan	[Plans] complete, realistic, consistent and executable	Review of document to ensure it reflects our ability to meet the requirements	Approval of document noted in meeting minutes	Project Management Team	
	[Reports] Consistent and truthful reflection of achievements, issues and lessons learned	Review of document	Approval of document noted in meeting minutes	Project Management Team	
	[All] Adherence with JISC Project Management, Guidelines, Standards and templates	Review of document	Approval noted in meeting minutes	Project Management Team	
	[All] Delivered on time to JISC	Date of delivery compared to Project Plan	Date of postage or e-mail	Project Manager or Director	
	[All] Documents are produced to the satisfaction of JISC VRE Programme Management Team	External review of document	Accepted by JISC	JISC VRE Programme Manager	
Output	**Software: third party for integration into GROWL**				
when needed	Appropriate functionality	Code provides required functionality in whole or in part without duplicating functionality of other modules and is of appropriate perceived quality	Developer report	Developers	
	Ready for incorporation into GROWL	Review based on developer report and its licence	Decision minuted at PMT meeting	PIs	
Output	**Software: library components, wrappers and user services/client code**				
as required	Appropriate design	Review of design, interface and standards used in team meeting	Discussion minuted	Development team	
	[C/C++] Code conforms to standards	Enable suitable warnings in C++ Compiler (minimum of `-std=ansi -W -Wall` for `g++` and `-std=c99 -W -Wall` for `gcc`). This will be enabled by default in build Makefiles	A “clean” compilation	Development team	`gcc` and `g++` compilers

Timing	Quality Criteria	QA Method(s)	Evidence of Compliance	Responsibility	Quality Tools)
	Ready for incorporation into GROWL	Review of interface, test cases and test results	Code and test script in CVS, results minuted	Development team	
Output	**GROWL Demonstrator**				
Monthly	Working OK	Regular checking that Demonstrator is working as intended (ideally an automatic test)	Report given to development team meeting	Development team	
Output	**Web site/Web pages**				
2005-02	JISC page and project Web site available within first month of project	Examination of JISC Web site	Web page is available on JISC Web site	Project Manager	Browser
Every 2 months	Web page content is up-to-date and relevent	Reviewed at progress meetings	Entry in minutes once agreed	Project Manager	
	Web pages conform to agreed W3C standards (XHTML, CSS, ECMAScript) and all links are valid	Regular use of online automatic testing tools, reported at meeting	Entry in minutes	All	Browser facilities (e.g. Firefox using Web Developer extension), W3C website or Altova XMLSpy (for XHTML)
Output	**Conference abstracts and papers**				
Conference deadlines	Timely delivery	Date of delivery compared to Conference deadlines	Date of postage or e-mail	Author	

18 Dissemination Plan

The project will deliver the following minimal items:

- Project Web page with links from JISC and collaborating institutions' Web sites.
- UK e-Science All-Hands Meeting 2005 (+ 2006?), talk and/or poster
- A scripted demo to demonstrate functionality and ease of use of GROWL
- Abstracts/papers submitted for peer-review to workshops or conferences as appropriate (e.g. an abstract for the First International Conference on e-Social Science has already been accepted for presentation in Manchester, June 2005).
- Articles submitted for publication in relevant peer-reviewed journals as appropriate
- Presentations on GROWL will be made at JISC VRE Programme events.

Demonstrations at conferences and other events will be treated as project milestones.

19 Exit/Sustainability Plan

The primary aim of the partners is to develop a demonstrator toolkit, which will have the potential for widespread use in the scientific community. The successful demonstration of this will lead to user-driven demand to further extend the functionality supported by GROWL and to provide user support. During the project the partners will create a log of newly arising user requirements and change requests, currently falling outside the scope of the project. The GROWL server is currently an institutional contribution from Daresbury. The partners aim to build a well founded case, providing justification for continued external funding of the work and services provided, thereby allowing GROWL to enter into a new development phase using Grid services middleware such as WSRF. The partners see this as the best way to conclude this project and to ensure sustainability of the achieved project outputs and services for the users and the assembled knowledge of the staff and institutions involved.

The exit plan below should therefore be viewed as a last resort for when our preferred take-up strategy fails to materialise, ensuring that the toolkit remains available to users for at least 3 years and documentation and the software repository remain available for a revival if the need were to arise and should the project obtain further funding at a later date.

See also table below.

Project Outputs	**Action for Take-up & Embedding**	**Action for Exit**
Toolkit software	The toolkit software will be freely available to academic users via the GROWL Web site which is currently hosted by CCLRC.	Without further funding, CCLRC will maintain this Web site for a further 3 years, but will not be able to commit to any software maintenance.
GROWL Toolkit server	The server used by GROWL clients will remain available and working for at least 3 years.	The GROWL server will not be maintained after 3 years.

Project Outputs	Action for Take-up & Embedding	Action for Exit
Project development documentation	The project development documentation will be partly available on the Web, as far as this is required for the users of the GROWL toolkit. This documentation will remain available on the Web site for a further 3 years and accurately reflect th ecurrent state of the project.	Documentation developed during the project will be stored and remain accessible, in accordance with the work packages at the (partner) institutions. This documentation will be stored for at least 3 years, but not updated.
Repository of development software	The latest, tested version of the software will be available on the Web site.	A repository of all code will be maintained for at least 5 years by CCLRC, available on request.
Further knowledge	Knowledge will be secured and remain available in publications, presentations and reports. Aim to undertake further related work.	Institutions and involved staff will continue to benefit from the experience and knowledge created during the project, but will be engaged in other projects.

Appendices

A Project Budget

Expense	Requested Funding		HICs	
	Year 1	Year 2	Year 1	Year 2
	10 months	8 months	10 months	8 months
Equipment				
3 PCs for the RAs (50%)	1,500		1,500	
3 laptops for the RAs (50%)	2,000		2,000	
Total	***3,500***	***0***	***3,500***	***0***
Staffing				
Project Management			12,500	11,000
Software Development (Cambridge)	25,253	22,222		
Software Development (Daresbury)	26,038	22,914		
Software Development (Lancaster)	25,116	22,102		
Overheads of staff on project (46%)			40,897	35,989
Total	***76,407***	***67,238***	***53,397***	***46,989***
Misc				
Travel and Subsistence	900	468		
Consumables	900	468		
Total		***3,500***		***3,500***
Grand Total	***81,707***	***68,174***	***56,897***	***46,989***

Summary over 18 months

Request from JISC	149,881
Host Institutional Contributions (HICs)	103,886

B Work-packages

Workpackages	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
	F	M	A	M	J	J	A	S	O	N	D	J	F	M	A	M	J	J
Project Management	●	●	♦	•	•	♦	•	•	♦	•	•	♦	•	•	♦	•	•	♦
WP1																		
Architecture and Design		●	●	♦														
Development			•	●	●	♦	●	●	●	●	●	●	●	●	●	●	●	♦
WP2																		
User Requirements		●	●	♦														
Development				•	●	♦	●	●	●	●	●	●	●	●	●	●	●	♦
WP3																		
User Requirements	●	●	●	♦														
Development					●	♦	●	●	●	●	●	●	●	●	●	●	●	♦
WP4																		
User Requirements			•	♦														
Development					●	♦	●	●	●	●	●	●	●	●	●	●	●	♦
Dissemination		♦		♦	♦		●	♦			●	♦					●	♦

Key	
♦	Major work and workpackage milestone
●	Major work
•	Minor work

Workpackages	Earliest Start	Latest Finish	Outputs	Milestone	Responsibility
Project Management					
0.1 Produce Draft Project Plan	2005-02	2005-03	**Draft Project Plan**		JK
0.2 Produce Final Project Plan	2005-03	2005-04	**Final Project Plan**	**1**	JK
0.3 Produce 1st Progress Report	2005-06	2005-07	**1st Progress Report**	**2**	JK
0.4 Produce 2nd Progress Report	2005-12	2006-01	**2nd Progress Report**	**4**	JK
0.5 Produce Draft Final Report	2006-05	2006-06	**Draft Final Report**		JK
0.6 Produce Final Report	2006-06	2006-07	**Final Report**	**6**	JK
0.7 Produce Completion Report	2006-07	2006-07	**Completion Report**	**6**	JK + MH
WP1: Toolkit and Services					
1.1 Produce Architectural Design	2005-03	2005-05	**Architectural Design**	**2**	JK
1.2 Update build system	2005-04	2005-07	for core **GROWL C Library**	**2**	JK
1.3 Integrate up-to-date version of gSOAP	2005-04	2005-07	for core **GROWL C Library**	**2**	JK
1.4 Integrate up-to-date version of R	2005-05	2005-07	**GROWL Wrapper for R**	**2**	JK
WP2: Bioinformatics					
2.1 Produce User Requirements for WP2	2005-02	2005-05	**User Requirements**	**2**	LM
2.2 Implement an example WP2 Use Case	2005-05	2005-07	**GROWL Demonstrator**	**2**	LM
WP3: Chemistry and Physics					
3.1 Produce User Requirements for WP3	2005-02	2005-05	**User Requirements**	**2**	JK
3.2 Implement an example WP3 Use Case	2005-05	2005-07	**GROWL Demonstrator**	**2**	JK
WP4: Social Science					
4.1 Produce User Requirements for WP4	2005-04	2005-05	**User Requirements**	**2**	DG
4.2 Implement an example WP4 Use Case	2005-06	2005-07	**GROWL Demonstrator**	**2**	DG
Dissemination					
5.1 Web Page on JISC site	2005-02	2005-02	**JISC Web Page**	**1**	RA
5.2 Project Web site	2005-02	2005-03	**Project Web site**	**1**	MH/RA
5.3 National e-Social Science Conference	2005-04	2005-05	Conference Paper		RC
5.4 National e-Social Science Conference	2005-05	2005-06	Conference Presentation		RC
5.5 All-Hands Meeting	2005-07	2005-09	presentation, poster, paper and/or demo	**3**	ALL